

Measuring the Diversification Effect of Private Assets in Public Portfolios: Implications for Defined-Contribution Plans¹

Hammad Qureshi²

September, 2026

Abstract

Should retail investors add private assets to a public portfolio? The question has gained relevance for households as policymakers debate private assets in defined contribution (DC) plans. The economic case rests in part on a diversification claim: private assets appear less volatile than, and weakly correlated with, public markets — properties that appraisal smoothing mechanically exaggerates. We correct reported private-market index returns for smoothing with a range of established methods and re-estimate constrained mean-variance frontiers. After correction, a private sleeve improves the estimated tangency Sharpe ratio by 19 to 23 percent over a public-only portfolio — 60 to 74 percent of the roughly 31 percent gain implied by reported returns. This improvement is statistically significant and persists under a zero-alpha prior and forward-looking return assumptions. Because private-market indexes are not directly investable, we extend the analysis to estimate the frontiers on simulated investable books of randomly drawn private funds, net of the cost of managing the private sleeve at levels anchored in the literature. The gain persists in these investable books and remains positive in the lower tail of simulated outcomes. Our findings suggest that a modest private-asset sleeve can offer a diversification benefit that survives measurement correction, fund-selection risk, and fees.

Keywords: private markets; appraisal smoothing; defined-contribution plans; target-date funds; fund-selection risk; fees

JEL classification: G11, G23, G51, J32

¹ I would like to thank Vikas Agarwal, Shelly Antoniewicz, Gregory Brown, Elena Chism, David Cohen, James Duvall, Sarah Holden, Lei Li, Christian Lundblad, Andrea Rossi, Laura Starks, and Shane Worner for valuable comments and suggestions on this analysis. Irina Atamanchuk and Alessandro Sitta provided excellent research assistance. I am responsible for any remaining errors or deficiencies. The views expressed in this paper are those of the author and do not necessarily reflect the views of the Investment Company Institute, its staff, or its member firms.

² Senior Economic Adviser, Investment Company Institute, 1401 H St NW, Washington, DC 20005, USA. Contact email: hammad.qureshi@ici.org.

1. Introduction

Over the past two decades, primary capital formation in the United States has shifted markedly toward private markets. Companies now raise more equity through private offerings than through initial public offerings. The number of publicly listed firms has fallen by nearly two-fifths since 2000 — from about 6,300 to about 3,900 — while the number of private-equity-backed firms has grown more than sixfold, and firms stay private far longer than they once did. Total private-market assets under management in the United States rose from about \$1.9 trillion at year-end 2010 to roughly \$9.1 trillion as of September 2025 (Figure 1). Institutional investors — public pension plans, endowments, and sovereign funds — have been responding for years by allocating increasing shares of their portfolios to private equity, private credit, real estate, and infrastructure. State and local defined-benefit (DB) pension plans, for example, raised their combined allocation to private equity and real estate from about 8 percent of investments in fiscal 2001 to about 23 percent in fiscal 2024, with further allocations to commodities, hedge funds, and other alternatives (Figure 2). Most retail investors, including the majority of defined-contribution (DC) plan participants, have not had comparable access.

The policy landscape is now moving to address this gap. The Department of Labor's recently proposed rule³ would establish an asset-class-agnostic, process-based safe harbor giving DC plan fiduciaries a clearer path for prudently selecting plan investment options, including options that incorporate private-market assets where appropriate. Whether and how a measured private-market allocation might be delivered to DC participants has thus moved to the center of retirement-policy debate. The economic case rests in part on a diversification claim: that private assets, because they appear less volatile than and weakly correlated with public markets, can expand the risk-return opportunity set available to long-horizon retirement savers. That claim, however, is clouded by a well-documented measurement problem. Reported private-fund net asset values are set through fair-value methodologies that blend information across reporting periods. The resulting appraisal smoothing mechanically dampens measured volatility, induces positive return autocorrelation, and weakens measured correlation with public indexes — precisely the three properties that, taken at face value, would overstate the diversification benefit of private assets. And even a correctly measured index benefit may not be attainable. In public markets an investor can hold essentially the entire market through a total-market index fund; private markets offer no counterpart — their indexes are constructed from reported fund results, not tradable portfolios, and investment vehicles do not hold every fund in every vintage. The relevant question is therefore what remains of the benefit for an investable book of funds, net of both funds' own fees and any additional fee charged by the vehicle that assembles them.

This paper addresses what remains of the diversification claim once the frictions a real-world allocator faces are priced in within a single framework: we estimate constrained mean-variance

³ US Department of Labor, Employee Benefits Security Administration. 2026. "Fiduciary Duties in Selecting Designated Investment Alternatives." Proposed rule, 91 Fed. Reg. 16088 (March 31). See also ICI's comment letter on the proposal (Investment Company Institute 2026) and the accompanying Investment Company Institute white paper (Antoniewicz et al. 2026). Parts of the unsmoothing and efficient-frontier analysis in Sections 3 and 4 were first developed by the author, with input from ICI colleagues, for that white paper. The ICI white paper also covers material beyond the scope of this paper. This paper extends the unsmoothing and frontier analysis with the fund-level simulation of investable books, formal statistical tests and bootstrap inference, forward-looking capital-market assumptions, and the private-sleeve fee analysis.

frontiers with and without a modest private-market allocation, and trace how the estimated gain changes as frictions are introduced. Three frictions organize the analysis. The first is measurement. We correct reported private-market index returns for appraisal smoothing using three established estimators and, as a sensitivity, repeat the analysis with expected returns drawn instead from forward-looking capital-market assumptions surveyed from major investment advisory firms — inputs commonly used in institutional asset allocation. The second is investability. Using a fund-level panel of more than 5,500 individual private funds, we simulate the frontier available to books of randomly selected funds, and thus quantify the fund-selection risk that index-level returns average away. The third is fees. The private-market returns in our data are net of the underlying funds' own fees. We subtract an additional “private-sleeve fee” — the cost of selecting, holding, and administering the underlying private funds, which may be paid to an intermediary such as a fund of funds or incurred by the manager of a target-date fund’s private sleeve — across a grid of levels anchored in the academic literature and ask whether the estimated gain survives. Every frontier is estimated on portfolios of the kind a DC plan investment option could hold — no short sales and a private sleeve capped at 20 percent — and statistical tests, sensitivity analyses, and simulation-based uncertainty bands are embedded throughout to gauge how sensitive the conclusions are to sampling error.

Our findings can be stated compactly. Unsmoothing materially raises measured private-asset volatility — annualized private-equity volatility rises from 6.0 percent to between 9.7 and 12.7 percent, depending on the estimator — and pulls first-order autocorrelation to levels statistically indistinguishable from zero. Even after this correction, a capped private sleeve shifts the estimated efficient frontier up and to the left across the full range of risk levels examined: the tangency Sharpe ratio rises from 0.85 for the public-only portfolio to between 1.01 and 1.05, an improvement of 19 to 23 percent, so that 60 to 74 percent of the apparent, smoothed diversification benefit survives.

The shift is statistically significant under three independent tests, is bounded above zero across 10,000 block-bootstrap simulations, survives a zero-alpha return prior and allocation caps ranging from 15 to 25 percent, and does not appear to worsen tail risk once volatility is equalized. Nor is the result an artifact of the estimation window: replacing realized returns with expected returns based on capital-market assumptions from more than 40 major asset managers and institutions lowers Sharpe ratios for public-only and public-plus-private portfolios alike, but leaves the gain essentially intact at +0.15 to +0.17.

The gain survives implementation as well. Following the framework of Brown, Gonçalves, and Hu (2026), we simulate books of nine funds per vintage; the median gain for a simulated book ranges from 0.12 to 0.14 under the corrected specifications and the 5th percentile remains positive even when fund-selection and sampling uncertainty are compounded. Finally, subtracting a private-sleeve fee of 50 to 150 basis points of sleeve assets — a range anchored in the academic literature, and approximately 10 to 30 basis points of total portfolio assets given the 20 percent cap — leaves the corrected index-level gain at 68 to 91 percent of its no-fee value and the fund-level median at 59 to 88 percent. The diversification benefit is thus smaller than reported index statistics imply, but it remains positive at every step, and the index-level gain is statistically significant.

This paper sits at the intersection of two strands of research. The first concerns the measurement of private-capital performance. Kaplan and Schoar (2005) introduced the public-

market-equivalent framework, and a substantial literature has since shown that estimated outperformance depends on data source, methodology, and the risk adjustment applied;⁴ Brown, Lundblad, and Volckmann (2025) survey the risk-adjusted evidence. Within this strand, a parallel line of work addresses the appraisal smoothing embedded in reported valuations. Geltner (1991, 1993) developed the autoregressive unsmoothing model for appraisal-based real-estate returns; Getmansky, Lo, and Makarov (2004) generalized it to a moving-average representation for hedge funds; and Dimson (1979), extended by Asness, Krail, and Liew (2001) and Anson (2013, 2016), recovers economic exposure from lagged public-market betas. The three estimators are not nested and rest on different assumptions about how new information enters reported valuations, which is why we estimate all three rather than selecting one: agreement among them is evidence that the correction captures smoothing rather than the idiosyncrasies of a particular filter. Estimating at the fund level, Coutts, Gonçalves, and Rossi (2024) find that a common smoothing component remains even after fund-by-fund correction; we therefore apply the correction not to individual funds but directly to aggregate index or investable-book returns, the level at which that common component operates.⁵

The second strand addresses the portfolio question directly. Markowitz (1952) and Tobin (1958) supply the framework, whose well-documented sensitivity to estimated inputs⁶ is blunted by the very constraints we impose on policy grounds (Jagannathan and Ma (2003)). Closest to our work, Brown, Gonçalves, and Hu (2026) — hereafter BGH — estimate private-capital alpha as the increase in the maximum Sharpe ratio when the investable set expands from public to private assets, simulating investable fund portfolios. We adapt BGH's sampling design to the DC question. On fees, Metrick and Yasuda (2010), Harris, Jenkinson, Kaplan, and Stucke (2018), and Brown, Hu, and Kuhn (2021) document the compensation structure of private funds and funds of funds; our contribution is to examine the private-sleeve fee directly, anchoring the examined fee range to their estimates.

Each of these literatures isolates one wedge between reported private-market performance and what an investor can actually earn, and each is typically studied on its own. Our contribution is to examine them within a single constrained-optimization framework: we subject the estimated diversification benefit to each wedge in turn — measurement, return outlook, investability, and fees — and ask how much survives at every step. Correcting for smoothing reduces the benefit but does not eliminate it; what remains persists under forward-looking return assumptions, carries over from the index to an investable book of funds, and stays positive after a private-sleeve fee. That the benefit survives each of these tests within the same framework suggests it reflects a real feature of the data rather than an artifact of measurement, implementation, or fees.

⁴ See, for example, Harris, Jenkinson, and Kaplan (2014) and Phalippou and Gottschalg (2009) on data source and methodology; Ang, Chen, Goetzmann, and Phalippou (2018) on inferring returns from irregular cash flows; Sorensen, Wang, and Yang (2014) on generalized PME with illiquidity; Korteweg and Nagel (2016, 2025) on stochastic-discount-factor methods; Cochrane (2005) on selection-corrected venture returns; Gupta and Van Nieuwerburgh (2021) on strip-by-strip valuation; Boyer, Nadauld, Vorkink, and Weisbach (2023) on secondary-market discount rates; Stafford (2022) on replicating private-equity risk with public securities.

⁵ Coutts, Gonçalves, and Rossi (2024) show that a common smoothing factor can survive fund-by-fund unsmoothing, so correcting each fund separately does not remove the aggregate distortion that matters for portfolio choice; Brown, Ghysels, and Gredil (2023) reach a similar conclusion from the nowcasting side. Applying the correction at the aggregate level targets this common component directly.

⁶ Michaud (1989); Best and Grauer (1991); Chopra and Ziemba (1993); DeMiguel, Garlappi, and Uppal (2009).

The rest of the paper proceeds as follows. Section 2 describes the data. Section 3 sets out the three unsmoothing estimators and their effects on volatilities, autocorrelations, and correlations with public markets. Section 4 estimates the constrained frontiers, tests whether the shift is distinguishable from zero, and reports simulation-based inference alongside the associated sensitivities. Section 5 simulates investable books of funds for each private asset class, pricing fund-selection risk. Section A re-estimates the frontiers under forward-looking capital-market assumptions surveyed from major investment advisory firms. Section 6 examines the effect of a fund-of-funds fee on the diversification gain, and Section 7 concludes.

2. Data

Our analysis uses aggregate and fund-level returns and cash flows from Preqin for five private-market asset classes: private equity (excluding venture capital), venture capital, private debt, private real estate, and private infrastructure. The sample covers 2010:Q1 through 2025:Q3. This choice reflects two considerations: Preqin's coverage has improved markedly since 2010, and post-global-financial-crisis (GFC) vintages better represent today's private-market opportunity set, avoiding the splice of a structurally different pre-2008 private universe onto the current one. We address any residual dependence on this window by re-estimating the analysis under forward-looking capital-market assumptions in Appendix A. We complement the private series with three public asset classes, sourced from Refinitiv and Morningstar: US equity (FT Wilshire 5000 Total Return Index), world equity ex-US (MSCI ACWI ex-USA Index, in US dollars), and US investment-grade bonds (FTSE US Broad Investment-Grade Bond Index).

For the asset-class analysis, we use Preqin's aggregate private-capital indexes.⁷ The risk-free rate corresponds to the three-month US Treasury-bill yield. Table 1 reports summary statistics, and the pattern that motivates the analysis is immediate. Four of the five private asset-class return series (all except infrastructure) exhibit strong, statistically significant first-order autocorrelation, and their reported volatilities are far below those of the public-equity benchmarks: reported annualized volatility is 5.8 percent for private equity and 8.2 percent for venture capital, against 15.8 percent for US equity. The Ljung–Box statistic rejects the no-autocorrelation null at the 1 percent level for private equity, venture capital, and private real estate, and at the 10 percent level for private debt.⁸ Private infrastructure is the exception: its reported returns show no significant serial dependence, so it is carried through the analysis at its reported values and is excluded from the unsmoothing step.

To construct investable portfolios, we turn to Preqin's fund-level cash-flow data: 338,190 dated capital calls, distributions, and quarterly net-asset-value observations for more than 5,500 US private-capital funds of vintages 1990 through 2025, mapped to the paper's five asset classes (Table 2). Fund-quarter returns enter the analysis over 2010:Q1–2025:Q3, the same window as the index-level results, so the fund-level and index-level analyses share a common sample period. Earlier vintages contribute the return history they generate within that window, with a

⁷ Preqin's aggregate private-capital indexes are size-weighted averages of constituent fund returns computed quarterly with the Modified Dietz formula and reported net of management fees, carried interest, and fund expenses. Because they aggregate reported net asset values, they inherit the valuation conventions of their constituents.

⁸ The Ljung–Box statistic (Ljung and Box 1978) tests the joint null that the first k autocorrelations are zero; we use $k = 4$ to span one year of quarterly data. With 63 quarters the 95 percent no-autocorrelation band for a single autocorrelation coefficient is approximately ± 0.247 .

useful consequence: the simulated books are seasoned from the outset, holding funds across the maturity spectrum at every quarter — younger vintages still calling capital alongside older vintages harvesting — as an ongoing private-markets program would, rather than beginning from newly raised funds and ramping up. Coverage naturally varies with the size and maturity of each market: private equity, the largest and longest-established class, contributes 2,531 funds with a median of 65 funds per vintage year; while infrastructure, a younger asset class, contributes 181 funds with a median of seven per vintage — a difference the simulation design in Section 5 accommodates directly. Our fund-level calculations follow Preqin's aggregation methodology in two respects: cash flows are dollarized from per-commitment figures, and fund-quarter returns are winsorized.⁹

3. Unsmoothing Reported Private-Market Returns

Public and private assets are valued through fundamentally different processes. A listed equity trades continuously at market-determined prices, so a fund striking its net asset value anchors to executed prices that update within seconds. The underlying holdings of a private fund do not trade frequently; their reported value is determined through fair-value methodologies that rely on unobservable inputs — recent transactions in comparable assets, discounted projected cash flows, multiples of trailing earnings, or some weighted combination. These methodologies are grounded in established fair-value principles, but they can involve a meaningful degree of judgment. In practice, valuation professionals tend to anchor estimates to the prior period's value and revise as new evidence accumulates, so each quarter's reported return is, in effect, a weighted average of the true economic return in the current and prior quarters rather than a clean measurement of the current period alone.

The empirical consequences all bias the diversification calculation in the same direction. Smoothing suppresses variance, so measured volatility falls. Carrying information across adjacent periods induces positive serial correlation, which additionally inflates the Sharpe ratio.¹⁰ And because the current-quarter public shock enters the private series only partially and with a lag, the measured contemporaneous correlation with public markets falls. All three distortions make private assets appear to be better diversifiers than they are, so correcting for them is a precondition for a reliable analysis.

3.1 Unsmoothing models

We apply three widely used estimators, each built on a different assumption about how new information enters reported valuations. Estimating all three lets us verify that the substantive conclusions do not hinge on any particular correction.

⁹ Cash flows are dollarized by multiplying Preqin's per-\$10 million figures by fund size divided by ten, to be consistent with Preqin's methodology, so the funds without reported size drop out of dollar aggregates and simulations. Fund-quarter returns are winsorized at the 2.5th and 97.5th percentiles within each category-quarter.

¹⁰ The Sharpe ratio is annualized by the square root of the number of periods, which can overstate it when returns are positively autocorrelated, because the variance of the annual return exceeds four times the quarterly variance (Lo 2002). This is a second, independent channel through which smoothing flatters private assets, and it is removed by the same correction: once the unsmoothed series has statistically insignificant autocorrelation, the square-root scaling is approximately valid.

Geltner (1991, 1993) models the reported return as a weighted average of current and past true returns, with non-negative weights that sum to one and decay at a roughly constant rate:

$$R_t^{\text{rep}} = \sum_{h=0}^{\infty} \psi(h) R_{t-h}^{\text{true}}, \quad \psi(h) \geq 0, \quad \sum_h \psi(h) = 1 \quad (1)$$

Under standard stationarity conditions this has an equivalent autoregressive representation, which can be estimated as an AR(p) regression of the reported return on its own p lags,

$$R_t^{\text{rep}} = \varphi_1 R_{t-1}^{\text{rep}} + \varphi_2 R_{t-2}^{\text{rep}} + \dots + \varphi_p R_{t-p}^{\text{rep}} + \varepsilon_t \quad (2)$$

and recover the unsmoothed series by inverting the estimated dynamics,

$$R_t^{\text{true}} = \frac{R_t^{\text{rep}} - \hat{\varphi}_1 R_{t-1}^{\text{rep}} - \hat{\varphi}_2 R_{t-2}^{\text{rep}} - \dots - \hat{\varphi}_p R_{t-p}^{\text{rep}}}{1 - \hat{\varphi}_1 - \hat{\varphi}_2 - \dots - \hat{\varphi}_p} \quad (3)$$

We set $p = 2$, following Coutts, Gonçalves, and Rossi (2024), who find that the autocorrelation structure of illiquid funds is well described by one or two lags; the AR(2) form also nests the AR(1) of Geltner's original work. Because the inversion divides by $1 - \hat{\varphi}_1 - \hat{\varphi}_2$, near-unit persistence would make it explosive, so we impose a stability cap.¹¹

Getmansky, Lo, and Makarov (2004) — GLM — relax the constant-decay assumption, truncating the weights after a finite horizon H :

$$R_t^{\text{rep}} = \sum_{h=0}^H \theta(h) R_{t-h}^{\text{true}}, \quad \sum_{h=0}^H \theta(h) = 1 \quad (4)$$

We estimate this by maximum likelihood with the lookback H selected by the Akaike Information Criterion (Akaike 1974) and the weights constrained to sum to one.¹²

Dimson (1979), extended by Asness, Krail, and Liew (2001) and Anson (2013, 2016), takes a different route. Rather than modeling the smoothing process itself, it recovers economic exposure by regressing the reported excess return on the contemporaneous and lagged excess returns of a matched public benchmark,

$$R_t^{\text{rep}} - r_t^f = \alpha + \sum_{\ell=0}^L \beta_{\ell} (R_{t-\ell}^{\text{mkt}} - r_{t-\ell}^f) + \varepsilon_t, \quad (5)$$

The asset's total economic beta is the sum of the contemporaneous and lagged loadings, and the unsmoothed series is reconstructed by placing that full systematic exposure on the contemporaneous factor:

$$\hat{\beta} = \sum_{\ell=0}^L \hat{\beta}_{\ell}, \quad R_t^{\text{true}} = \hat{\alpha} + r_t^f + \hat{\beta} (R_t^{\text{mkt}} - r_t^f) + \hat{\varepsilon}_t \quad (6)$$

¹¹ The inversion in equation (3) divides by $1 - \hat{\varphi}_1 - \hat{\varphi}_2$, so coefficients summing close to one make the recovered series explosive. We therefore impose a stability cap that proportionally rescales $\hat{\varphi}_1$ and $\hat{\varphi}_2$ whenever their sum exceeds 0.95. The cap never binds at the index level; it binds in a small minority of fund-level draws in Section 5.

¹² The moving-average weights are identified only up to scale, which is why they are normalized to sum to one: writing the fitted coefficients as $\hat{\beta}_j$, we set $\hat{\theta}_j = \hat{\beta}_j / (1 + \sum_j \hat{\beta}_j)$ with $\hat{\theta}_0 = 1 / (1 + \sum_j \hat{\beta}_j)$. The recursion $R_t^{\text{true}} = (R_t^{\text{rep}} - \sum_{j=1}^H \hat{\theta}_j R_{t-j}^{\text{true}}) / \hat{\theta}_0$ is seeded with demeaned observed values over the first H quarters, so the first H unsmoothed observations are approximations; dropping them changes the estimated volatilities by less than 0.2 percentage points.

The approach requires a public proxy for each private class, and it attributes to the private asset only the risk its benchmark can explain, together with the asset's own residual. We pair each private index with an exchange-traded proxy: the Russell 2000 for private equity, the Russell 2000 Growth for venture capital, an iBoxx US high-yield corporate bond index for private debt, an MSCI US real-estate index for private real estate, and an S&P Global Infrastructure index for infrastructure.¹³

3.2 Estimated smoothing and its effects: volatilities, autocorrelations, and correlations

Table 3 reports the estimated smoothing parameters and the resulting unsmoothed volatilities. All three estimators detect significant smoothing. The Geltner first-lag coefficient is significant at the 1 percent level for every class and ranges from 0.48 to 0.58, implying that roughly half of each quarter's reported return is an echo of the previous quarter. The GLM lookback selected by AIC ranges from one quarter for private debt to four for private real estate. And the Dimson F-test rejects the joint hypothesis that all lagged betas are zero at the 1 percent level in every case, with individually significant loadings extending to the fourth lag for most classes — direct evidence that public-market information impacts reported private valuations with a lag of up to a year.

Placing reported returns alongside their smoothing-adjusted counterparts illustrates the economic content of these parameters. Figure 3 plots the reported and smoothing-adjusted quarterly return series for each of the four smoothed asset classes under all three estimators. Reported returns stay close to their average, while the adjusted returns fluctuate considerably more — most visibly in 2020:Q1, when the reported series registered a moderate decline of about 6 percent against an adjusted decline of about 15 percent. The COVID-19 quarter is the clearest available natural experiment on smoothing, because neither the size nor the timing of the public-market shock is in dispute.

The effect on second moments is substantial (Figure 4, Panel A). Across estimators, annualized private-equity volatility rises from a reported 6.0 percent to between 9.7 and 12.7 percent; venture-capital volatility rises from 8.5 percent to between 15.6 and 22.4 percent; private-real-estate volatility rises from 4.5 percent to between 7.0 and 12.4 percent; and private debt, the least heavily smoothed of the four, rises only from 3.8 percent to between 5.4 and 5.8 percent.

The autocorrelation evidence points the same way (Figure 4, Panel B). Reported first-order autocorrelations of 0.37 to 0.68 lie well outside the band statistically indistinguishable from zero at the 95 percent level, while all three corrections pull every series inside it, to between -0.20 and $+0.11$. This is a genuine specification check rather than a mechanical result: none of the three estimators targets the autocorrelation of the corrected series directly, so the fact that all three deliver serially uncorrelated output is consistent with a smoothing process of roughly the assumed form having generated the reported data.

¹³ Because the reconstruction in equation (6) loads the asset's full systematic exposure onto the contemporaneous factor, a non-positive estimated total beta would imply a degenerate series; in that case we revert to reported returns. However, the reversion never triggers in the index analysis and never triggers in the fund-level draws.

Unsmoothing also restores the correlation with public markets that smoothing had suppressed (Figure 5). The correlation of each unsmoothed private series with the US public-equity benchmark rises under all three estimators — for private equity from 0.61 to between 0.74 and 0.91, and for private real estate, the most heavily smoothed class, from 0.20 to between 0.50 and 0.65. The increases are largest for the two classes whose volatilities moved most, as a common smoothing explanation requires. Because these correlations enter the diversification calculation directly, raising them is exactly what tempers the inflated benefit implied by the reported data.

4. The Efficient Mean-Variance Frontier

The natural framework is Markowitz's (1952) mean-variance efficient frontier, which traces, for each level of portfolio volatility, the maximum attainable expected return. A new asset class confers a diversification benefit if and only if it shifts that frontier up or to the left. For each of a sequence of target returns we solve the minimum-variance problem

$$\min_w \quad w' \hat{\Sigma} w \quad \text{subject to} \quad \hat{\mu}' w = \mu, \quad 1' w = 1, \quad w \geq 0 \quad (7)$$

subject to two additional constraints on the private sleeve P:

$$\sum_{i \in P} w_i \leq 0.20, \quad w_i \leq 0.10 \quad \text{for all } i \in P \quad (8)$$

Rather than searching over the unconstrained mean-variance solution, we restrict the optimization to portfolios of the kind relevant to the inclusion of private assets in DC plan investment alternatives—those a DC plan investment option, such as a target-date fund (TDF), could hold. The portfolio is fully invested and long-only; the aggregate private weight is capped at 20 percent, broadly consistent with the private allocations of large institutional investors shown in Figure 2; and no single private class may exceed 10 percent, which forces diversification within the sleeve.¹⁴ We summarize each frontier by its tangency portfolio, the point of maximum Sharpe ratio against the three-month Treasury-bill yield:

$$S^* = \max_{w \in F} \frac{\hat{\mu}' w - r^f}{\sqrt{w' \hat{\Sigma} w}} \quad (9)$$

where F denotes the feasible set defined by (7) and (8). We compute five frontiers — a public-only baseline and four public-plus-private frontiers using, respectively, reported (smoothed) returns and the three unsmoothed series.

4.1 Baseline results

Adding private-market assets shifts the estimated efficient frontier up and to the left of the public-only frontier across the full range of risk levels shown, consistently across the three unsmoothing models (Figure 6). At an annualized expected return of 7 percent, the public-only

¹⁴ The frontier is traced over a grid of 150 target returns spanning the minimum-variance and maximum-return endpoints, solved with an active-set quadratic-programming routine. The long-only and cap constraints are not innocuous: Jagannathan and Ma (2003) show that such constraints act like shrinkage on the covariance matrix and can improve out-of-sample performance even when they are economically wrong, and DeMiguel, Garlappi, and Uppal (2009) document how estimation error degrades unconstrained mean-variance solutions. Here the constraints are policy primitives rather than statistical devices, but the same logic means our estimates are less exposed to estimation error than an unconstrained optimization would be.

frontier requires an estimated 6.7 percent annualized volatility, whereas the three unsmoothed frontiers reach the same return at 5.3 to 5.5 percent — a reduction of 1.2 to 1.4 percentage points in required risk. The reported (smoothed) frontier reaches 7 percent at just 5.0 percent volatility, which is a direct measure of how much of the apparent benefit is a smoothing artifact.

The frontier comparison is summarized by the tangency Sharpe ratio (Figure 7). Each unsmoothed model yields a higher estimated Sharpe ratio than the public-only baseline of 0.85, and the three cluster tightly — 1.04 (Geltner), 1.05 (GLM), and 1.01 (Dimson) — an estimated improvement of 19 to 23 percent. All three lie well below the 1.12 implied by reported returns. To separate the part of the raw benefit that survives correction from the part attributable to smoothing, we compute the surviving share

$$s = \frac{s^{\text{unsm}} - s^{\text{pub}}}{s^{\text{rep}} - s^{\text{pub}}} \times 100\% \quad (10)$$

Figure 8 presents the decomposition: the Geltner model preserves about 70 percent of the benefit, attributing 30 percent to smoothing; GLM gives a very similar 74 percent; and Dimson preserves about 60 percent, the most conservative of the three. That all three models — built on quite different assumptions about how smoothing arises — preserve a majority of the benefit confirms that the diversification gain is not an artifact of any one correction. The bootstrap sharpens the point: in none of the simulations does the surviving share fall to zero or reach 100 percent, so the surviving benefit and the smoothing overstatement are each statistically significant (Figure 8).

To translate Sharpe ratios into cumulative-return terms, we follow Tobin's (1958) two-fund separation theorem and blend each tangency portfolio with cash so that all portfolios share a common volatility of 5.04 percent annualized, the lowest tangency volatility under the log-return convention; because cash carries no volatility, this leaves each Sharpe ratio unchanged. On \$1 invested at the start of 2010, the public-only portfolio grows to an estimated \$2.41 by 2025:Q3, while the three unsmoothed portfolios — the appropriate basis for comparison — end between \$2.73 and \$2.81, roughly 13 to 17 percent above the public-only outcome (Figure 9). The reported-returns portfolio ends highest, at \$2.97, but overstates the achievable outcome because the smoothed series understates the volatility against which the cash blend is calibrated.

4.2 Statistical tests of the frontier shift

Is the public-plus-private frontier statistically distinguishable from the public-only frontier at all? We answer with three tests that rest on very different assumptions. The first is the Gibbons, Ross, and Shanken (1989) intercept test of whether the five private assets improve the unconstrained tangency of the three public classes — a necessary but not sufficient condition for the constrained gain. Regressing each private asset's excess return on the public classes, the test of the joint null that all five intercepts are zero is

$$J = \frac{T - N - K}{N} \cdot \left(1 + \hat{\mu}' \hat{\Sigma}_F^{-1} \hat{\mu}\right)^{-1} \cdot \hat{\alpha}' \hat{\Sigma}_e^{-1} \hat{\alpha} \sim F(N, T - N - K) \quad (11)$$

All four return specifications reject the null decisively, with F-statistics between 15.1 and 18.2 and p-values below 0.0001 (Table 4, Panel A); the four statistics cluster tightly, so the

conclusion does not depend on the choice of correction.¹⁵ The per-asset intercepts (Panel B) show where the abnormal return concentrates: private equity, private debt, and infrastructure carry positive, significant alphas under every specification, while venture capital's alpha is insignificant under every correction — negative under Geltner and Dimson, essentially zero under GLM — and real estate's fades under correction — significant only under Geltner, marginal under Dimson, insignificant under GLM.

The second test is a nonparametric comparison of realized returns at common volatility. We cash-blend the with-private and public-only tangencies to the common target and apply a paired-difference test to

$$d_t = \lambda_{\text{pvt}} r_{\text{pvt},t} + (1 - \lambda_{\text{pvt}}) r_{f,t} - [\lambda_{\text{pub}} r_{\text{pub},t} + (1 - \lambda_{\text{pub}}) r_{f,t}] \quad (12)$$

using Newey–West (1987) standard errors with four lags and a Bartlett kernel,

$$\hat{\sigma}_{\text{HAC}}^2 = \hat{\gamma}_0 + 2 \sum_{j=1}^q \left(1 - \frac{j}{q+1}\right) \hat{\gamma}_j, \quad q = 4 \quad (13)$$

a formulation robust to high correlation between the two tangencies.¹⁶ All four specifications reject equality at the 5 percent one-sided level, with z-statistics between +2.22 and +2.47 (Table 5).

Third, the block-bootstrap replicates described in Section 4.3 bound the sampling variability of the shift: the 5th percentile of the gain distribution remains positive under every specification, at +0.08 (Geltner) to +0.10 (reported). Because the public-only portfolio is feasible within the with-private problem, the estimated gain cannot be negative in any simulation by construction; we therefore read the bootstrap as an interval on the size of the shift rather than as an independent test of its existence, and rely on the first two tests for formal significance. Table 6 summarizes all three.

4.3 Simulation-based uncertainty: bootstrap distributions and frontier bands

A point estimate computed on 63 quarterly observations requires an explicit account of sampling uncertainty. We use the circular moving-block bootstrap of Künsch (1989) and Politis and Romano (1992), which preserves short-run dependence by resampling in contiguous blocks: 10,000 simulations, each formed by resampling the 63 quarters in four-quarter blocks with circular wrap-around.¹⁷ On each simulation we recompute the mean vector and covariance

¹⁵ The Gibbons, Ross, and Shanken (1989) statistic is exact under multivariate normal, homoskedastic, serially uncorrelated residuals — conditions that quarterly private-asset returns are unlikely to meet fully. We therefore pair the test with the common-volatility comparison in equation (12), which makes minimal distributional assumptions, and with the nonparametric bootstrap. Because the three tests rest on different assumptions and are sensitive to different departures from them, agreement among all three is more informative than any single p-value.

¹⁶ The Jobson–Korkie (1981) statistic as corrected by Memmel (2003) is the standard test of equal Sharpe ratios, but its estimated variance can turn negative when the two portfolios are highly correlated. We therefore compare realized returns at a common volatility and apply Newey–West (1987) standard errors, in the spirit of the robust Sharpe-ratio inference advocated by Ledoit and Wolf (2008).

¹⁷ Four-quarter blocks are chosen to span the longest smoothing lag selected by any estimator, so that a simulation preserves the serial dependence the correction is designed to remove. The circular construction of Politis and

matrix from the already-unsmoothed series — holding the smoothing-model coefficients fixed at their full-sample values, which isolates sampling uncertainty in the moments — and re-solve the constrained tangency for all five specifications (Figure 10).

Every estimated gain distribution is bounded above zero. The 5th-percentile gain ranges from +0.08 under the Geltner correction, the least favorable of the four specifications on this measure, to +0.10 under reported returns, so even under a pessimistic one-in-twenty draw the tangency Sharpe ratio still improves. Sharpe levels are far noisier — their 5th-to-95th-percentile bands span roughly 0.6 to 2.2 — but the public-only and with-private level draws are correlated across simulations, so uncertainty as to the improvement is an order of magnitude tighter than the level bands would suggest.¹⁸

The same simulations allow uncertainty bands on the whole frontier, not just its tangency point (Figure 11). At each level of annualized volatility we compute the distribution across simulations of the attainable expected return and of the with-private improvement, restricting attention to volatility levels reached by at least half of the simulations. Two findings stand out. First, over the volatility range attained by at least 50 percent of simulations, the improvement is positive with one-sided $p \leq 0.05$ at every volatility level, for every specification. Second, private assets expand the opportunity set itself: the global minimum-variance volatility falls from an estimated 4.6 percent (public only) to 3.7 to 3.8 percent under the corrected specifications, a paired within-simulation reduction of 0.8 to 0.9 percentage points whose bootstrap 5th percentile remains positive in every specification. The with-private investor can, in other words, reach volatility levels the public-only investor cannot reach at all — a form of benefit that a Sharpe-ratio comparison alone does not capture.

4.4 Sensitivities: allocation caps, return priors, estimation window, and tail risk

The 20 percent aggregate cap is intended to capture the allocation a DC plan option may hold, but that number could reasonably be set higher or lower, so we re-solve the optimization at caps of 15 and 25 percent (Figure 12). The improvement over the public-only baseline is present at every cap and rises as the optimizer is permitted more of the diversifying private asset classes: the corrected tangency Sharpe ratio runs from 0.98 to 1.00 at a 15 percent cap, 1.01 to 1.05 at 20 percent, and 1.05 to 1.10 at 25 percent, against the invariant public-only 0.85 (Panel A). The share of the benefit attributable to smoothing also rises with the cap — for the Dimson model from 33 percent at a 15 percent cap to 46 percent at a 25 percent cap (Panel B) — which is precisely why unsmoothing matters more, not less, as allocations to private assets grow. Even at a 25 percent cap, the majority of the estimated benefit is retained under every model. Panel C translates these Sharpe-ratio gains into cumulative-return terms. Holding volatility constant across portfolios, a portfolio with a private sleeve ends the January 2010–September 2025 period with a higher estimated value than the public-only portfolio, and the margin widens with the

Romano (1992) wraps the sample so every quarter has equal resampling probability, avoiding the end-of-sample underweighting of the original moving-block bootstrap of Künsch (1989).

¹⁸ Two conventions govern the display of bootstrap uncertainty. First, whiskers and bands are drawn on the gain rather than the level, because the public-only and with-private Sharpe levels are correlated across simulations, so overlapping level intervals would badly understate the precision of the difference (Schenker and Gentleman 2001). Second, displayed medians are aligned to the full-sample point estimates for presentation only; every p-value and percentile quoted in the text is computed from the raw draws.

allocation: even at the lowest cap of 15 percent, estimated end values are 10 to 12 percent higher than the public-only benchmark.

The expected-return vector in the baseline comes entirely from realized 2010–2025 averages — a window of unusually strong returns in which public equity returns outpaced those of every private class. Holding each model's covariance fixed, we re-solve the constrained tangency under a zero-alpha prior that forces each private class to earn the historical mean of its public comparable (Table 7). The gain is +0.19, +0.20, and +0.16 under the historical prior (Geltner, GLM, Dimson) and +0.24, +0.20, and +0.17 under zero alpha — essentially unchanged under GLM and larger under the other two. The pattern follows from the window: because the US-equity mean exceeds every private mean,¹⁹ the prior raises expected returns for the equity-mapped classes and lowers them for the bond-mapped pair, and the optimizer rotates the sleeve toward the former while the covariance structure is unchanged. The benefit under this prior therefore requires no return premium over the public comparables, which is exactly what the zero-alpha experiment is designed to isolate. Appendix A examines the same question from the opposite direction, replacing realized means with forward-looking assumptions, published by 41 investment advisory firms. Sharpe levels fall by roughly half — the survey compresses the equity premium and more than doubles the cash rate — but the estimated gain survives essentially intact at +0.15 to +0.17, remains positive at +0.08 to +0.10 when every private class is assigned the survey median of its closest public comparable, and stays positive in all 60,000 Monte Carlo draws over the cross-adviser distribution of published assumptions, with a 5th percentile of +0.09 to +0.12.

An additional sensitivity concerns the length of the estimation window rather than its character. The baseline period starts in 2010:Q1 by design: post-crisis vintages better represent today's private-market opportunity set, and the post-GFC span avoids blending two structurally different private-market regimes into one covariance matrix. Additionally, Preqin's coverage is materially broader over the post-2010 span, with earlier years built from thinner constituent-fund samples — another reason the longer window enters as a sensitivity rather than the baseline. Nonetheless, 63 quarters is a short sample on which to estimate an eight-asset covariance matrix, and a window opening in 2010:Q1 excludes the 2008 global financial crisis. We therefore re-estimate the frontier on Preqin index history extending back to 2002:Q1, paired with the same public series (Table 8). Because Preqin's infrastructure index begins only in 2006:Q4 the exercise uses a seven-asset universe, with the baseline re-estimated on the same seven assets for comparability.²⁰ Measured against each window's own realized bill rate, the estimated gain is larger on the longer sample — +0.17 under Geltner and +0.24 under GLM, against +0.14 and +0.15 on the re-estimated baseline — and the frontier-expansion tests reported below reject the no-expansion null in both windows. However, the comparison is shaped by the environment of the added years: the extended window includes two major public-equity drawdowns as well as a higher pre-2008 risk-free rate, so the public-only tangency Sharpe ratio falls from 0.85 to 0.72 across windows and the larger gain partly reflects a weaker public benchmark rather than a

¹⁹ These comparisons are based on arithmetic means of the optimization inputs (unsmoothed private series), under which US equity averages 14.5 percent versus 13.2 percent or less for every private class. On Table 1's log-return basis the reported private-equity and US-equity means are nearly equal; the two bases differ by roughly half the return variance, which is far larger for the more volatile public-equity series.

²⁰ Because the private-debt benchmark used for the Dimson correction is not available prior to 2008, the extended results cover the reported, Geltner, and GLM specifications.

stronger private sleeve. Taken together, the exercise indicates that the central result is not an artifact of the estimation window: the estimated diversification benefit remains positive, and the frontier expansion statistically significant, whether or not the sample includes the global financial crisis.

Mean-variance optimization is silent on tail behavior, yet the failure mode that should most concern a retirement saver is diversification that vanishes in a drawdown. We therefore scale every portfolio to a common 5.0 percent target volatility — the lowest tangency volatility under the arithmetic convention, so the comparison is like-for-like — and compute conditional value-at-risk (Table 9).²¹

$$\text{CVaR}_\alpha = E[r_p \mid r_p \leq q_\alpha], \quad \alpha = 0.05 \quad (14)$$

At common volatility, single-quarter tail losses are modestly better for the with-private portfolios, with 95 percent conditional value-at-risk of -4.3 to -4.4 percent for the corrected specifications against -4.7 percent public-only. The improvement is genuine but should be read for what it is: the co-movement between the private and public portions in the public portfolio's worst-decile quarters is positive in every specification (roughly $+0.3$ to $+0.5$), so the private sleeve diversifies rather than hedges — it spreads risk, including in the tail, but does not offset public losses in the worst quarters.

5. The Investable Universe: A Fund-Level Simulation Analysis

The index-level analysis implicitly assumes an investor holds the fully diversified index of every fund in every vintage — a breadth of diversification no single vehicle currently assembles. What matters for an institutional investor or a DC plan fiduciary is the frontier available to an investable book of funds. Following the sampling design of Brown, Gonçalves, and Hu (2026), we replace each private index with simulated investable series built from randomly selected funds and re-estimate the entire pipeline on each draw. A book holding roughly nine randomly selected funds from each of some thirty vintages, value-weighted and measured only once the book has seasoned, is broadly similar in structure to the underlying portfolio of a mature multi-manager evergreen vehicle — the kind of access structure that may be used in TDF-private partnerships. The simulation therefore isolates a single new friction — underdiversification, or fund-selection risk — while holding everything else identical to the index analysis: the eight-asset universe, the three unsmoothing estimators applied within each draw, the constraints, and the 2010:Q1–2025:Q3 measurement window.

5.1 Simulation design

Each draw proceeds in four steps. First, select: for each private category and each vintage year from 1990 through 2025, draw n funds at random with equal probability and without replacement, taking all funds where fewer than n exist. The vintage range opens twenty years before the measurement window²² so that the book is fully seasoned by 2010:Q1, holding live

²¹ Conditional value-at-risk is estimated as the sample average of the returns at or below the empirical 5th percentile, following the formulation of Rockafellar and Uryasev (2000).

²² The twenty-year horizon follows BGH, whose steady-state allocation aggregates the trailing twenty vintage years.

net asset value across the full distribution of fund ages rather than ramping up within the window. Second, weight: value-weight commitments across the selected funds within each vintage, and scale each vintage's total commitment in proportion to the capital actually raised by that category-vintage, mirroring BGH's cross-vintage scaling. Third, pool: each quarter, sum net asset values, capital calls, and distributions across all holdings alive that quarter and compute one pooled return with the day-weighted Modified Dietz formula, the convention Preqin uses for its published indexes,

$$R_t^{\text{pool}} = \frac{\text{NAV}_t - \text{NAV}_{t-1} - C_t + D_t}{\text{NAV}_{t-1} + \sum_i \omega_i C_{i,t} - \sum_i \omega_i D_{i,t}} \quad (15)$$

where C and D are the quarter's pooled capital calls and distributions and the day weights measure the fraction of the quarter each flow was invested.²³ Pre-window quarters are retained to supply lags for the unsmoothing models. Because the pooled book holds some thirty vintages simultaneously, young funds' J-curve quarters are a small NAV-weighted sliver in a pot dominated by mature vintages, so life-cycle effects average out cross-sectionally. Fourth, unsmooth and optimize: apply the three unsmoothing estimators to each draw's pooled category series, then solve the identical constrained tangency problem. The Geltner and Dimson corrections are re-estimated within each draw, under the same safeguards as the index analysis; the GLM correction is instead calibrated once per asset class and applied as a fixed filter.²⁴

Following BGH, we run 10,000 draws per scenario at $n = 9$ (representing a mature multi-manager platform book holding), $n = 3$ (representing a concentrated or single-manager vehicle), and $n = \text{all}$ (the value-weighted universe, a validation case). A second set of simulations compounds fund selection with the four-quarter block bootstrap of Section 4.3, re-centered so both layers share the same median.

5.2 Simulation results

Table 10 reports the headline comparison; Figures 13 and 14 add the Sharpe levels with uncertainty bands and the full gain distributions, respectively. Two findings emerge. First, moving from the universe of all funds²⁵ to a nine-fund-per-vintage book changes the median gain remarkably little: +0.125 (Geltner), +0.141 (GLM), and +0.119 (Dimson). The structural reason is that the 20 percent aggregate and 10 percent per-class caps mean the optimizer never fully exploited the index-level moments in the first place, so the gain is far less sensitive to deterioration in the private sleeve's standalone Sharpe ratio than an unconstrained analysis would be. Second, what underdiversification changes is certainty rather than central tendency: one frontier becomes a distribution of frontiers. The 5th percentile across fund-selection draws

²³ The optimizer never sees individual funds. Letting it choose weights across funds would load onto ex-post winners and measure lookahead bias rather than an investable frontier.

²⁴ Maximum-likelihood estimation of a moving-average process on a short, noisy pooled series is fragile, and re-fitting it inside every draw produces occasional near-explosive inversions that contaminate the simulated tails. We therefore calibrate the GLM filter once per asset class on the value-weighted all-funds book and apply it as a fixed filter in every draw, in the spirit of the common-filter result of Coutts, Gonçalves, and Rossi (2024).

²⁵ The value-weighted all-funds book delivers gains of +0.12 to +0.14 under the corrected specifications, modestly below the corresponding index-level point estimates of +0.16 to +0.20 under the same arithmetic convention. The gap is a construction difference rather than sampling noise: the Preqin index draws on a broader universe than the cash-flow fund panel — including open-end and NAV-only-reporting vehicles that the fund-level book cannot replicate — with weighting differences contributing at the margin.

remains clearly positive at $n = 9$, at +0.10 to +0.12 for the corrected specifications, and even when fund selection is compounded with block-bootstrap sampling uncertainty the 5th percentile stays above zero, at about +0.05 both at $n = 9$ and at $n = 3$. The concentrated $n = 3$ case widens the tails without moving the median much; Figure 15 overlays the two book sizes against the all-funds reference and shows the pattern directly.

6. The Diversification Gain Net of Private-Sleeve Fees

A recurring theme in the literature on retail and DC access to private markets is that the delivery vehicle, not the asset class, is where the economics can change. Funds of funds and other delegated intermediaries charge an additional fee on top of the underlying funds' net-of-fee terms (Harris, Jenkinson, Kaplan, and Stucke 2018; Andonov 2024), and in the DC setting specifically, Brown et al. (2022) weigh the potential diversification and return benefits of private equity against the costs of providing access, noting that the net benefit will depend both on how private funds perform and on how the institutional features of DC plans evolve.^{26,27}

This private-sleeve fee is relevant to interpreting our estimates. The return series used in this paper are net of fund-level fees — the Preqin indexes and the fund-level cash flows are both net of the underlying funds' management fees and carried interest — but gross of any private-sleeve fee. Whether, and how much, a defined-contribution implementation would add such a charge is difficult to predict: target-date funds are typically structured as funds of funds, and a private sleeve could arrive through an underlying vehicle performing the functions the funds-of-funds literature describes — selecting funds, pacing commitments, assembling and rebalancing the portfolio — but the pricing of those services varies with the implementation route, the allocator's scale and fee tier, and the fee structure itself (Andonov 2024; Begenau and Siriwardane 2024),²⁸ and we take no position on what any vehicle may charge. Nonetheless, we examine, as a sensitivity exercise, how an additional fee at levels documented in the academic literature would affect the diversification gain estimated in Sections 4 and 5.

At each annual fee level we subtract the fee, converted to quarterly, from every private return series and re-solve the constrained optimization:

$$g(f) = S^*(\tilde{\mu}(f), \hat{\Sigma}) - S^{\text{pub}}, \quad \tilde{\mu}_i(f) = \hat{\mu}_i - f \cdot 1\{i \in P\} \quad (16)$$

perturbing the mean vector alone — a fee reduces expected return without altering risk — and charging the fee on sleeve assets, following Brown, Hu, and Kuhn (2021) and matching how

²⁶ Fees on the investment options most widely used in DC plans, such as target-date funds, have continued to fall in recent years, reaching an asset-weighted average of 0.27 percent for target date mutual funds in 2025 (Li 2026). A similar decline for vehicles that hold private assets would make the lower end of the fee range examined below the more relevant.

²⁷ Pegoraro, Shive, and Zambrana (2025) similarly show that the design and pricing of semi-liquid retail vehicles materially shape realized outcomes. We do not calibrate to the expense ratios of such vehicles, or to all-in cost estimates for private-markets implementation, because measures of that kind embed the fund-level economics our returns are already net of.

²⁸ Delegated routes cost materially more than internal management or co-investment (Andonov 2024); management fees on identical positions in the same fund differ by 91 basis points across investor fee tiers (Begenau and Siriwardane 2024); charges may be flat or performance-contingent, with step-downs over the fund's life (Harris et al. 2018); and some costs — notably the drag from semi-liquid vehicles' liquidity buffers — sit outside the fee schedule altogether (Ewens and Faber 2026).

fund-of-funds schedules and underlying-fund expenses apply. Because the 20 percent cap binds in most cases, a fee expressed on sleeve assets translates to roughly one-fifth of that on the total portfolio — 100 basis points of sleeve net asset value is approximately 20 basis points of portfolio assets. We report results at 0, 50, 100, and 150 basis points of sleeve net asset value per year, a grid spanning the range documented in the literature — from the near-zero cost of a large in-house program (Fang, Ivashina, and Lerner 2015) through the contractual management-fee-plus-carry terms of the funds of funds studied by Harris et al. (2018).²⁹

Figure 16 plots the index-level gain at each fee level; Tables 11 and 12 report the full grid at index and fund level. The gain declines close to linearly in the fee and remains positive in every cell. At 150 basis points the corrected index-level estimates are +0.11 to +0.14 — 68 to 73 percent of their no-fee values — with block-bootstrap 5th percentiles positive under every specification. At fund level, the median nine-fund book retains 59 to 64 percent of its no-fee gain, and the gain stays positive at the 5th percentile of both the fund-selection and the compounded selection-and-sampling distributions.

Across the fee levels considered, the smoothing treatments, and both the index and simulated-book implementations, the estimated gain is reduced but remains positive. Within the documented range, a meaningful share of the estimated benefit survives the fee — smaller at higher fee levels, but positive throughout, including at the 5th percentile of the selection-and-sampling distribution.

7. Conclusion

This paper examines whether the diversification case for a measured private-market sleeve in DC portfolios holds up under three successive refinements toward realism. The claim rests on private assets appearing less volatile than, and weakly correlated with, public markets — properties that measurement, implementation, and fees could each overstate — so we test whether the estimated gain survives all three. We correct reported returns for the appraisal smoothing that inflates those properties; we replace historical means estimated over a period of strong performance with published forward-looking assumptions; we move from the index no investor holds to the finite fund books a vehicle would hold; and we subtract fees spanning the range the funds-of-funds literature reports. The answer at each stage is yes: the estimated benefit is attenuated by some refinements and eliminated by none.

Correcting for smoothing raises measured private-asset volatilities to between roughly one-and-a-half and nearly three times their reported levels, removes significant return autocorrelation, and reveals that between a quarter and two-fifths of the naive diversification benefit is a measurement artifact. The majority survives: the corrected tangency Sharpe ratio rises from 0.85

²⁹ Harris et al. report a median 1 percent management fee plus 5 percent carried interest for primary funds of funds; the 150 basis-point upper gridpoint spans an annual-equivalent allowance for that carry, a conversion that is ours rather than theirs. Zero corresponds to the no-additional-cost baseline of Brown, Hu, and Kuhn (2021), whose robustness case assumes 50 basis points. Realized evidence suggests the net cost of the layer has historically run below contractual terms, since intermediaries recover part of the fee through selection and access — Harris et al. find excess public-market equivalents of -0.06 (buyout) and -0.10 (generalist) against simulated random direct-fund portfolios, and an insignificant $+0.02$ for venture.

to between 1.01 and 1.05, the frontier shift is statistically significant under three independent tests, and the gain remains positive across bootstrap simulations, tail regimes, allocation caps from 15 to 25 percent, and an estimation window extended to include the global financial crisis. Testing the gain against the published assumptions of 41 investment advisory firms lowers Sharpe levels throughout but leaves it intact — at the survey medians, at the least-favorable published combination, and in all 60,000 Monte Carlo draws over the cross-adviser distribution. Within the investable universe, moving from the all-funds book to simulated books of randomly selected funds changes the median gain remarkably little; what changes is certainty, one frontier becoming a distribution of frontiers, and even the 5th percentile of the compounded selection-and-sampling distribution remains positive. Finally, subtracting a private-sleeve fee of 50 to 150 basis points of sleeve assets — a range anchored in the academic literature — leaves the corrected index-level gain at 68 to 91 percent of its no-fee value and the fund-level median at 59 to 88 percent. The diversification benefit is thus smaller than reported index statistics imply, but it remains economically and statistically significant.

References

- Akaike, Hirotugu. 1974. "A New Look at the Statistical Model Identification." *IEEE Transactions on Automatic Control* 19 (6): 716–723.
- Andonov, Aleksandar. 2024. "Delegated Investment Management in Alternative Assets." *Review of Corporate Finance Studies* 13 (1): 264–301.
- Ang, Andrew, Bingxu Chen, William N. Goetzmann, and Ludovic Phalippou. 2018. "Estimating Private Equity Returns from Limited Partner Cash Flows." *Journal of Finance* 73 (4): 1751–1783.
- Anson, Mark. 2013. "Performance Measurement in Private Equity: Another Look at the Lagged Beta Effect." *Journal of Private Equity* 17 (1): 29–44.
- Anson, Mark. 2016. "Asset Allocation with Private Equity." *Journal of Investment Consulting* 17 (1): 30–36.
- Antoniewicz, Shelly, James Duvall, Sarah Holden, Lei Li, Hammad Qureshi, and Shane Worner. 2026. "ICI Economic Analysis of the Department of Labor's Designated Investment Alternatives Proposal," Investment Company Institute White Paper.
- Asness, Clifford, Robert Krail, and John Liew. 2001. "Do Hedge Funds Hedge?" *Journal of Portfolio Management* 28 (1): 6–19.
- Begenau, Juliane, and Emil Siriwardane. 2024. "Fee Variation in Private Equity." *Journal of Finance* 79 (2): 1199–1247.
- Best, Michael J., and Robert R. Grauer. 1991. "On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results." *Review of Financial Studies* 4 (2): 315–342.
- Boyer, Brian, Taylor D. Nadauld, Keith P. Vorkink, and Michael S. Weisbach. 2023. "Discount-Rate Risk in Private Equity: Evidence from Secondary Market Transactions." *Journal of Finance* 78 (2): 835–885.
- Brown, G., K. Crouch, A. Ghent, R. Harris, Y. Hochberg, T. Jenkinson, S. Kaplan, R. Maxwell, and D. Robinson. 2022. "Should Defined Contribution Plans Include Private Equity Investments?" *Financial Analysts Journal* 78 (4): 5–17.
- Brown, Gregory W., Eric Ghysels, and Oleg R. Gredil. 2023. "Nowcasting Net Asset Values: The Case of Private Equity." *Review of Financial Studies* 36 (3): 945–986.
- Brown, Gregory W., Andrei S. Gonçalves, and Wendy Hu. 2026. "The Private Capital Alpha." Working paper, Ohio State University.
- Brown, Gregory W., Wendy Hu, and Bert-Klemens Kuhn. 2021. "Private Investments in Diversified Portfolios." Working paper, Institute for Private Capital, University of North Carolina at Chapel Hill, January 29.
- Brown, Gregory W., Christian T. Lundblad, and William Volckmann. 2025. "Risk-Adjusted Performance of Private Funds: What Do We Know?" Working paper, Institute for Private Capital.
- Chopra, Vijay K., and William T. Ziemba. 1993. "The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice." *Journal of Portfolio Management* 19 (2): 6–11.
- Cochrane, John H. 2005. "The Risk and Return of Venture Capital." *Journal of Financial Economics* 75 (1): 3–52.
- Couts, Spencer J., Andrei S. Gonçalves, and Andrea Rossi. 2024. "Unsmoothing Returns of Illiquid Funds." *Review of Financial Studies* 37 (7): 2110–2155.

- DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal. 2009. "Optimal Versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy?" *Review of Financial Studies* 22 (5): 1915–1953.
- Dimson, Elroy. 1979. "Risk Measurement When Shares Are Subject to Infrequent Trading." *Journal of Financial Economics* 7 (2): 197–226.
- Ewens, Michael, and Jacob W. Faber. 2026. "Liquid Claims on Illiquid Assets: The Economics of Retail Access to Private Markets." Working paper, February.
- Fang, Lily H., Victoria Ivashina, and Josh Lerner. 2015. "The Disintermediation of Financial Markets: Direct Investing in Private Equity." *Journal of Financial Economics* 116 (1): 160–178.
- Geltner, David. 1991. "Smoothing in Appraisal-Based Returns." *Journal of Real Estate Finance and Economics* 4 (3): 327–345.
- Geltner, David. 1993. "Estimating Market Values from Appraised Values without Assuming an Efficient Market." *Journal of Real Estate Research* 8 (3): 325–345.
- Getmansky, Mila, Andrew W. Lo, and Igor Makarov. 2004. "An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns." *Journal of Financial Economics* 74 (3): 529–609.
- Gibbons, Michael R., Stephen A. Ross, and Jay Shanken. 1989. "A Test of the Efficiency of a Given Portfolio." *Econometrica* 57 (5): 1121–1152.
- Gupta, Arpit, and Stijn Van Nieuwerburgh. 2021. "Valuing Private Equity Investments Strip by Strip." *Journal of Finance* 76 (6): 3255–3307.
- Harris, Robert S., Tim Jenkinson, and Steven N. Kaplan. 2014. "Private Equity Performance: What Do We Know?" *Journal of Finance* 69 (5): 1851–1882.
- Harris, Robert S., Tim Jenkinson, Steven N. Kaplan, and Rüdiger Stucke. 2018. "Financial Intermediation in Private Equity: How Well Do Funds of Funds Perform?" *Journal of Financial Economics* 129 (2): 287–305.
- Horizon Actuarial Services. 2024. *Survey of Capital Market Assumptions: 2024 Edition*. Silver Spring, MD: Horizon Actuarial Services, LLC.
- Horizon Actuarial Services. 2025. *Survey of Capital Market Assumptions: 2025 Edition*. Silver Spring, MD: Horizon Actuarial Services, LLC.
- Investment Company Institute. 2026. Comment letter on the Department of Labor's proposed rule, "Fiduciary Duties in Selecting Designated Investment Alternatives" (RIN 1210-AC38), filed June 1, 2026. Washington, DC: Investment Company Institute.
- Jagannathan, Ravi, and Tongshu Ma. 2003. "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps." *Journal of Finance* 58 (4): 1651–1683.
- Jobson, J. D., and Bob M. Korkie. 1981. "Performance Hypothesis Testing with the Sharpe and Treynor Measures." *Journal of Finance* 36 (4): 889–908.
- Kaplan, Steven N., and Antoinette Schoar. 2005. "Private Equity Performance: Returns, Persistence, and Capital Flows." *Journal of Finance* 60 (4): 1791–1823.
- Korteweg, Arthur, and Stefan Nagel. 2016. "Risk-Adjusting the Returns to Venture Capital." *Journal of Finance* 71 (3): 1437–1470.
- Korteweg, Arthur, and Stefan Nagel. 2025. "Risk-Adjusted Returns of Private Equity Funds: A New Approach." *Review of Financial Studies* 38 (9): 2557–2601.
- Künsch, Hans R. 1989. "The Jackknife and the Bootstrap for General Stationary Observations." *Annals of Statistics* 17 (3): 1217–1241.

- Ledoit, Olivier, and Michael Wolf. 2008. "Robust Performance Hypothesis Testing with the Sharpe Ratio." *Journal of Empirical Finance* 15 (5): 850–859.
- Li, Lei. 2026. "Trends in the Expenses and Fees of Funds, 2025." *ICI Research Perspective* 32, no. 1
- Ljung, Greta M., and George E. P. Box. 1978. "On a Measure of Lack of Fit in Time Series Models." *Biometrika* 65 (2): 297–303.
- Lo, Andrew W. 2002. "The Statistics of Sharpe Ratios." *Financial Analysts Journal* 58 (4): 36–52.
- Markowitz, Harry. 1952. "Portfolio Selection." *Journal of Finance* 7 (1): 77–91.
- Mommel, Christoph. 2003. "Performance Hypothesis Testing with the Sharpe Ratio." *Finance Letters* 1: 21–23.
- Metrick, Andrew, and Ayako Yasuda. 2010. "The Economics of Private Equity Funds." *Review of Financial Studies* 23 (6): 2303–2341.
- Michaud, Richard O. 1989. "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal* 45 (1): 31–42.
- Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–708.
- Pegoraro, Stefano, Sophie Shive, and Rafael Zambrana. 2025. "Democratizing Illiquid Assets: Liquidity Transformation and Performance in Interval Funds." Working paper, University of Notre Dame.
- Phalippou, Ludovic, and Oliver Gottschalg. 2009. "The Performance of Private Equity Funds." *Review of Financial Studies* 22 (4): 1747–1776.
- Politis, Dimitris N., and Joseph P. Romano. 1992. "A Circular Block-Resampling Procedure for Stationary Data." In *Exploring the Limits of Bootstrap*, edited by Raoul LePage and Lynne Billard, 263–270. New York: Wiley.
- Rockafellar, R. Tyrrell, and Stanislav Uryasev. 2000. "Optimization of Conditional Value-at-Risk." *Journal of Risk* 2 (3): 21–41.
- Schenker, Nathaniel, and Jane F. Gentleman. 2001. "On Judging the Significance of Differences by Examining the Overlap Between Confidence Intervals." *The American Statistician* 55 (3): 182–186.
- Sorensen, Morten, Neng Wang, and Jinqiang Yang. 2014. "Valuing Private Equity." *Review of Financial Studies* 27 (7): 1977–2021.
- Stafford, Erik. 2022. "Replicating Private Equity with Value Investing, Homemade Leverage, and Hold-to-Maturity Accounting." *Review of Financial Studies* 35 (1): 299–342.
- Tobin, James. 1958. "Liquidity Preference as Behavior towards Risk." *Review of Economic Studies* 25 (2): 65–86.

Exhibits

Table 1. Summary Statistics for Private and Public Market Benchmarks

Quarterly total returns, 2010:Q1–2025:Q3; means and volatilities annualized

	Private market returns					Public market returns		
	Private equity	Venture capital	Private debt	Private real estate	Private infrastructure	US equity	World equity	US bond
Annualized mean return	13.1%	11.0%	9.1%	8.7%	9.0%	13.1%	6.8%	2.5%
Annualized volatility	5.8%	8.2%	3.8%	4.4%	2.8%	15.8%	16.6%	4.7%
Autocorrelation	0.49	0.67	0.36	0.58	0.11	−0.17	−0.14	0.11
Ljung-Box Test	15.8***	29.2***	8.1*	21.0***	0.77	1.95	1.29	0.85

Notes: All returns are adjusted for expenses and are expressed as log returns. Autocorrelation is the first-order sample autocorrelation of quarterly returns. The Ljung–Box statistic tests the joint null that the first four autocorrelations are zero. *** = significant at the 1% level; ** = 5%; * = 10%.

Table 2. Preqin Private-Capital Funds in the Fund-Level Sample

	Private equity (ex VC)	Venture capital	Private debt	Real estate	Infrastructure	All classes
Panel A — Universe and coverage (vintages 1990–2025)						
Number of funds	2,531	1,053	676	1,064	181	5,505
First vintage year	1990	1990	1990	1992	2003	1990
Last vintage year	2025	2025	2025	2025	2025	2025
Median funds per vintage year	64.5	27.5	15.0	39.0	7.0	148.5
Avg funds reporting per quarter (2010:Q1–2025:Q3)	815	310	184	297	50	1,656
Panel B — Fund characteristics						
Fund size, mean (\$mn)	1,560	414	1,486	943	2,635	1,242
Fund size, median (\$mn)	615	260	793	500	1,000	500
Fund size, p10 (\$mn)	122	57	162	100	185	101
Fund size, p90 (\$mn)	3,592	800	3,400	1,975	6,000	2,800
Median life, liquidated funds (years)	14	14	12	8	10	12
Panel C — Cash-flow						
Median capital calls per fund	14	13	10	9	11	12
Median distributions per fund	18	15	15	14	12	16

Table 3. Smoothing-Parameter Estimates and Implied Unsmoothed Volatilities*Quarterly returns, 2010:Q1–2025:Q3; volatilities annualized, percent*

Panel A. Geltner autoregressive model			
Asset class	φ_1	φ_2	Unsmoothed volatility
Private equity	0.52***	-0.03	10.2%
Venture capital	0.58***	0.14	22.4%
Private debt	0.55***	-0.22*	5.4%
Private real estate	0.48***	0.08	8.3%

Panel B. GLM moving-average model							
Asset class	$\theta(0)$	$\theta(1)$	$\theta(2)$	$\theta(3)$	$\theta(4)$	Selected lag (AIC)	Unsmoothed volatility
Private equity	0.53	0.28	0.19	—	—	2	9.7%
Venture capital	0.37	0.22	0.28	0.13	—	3	15.6%
Private debt	0.66	0.34	—	—	—	1	5.4%
Private real estate	0.31	0.24	0.25	0.11	0.09	4	12.4%

Panel C. Dimson lagged-beta model							
Asset class	β_0	β_1	β_2	β_3	β_4	F-stat	Unsmoothed volatility
Private equity	0.23***	0.15***	0.09***	0.07***	0.07***	16.91***	12.7%
Venture capital	0.26***	0.19***	0.15***	0.12***	0.09***	15.18***	18.1%
Private debt	0.29***	0.23***	0.09*	0.08**	-0.00	8.02***	5.8%
Private real estate	0.11***	0.11***	0.05*	0.02	0.06**	4.48***	7.0%

Notes: Panel A reports the AR(2) coefficients of equation (2); Panel B the normalized moving-average weights of equation (4) and the lookback selected by the Akaike Information Criterion; Panel C the contemporaneous and lagged betas of equation (5) and the F-statistic for the joint null that all lagged betas are zero. *** = significant at the 1% level; ** = 5%; * = 10%. Smoothing parameters are estimated on quarterly returns from 2009:Q1, with the 2009 quarters supplying initial lags; all reported moments are computed over 2010:Q1–2025:Q3. The unsmoothed GLM series is recovered by recursive inversion of the normalized weights, with the recursion initialized at the demeaned reported returns

Table 4. Gibbons–Ross–Shanken Test of Frontier Expansion, by Specification*Arithmetic quarterly returns, 2010:Q1–2025:Q3***Panel A. Joint test of zero intercepts**

Specification	F-statistic	p-value	Reject H_0 at 5%?
Reported	16.92	<0.0001	Yes
Geltner	16.02	<0.0001	Yes
GLM	15.12	<0.0001	Yes
Dimson	18.19	<0.0001	Yes

Panel B. Per-asset annualized alpha (percent)

Asset	Reported	Geltner	GLM	Dimson
Private equity	+9.35***	+6.02***	+6.86***	+3.21**
Venture capital	+5.73**	−3.80	+0.05	−3.48
Private debt	+6.40***	+5.15***	+5.09***	+4.38***
Private real estate	+6.67***	+4.85**	+3.01	+3.02*
Private infrastructure	+7.28***	+7.28***	+7.28***	+7.28***

Notes: Panel A reports the statistic of equation (11) based on a reference distribution of $F(5, 55)$. Panel B reports the intercepts from regressing each private asset's excess return on the three public classes. Infrastructure is carried at reported returns in every specification, so its alpha is identical across columns. *** = significant at the 1% level; ** = 5%; * = 10%.

Table 5. Common-Volatility HAC Paired-Difference Test of Equal Sharpe Ratios

Specification	SR private	SR public	Difference	z (HAC)	p (one-sided)	Reject?
Reported	1.12	0.851	0.268	2.47	0.007	Yes
Geltner	1.04	0.851	0.187	2.38	0.009	Yes
GLM	1.05	0.851	0.197	2.22	0.013	Yes
Dimson	1.01	0.851	0.161	2.47	0.007	Yes

Notes: Both tangency portfolios are blended with cash to the common target volatility of equation (12) before the comparison. The z-statistic uses the heteroskedasticity- and autocorrelation-consistent variance of equation (13).

Table 6. Summary of the Three Frontier-Shift Tests, by Specification

Specification	GRS F	GRS p	HAC z	HAC p (one-sided)	Bootstrap (5 th percentile)
Reported	16.92	0.00	2.47	0.007	0.10
Geltner	16.02	0.00	2.38	0.009	0.08
GLM	15.12	0.00	2.22	0.013	0.09
Dimson	18.19	0.00	2.47	0.007	0.08

Notes: The bootstrap column reports the 5th percentile, across the 10,000 circular moving-block simulations of Section 4.3, of the gain in the tangency Sharpe ratio from adding the private sleeve.

Table 7. Tangency Sharpe Ratio Under Alternative Expected-Return Scenarios

Model	Scenario	Public Sharpe	+Private Sharpe	Sharpe gain
Geltner	Historical	0.85	1.04	0.19
Geltner	Zero-alpha	0.85	1.09	0.24
GLM	Historical	0.85	1.05	0.20
GLM	Zero-alpha	0.85	1.05	0.20
Dimson	Historical	0.85	1.01	0.16
Dimson	Zero-alpha	0.85	1.03	0.17

Notes: The zero-alpha scenario assigns each private class the historical mean of its closest public comparable, leaving the covariance structure unchanged.

Table 8. The Diversification Gain on a Longer Estimation Window

Specification	Sharpe-ratio gain	GRS p
Panel A. Baseline: 2010:Q1–2025:Q3		
Reported (smoothed)	0.261	<0.0001
Geltner	0.138	0.0002
GLM	0.15	<0.0001
Panel B. Extended: 2002:Q1–2025:Q3		
Reported (smoothed)	0.352	<0.0001
Geltner	0.166	0.0001
GLM	0.241	<0.0001

Notes: Both panels use a seven-asset universe excluding private infrastructure, whose Preqin index begins only in 2006:Q4; the Dimson correction requires a matched high-yield benchmark unavailable before 2008 and is omitted. p-values are for the Gibbons–Ross–Shanken test of the joint null that all private-asset intercepts are zero, $F(4, T-7)$.

Table 9. Tail-Risk Measures at a Common Target Volatility

Specification	λ (vol scalar)	CVaR-95%
Public only	0.51	−4.71%
Reported	1.00	−4.43%
Geltner	0.93	−4.31%
GLM	0.91	−4.39%
Dimson	0.87	−4.28%

Notes: Every portfolio is blended with cash to the lowest tangency volatility across the five specifications under the arithmetic convention — 5.0 percent annualized, attained by the reported-returns portfolio, whose λ is therefore 1.00. Conditional value-at-risk follows equation (14) and is the average of the returns at or below the empirical 5th percentile.

Table 10. The Diversification Gain Survives at Fund Level: Index Versus Simulated Investable Books

Specification	All funds	n = 9, median	n = 9, p5 (selection)	n = 9, p5 (selection + sampling)	n = 3, median	n = 3, p5 (selection + sampling)
Reported	0.240	0.237	0.198	0.071	0.248	0.075
Geltner	0.137	0.125	0.102	0.053	0.151	0.049
GLM	0.135	0.141	0.116	0.055	0.149	0.055
Dimson	0.119	0.119	0.099	0.055	0.124	0.046

Notes: Fund selection varies which funds are drawn; fund selection and sampling compounds fund selection with the four-quarter block bootstrap, re-centered so both layers share the same median. n is the number of funds drawn per category-vintage; books are value-weighted within vintage and scaled across vintages by capital raised.

Table 11. Index-Level Gain Net of the Private-Sleeve Fee

Gain in annualized Sharpe ratio over the public-only tangency portfolio

Specification	Public-only Sharpe ratio	Private-sleeve fee (basis points of sleeve assets per year)			
		0 (no fee)	50	100	150
Panel A. Deterministic gain					
Reported (smoothed)	0.851	+0.268	+0.248	+0.229	+0.210
Geltner	0.851	+0.187	+0.169	+0.151	+0.133
GLM	0.851	+0.197	+0.179	+0.162	+0.144
Dimson	0.851	+0.161	+0.144	+0.127	+0.110
Panel B. Block-bootstrap 5th percentile					
Reported (smoothed)	—	+0.103	+0.095	+0.086	+0.079
Geltner	—	+0.077	+0.068	+0.061	+0.055
GLM	—	+0.085	+0.077	+0.071	+0.063
Dimson	—	+0.081	+0.074	+0.063	+0.053

Notes: The private-sleeve fee is the direct fee charged on sleeve assets for selecting, holding, and administering the underlying private funds. It is in addition to the underlying funds' own fees (indirect fees), which are already reflected in the net-of-fee returns used throughout. With the 20 percent sleeve cap binding, 50, 100, and 150 basis points of sleeve assets correspond to approximately 10, 20, and 30 basis points of total portfolio assets. The public-only Sharpe ratio is the full-sample estimate; it is the same at every fee level and under every specification because the public-only portfolio holds no private assets. Panel B reports the 5th percentile of 10,000 block-bootstrap simulations, re-centered on the deterministic point estimate; the public-only Sharpe ratio varies across bootstrap draws, so no single value applies (—).


Table 12. Fund-Level Gain Net of the Private-Sleeve Fee*Gain in annualized Sharpe ratio over the public-only tangency portfolio*

Specification	Public-only Sharpe ratio	Private-sleeve fee (basis points of sleeve assets per year)			
		0 (no fee)	50	100	150
Panel A. Median simulated book					
Reported (smoothed)	0.851	+0.237	+0.218	+0.200	+0.182
Geltner	0.851	+0.125	+0.108	+0.091	+0.075
GLM	0.851	+0.141	+0.123	+0.106	+0.090
Dimson	0.851	+0.119	+0.102	+0.086	+0.071
Panel B. 5th percentile — fund selection only					
Reported (smoothed)	0.851	+0.198	+0.180	+0.163	+0.146
Geltner	0.851	+0.102	+0.086	+0.070	+0.055
GLM	0.851	+0.116	+0.100	+0.083	+0.068
Dimson	0.851	+0.099	+0.084	+0.068	+0.053
Panel C. 5th percentile — fund selection compounded with sampling					
Reported (smoothed)	—	+0.071	+0.064	+0.058	+0.049
Geltner	—	+0.053	+0.042	+0.030	+0.015
GLM	—	+0.055	+0.047	+0.038	+0.025
Dimson	—	+0.055	+0.042	+0.029	+0.020

Notes: Results are for simulated books of nine funds per category-vintage. The private-sleeve fee is the direct fee charged on sleeve assets for selecting, holding, and administering the underlying private funds. It is in addition to the underlying funds' own fees (indirect fees), which are already reflected in the net-of-fee returns used throughout. With the 20 percent sleeve cap binding, 50, 100, and 150 basis points of sleeve assets correspond to approximately 10, 20, and 30 basis points of total portfolio assets. The public-only Sharpe ratio is the full-sample estimate; it is the same at every fee level and under every specification because the public-only portfolio holds no private assets. Panel B varies the identity of the funds drawn; Panel C additionally applies the block bootstrap to the time dimension, under which the public-only Sharpe ratio varies across draws, so no single value applies (—).

Figure 1. US Private Markets Have Grown Substantially in Recent Years


Trillions of US dollars, 2010–September 2025



Source: Preqin (a part of BlackRock). Adapted from the Investment Company Institute white paper (Antoniewicz et al. 2026), Figure 2.

Figure 2. State and Local DB Pension Plan Allocations to Private and Alternative Assets

Percentage of investments, fiscal years 2001–2024




Source: Public Plans Database. Adapted from the Investment Company Institute white paper (Antoniewicz et al. 2026), Figure 3.

Figure 3. Smoothing-Adjusted Private Market Returns Are More Volatile Than Reported Returns

Quarterly returns, percent, 2010:Q1–2025:Q3; four asset classes by unsmoothing model

Panel A. GLM model



Panel B. Geltner model




Panel C. Dimson model



Notes: Each panel plots quarterly index returns, in percent, for one private asset class over 2010:Q1–2025:Q3. Dashed lines are reported (smoothed) index returns; solid lines are smoothing-adjusted returns under the estimator named in the panel-block heading — the GLM moving-average filter in Panel A, the Geltner autoregressive inversion in Panel B, and the Dimson lagged-beta reconstruction in Panel C, each estimated as described in Section 3.

Figure 4. Unsmoothing Raises Measured Volatility and Removes Return Autocorrelation


Panel A: annualized volatility, percent. Panel B: first-order autocorrelation of quarterly returns



Notes: The no-autocorrelation band in Panel B is the 95 percent interval for a single autocorrelation coefficient, approximately ± 0.247 at this sample size.

Figure 5. Unsmoothing Increases Correlations with Public Markets


Correlation of quarterly private-asset returns with US public equity, 2010:Q1–2025:Q3



Notes: The public-equity benchmark is the FT Wilshire 5000 Total Return Index.

Figure 6. Private Market Assets Can Provide Diversification Benefits and Expand the Opportunity Set


Annualized expected return and volatility, percent



Notes: The public-plus-private frontiers are calculated from smoothed (reported) returns and from each of the three unsmoothing corrections. All frontiers impose the constraints of equations (7) and (8).

Figure 7. Adding Private Market Assets Raises the Estimated Sharpe Ratio of the Tangency Portfolio


Annualized Sharpe ratio of the constrained tangency portfolio



Notes: The tangency portfolio is the frontier point of maximum Sharpe ratio against the three-month US Treasury-bill yield, subject to equations (7) and (8).

Figure 8. Decomposition of the Estimated Diversification Benefit


Percent of the reported (smoothed) diversification benefit



Notes: The benefit is decomposed with equation (10): the lower segment is the share of the reported gain in the tangency Sharpe ratio that survives the smoothing correction, and the upper segment the share attributable to appraisal smoothing. Whiskers are 5th-to-95th-percentile intervals of the surviving share, computed within each block-bootstrap simulation described in Section 4.3 as the ratio of the corrected gain to the reported gain, from raw draws without re-centering — whiskers therefore need not be centered on the point-estimate shares. In no simulation does the share fall to zero or reach 100 percent, so both components of the decomposition are statistically distinguishable from zero at well below the 1 percent one-sided level.

Figure 9. Growth of Efficient Frontier Portfolios at the Same Level of Risk


Value of \$1 invested on January 1, 2010 through September 30, 2025



Notes: Each tangency portfolio is blended with cash so that all portfolios share a common annualized volatility of 5.04 percent, the lowest tangency volatility under the log-return convention. Because cash carries no volatility, the blend leaves each Sharpe ratio unchanged.

Figure 10. Bootstrap Uncertainty in the Tangency Sharpe Ratio


Annualized Sharpe ratio




Notes: Bars correspond to full-sample point estimates; whiskers are bootstrap 5th-to-95th-percentile intervals for the improvement. Whiskers are drawn on the gain rather than the level and are anchored at the public-only point estimate, because the public-only and with-private Sharpe levels are correlated at 0.95 to 0.98 across simulations. Displayed medians are aligned to the full-sample point estimates for presentation; percentiles and p-values quoted in the text are computed from the raw draws.

Figure 11. The Frontier Improvement Is Positive Across the Full Range of Attainable Risk Levels

Annualized expected return and volatility, percent






Notes: Bands are drawn only over volatility levels attained by at least half of the 10,000 simulations.


Figure 12. The Diversification Benefit Is Present at Every Private-Allocation Cap

Panel A. Sharpe-ratio gain over the public-only portfolio by private-allocation cap




Panel B. Decomposition of the benefit by private-allocation cap

- Estimated overstated benefit from appraisal smoothing
 ■ Estimated resulting benefit after unsmoothing



Panel C. Public-plus-private investment values relative to public-only at a common risk level




Notes: The 10 percent per-class cap of equation (8) is unchanged throughout. Panel B decomposes the reported benefit with equation (10).


Figure 13. Fund-Level Sharpe Ratios Under Fund-Selection and Sampling Uncertainty

Annualized Sharpe ratio, nine funds per vintage, 10,000 simulated books

Panel A. Fund-selection dispersion



Panel B. Fund selection compounded with sampling




Notes: Bars are medians in both panels; only the whiskers differ. Panel A whiskers are the fund-selection 5th-to-95th-percentile gain interval added to the invariant public-only point Sharpe ratio; Panel B whiskers are the re-centered compound interval.

Figure 14. Distribution of the Fund-Level Diversification Gain


Gain in annualized Sharpe ratio over public only

Panel A: Nine funds per vintage

Fund selection only




Fund selection compounded with sampling




Panel B: Three funds per vintage

Fund selection only




Fund selection compounded with sampling



Notes: Panel A shows books of nine funds per vintage and Panel B books of three funds per vintage; within each panel the upper chart varies only which funds are drawn and the lower chart compounds fund selection with the block bootstrap of Section 4.3. The concentrated three-fund case widens the tails without moving the median much, the pattern that right-skewed fund returns under value weighting imply.

Figure 15. Fund Concentration Widens the Tails of the Estimated Gain Without Moving Its Median


Gain in annualized Sharpe ratio over public only, by number of funds drawn per vintage; fund-selection draws



Notes: In each panel the solid curve is the nine-fund-per-vintage book, the dashed curve the three-fund book, and the vertical line the deterministic all-funds book.

Figure 16. Index-Level Diversification Gain at Alternative private-Sleeve Fee Levels

Gain in annualized Sharpe ratio over the public-only tangency portfolio



Notes: The fee is charged on sleeve net asset value following equation (16). With the 20 percent sleeve cap binding, 50, 100, and 150 basis points of sleeve assets correspond to approximately 10, 20, and 30 basis points of total portfolio assets.

Appendix A: The Frontier Shift Under Forward-Looking Capital-Market Assumptions

The results in Sections 4 and 5 are estimated on realized 2010–2025 returns, and a natural question is whether they would carry over to an allocation decision that is necessarily a bet on future returns rather than past ones — particularly given that the estimation window had US equities compounding well above their long-run average and cash returning close to zero for more than half the sample. Optimized portfolios, moreover, can be sensitive to expected-return inputs (Michaud 1989; Best and Grauer 1991), with errors in means mattering roughly an order of magnitude more than errors in variances (Chopra and Ziemba 1993). And institutional practice is forward-looking: plan investment fiduciaries set policy allocations against published capital-market assumptions, not sample means. This appendix asks whether the frontier shift survives when historical expected returns are replaced wholesale by the assumptions used in that practice.

The forward-looking inputs come from the Horizon Actuarial Survey of Capital Market Assumptions, which compiles the published long-run assumptions of 41 investment advisory firms (Horizon Actuarial Services 2024, 2025). Two features of the survey suit the exercise. First, it reflects no single firm's view: the cross-adviser median reduces the dependence on any one source's assumptions that can make single-source exercises less persuasive. Second, it reports the dispersion of assumptions across respondents, which lets expected returns enter as a measured range of professional views rather than a single point. The design uses both in three steps, each addressing a distinct question: a central case under survey-median means (is the gain an artifact of the realized sample?); a zero-alpha variant stripping away all private return premia (does anything survive on diversification alone?); and a least-favorable bound plus a Monte Carlo over the cross-adviser distribution (does any published set of assumptions overturn the sign?).

A.1 Empirical design

The survey and our own estimates play different roles because each measures a different thing well. Expected returns are not distorted by appraisal smoothing, so the survey consensus is a natural forward-looking mean vector. Its second moments are less suitable, because the survey does not standardize how respondents treat smoothing in private-asset volatilities and correlations, whereas measuring those second moments is precisely what the machinery of Section 3 is built to do. The core specification therefore pairs survey means with our own unsmoothed covariances and is run separately for all three unsmoothing models.³⁰

Table A.1 sets the two sets of inputs side by side, and the mean comparison is itself informative. Forward-looking expected returns are markedly lower than realized ones for US

³⁰ Four mapping conventions are fixed once and applied everywhere in the calculations. The survey publishes no venture-capital class, so venture capital receives the survey private-equity assumption less the fixed historical venture-buyout spread estimated on our own panel, 2.0 percentage points annualized, at every percentile. Survey returns are geometric ten-year expectations and are converted to the arithmetic means the optimizer requires as $\mu_a = \mu_g + \sigma^2/2$, always using the volatility of the covariance source active in the given specification. Tangencies are computed against the survey's own cash assumption of 3.6 percent, since pairing forward-looking means with a backward-looking risk-free rate would overstate forward Sharpe ratios. And the constraint set, sample conventions, and solver are identical to the baseline.

equity — 6.5 percent against 13.9 percent — and higher for the bond class, 5.0 percent against 2.6 percent, so the survey compresses the realized equity premium substantially; the private-class means sit closer to their realized counterparts, which is why the relative attractiveness of the private sleeve can rise even as the equity premium falls.³¹

A.2 The frontier shift under survey-consensus means

Table A.2 reports the constrained tangency under the survey-median vector with each model's own unsmoothed covariance. Forward-looking Sharpe levels are much lower than historical ones — the public-only tangency Sharpe falls from 0.85 to 0.41 — because the survey compresses the equity premium relative to the exceptional realized window and more than doubles the cash rate. The gain from the private sleeve, however, survives essentially intact: 0.17 under Geltner, 0.17 under GLM, and 0.15 under Dimson, against 0.16 to 0.20 under historical means.³²

The zero-alpha variant addresses the skeptical objection, that reported private premia reflect fees, selection, and vintage luck, by assigning each private class the survey median of its closest public comparable so that the sleeve retains only its covariance structure. Even then the gain is +0.08 to +0.10. This is the paper's diversification floor: a positive gain with the private return premium assumed away entirely, on forward-looking means, at a cash rate more than double the realized sample average.

A.3 Does professional disagreement overturn the result?

The survey's dispersion across its 41 respondents is itself usable information, and we use it in two ways (Table A.3). The first is a least-favorable bound: every private class at its cross-adviser 25th percentile and every public class at its 75th — the combination of published assumptions least favorable to the diversification claim. The gain remains positive under all three models, at +0.14, +0.14, and +0.12; the symmetric most-favorable bound of +0.19 to +0.21 completes the bracket. The second is a Monte Carlo over assumption uncertainty: 10,000 expected-return vectors per mode are drawn from each class's five-point cross-adviser distribution — minimum, 25th percentile, median, 75th, and maximum — by piecewise-linear inverse-CDF sampling, under two views of how assumptions relate across classes. In the blended view, each class's assumption is drawn separately, mixing positions across advisers as a committee synthesizing several sources might. In the single-adviser view, one uniform draw governs all classes, representing an adviser who is consistently at the same point of the published range everywhere — an internally coherent optimist or pessimist. Across all 60,000 draws the gain is never negative, and the 5th percentile ranges from +0.09 to +0.12.³³ Together with the bootstrap treatment of sampling uncertainty in the covariance estimates, the two layers of input uncertainty

³¹ The survey's second moments are broadly consistent with our unsmoothed treatment, suggesting that at least some estimates likely already embed an implicit unsmoothing.

³² In proportional terms the forward-looking gain is mechanically larger not because the sleeve contributes more, but because the survey compresses the denominator.

³³ The 2025 edition publishes its cross-adviser percentiles in chart form only, so the five-point distributions used here are the 2024-edition percentile table re-centered on 2025 averages.

— means across the published range of professional views, covariances within sampling error — independently bound the gain away from zero.

Table A1. Survey-Consensus Assumptions Mapped to the Eight-Asset Universe


Percent per year

Asset class	Average return	
	Historical*	Forward-looking
US equity	13.9	6.5
World equity (ex-US)	7.0	7.5
US investment-grade bond	2.6	5.0
Private equity (ex VC)	14.0	9.0
Venture capital	11.6	7.0
Private debt	9.5	7.9
Private real estate	9.1	6.2
Private infrastructure	9.4	7.2

Notes: Returns are shown on the annualized geometric (compound) basis for consistency with the survey, whose published expected returns are geometric; the survey figures are cross-adviser medians of 10-year assumptions. *Historical figures are compound annual returns of the reported series over 2010:Q1–2025:Q3; they differ slightly from Table 1, which reports annualized log returns. The survey publishes no venture-capital class; venture capital receives the private-equity median less the historical venture-minus-buyout spread of 2.0 percentage points. Sources: Horizon Actuarial Services (2024, 2025); Investment Company Institute calculations of Prequin (a part of BlackRock), Morningstar, and Refinitiv data.

Table A2. Tangency Sharpe Ratios Under Survey-Median Expected Returns

Model	Public only	+Private (survey-median μ)	+Private (zero-alpha)	Sharpe gain	Zero-alpha gain
Geltner	0.408	0.574	0.507	+0.166	+0.099
GLM	0.408	0.574	0.492	+0.166	+0.084
Dimson	0.408	0.557	0.494	+0.149	+0.086




Notes: Survey results are based on Horizon Actuarial Services (2024, 2025).

Table A3. Dispersion Tests: Cross-Adviser Percentile Bounds and Monte Carlo Simulations


Panel A. Least and most favorable cross-adviser percentile bounds

Model	Scenario	Public only	+Private	Sharpe gain
Geltner	Least favorable (p25 pvt / p75 pub)	0.466	0.605	+0.139
Geltner	Most favorable (p75 pvt / p25 pub)	0.310	0.513	+0.203
GLM	Least favorable (p25 pvt / p75 pub)	0.466	0.607	+0.141
GLM	Most favorable (p75 pvt / p25 pub)	0.310	0.517	+0.207
Dimson	Least favorable (p25 pvt / p75 pub)	0.466	0.587	+0.122
Dimson	Most favorable (p75 pvt / p25 pub)	0.310	0.504	+0.194



Panel B. Monte Carlo over the cross-adviser distribution

Model	Mode	Draws	Median gain	5th pct.	95th pct.	Pr(gain ≤ 0)
Geltner	independent	10,000	+0.175	+0.103	+0.257	0.0%
Geltner	rank-aligned	10,000	+0.168	+0.117	+0.274	0.0%
GLM	independent	10,000	+0.181	+0.105	+0.260	0.0%
GLM	rank-aligned	10,000	+0.165	+0.116	+0.269	0.0%
Dimson	independent	10,000	+0.155	+0.087	+0.238	0.0%
Dimson	rank-aligned	10,000	+0.149	+0.109	+0.246	0.0%



Notes: Five-point cross-adviser distributions are the 2024-edition percentile table re-centered on 2025 averages, with venture capital carried at the private-equity distribution less the historical spread. Independent draws (the blended view) sample each class separately; rank-aligned draws (the single-adviser view) use one uniform for all classes, representing an internally consistent optimist or pessimist adviser. Survey results are based on Horizon Actuarial Services (2024, 2025).