Insider Alpha: Evidence from Private Foundation Portfolios *

Zhongjin Lu† Miguel I. Puertas‡

March 22, 2026

Abstract

Using a novel dataset of private foundation returns and security-level holdings, we show that size and alternative-asset tilt generate no alpha once return smoothing is accounted for. Alpha is concentrated among foundations founded by hedge fund (HF) principals. Consistent with insider advantage rather than transferable skill, HF-affiliated foundation's outperformance arises primarily from highly concentrated portfolios; across HF- and private-equity-affiliated foundations, holdings concealed from regulators predict alpha more strongly than disclosed ones, related investments predict alpha whereas unrelated investments do not, and alpha persists after founders retire. These findings challenge the view that alternative alpha is available to investors with sufficient capital and patience.

JEL Classification: G11, G23, G24
Keywords: Private Foundations; Alternative Assets; Hedge Funds; Private Equity; Portfolio Choice

*The authors have no conflicts of interest to disclose.

†Terry College of Business, University of Georgia. Email: zlu15@uga.edu.
‡Terry College of Business, University of Georgia. Email: miguel.puertas@uga.edu

1 Introduction

Once exclusive to large institutions and ultra-wealthy investors, alternative assets such as private equity, hedge funds, and venture capital are increasingly positioned as a mainstream asset class. Policymakers and practitioners advocate "democratizing" access to this $30.9 trillion industry, arguing that broader participation can improve household portfolio efficiency (Council of Economic Advisers, 2025). This view presumes that alternative alpha is a scalable asset-class characteristic—available to any investor with sufficient capital and patience.

Yet private markets are characterized by capital rationing and two-sided matching, so access to high-performing funds is selective rather than proportional to capital. Consequently, despite extensive research on private fund performance, it remains unclear whether alternative alpha—if it exists—can be systematically captured by investors with sufficient long-term capital. Answering this question requires investor-level portfolio holdings data, which are scarce due to limited investor-side transparency.

To address this gap, we examine realized outcomes of private market participation in a unique laboratory: U.S. private foundations. These grant-making entities, endowed with permanent capital, typically have broad access to alternative assets. They control more than $1 trillion in assets, vary widely in size and founder background, and disclose granular portfolio holdings through IRS Form 990-PF, a rare source of transparency in private markets. By digitizing these filings and integrating them with existing databases, we construct a comprehensive, survivor-bias-free dataset of foundation returns and full portfolio holdings. This setting allows us to directly observe investor-level allocations and realized returns and, in turn, to identify which limited partners (LPs) capture alternative alpha and through what mechanism. [1]

We examine which LPs capture alternative alpha by evaluating three competing explanations: (i) scale-related advantages, (ii) broad portfolio tilts toward private funds, and (iii) non-replicable advantages tied to the LP's identity. The first two explanations, if confirmed,

[1] Earlier work such as Lerner et al. (2007) and Sensoy et al. (2014) documents performance differences across LP types within private equity funds. We extend this line of inquiry by examining performance differences across LP's professional background using their complete alternative asset holdings.

have straightforward policy implications: pooling capital or broadening access to alternative asset classes could close the performance gap. The third implies that neither pooling nor tilting would be effective, as the advantage is inherently non-replicable.

Our evidence consistently supports the third explanation. We first show that previously documented premia associated with size and private-market exposure largely reflect return smoothing. Consistent with Binfare and Zimmerschied (2026), larger foundations appear to outperform smaller ones by approximately 80 basis points per year in unadjusted CAPM alpha and allocate a greater share of their portfolios to alternative assets. However, once we account for smoothing by augmenting standard factor models with lagged market returns (Dimson, 1979), the alpha associated with both foundation size and private-market tilts disappears.

In contrast, we find substantial heterogeneity across foundations when segmenting by founders' professional backgrounds. Risk-adjusted outperformance is concentrated among foundations funded by hedge fund (HF) principals. These HF-affiliated foundations earn annual alphas approximately 2.6% to 3.2% higher than non-insider foundations over the full 1995–2023 sample. All other foundations have alphas statistically indistinguishable from zero, including those established by private equity and venture capital (PEVC) general partners and those whose founders come from other financial backgrounds. We refer to HF- and PEVC-affiliated foundations as *fund-insider* foundations.[2]

These patterns raise a central question: does the outperformance of HF-affiliated foundations reflect a replicable investment technology or a non-replicable advantage tied to LP–GP relationships? Under a *replicable-technology* view, HF-affiliated foundations are sophisticated investors, earning alpha through superior manager selection across diversified private fund portfolios. Under a *non-replicable* view, their edge arises from their role as GPs of the private funds in which their foundations invest, conferring privileged access or informational advantages unavailable to outside investors.

We test four implications of the non-replicable view. First, alpha should be concentrated among HF-affiliated foundations with highly concentrated alternative-asset portfolios rather

[2]A "fund insider" simultaneously serves as a general partner of a private investment fund and as the capital provider of a foundation that invests in that fund.

than those broadly diversified across managers. Second, alternative-asset holdings misclassified as public equity should be more strongly associated with alpha than correctly disclosed holdings, consistent with strategic concealment of relationship-specific advantages. Third, alpha should be present in affiliated fund holdings and absent in unaffiliated holdings, even within the same foundations. Fourth, alpha should persist after the founder retires from day-to-day management, consistent with a structural advantage rather than a replicable investment process.

The evidence strongly supports the non-replicable view. Alpha is concentrated among HF-affiliated foundations with highly concentrated alternative portfolios rather than diversified ones. Among fund-insider foundations, the concealed alternative-asset weight—the gap between our LLM-based measure and the SOI classification—strongly predicts alpha. Decomposing holdings into affiliated and unaffiliated investments, we find that excess returns among fund-insider foundations load entirely on affiliated positions and are absent in unaffiliated ones. In contrast, neither the concealed alternative-asset weight nor affiliated investment exposure predicts alpha among non-insider foundations. Finally, we find little evidence of performance decline following the retirement of foundation founders from their investment firms. A replicable-technology explanation cannot jointly account for these patterns.

Taken together, our findings suggest that alternative alpha is unlikely to be a replicable or scalable asset-class characteristic. The two commonly cited channels—size and broad private-market exposure—generate no alpha once return smoothing is properly accounted for. In contrast, the one channel that consistently generates alpha—LP-GP identity overlap—is by construction non-replicable. Expanding retail access to private markets will not mechanically improve portfolio efficiency if the alpha resides in such relationships rather than in capital itself. As the architect of the Yale endowment model that pioneered alternative investing, David Swensen cautioned, "unless an investor has access to incredibly high-quality professionals, they should be 100 percent passive... that includes all individual investors and most institutional investors."

1.1 Related Literature

Our main contribution is to provide systematic empirical evidence against the view that alternative alpha is broadly accessible to investors with sufficient capital and patience. Instead, our results suggest that alternative alpha is largely insider alpha. Our paper relates to three strands of the private markets literature.

First, our paper contributes to the literature on LP outcomes. A large body of work documents heterogeneity and persistence in fund performance at the GP level (e.g., Kaplan and Sensoy (2015); Agarwal et al. (2015)). However, because private markets are characterized by capital rationing and two-sided matching, realized LP returns depend not only on fund performance but also on which funds investors can access (Sørensen, 2007). An emerging LP-centric literature therefore studies investor outcomes using administrative data on households, pension funds, endowments, and private foundations (see, e.g., Bach et al. (2020); Dahiya and Yermack (2018); Lo et al. (2019); Binfare and Zimmerschied (2026); Balloch et al. (2025); Gabaix et al. (2025); Begenau et al. (2025)). These studies generally find that larger investors allocate more to alternatives and earn higher unadjusted returns, although whether scale and private-asset tilts translate into superior risk-adjusted performance remains unsettled. Both Bach et al. (2020) and Binfare and Zimmerschied (2026) study private foundation performance and reach different conclusions.

Progress in this literature has been constrained by limited and selective investor-side data: both high- and low-performing private funds have incentives to withhold information, and LPs likewise have limited incentives to disclose their allocations. We address this data limitation by digitizing Form 990-PF filings to recover three key dimensions simultaneously: investor identity, security-level holdings, and investor–investment relationship mappings. These dimensions are not available in existing Form 990-PF databases, and to our knowledge no existing private-markets dataset contains all three. Using these data, we show that the previously documented premia associated with size and alternative-asset tilts disappear once return smoothing is accounted for, and that alternative alpha is instead closely tied to LP identity. Our dataset is also substantially more comprehensive than those used in prior studies of private foundations (see Section 2).

Second, our paper contributes to the literature on cross-LP performance differences. Lerner et al. (2007) show that endowments outperform other LP classes in private equity, while Sensoy et al. (2014) find that this gap narrowed substantially as the industry matured. We extend this literature by showing, in the private foundation context, that the performance gap is not explained by institutional scale, broad alternative exposure, or general financial sophistication, but instead reflects fund-insider advantage. Our holding-level relationship data allow us to identify this mechanism through four tests that distinguish replicable investment technology from non-replicable insider advantage.

Finally, our findings contribute to the literature on information and disclosure in private markets. Abuzov et al. (2025) show that top-performing venture capital funds responded to FOIA rulings by excluding public pension LPs rather than allowing their LP relationships to become publicly disclosed. We document a complementary pattern on the LP side: fund-insider investors strategically misclassify their highest-performing affiliated holdings in public regulatory filings. Together, these results suggest that informational asymmetries in private markets operate on both sides of the LP–GP relationship and help preserve the rents associated with insider positions.

2 Institutional Background

Private foundations (PFs) are tax-exempt charitable organizations under Section 501(c)(3) of the Internal Revenue Code that do not qualify as public charities, because they are primarily funded by a single individual, family, or corporation. Their favorable tax treatment includes income-tax deductions for donors and exemption from most income taxes, subject only to a 1.39% excise tax on net investment income.

Due to the preferential tax status and the self-dealing risk arising from concentrated control, private foundations face significantly stricter reporting requirements than other institutional investors. They must file IRS Form 990-PF annually, disclosing detailed balance-sheet and income-statement data, officer and trustee compensation, grant recipients, and granular security-level risky investment holdings. These disclosures are more detailed than those required of most institutional investors—including public pension funds and university en-

dowments—and provide a rare window into private-market participation that is unavailable through any other regulatory source.

Private foundations are well positioned to access private investment vehicles. Most in our main sample exceed the asset thresholds for accredited investor status under U.S. securities law, and many also meet the higher qualified purchaser threshold—at least $5 million in investments—required for funds relying on the Section 3(c)(7) exemption, including private equity and hedge funds. Foundations with assets above $100 million, the subsample for which we collect founder background information, are typically large enough to satisfy the minimum commitments of institutional-quality alternative asset managers. Together, these features make private foundations an unusually transparent and economically meaningful setting for studying investor heterogeneity in private markets.

3 Data Construction

Studying LP-level outcomes in private markets is empirically challenging due to incomplete coverage of investor returns and the inability to reconstruct complete LP-level private portfolios. Commercial private-market databases rely on voluntary disclosure by GPs and LPs, raising concerns about systematic selection biases (Phalippou, 2026). Standard PE/VC databases such as Preqin omit important vehicles such as co-investments, which account for roughly 40% of private equity capital commitments by 2017 (Lerner et al., 2022). More fundamentally, these datasets cannot be combined to recover the complete private holdings of any LP, as investor identities and allocation amounts are incomplete and selectively reported. Existing studies partially address these limitations by focusing on public pension funds subject to FOIA disclosures, though results may not generalize to other LPs given their distinct political mandates and disclosure constraints (Andonov et al., 2017; Abuzov et al., 2025).

U.S. private foundations provide an alternative, transparent, and economically meaningful laboratory for studying LP-level outcomes. These grant-making entities, endowed with permanent capital, have long had broad access to alternative assets. They control more than $1 trillion in assets, vary widely in size and founder background, and are required to

disclose detailed portfolio information through IRS Form 990-PF. We digitize the relevant sections of these filings and integrate them with IRS administrative data to construct a comprehensive, survivor-bias-free panel of foundation returns and security-level private holdings. Unlike existing datasets, our data allow us to observe investor identity, security-level investments, and investor–fund linkages that distinguish affiliated from unaffiliated exposure—dimensions essential for testing whether outperformance reflects scalable allocation choices or relationship-driven access.

3.1 Addressing Missing Observations

Existing administrative datasets on foundations—the IRS Statistics of Income (SOI) microdata and the NCCS Core files—are valuable but incomplete. SOI microdata are designed to provide a cross-sectional representation of the sector: all foundations with fair market assets above $10 million are included, while smaller foundations are sampled at stratified rates, limiting the ability to construct continuous return panels. The NCCS Core files cover the full population but are available only for selected years and contain fewer variables than SOI.

To construct a comprehensive return panel, we augment SOI with four additional sources: NCCS Core files, IRS Annual Extract files, XML Form 990-PF filings (post-2013), and historical PDF filings (post-2001).[3] A foundation enters our sample when it first appears in SOI with assets above $1 million and positive operating expenses. We then construct a continuous return series through its termination year. To ensure that each foundation in the estimation sample has at least five annual observations for alpha estimation, we stop admitting new foundations after 2019 and extend all return series through 2023.

Our merged dataset allows us to quantify the extent of missing observations in SOI, the most commonly used source for private-foundation research. Nearly 30% of foundation-year returns would be missing if we relied solely on SOI, with missingness particularly severe among smaller foundations near the sampling threshold. Table 1 reports, for foundations

[3]Electronic filings of Form 990-PF became publicly available in 2001 and were initially mostly PDFs; since 2013 filings increasingly appear in XML, and XML filing became mandatory for tax years ending July 31, 2020 and later under the Taxpayer First Act.

sorted by size in a given year, the probability of at least one missing annual return over the subsequent 15-year window. In SOI, this probability ranges from 57% to 77% for foundations smaller than $10 million and from 3% to 7% even for the largest 100. In our merged dataset, these probabilities fall to effectively zero for the largest 100 and below 1% (2.5%) for foundations larger than $100 million (smaller than $10 million) in sorting years after 2000, when PDF filings become available.

This improvement is economically meaningful. Even if returns are missing at random, shorter histories reduce the precision of alpha estimates. If missingness correlates with foundation size or reporting behavior, completeness is essential to avoid biased inference about performance differences across foundations.

Insert Table 1 here.

3.2 LLM-Based Classification of Alternative Asset Holdings

The second key limitation of existing datasets is the absence of security-level information on alternative asset holdings. To support analysis of alternative-asset exposure—and, critically, to construct investor–fund relationship maps that distinguish related from unrelated holdings—we develop a structured, multi-stage classification pipeline powered by a large language model (LLM). Without this pipeline, it would be impossible to test the insider-access mechanisms central to our analysis, since a foundation concentrating in affiliated funds is observationally equivalent to a fund-of-funds style foundation when only aggregate asset-class exposure is available.

The key design principle is to avoid asking the LLM to directly classify a line item as an alternative asset or not. Although simple, that approach produces a classification regardless of whether the holding description contains sufficient information to support one—reliable and unreliable outputs are then indistinguishable. Instead, we design the pipeline to replicate how a human analyst would approach the problem: extracting verifiable attributes first, applying a classification rule second, and resolving residual ambiguity through a rule-based fallback third. At each stage, uncertainty is expressed explicitly rather than suppressed.

In the first stage, the LLM extracts five attributes for each investment line item: (i) *In-*

termediation—whether investment decisions are delegated to an intermediary; (ii) *Access*—whether participation is subject to discretionary approval; (iii) *Specificity*—whether the description is precise enough to identify a specific investment; (iv) *Investment Type* (e.g., hedge fund, PE, VC, bond, equity); and (v) *Broad Asset Type* (e.g., private fund, equity, fixed income, real asset, cash). For all five attributes, the model may return "TBD" to express uncertainty rather than commit to an unsupported classification.

We define an *alternative asset* using attributes (i) and (ii): a holding is classified as alternative if it is both intermediated and access-restricted. When the LLM returns "TBD" for either attribute, we use attributes (iii)–(v) to resolve the ambiguity according to a prespecified, deterministic mapping (Table A1). For robustness, we construct both a more restrictive and a less restrictive definition; our results are consistent across all three specifications. When the extracted attributes indicate that a holding description is too vague—for example, a line item labeled simply "investments"—we apply a schedule-based default rule that requires no model judgment: holdings reported on the corporate stock schedule are classified as non-alternative, whereas holdings reported on the other-investments schedule are classified as alternative. This rule is conservative as it likely understates alternative-asset exposure for foundations that misclassify private funds under public-asset schedules. For holdings deemed sufficiently specific by attribute (iii), we further conduct Gemini-grounded web searches to identify relationships between each founder and the underlying fund investment, allowing us to classify holdings as related, unrelated, or of unknown relation.

The pipeline offers two advantages over SOI aggregates. First, security-level granularity enables us to map specific founders to specific fund investments and test insider-access mechanisms that are unidentifiable from coarser data. Second, it uncovers systematic misclassification that is invisible in SOI. A representative example is illustrated in Appendix Figure A3: the Overdeck Family Foundation reports both an affiliated hedge fund and a Vanguard index fund under the corporate stock schedule; SOI treats both as corporate equity.

3.3 Foundation Founder Background

A central element of our analysis is classifying each foundation by the professional background of the individual most responsible for providing its capital. Because this classification requires intensive manual research, we restrict it to foundations large enough that regulatory eligibility and minimum-commitment requirements are unlikely to constrain access to alternative assets. A foundation enters this subsample in the first year its fair market assets exceed $100 million; we then track it continuously until termination. To ensure that every foundation in the subsample has sufficient return history for alpha estimation, we stop admitting new entrants after 2017.[4] We refer to this sample as the $100M sample, which contains 1,655 unique foundations.

We manually review each foundation to identify the individual primarily responsible for financing it. In most cases this is the founding donor; when capital originates from someone other than the nominal founder, we attribute the foundation to the true capital provider, whom we refer to as the *founder*. We classify founders by wealth source: Hedge Fund (HF), Private Equity or Venture Capital (PEVC), other financial industries (FIN; e.g., mutual funds, banking, real estate), and Non-Financial industries (NONFIN). Foundations without an identifiable individual founder—including corporate, community, and hospital-conversion foundations—are labeled Corp, Community, and Hosp, respectively. For example, the Overdeck Family Foundation is classified as HF (John Overdeck, Two Sigma), and Beneficus Foundation as PEVC (John Doerr, Kleiner Perkins).[5] As defined earlier in the paper, we refer to HF- and PEVC-affiliated foundations as *fund-insider foundations*. FIN and NONFIN foundations are collectively termed *non-insider foundations*. Because extracting security-level holdings from PDF filings is resource-intensive, the LLM pipeline is currently applied only to this $100M subsample; accordingly, all analyses involving foundation type and LLM-based alternative-asset weights are conducted within it.

[4]The subsample is constructed using only information available at each point in time, avoiding look-ahead bias. Once a foundation crosses the $100 million threshold, it remains in the sample even if assets subsequently fall below that level.

[5]To ensure consistency and flag potential errors, all manual classifications are reviewed by an LLM verifier; disagreements are subject to a second round of manual review.

3.4 Return Construction

Annual investment return (R_{it}) is computed as the change in net assets net of non-investment flows, scaled by the mid-year asset base, following standard practice in the literature (e.g., (Dahiya and Yermack, 2018; Binfare and Zimmerschied, 2026)).

$$R_{it} = \frac{\Delta A_{it} - \Delta L_{it} - F_{it}}{A_{i,t-1} + \frac{1}{2}F_{it}}$$

where A_{it} is the end-of-year fair market value of total assets for foundation i in year t, L_{it} the end-of-year book value of liabilities, and $F_{it} \equiv C_{it} - D_{it}$ denotes net non-investment cash flows—contributions received (C_{it}) net of grants and expenses paid (D_{it}).

Returns are gross of excise taxes. To ensure comparability across fiscal years, each return is measured over exactly twelve calendar months, with the period end date taken from Form 990-PF field E010. We manually review all extreme return observations in the $100 million subsample and retain those attributable to verified economic events—such as large investment gains or the realization of alternative assets upon a liquidity event—trimming only observations associated with measurement error from reporting mistakes or asset transfers not fully captured by existing data sources. This approach preserves genuine extreme returns that standard trimming rules would otherwise discard. Full details are provided in Appendix B.

4 Analysis of Performance Difference

4.1 Summary Statistics

Table 2 reports the descriptive statistics for both the full sample and the $100M sample. Our primary measure of risk-adjusted performance, hereafter referred to as CAPM alpha, is the intercept from a regression of excess returns on the contemporaneous market factor and its one-year lag. This Dimson (1979) correction accounts for return smoothing induced by illiquid holdings. We define the *unadjusted alpha* as the intercept from a regression that excludes the lagged market factor.

The full sample comprises 20,118 unique foundations with 18,498 of them having at

least five annual observations for computing performance statistics. These foundations earn an average excess return of 6.09% per year and an average alpha of 0.39%. The alpha is virtually indistinguishable from the unadjusted alpha for the typical foundation, indicating that illiquidity-related return smoothing is not economically important in the broad cross-section. The median alpha is approximately zero, suggesting that the typical foundation does not earn returns beyond what its market exposure would predict.

Performance dispersion is substantial. The 5th-to-95th percentile range of alpha spans 11.07 percentage points. Risk-taking also varies widely: return volatility ranges from 6.08% at the 5th percentile to 20.99% at the 95th percentile, and market betas range from 0.06 to 0.88. Portfolio composition is equally heterogeneous. More than one quarter of foundations allocate essentially nothing to alternative assets, and the median alternative-asset weight is only 4.09%, yet foundations at the 95th percentile allocate nearly 70% of assets to alternatives.

The $100M subsample exhibits similar heterogeneity in performance, risk, and alternative-asset allocation, with slightly higher average excess returns and alphas and substantially greater average alternative-asset weights. Notably, the gap between alpha and unadjusted alpha is 67 basis points per year in this subsample, which is substantially larger than in the full sample. This pattern provides an early indication that the Dimson correction is economically important for larger foundations, a point we examine directly in Section 4.2.1.

The total assets controlled by the private foundations in our sample have increased ten-fold, from $100 billion in 1995 to more than $1 trillion in 2023 (see Appendix Figure A1). Over the same period, the average foundation gradually increased its allocation to alternative assets from 6% to nearly 20% (see Appendix Figure A2).

This substantial rise in alternative asset allocations appears at odds with the (virtually) zero CAPM alpha earned by the median foundation. However, it is possible that foundations with more substantial allocations to alternative assets earn positive CAPM alphas. We now turn to a more detailed analysis of alpha.

Insert Table 2 here.

4.2 Test Competing Hypotheses

The substantial cross-sectional dispersion in alpha documented above raises a central question: what determines whether a private foundation earns positive alpha? We evaluate three competing explanations. The first is *scale*: larger foundations may negotiate better fee terms, deploy more sophisticated due diligence, or achieve economies of scale in manager monitoring. The second is *alternative-asset tilt*: allocating a greater share of the portfolio to alternative assets may itself capture an asset-class premium, regardless of who is making the allocation decision. The third is *fund-insider status*: alpha may reflect a non-replicable advantage rooted in the investor's identity, rather than in the scale of capital or the composition of the portfolio.

These three explanations carry distinct implications for policy and for the broader literature on investor heterogeneity. If scale drives superior outcomes, access could in principle be democratized by replicating institutional size through capital pooling. If portfolio construction is the mechanism, broadening access to alternative asset classes becomes the relevant policy lever. But if alpha is tied to the investor's identity, neither pooling capital nor tilting toward alternatives will close the gap.

4.2.1 Scale and Alternative-Asset Tilt

We begin by examining the scale hypotheses. To facilitate comparison with prior studies, we first construct equal-weighted portfolios of foundations sorted into four size groups: below \$50 million, \$50–\$100 million, above \$100 million (excluding the top 100), and the top 100 largest foundations. The sort is restricted to foundations with December fiscal year ends (approximately 75% of the sample), and size is measured by CPI-adjusted fair market assets at the end of the prior year.

Table A3 reveals a size gradient in both excess returns and unadjusted alphas. Average excess returns rise monotonically from 6.19% for the smallest foundations to 7.06% for the top 100, and unadjusted alphas follow the same pattern, increasing from 0.91% to 1.35%. Alternative-asset allocation also rises monotonically with size, from 12.80% for the smallest group to 33.04% for the top 100. Using a more comprehensive dataset with substantially

fewer missing observations, our results are consistent with those in Binfare and Zimmerschied (2026). Based on these patterns, Binfare and Zimmerschied (2026) conclude that both scale and an alternative-asset tilt are associated with superior risk-adjusted returns in the private foundation sample.

However, the apparent size gradient in unadjusted alpha does not survive the Dimson correction. After adjusting for return smoothing, alphas are nearly flat across the three smaller size groups—1.09%, 1.09%, and 1.03%—and actually decline to 0.89% for the top 100 largest foundations. The higher unadjusted alpha of the largest foundations therefore reflects not genuine risk-adjusted outperformance, but rather greater exposure to the lagged market return, consistent with the more pronounced return smoothing induced by their larger alternative-asset allocations. The lagged-factor loading for the top 100 reaches 0.05, approximately 10% of their contemporaneous market beta, compared with nearly zero lagged-factor loadings for the other three size groups.

As a robustness check, we replicate the specification of Binfare and Zimmerschied (2026) using their four-factor model and size-bucket definitions in Appendix Table A4. The same pattern emerges: the four-factor alpha for their largest size group declines from 0.73% to −0.03% once the lagged market factor is included. Thus, whether we use a one-factor market model or a four-factor model that additionally includes international equity, U.S. bond, and hedge fund factors, large foundations do not earn higher alphas than their smaller counterparts once lagged market exposure is properly accounted for.

We find a parallel null result for alternative-asset tilt when sorting foundations on lagged alternative-asset weights. We find no significant evidence that a higher allocation to alternative assets, by itself, produces risk-adjusted outperformance.

Because univariate portfolio sorts can evaluate only one dimension at a time, we next estimate a multivariate panel regression:

$$R_{i,t} - R_{f,t} = \left(\alpha_0 + \boldsymbol{\gamma}'\mathbf{X}_{i,t-1}\right) + \beta_{i,0}, \mathrm{MKT}_t + \beta_{i,1}, \mathrm{MKT}_{t-1} + \varepsilon_{i,t}, \tag{1}$$

where $\mathbf{X}_{i,t-1}$ is a vector of lagged foundation characteristics, and contemporaneous plus

one-year-lagged market factor controls for the market exposure. This panel regression approach also allows us to retain all observations without restricting the sample to December fiscal year ends.

We use two size proxies based on lagged assets: *Top100*, an indicator for the 100 largest foundations, and *Size(rank)*, a normalized cross-sectional size rank scaled from 0 (smallest) to 1 (largest). We also include the lagged alternative-asset weight. For each specification, we report results both with and without year fixed effects, and both in the full sample and the $100M sample.

Table 3 shows that the regression coefficients on both size measures and alternative-asset weights are consistently statistically indistinguishable from zero across all specifications, with the exception of Column (4), where the coefficient on private weight is positive and significant at the 10% level. As we show in subsequent analysis, however, this significance is completely driven by HF- and PEVC-affiliated foundations. Once these insider-foundations are excluded, alternative-asset weights are no longer correlated with alphas. Overall, the panel regression results corroborate the portfolio-sort evidence and indicate that neither the scale hypothesis nor the alternative-asset-tilt hypothesis receives empirical support in the data.

Insert Table 3 here.

4.2.2 Founder Background

We next turn to our third hypothesis, which posits a non-replicable advantage tied to investor identity. Specifically, we examine whether alpha is concentrated among fund-insider foundations whose founders are themselves private fund providers.

Table 4 reports portfolio characteristics by founder type for the $100M subsample. The sample comprises 1,655 unique foundations, of which 1,463 have at least five annual observations for alpha estimation: 86 HF-affiliated foundations, 42 PEVC-affiliated foundations, 285 FIN-affiliated foundations, and 1,050 NONFIN-affiliated foundations.

Performance differences across groups are stark. HF- and PEVC-affiliated foundations earn average alphas of 1.52% and 1.12% per year, respectively, compared with −1.13% for

FIN foundations and 0.40% for NONFIN foundations. The Dimson correction is largest for HF-affiliated foundations, at 2.17 percentage points per year, consistent with their substantially higher alternative-asset allocations (66.25% on average). Importantly, these performance differences do not reflect compensation for higher systematic risk: market betas cluster tightly between 0.45 and 0.58 across all four groups, and return volatility ranges between 13% and 16%.

For statistical inference, we estimate the panel regression in Equation (1) augmented with founder-type indicators in Table 5. Columns (1) and (2) show that HF affiliation is the strongest predictor of alpha: the HF indicator is associated with annual alpha that is 2.55%–3.17% higher than the baseline NONFIN group, depending on whether year fixed effects are included. The PEVC coefficient is positive but smaller and statistically insignificant. The FIN coefficient is neither consistently positive nor statistically significant, indicating that general financial sophistication does not translate into an alpha advantage in private markets.

Figure 2 illustrates both the economic magnitude and persistence of the fund-insider alpha premium. Panel A plots the cumulative sum of cross-sectional average market-adjusted returns by founder type, showing that the outperformance of HF-affiliated foundations is both large and persistent over time. Panel B places this premium in the context of the size gradient, demonstrating that the fund-insider alpha premium substantially exceeds any size effect in the data.

Columns (3) and (4) further add controls for alternative-asset weights. Controlling for the high alternative-asset allocations of HF-affiliated foundations explains little of the positive coefficient on the HF indicator. Columns (5) and (6) additionally include interactions between founder type and alternative-asset allocation, allowing the effect of alternative-asset tilts on CAPM alpha to vary across groups. We find that a higher alternative-asset weight is a strong and statistically significant predictor of alpha for fund-insider foundations. The interaction coefficients are 4.76 for HF-affiliated and 3.97 for PEVC-affiliated foundations without year fixed effects, and 3.90 and 2.62, respectively, with year fixed effects. These coefficients are economically large. For example, a coefficient of 4.76 implies that a 10-percentage-point increase in alternative-asset allocation is associated with an additional 47.6

basis points of annual alpha.

In contrast, the interaction coefficient for FIN foundations is negative and statistically insignificant, as is the coefficient on alternative-asset weights for the baseline NONFIN group. These results indicate that alternative-asset exposure per se does not generate alpha; rather, it is alternative-asset exposure deployed by fund insiders that is associated with positive CAPM alpha.

Columns (7) and (8) confirm this pattern by regressing excess returns on lagged alternative-asset weights after excluding HF- and PEVC-affiliated foundations. The coefficient on the alternative-asset weight is -0.99 without year fixed effects and 0.28 with year fixed effects, and is statistically insignificant in both cases.

Comparing specifications with and without year fixed effects reveals an additional asymmetry between fund-insider and non-insider foundations. Adding year fixed effects increases the alternative-asset-weight coefficient for non-insider foundations but decreases it for HF- and PEVC-affiliated foundations. This divergence suggests that non-insider foundations tend to increase their alternative-asset allocations in periods when subsequent alpha is low, whereas fund-insider foundations do the opposite. The pattern is consistent with fund insiders possessing a form of market-timing ability that their non-insider peers lack.

Overall, the evidence in this section demonstrates that neither scale nor alternative-asset tilts are associated with foundations' CAPM alphas. Instead, alpha is concentrated among fund-insider foundations, most notably those established by hedge fund principals. These findings suggest that CAPM alphas in alternative assets are not a scalable asset-class characteristic, but rather accrue to fund insiders who can identify and access them. We next examine this hypothesis in greater detail.

5 The Mechanism of Insider Alpha

The preceding evidence raises a central question for both theory and policy: does alpha among fund-insider foundations reflect generalizable fund-selection skill that sophisticated investors could replicate, or a non-replicable advantage tied to the founder's identity as a private fund provider? Under a *selection-skill* view, fund-insider foundations should resemble

sophisticated fund-of-funds investors—broadly diversified across high-quality private funds and earning alpha through superior manager screening. Under an *insider-advantage* view, outperformance stems from informational asymmetries inherent in the founder's role as a fund provider: direct knowledge of which vehicles are highest performing, privileged access to fund capacity, and fee terms unavailable to outside LPs. This view predicts (i) concentrated portfolios focused on affiliated flagship funds; (ii) strategic opacity in regulatory filings to obscure affiliated positions; (iii) no alpha where insider linkages are absent, even for the same sophisticated founders; and (iv) persistence of alpha after the founder steps back from active investment decisions. We examine each prediction in turn.

5.1 Concentration

We begin with prediction (i). Table 6 augments the baseline specification with a *SmallN* indicator equal to one if a foundation holds three or fewer alternative-asset investments. Columns (1) and (2) show that concentration is associated with higher alpha, although the effect is statistically significant only in the specification without year fixed effects.

Columns (3) and (4) introduce interactions between *SmallN* and the HF and PEVC indicators. The HF interaction is the only term that is both economically large and statistically significant: 3.35 without year fixed effects and 3.36 with year fixed effects, both significant at the 1% level. In the preferred specification with year fixed effects, the HF coefficient is 1.21 and the *SmallN* coefficient is 0.10. A HF-affiliated foundation with a concentrated portfolio therefore earns an annual alpha of 4.57% (1.21 + 3.36), whereas one with a more diversified portfolio earns only 1.21%. A non-HF-affiliated foundation with a similarly concentrated portfolio earns an economically small and statistically insignificant 0.10%. These results are robust to additional controls interacting founder type with alternative-asset weights (Columns (5) and (6)).

These findings are difficult to reconcile with a selection-skill explanation. A skilled fund selector seeking diversification benefits is unlikely to concentrate the portfolio in three or fewer positions when multiple outperforming managers can be identified. Concentration is rational only if the information advantage is specific to a single vehicle. The evidence thus provides an initial indication that a small number of high-conviction affiliated flagship

funds are the primary source of alpha for HF-affiliated foundations—an advantage that is inherently neither scalable nor replicable.

5.2 Strategic Opacity

Prediction (ii) tests whether the alternative-asset exposure that HF- and PEVC-affiliated foundations conceal through reporting discretion is more strongly associated with alpha than the exposure they transparently disclose.

5.2.1 Disclosure Incentives

Form 990-PF, Part II, requires foundations to report risky holdings under three schedule categories: corporate stock, other investments, and other assets. The IRS defines these categories at a high level without mandating a uniform classification rule, so a limited partnership interest in a hedge fund may appear under either corporate stock or other investments depending on the preparer's judgment. The widely used SOI microdata largely preserves foundations' self-reported classifications, subject to limited adjustments.[6] Furthermore, while foundations are required to report a description, book value, and fair market value for each holding, the instructions specify no minimum granularity for the description field, leaving foundations free to report a position as, e.g., "limited partnership interest" without naming the issuer.

The incentive to exploit this discretion is not absolute. Because Form 990-PF is a public document, foundations face reputational pressure to file with at least some degree of transparency: systematically vague disclosures may invite regulatory scrutiny and undermine the charitable stewardship narrative that supports tax-exempt status. For prominent HF and PEVC founders, disclosing the name of a well-known flagship fund already closed to outside investors carries little competitive cost and may even yield branding or reputational benefits. These countervailing pressures imply that opacity is partial rather than complete, concentrated in positions where transparency carries the highest competitive cost. Foundations can adopt different opacity strategies. For HF-affiliated founders, flagship funds are often already

[6]Based on our communications with IRS staff, SOI adjustments are narrow in scope and rely on a short keyword list to reclassify certain items reported under "other investments." These adjustments do not extend to assets classified as corporate stock. Consequently, private fund interests reported under the corporate-stock category are generally not reclassified in the SOI data.

closed to outside investors, making name disclosure relatively harmless. Strategic opacity thus operates primarily at the schedule level: classifying a private fund under the corporate-stock schedule—rather than under "other investments"—prevents systematic screening by data vendors and automated parsers that rely on schedule codes to detect private-fund exposure. As a result, a database that reports aggregate holdings by reported asset class will treat the position as public equity, even if the private fund name appears in the description field. For PEVC-affiliated founders, the incentives are stronger and more granular. Affiliated funds may still be raising capital or managing portfolios with co-investment opportunities. Publicly disclosing the fund name can signal available capacity to competing LPs, strain relationships with excluded investors, compromise early-stage valuations, or intensify deal-level competition. PEVC-affiliated foundations therefore have stronger incentives to obscure fund identity itself, rather than misclassify the asset category.

Appendix Figure A3 illustrates the HF case. The Overdeck Family Foundation reports a multi-million-dollar allocation to an affiliated Two Sigma vehicle alongside a Vanguard index fund under the corporate-stock schedule, so that both positions appear as corporate stock in the SOI data despite their fundamentally different economic exposures.

5.2.2 Summary Statistics

Comparing our LLM-based alternative-asset weights with SOI-based weights reveals systematic differences across founder types that are consistent with this incentive framework. Table 7 shows that SOI-based weights understate true alternative-asset exposure for HF-affiliated foundations by 7.44 percentage points on average, while overstating it for NONFIN (1.79 pp), FIN (3.46 pp), and PEVC (3.28 pp) foundations.[7] Within-group dispersion is substantial, with standard deviations ranging from 16.25 pp (FIN) to 21.17 pp (HF). Although the overall correlation between LLM-based and SOI-based weights is 0.70, offsetting biases cause cross-sectional averages to align closely (Appendix Figure A4), masking large EIN-level discrepancies.

The two opacity strategies also leave distinct footprints in the composition of each group's

[7]For non-insider foundations, overstatement typically arises when mutual funds and brokerage accounts are classified as "other investments" alongside genuine private funds—an error consistent with inadvertent rather than strategic misreporting.

alternative-asset holdings. Table 7 shows that PEVC-affiliated foundations allocate 48.86% of their alternative assets to holdings classified as *unknown*—positions whose descriptions are too vague to identify either the fund or the investor–fund relationship—compared with 31.86% for HF-affiliated foundations. As we show below, this elevated share of unknown holdings among PEVC foundations is itself a stable predictor of alpha, consistent with these undisclosed positions being intentional obfuscation.

5.2.3 Strategic versus Nonstrategic Misclassification

If misclassification were random or inadvertent, it should not predict future returns. We test this using *L1PrivDiffW*, defined as the lagged difference between a foundation's LLM-based alternative-asset weight and its SOI-reported weight.

Columns (1)–(4) of Table 8 compare the predictive content of LLM-based and SOI-based weights over 2002–2023, when both measures are available. LLM-based weights are strongly positively associated with CAPM alpha for fund insiders; SOI-based weights are not. Without year fixed effects, the coefficient on the HF (PEVC) interaction with LLM-based weights is 4.85 (4.05), compared with 0.04 (2.57) for the corresponding SOI interaction. Among HF-affiliated foundations, the association between alternative-asset exposure and alpha is more than an order of magnitude stronger using our measure. With year fixed effects, the coefficient on the HF (PEVC) interaction with LLM-based weights is no longer statistically significant but remains economically large, whereas the coefficient on the HF interaction with SOI-based weights turns negative.

Columns (5) and (6) directly test the strategic-opacity hypothesis by including both the SOI-based weight and *L1PrivDiffW*. The coefficient on *L1PrivDiffW* for the average foundation is small and statistically insignificant. In contrast, the coefficients on the interaction between *L1PrivDiffW* and fund-insider indicators are large: 5.34 (3.79) for PEVC and 7.52 (6.67) for HF, without (with) year fixed effects. Economically, a concealed 10-percentage-point allocation to a private fund is associated with 37.9 (66.7) additional basis points of annual alpha for PEVC- (HF-) affiliated foundations in the year-fixed-effects specification. Despite the large magnitude, these coefficients are statistically significant for the HF foundations and insignificant for the PEVC foundations.

Taken together, the evidence indicates that for fund-insider foundations, the portion of alternative-asset exposure that is concealed is precisely the portion most strongly associated with subsequent alpha—consistent with strategic opacity concentrated in the highest-performing positions, where the competitive cost of transparency is greatest.

5.3 Related versus Unrelated Holdings

Prediction (iii) provides another sharp test. If fund-insider foundations outperform because their principals are skilled evaluators of alternative assets, that skill should apply to all alternative-asset positions. By contrast, under the insider-advantage explanation, alpha should be concentrated in holdings offered by the investment firm connected to the founder (e.g., as a GP or senior executive).

We test this by using Gemini to conduct structured Google searches on the relationship between each foundation's founder and each alternative-asset holding exceeding 0.1% of total assets. We classify each holding as **related** (the founder or spouse is a director or senior executive of the fund's parent firm, or a documented secondary business or family connection exists), **unrelated** (the fund is identifiable but no documented connection to the founder exists), or **unknown** (the description is insufficient to identify the relationship, including all vague descriptions). We separately flag personal investment vehicles and family offices and assign them to category **FAM**.

We first examine allocation differences across these categories. Table 7 shows that HF-affiliated foundations allocate 42.83% of alternative assets to related holdings on average; the corresponding shares are 15.83% for PEVC, 2.36% for FIN, and 1.01% for NONFIN.

Table 9 shows that this stark allocation difference is strongly associated with differences in returns. Columns (1) and (2) indicate that for non-insider foundations, neither related nor unknown holdings significantly predict alpha. In contrast, Columns (3) and (4) show that among HF-affiliated foundations, **related** holdings are the strongest predictor of alpha, with coefficients of 7.16 without year fixed effects and 6.06 with year fixed effects. Coefficients on **unknown** holdings are also positive and economically large (4.93 and 4.66), consistent with some unknown positions representing undisclosed affiliated funds. By contrast, **unrelated** holdings exhibit mixed signs and are uniformly statistically insignificant. Family-office hold-

ings carry negative and significant coefficients, indicating that managing personal capital is economically distinct from allocating to affiliated institutional funds.

Among PEVC-affiliated foundations, alpha loads on both related and unknown holdings; however, the coefficient on unknown holdings has a substantially smaller standard error, making it statistically more significant. The stronger and more stable relationship between unknown holdings and alpha is consistent with greater incentives to obscure PE and VC positions, given competitive pressures surrounding access to oversubscribed funds.

Taken together, these results strongly favor the insider-advantage explanation over the skill hypothesis. Alternative-asset investments generate no excess return outside affiliated positions. Alpha appears to be a property of the LP–GP identity overlap, not of the LP's general investment capability.

5.4 Retirement

Prediction (iv) examines whether insider alpha persists after the founder steps back from active management. If insider-foundation alpha reflects the founder's position as owner, partner, or senior executive of the investment firm providing the affiliated private funds—rather than superior personal investment skill—then alpha should persist even after the founder steps back from day-to-day management.

We manually collect the retirement dates of founders of insider-affiliated foundations. Because these foundations are large and their founders are public figures, retirement dates are observable for all cases in our sample. We re-estimate the baseline panel regression in Equation (1), augmented with an indicator for the post-retirement period and event-time dummies for the retirement year (t_0) and the three preceding years ($t_0 - 1$, $t_0 - 2$, $t_0 - 3$). These lead indicators account for the possibility that retirement timing is endogenous to recent performance or portfolio adjustments. We include foundation fixed effects to isolate within-foundation changes around retirement and year fixed effects to control for aggregate trends in insider alpha. To strengthen identification, we match each insider-affiliated foundation to a NONFIN foundation on lagged size and lagged alternative-asset weights using both propensity score matching (logit specification) and Mahalanobis distance, then estimate difference-in-differences specifications comparing insider foundations to their matched

counterparts before and after retirement.

Across specifications, we find no evidence that insider alpha declines following retirement. For HF-affiliated foundations, alpha increases by 24 basis points under propensity score matching and by 12 basis points under Mahalanobis matching—both economically small relative to average HF alpha and statistically insignificant—alongside modest reductions of 2–8 percentage points in alternative-asset allocations. For PEVC-affiliated foundations, the point estimates suggest alpha increases of 8 to 26 basis points post-retirement across the two matching methods, also statistically insignificant, while alternative-asset allocations increase by 5 percentage points under propensity score matching and remain unchanged under Mahalanobis matching.

The persistence of insider alpha after retirement supports the insider-advantage interpretation. Performance appears tied to the founder's ownership stake and continuing structural relationship with the affiliated investment firm, rather than to direct involvement in day-to-day investment decisions—an advantage that outside LPs cannot replicate.

5.5 Time-Series Evolution of Insider Alpha

We next examine whether the association between insider alternative-asset allocations and alpha has been stable over time. This question is motivated by two trends documented in the literature. First, average net-of-fee private fund alpha has declined substantially since the financial crisis (Bollen et al., 2021; Sullivan, 2020; Phalippou, 2020). Second, the systematic performance differences across LP types documented by **?** and Sensoy et al. (2014), as well as buyout fund performance persistence Harris et al. (2023), have weakened since the early 2000s. We therefore ask whether insider alpha has eroded alongside these broader industry trends.

To examine this question, we re-estimate Equation (1) in 15-year rolling windows, using lagged alternative-asset weight as the key explanatory variable for alpha. We also estimate two specifications that isolate the components most strongly associated with insider advantage: one replacing total weight with misclassified weight (the difference between our LLM-based and SOI-based measures), and one decomposing total weight into related, unrelated, family-office, and unknown components.

Figure 3 reports the estimated coefficients across eight rolling windows, with the first spanning 2002–2016 and the last spanning 2009–2023. For HF-affiliated foundations, every measure of insider alternative-asset exposure—total weight, misclassified weight, related weight, and unknown weight—predicts alpha with remarkably stable coefficients across windows ending between 2016 and 2022. There is no evidence of a gradual decline in the relationship between insider exposure and CAPM alpha. In the final window, however, the coefficients fall across all measures, and in three of the four cases they are no longer statistically significant at the 10% level. This weakening coincides with the industry-wide underperformance of hedge funds in 2023 (HFRI hedge fund index return of 7.5% versus 26.3% for the S&P 500).

For PEVC-affiliated foundations, the pattern is broadly similar but shows an earlier decline in total alternative-asset weight beginning in the fifth window (ending in 2020), a trend largely mirrored by the unknown-weight component. The coefficients on misclassified and related weights remain relatively stable until the final window ending in 2023, when all four measures fall below earlier estimates and none remains statistically significant at the 10% level. As with HF foundations, this timing coincides with weak private-market performance relative to public equities in 2023 (Cambridge Associates U.S. PE return of 9.3% and VC return of -3.4% versus 26.3% for the S&P 500).

Overall, the relationship between insider alternative-asset allocations and alpha is stable for most of the sample period and shows little evidence of gradual structural decay, consistent with an advantage tied to the GP-founder relationship rather than to industry-wide conditions. Whether the weakening in the final window reflects a temporary cyclical downturn or the onset of a more persistent structural shift requires future data to assess.

Taken together, the evidence across our four tests tilts the weight of evidence toward insider advantage rather than selection skill, though we do not claim to rule out the skill hypothesis entirely. A skilled investor could plausibly maintain a concentrated portfolio, prefer some degree of opacity, or invest more successfully in funds where she has closer relationships. Nevertheless, the joint appearance of all four patterns—concentration, targeted misclassification of the highest-returning positions, the absence of alpha in unrelated alternative holdings of the same sophisticated founders, and the persistence of alpha after retirement—arises nat-

urally under insider advantage but requires increasingly strong auxiliary assumptions under a selection-skill account. We also acknowledge that insider advantage and selection skill need not be mutually exclusive: these founders often become GPs precisely because they possess investment ability. What our results indicate is that the alternative alpha earned by insider foundations is not transferable to outside investors—it reflects the structural advantage of the GP-founder relationship rather than a replicable investment capability.

6 Discussion

Berk and Green (2004) predict that competitive capital flows drive the net-of-fee CAPM alpha of active mutual managers to zero in equilibrium. Whether this logic applies to private markets is theoretically unclear. Illiquid assets may command an illiquidity premium (**?**), LPs are more resourceful and sophisticated than average mutual fund investors, and GPs may share inframarginal rents with LPs whose non-capital contributions they value (Hochberg et al., 2014). Empirically, whether the average private fund generates positive net-of-fee alpha also remains actively contested (**?**Harris et al., 2014; **?**). We take no position on this question. Instead, conditional on alpha existing in some private funds, we ask who captures it.

The answer is more pessimistic for outside investors than Berk–Green would imply. Non-insider foundations in our sample—including large, well-staffed institutional investors with substantial assets and, in many cases, professional financial expertise—earn zero CAPM alpha from alternative investments despite their ability to deploy sizable capital over long horizons.

7 Conclusion

We study the sources of alternative alpha in private markets using a novel dataset of private foundation returns and security-level holdings recovered from IRS Form 990-PF filings. We document three main findings. First, the size and alternative-asset-tilt premia documented in prior work disappear once return smoothing is properly accounted for using the Dimson

(1979) correction: neither scale nor broad private-market exposure generates risk-adjusted outperformance. Second, alpha is concentrated among foundations established by hedge fund principals; foundations established by private equity and venture capital managers, other financial professionals, and nonfinancial founders earn alphas statistically indistinguishable from zero. Third, four mechanism tests—portfolio concentration, strategic misclassification of the highest-returning positions, the absence of alpha in unaffiliated holdings of the same sophisticated founders, and persistence after founder retirement—jointly indicate that this outperformance reflects LP–GP identity overlap rather than transferable investment skill.

These findings have implications for the literature on LP heterogeneity and for policy arguments advocating expanded retail access to alternative assets. The LP performance differences documented by Lerner et al. (2007) and Sensoy et al. (2014) have narrowed over time, yet we show that performance differences tied to investor identity persist. Recent proposals cite evidence that private equity buyout funds outperform public equities by roughly 3% per year as justification for democratizing access (Council of Economic Advisers, 2025). Even accepting this estimate, the relevant benchmark is not the average GP return but the return available to outside LPs—and the institutional investors in our sample, whose resources far exceed those of any plausible retail entrant, already fail to earn it. Retail investors, lacking the relationship capital, co-investment networks, and negotiating leverage available to institutional LPs, therefore face an even more disadvantaged starting point. Expanding access without altering the underlying market structure exposes investors to the fees and illiquidity of the asset class without offering a credible path to the alpha that justifies those costs.

References

Abuzov, R., Gornall, W., Strebulaev, I. A., 2025. The value of privacy and the choice of limited partners by venture capitalists. Journal of Financial Economics 169, 104063.

Agarwal, V., Mullally, K. A., Naik, N. Y., 2015. The Economics and Finance of Hedge Funds: A Review of the Academic Literature. Foundations and Trends(R) in Finance 10, 1–111.

Andonov, A., Bauer, R. M. M. J., Cremers, K. J. M., 2017. Pension Fund Asset Allocation and Liability Discount Rates. The Review of Financial Studies 30, 2555–2595.

Bach, L., Calvet, L. E., Sodini, P., 2020. Rich Pickings? Risk, Return, and Skill in Household Wealth. American Economic Review 110, 2703–2747.

Balloch, C., Mainardi, F., Oh, S., Vokata, P., 2025. Democratizing Private Markets: Private Equity Performance of Individual Investors.

Begenau, J., Liang, P., Siriwardane, E., 2025. The Rise of Alternatives.

Berk, J. B., Green, R. C., 2004. Mutual Fund Flows and Performance in Rational Markets. Journal of Political Economy 112, 1269–1295.

Binfare, M., Zimmerschied, K., 2026. Investing with Purpose: Evidence from Private Foundations. Journal of Finance, forthcoming .

Bollen, N. P., Joenväärä, J., Kauppila, M., 2021. Hedge Fund Performance: End of an Era? Financial Analysts Journal 77, 109–132.

Council of Economic Advisers, 2025. Retail Access to Alternative Investments Via Defined Contribution Plans. In: *The White House*, The White House.

Dahiya, S., Yermack, D., 2018. Investment Returns and Distribution Policies of Non-Profit Endowment Funds.

Dimson, E., 1979. Risk measurement when shares are subject to infrequent trading. Journal of Financial Economics 7, 197–226.

Gabaix, X., Koijen, R. S. J., Mainardi, F., Oh, S., Yogo, M., 2025. Asset Demand of U.S. Households.

Harris, R. S., Jenkinson, T., Kaplan, S. N., 2014. Private Equity Performance: What Do We Know? The Journal of Finance 69, 1851–1882.

Harris, R. S., Jenkinson, T., Kaplan, S. N., Stucke, R., 2023. Has persistence persisted in private equity? Evidence from buyout and venture capital funds. Journal of Corporate Finance 81, 102361.

Hochberg, Y. V., Ljungqvist, A., Vissing-Jørgensen, A., 2014. Informational Holdup and Performance Persistence in Venture Capital. The Review of Financial Studies 27, 102–152.

Kaplan, S. N., Sensoy, B. A., 2015. Private Equity Performance: A Survey. Annual Review of Financial Economics 7, 597–614.

Lerner, J., Mao, J., Schoar, A., Zhang, N. R., 2022. Investing outside the box: Evidence from alternative vehicles in private equity. Journal of Financial Economics 143, 359–380.

Lerner, J., Schoar, A., Wongsunwai, W., 2007. Smart Institutions, Foolish Choices: The Limited Partner Performance Puzzle. The Journal of Finance 62, 731–764.

Lo, A. W., Matveyev, E., Zeume, S., 2019. The Risk, Reward, and Asset Allocation of Nonprofit Endowment Funds.

Phalippou, L., 2020. An Inconvenient Fact: Private Equity Returns and the Billionaire Factory. The Journal of Investing 30, 11–39.

Phalippou, L., 2026. Navigating Private Equity Data: A Critical Review of Key Sources.

Sensoy, B. A., Wang, Y., Weisbach, M. S., 2014. Limited partner performance and the maturing of the private equity industry. Journal of Financial Economics 112, 320–343.

Sørensen, M., 2007. How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital. The Journal of Finance 62, 2725–2762.

Sullivan, R. N., 2020. Hedge Fund Alpha: Cycle or Sunset? The Journal of Alternative Investments 23, 55–79.

Figure 1: Equal-Weighted Size Sorted Portfolios of All Foundations (1995–2023)
This figure plots the performance of equal-weighted size-sorted portfolios, based on CPI-adjusted one-year-lagged total assets. $\alpha_i^{unadjusted}$ is estimated from the time-series regression:

$$R_{i,t} - R_{f,t} = \alpha_i^{unadjusted} + \beta_i MKT_t + \epsilon_{i,t}$$

The Dimson-adjusted α_i is estimated as:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{i,0} MKT_t + \beta_{i,1} MKT_{t-1} + \varepsilon_{i,t}$$

where $R_{i,t}$ is portfolio i's annual return, $R_{f,t}$ is the risk-free rate, and MKT represents the market factor. Size groups are: Below 50M, 50M to 100M, Above 100M (Excludes top 100), and Top 100.

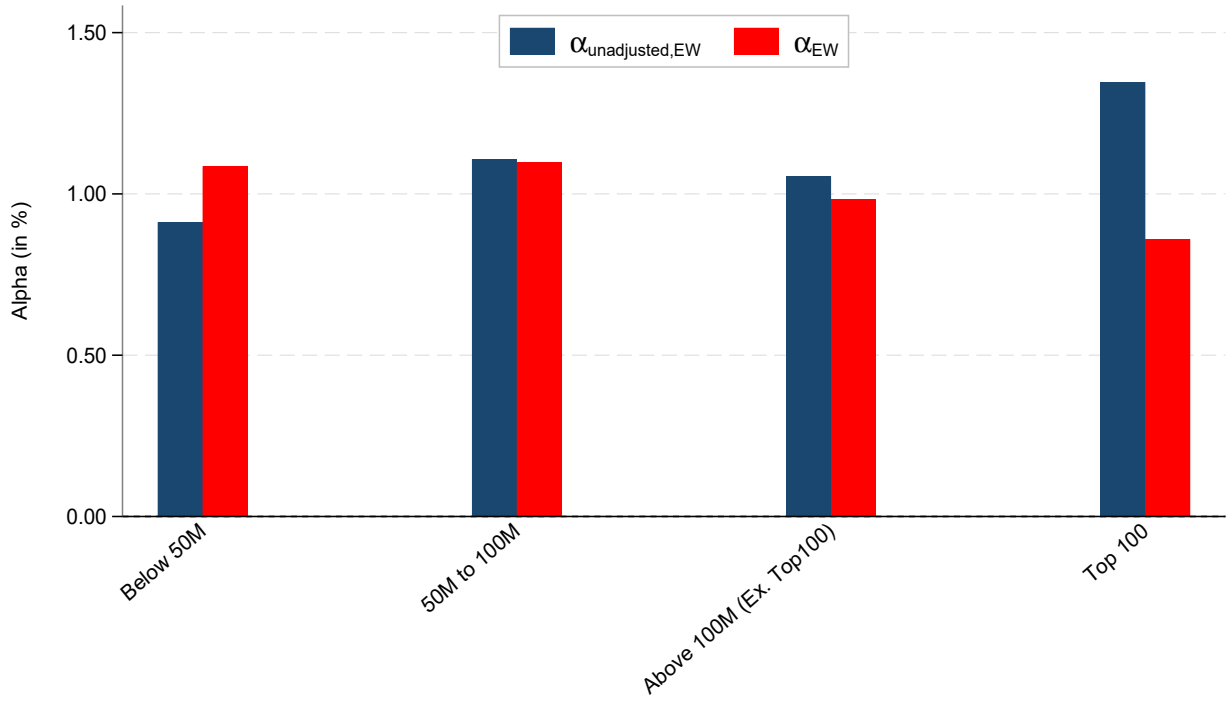
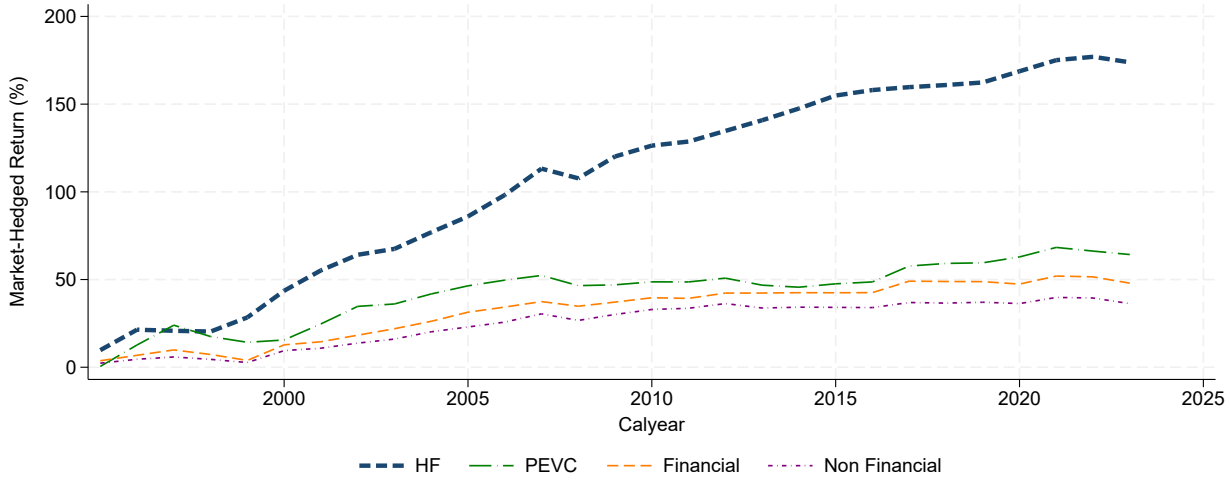


Figure 2: Cumulative Market-Hedged Returns (1995–2023)
This figure plots the cumulative market-hedged returns estimated as:

$$R_{i,t}^{\text{hedged}} = R_{i,t} - \beta_{i,0} MKT_t - \beta_{i,1} MKT_{t-1}$$

where $R_{i,t}$ is foundation i's annual return, and MKT represents the market factor. Panel A plots the cumulative sum of cross-sectional average market-adjusted returns by founder type. Panel B illustrates HF premium in the context of the size gradient by plotting HF and equal-weighted size-sorted portfolios.

Panel A: LP's Performance Difference Across Investor Identity



Panel B: Size Premium

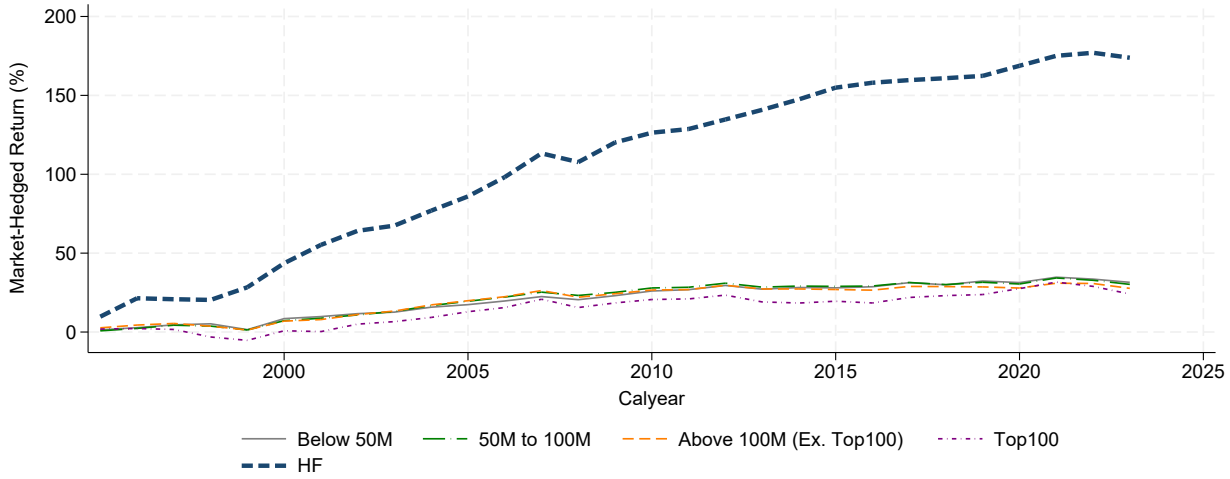


Figure 3: Affiliated Private Weight and Alpha
This figure plots the estimated 15-year rolling coefficients and 10% confidence intervals for the interactions of the indicated variable with HF and PEVC indicators for the $100M sample. Coefficients are computed following the specifications in Table 5 column (6), Table 8 column (6), and Table 9 column (4).

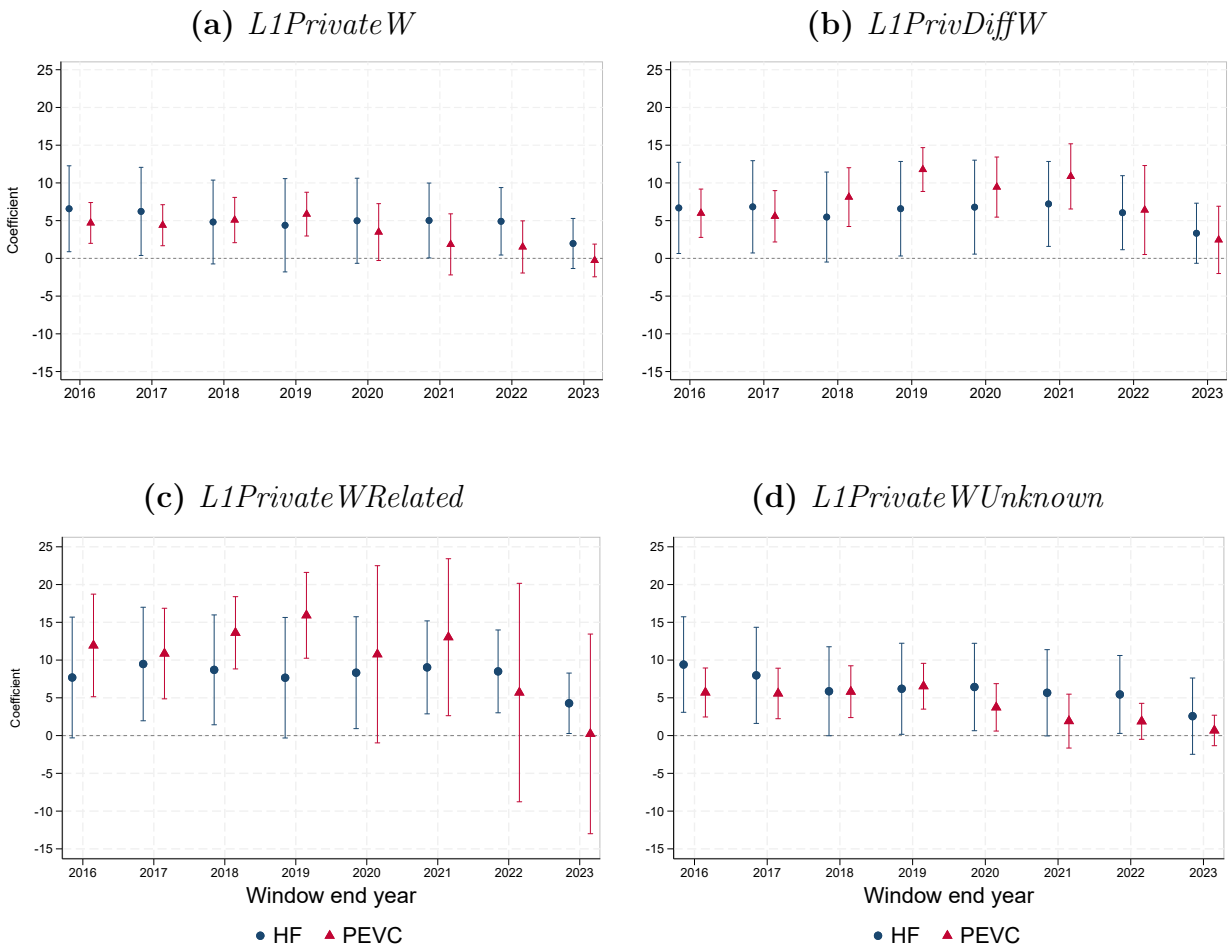


Table 1: Missing Returns Analysis

This table reports, for foundations sorted by size in a given year, the probability of having at least one missing annual return over the next 15-year window. *Rank Yr* is the year in which foundations are sorted into size groups. *Return Yrs* is the 15-year window over which missing returns are assessed. *Nobs* is the number of foundations in each size group. *% CS* and *% SOI* report the probability of at least one missing return in our comprehensive sample and the SOI dataset, respectively.

		Below $10 M			[$10 M, $100 M]			Above $100 M			Top 100		
Rank Yr	Return Yrs	Nobs	% CS	% SOI	Nobs	% CS	% SOI	Nobs	% CS	% SOI	Nobs	% CS	% SOI
1994	1996-2009	1006	7.3	56.9	1599	1.8	20.4	255	0.8	9.0	100	0.0	7.0
1995	1996-2010	1333	7.4	62.9	2025	2.1	25.3	314	1.0	8.9	100	0.0	6.0
1996	1997-2011	1701	6.3	61.6	2317	1.5	26.8	355	0.0	8.2	100	0.0	5.0
1997	1998-2012	2136	5.2	65.1	2777	1.4	32.0	462	0.2	8.0	100	0.0	5.0
1998	1999-2013	2752	3.7	72.6	3406	0.9	38.8	535	0.0	8.2	100	0.0	4.0
1999	2000-2014	3059	1.9	75.7	3891	0.7	43.7	609	0.2	8.0	100	0.0	4.0
2000	2001-2015	3417	1.5	76.9	4162	0.5	44.2	650	0.2	7.7	100	0.0	3.0
2001	2002-2016	3809	1.7	76.9	4280	0.5	41.8	602	0.5	6.5	100	0.0	5.0
2002	2003-2017	4262	1.6	77.0	4195	0.7	37.4	533	0.2	5.4	100	0.0	3.0
2003	2004-2018	4405	1.8	69.3	4455	0.9	37.4	610	0.5	5.2	100	0.0	3.0
2004	2005-2019	4547	1.8	66.4	4717	0.8	37.1	633	0.3	4.3	100	0.0	5.0
2005	2006-2020	4697	2.5	66.3	4908	1.3	37.4	671	0.7	5.2	100	0.0	4.0

Table 2: Foundation Performance (1995–2023)

This table reports the performance summary statistics for foundations in the full sample (Panel A) and the $100M sample (Panel B), requiring at least five annual EIN-return observations. $\alpha_i^{unadjusted}$ is estimated from the time-series regression:

$$R_{i,t} - R_{f,t} = \alpha_i^{unadjusted} + \beta_i MKT_t + \epsilon_{i,t}$$

The Dimson-adjusted α_i is estimated as:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{i,0} MKT_t + \beta_{i,1} MKT_{t-1} + \varepsilon_{i,t}$$

where $R_{i,t}$ is foundation i's annual return, $R_{f,t}$ is the risk-free rate, and MKT represents the market factor. *Avg Private Weight* is the mean fraction of assets allocated to alternative investments, measured using our LLM-based weights and complemented with SOI data when LLM weights are missing.

Panel A: Full Sample

	N	Mean	SD	P5	P25	P50	P75	P95
Avg Excess Ret (in %)	18498	6.09	3.46	1.08	4.45	5.85	7.46	11.59
$\alpha^{unadjusted}$ (in %)	18498	0.39	3.63	-3.42	-1.56	-0.12	1.51	6.46
α (in %)	18498	0.39	5.25	-4.12	-1.44	-0.00	1.67	6.95
β	18498	0.55	0.26	0.06	0.44	0.59	0.69	0.88
β_0	18498	0.55	0.28	0.07	0.44	0.59	0.69	0.88
β_1	18498	-0.01	0.17	-0.20	-0.07	-0.01	0.05	0.22
Volatility (in %)	18498	12.63	4.88	6.08	10.09	12.07	14.53	20.99
Avg Private Weight (in %)	17646	15.15	22.84	0.00	0.00	4.09	21.05	69.66

Panel B: $100M Sample

	N	Mean	SD	P5	P25	P50	P75	P95
Avg Excess Ret (in %)	1463	6.41	5.22	0.99	4.65	5.99	7.50	12.33
$\alpha^{unadjusted}$ (in %)	1463	0.86	4.32	-3.43	-1.04	0.30	1.83	7.56
α (in %)	1463	0.19	10.92	-5.08	-1.36	0.16	1.83	7.00
β	1463	0.56	0.33	0.09	0.45	0.58	0.67	0.90
β_0	1463	0.57	0.49	0.09	0.46	0.58	0.68	0.92
β_1	1463	0.03	0.25	-0.19	-0.05	0.01	0.07	0.31
Volatility (in %)	1463	13.67	10.25	6.84	10.50	12.39	14.87	22.92
Avg Private Weight (in %)	1609	23.49	25.38	0.00	2.06	14.48	36.96	79.34

Table 3: Foundation-Level Panel Regression – Size and Private Tilt (2002–2023)

This table reports the results from a panel regression at the foundation level:

$$R_{i,t} - R_{f,t} = \left(\alpha_0 + \boldsymbol{\gamma}'\mathbf{X}_{i,t-1}\right) + \beta_{i,0}, \mathrm{MKT}_t + \beta_{i,1}, \mathrm{MKT}_{t-1} + \varepsilon_{i,t},$$

$\beta_{i,0}$ and $\beta_{i,1}$ are foundation i's exposures to the contemporaneous and one-year-lagged market excess return (*MKT*), estimated but not reported. $\mathbf{X}_{i,t-1}$ is a vector of foundation characteristics. *Top 100* equals one if the foundation is in the top 100 in year t based on CPI-adjusted total assets at $t-1$. *Size(rank)* represents a normalized cross-sectional size rank scaled from 0 (smallest) to 1 (largest). *L1PrivateW* is foundation i's one-year-lagged LLM-based allocation to alternative investments.

	(1)	(2)	(3)	(4)
Top100	0.37	-0.13		
	(0.4)	(0.3)		
Size(rank)			-0.08	-0.34
			(0.6)	(0.7)
L1PrivateW	-0.18	1.06	-0.13	1.09*
	(0.7)	(0.6)	(0.7)	(0.6)
Constant	0.49		0.56	
	(0.5)		(0.5)	
Sample	100M	100M	100M	100M
Year FE	No	Yes	No	Yes
R-squared	0.57	0.59	0.57	0.59
N	21850	21850	21850	21850

Standard errors in parentheses

Standard errors clustered by calendar year. Fixed effects when indicated.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Foundation Performance by Category (1995–2023)

This table reports performance summary statistics for the 100M sample by foundation category, following the same approach as in Table 2. Each foundation is classified into a category based on the professional background of the individual most responsible for providing its capital. There are 1,463 foundations that have a minimum of five annual return observations to estimate alpha, consisting of 42 PEVC, 86 hedge fund (HF), 285 financial firm (FIN), and 1,050 non-financial (NONFIN) foundations.

	N	Mean	SD	P5	P25	P50	P75	P95
PEVC								
Avg Excess Ret (in %)	42	7.66	7.09	-0.16	4.12	6.68	8.21	18.54
$\alpha^{unadjusted}$ (in %)	42	2.11	5.72	-3.49	-1.50	0.40	2.97	14.56
α (in %)	42	1.12	4.08	-4.81	-1.19	1.41	3.50	8.72
β	42	0.54	0.34	-0.03	0.35	0.53	0.72	1.11
β_0	42	0.56	0.37	0.06	0.35	0.53	0.77	1.11
β_1	42	0.06	0.30	-0.33	-0.12	0.05	0.16	0.51
Volatility (in %)	42	16.20	11.13	7.14	10.91	13.77	17.57	34.77
Hedge Fund								
Avg Excess Ret (in %)	86	8.90	7.54	1.86	4.67	7.23	10.97	21.39
$\alpha^{unadjusted}$ (in %)	86	3.69	6.65	-4.92	-0.59	2.92	7.06	16.41
α (in %)	86	1.52	7.56	-9.69	-2.19	1.02	5.27	13.42
β	86	0.45	0.41	-0.14	0.22	0.43	0.60	1.17
β_0	86	0.49	0.49	-0.23	0.23	0.44	0.64	1.24
β_1	86	0.14	0.33	-0.22	-0.03	0.06	0.20	0.62
Volatility (in %)	86	14.04	8.82	4.88	9.36	11.98	15.92	33.08
Financial Firm								
Avg Excess Ret (in %)	285	6.79	7.27	1.93	4.65	6.32	7.73	11.22
$\alpha^{unadjusted}$ (in %)	285	0.80	4.74	-3.91	-0.95	0.53	2.05	6.86
α (in %)	285	-1.13	22.70	-5.41	-1.36	0.33	2.02	6.58
β	285	0.58	0.48	0.09	0.40	0.56	0.69	0.92
β_0	285	0.63	0.92	0.09	0.42	0.58	0.69	0.96
β_1	285	0.06	0.40	-0.19	-0.06	0.01	0.08	0.40
Volatility (in %)	285	14.65	19.15	6.53	9.94	12.43	15.00	24.88
Non-financial								
Avg Excess Ret (in %)	1050	6.05	4.04	0.83	4.67	5.86	7.29	11.53
$\alpha^{unadjusted}$ (in %)	1050	0.60	3.78	-3.24	-1.04	0.19	1.57	5.68
α (in %)	1050	0.40	4.56	-4.15	-1.35	0.10	1.62	5.81
β	1050	0.56	0.27	0.13	0.48	0.59	0.67	0.87
β_0	1050	0.57	0.27	0.13	0.48	0.59	0.67	0.88
β_1	1050	0.01	0.17	-0.17	-0.05	0.01	0.06	0.24
Volatility (in %)	1050	13.27	5.95	7.00	10.57	12.39	14.78	22.07

Table 5: Foundation-Level Panel Regression – Size, Private Tilt, and Foundation Characteristics (2002–2023)

This table reports the results from a panel regression at the foundation level in the 100M sample, following the same specification as in Table 3. $\mathbf{X}_{i,t-1}$ is augmented with indicator variables for corporate, hospital conversion, and community foundations, and with founder-type indicators for hedge fund (HF), private equity and venture capital (PEVC), and other financial foundations (FIN). *L1PrivateW(HF)*, *L1PrivateW(PEVC)*, and *L1PrivateW(FIN)* are interactions of *L1PrivateW* with the corresponding founder-type indicators. Models (7) and (8) exclude HF, PE, and VC foundations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Top100	0.34	-0.03	0.41	-0.09	0.40	-0.10	0.50	-0.04
	(0.4)	(0.3)	(0.4)	(0.3)	(0.4)	(0.3)	(0.4)	(0.3)
HospitalConv	0.10	0.12						
	(0.2)	(0.2)						
Corporate	-0.25	-0.23						
	(0.3)	(0.3)						
Community	-0.30	-0.04						
	(0.6)	(0.7)						
HF	2.55***	3.17***	2.82***	2.94***	-0.08	0.56		
	(0.7)	(0.7)	(0.6)	(0.6)	(1.3)	(1.4)		
PEVC	0.16	0.59	0.27	0.51	-1.18	-0.44		
	(0.6)	(0.6)	(0.6)	(0.6)	(0.7)	(0.7)		
FIN	-0.08	0.05	-0.09	0.09	-0.01	0.15		
	(0.2)	(0.2)	(0.2)	(0.2)	(0.3)	(0.3)		
L1PrivateW			-0.62	0.62	-0.90	0.39	-0.99	0.28
			(0.7)	(0.6)	(0.7)	(0.6)	(0.7)	(0.6)
L1PrivateW(HF)					4.76*	3.90		
					(2.5)	(2.5)		
L1PrivateW(PEVC)					3.97**	2.62*		
					(1.6)	(1.5)		
L1PrivateW(FIN)					-0.51	-0.40		
					(0.9)	(0.8)		
Constant	0.38		0.50		0.57		0.56	
	(0.6)		(0.5)		(0.5)		(0.5)	
Ex. HF/PE/VC	No	No	No	No	No	No	Yes	Yes
Year FE	No	Yes	No	Yes	No	Yes	No	Yes
R-squared	0.57	0.59	0.57	0.59	0.57	0.59	0.57	0.59
N	21850	21850	21850	21850	21850	21850	20371	20371

Standard errors in parentheses
Standard errors clustered by calendar year. Fixed effects when indicated.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Foundation-Level Panel Regression – Concentration (2002–2023)

This table reports the results from a panel regression in the 100M sample, following the main specification in Table 3. *ZeroPri* equals one if the foundation's private weight is zero. *SmallN* equals one if the foundation holds three or fewer alternative-asset investments. *SmallN(HF)* and *SmallN(PEVC)* are interactions of *SmallN* with the HF and PEVC indicators.

	(1)	(2)	(3)	(4)	(5)	(6)
Top100	0.48	-0.08	0.47	-0.09	0.45	-0.11
	(0.4)	(0.3)	(0.4)	(0.3)	(0.4)	(0.3)
HF	2.69***	2.90***	1.01	1.21**	-1.80	-1.46
	(0.6)	(0.6)	(0.6)	(0.6)	(1.6)	(1.5)
PEVC	0.33	0.49	0.37	0.45	-1.24	-0.75
	(0.6)	(0.6)	(1.0)	(1.0)	(1.1)	(1.0)
L1PrivateWeight	0.10	0.63	0.05	0.58	-0.35	0.22
	(0.5)	(0.6)	(0.5)	(0.5)	(0.4)	(0.5)
ZeroPri	-0.04	-0.18	-0.01	-0.15	-0.05	-0.18
	(0.3)	(0.3)	(0.3)	(0.3)	(0.3)	(0.3)
SmallN	1.36**	0.25	1.21*	0.10	1.15*	0.04
	(0.6)	(0.2)	(0.6)	(0.2)	(0.7)	(0.2)
SmallN(HF)			3.35***	3.36***	3.55***	3.54***
			(1.1)	(1.0)	(1.1)	(1.0)
SmallN(PEVC)			-0.09	0.07	0.38	0.42
			(1.4)	(1.3)	(1.4)	(1.3)
L1PrivateW(HF)					4.53	4.28
					(2.7)	(2.6)
L1PrivateW(PEVC)					3.79**	2.85*
					(1.5)	(1.5)
Constant	-0.52		-0.42		-0.28	
	(0.8)		(0.8)		(0.8)	
Year FE	No	Yes	No	Yes	No	Yes
R-squared	0.57	0.59	0.57	0.59	0.57	0.59
N	21850	21850	21850	21850	21850	21850

Standard errors in parentheses
Standard errors clustered by calendar year. Fixed effects when indicated.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Schedule Reporting Summary Statistics by Category

This table reports summary statistics on alternative investment allocations by foundation category in the 100M sample. *Avg Priv Weight* represents the LLM-based private weight, and *Avg Priv Weight Diff* is the difference between the LLM and SOI weights. Avg percentage of alternative assets in *related investments*, *unrelated investments*, *unknown investments* and *family office*, decompose the foundation's total LLM-based private weight into the four components described in Table 9, expressed as a percentage of total private weight. The sample requires at least five annual foundation returns to estimate alpha. Relatedness is assessed only for holdings exceeding 0.1% of total assets.

	N	Mean	SD	P5	P25	P50	P75	P95
PEVC								
Avg Priv Weight (in %)	42	37.89	28.48	0.07	12.06	34.85	61.39	85.82
Avg Priv Weight Diff: LLM vs SOI (in %)	42	-3.28	20.59	-23.46	-8.90	-0.19	1.88	28.72
Avg % of Alternative Assets in Related Investments	38	15.83	28.37	0.00	0.00	0.09	15.40	92.98
Avg % of Alternative Assets in Unrelated Investments	38	24.57	28.08	0.00	0.00	8.47	47.39	77.29
Avg % of Alternative Assets in Unknown Investments	38	48.86	36.78	2.33	12.73	47.07	92.74	100.00
Avg % of Alternative Assets in Family Office	42	1.29	5.53	0.00	0.00	0.00	0.00	2.48
Hedge Fund								
Avg Priv Weight (in %)	86	66.25	29.49	0.95	50.52	74.26	93.14	99.26
Avg Priv Weight Diff: LLM vs SOI (in %)	86	7.44	21.17	-9.18	-0.80	0.00	5.94	59.97
Avg % of Alternative Assets in Related Investments	72	42.83	40.22	0.00	0.00	33.03	87.74	100.00
Avg % of Alternative Assets in Unrelated Investments	72	20.64	27.65	0.00	0.00	4.89	31.99	74.29
Avg % of Alternative Assets in Unknown Investments	72	31.86	35.69	0.00	0.39	17.80	59.50	100.00
Avg % of Alternative Assets in Family Office	86	6.26	20.92	0.00	0.00	0.00	0.00	60.13
Financial Firm								
Avg Priv Weight (in %)	282	19.59	22.52	0.00	1.35	10.34	30.29	66.89
Avg Priv Weight Diff: LLM vs SOI (in %)	282	-3.46	16.25	-34.16	-4.12	-0.05	0.78	12.72
Avg % of Alternative Assets in Related Investments	264	2.36	10.36	0.00	0.00	0.00	0.00	13.13
Avg % of Alternative Assets in Unrelated Investments	264	26.19	30.32	0.00	0.00	11.87	48.04	83.95
Avg % of Alternative Assets in Unknown Investments	264	48.74	34.65	0.00	18.76	43.65	80.02	100.00
Avg % of Alternative Assets in Family Office	282	2.54	12.58	0.00	0.00	0.00	0.00	3.78
Nonfinancial								
Avg Priv Weight (in %)	1042	22.28	23.86	0.00	1.70	14.01	35.82	71.00
Avg Priv Weight Diff: LLM vs SOI (in %)	1042	-1.79	17.93	-26.91	-3.33	0.00	2.11	19.04
Avg % of Alternative Assets in Related Investments	924	1.01	6.79	0.00	0.00	0.00	0.00	0.96
Avg % of Alternative Assets in Unrelated Investments	924	22.23	28.59	0.00	0.00	5.19	39.99	82.54
Avg % of Alternative Assets in Unknown Investments	924	44.15	35.28	0.00	9.55	41.04	74.40	100.00
Avg % of Alternative Assets in Family Office	1042	1.94	10.98	0.00	0.00	0.00	0.00	3.20

Table 8: Foundation-Level Panel Regression – LLM vs SOI Weights (2002–2023)

This table reports the results from a panel regression in the 100M sample following the main specification in Table 3. The sample is restricted to observations with non-missing LLM and SOI lagged private weights. *L1PrivateW_SOI* is the foundation's one-year-lagged alternative asset allocation as reported in SOI. *L1PrivDiffW* is the difference between the LLM-based weight (*L1PrivateW*) and the SOI-based weight (*L1PrivateW_SOI*). Interactions of these variables with the HF and PEVC indicators are also included.

	(1)	(2)	(3)	(4)	(5)	(6)
Top100	0.40	-0.10	0.41	-0.04	0.37	-0.11
	(0.4)	(0.3)	(0.4)	(0.3)	(0.4)	(0.3)
HF	-0.07	0.53	2.91***	3.33***	0.34	0.90
	(1.4)	(1.4)	(1.0)	(1.0)	(1.4)	(1.4)
PEVC	-1.18	-0.47	-0.61	-0.06	-0.98	-0.34
	(0.7)	(0.8)	(0.8)	(0.8)	(0.8)	(0.8)
L1PrivateW	-0.99	0.33				
	(0.7)	(0.6)				
L1PrivateW_SOI			-1.16	0.26	-1.24	0.35
			(0.7)	(0.6)	(0.8)	(0.7)
L1PrivDiffW					-0.29	0.27
					(0.5)	(0.5)
L1PrivateW(HF)	4.85*	3.97				
	(2.6)	(2.6)				
L1PrivateW(PEVC)	4.05**	2.68				
	(1.6)	(1.6)				
L1PrivateW(HF)_SOI			0.04	-0.44	3.49	2.69
			(1.9)	(1.9)	(2.6)	(2.6)
L1PrivateW(PEVC)_SOI			2.57	1.58	3.70**	2.34
			(1.6)	(1.5)	(1.6)	(1.6)
L1PrivDiffW(HF)					7.52**	6.67**
					(2.7)	(2.7)
L1PrivDiffW(PEVC)					5.34*	3.79
					(2.9)	(2.8)
Year FE	No	Yes	No	Yes	No	Yes
R-squared	0.57	0.59	0.57	0.59	0.57	0.59
N	21802	21802	21802	21802	21802	21802

Standard errors in parentheses
Standard errors clustered by calendar year. Fixed effects when indicated.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Foundation-Level Panel Regression – Relatedness (2002–2023)

This table reports results from a panel regression in the 100M sample, following the main specification in Table 3. *L1PrivateW* is the LLM-based weight decomposed into four components: *L1PrivateWRelated*, allocations where a personal or professional relationship between the fund and the foundation founder is identified; *L1PrivateWUnRelated*, allocations to identifiable holdings with no such relationship; *L1PrivateWUnknown*, the residual after removing related and unrelated weights from total private weight; and *L1PriWFAM*, allocations to family office-type investments. Interactions of these variables with the HF and PEVC indicators are also included. Models (1) and (2) exclude HF, and PEVC.

	(1)	(2)	(3)	(4)	(5)	(6)
Top100	0.37	-0.10	0.26	-0.17	0.25	-0.17
	(0.4)	(0.4)	(0.4)	(0.3)	(0.4)	(0.4)
HF			0.12	0.72	0.09	0.70
			(1.4)	(1.4)	(1.4)	(1.4)
PEVC			-1.10	-0.42	-1.13	-0.43
			(0.8)	(0.8)	(0.8)	(0.8)
L1PriWRelated	-1.82	-0.38	-1.82	-0.36		
	(1.7)	(1.7)	(1.7)	(1.7)		
L1PriWUnRelated	-2.32**	-0.67	-2.32**	-0.61	-2.32**	-0.61
	(1.1)	(0.9)	(1.1)	(0.9)	(1.1)	(0.9)
L1PriWUnknown	-0.68	0.32	-0.66	0.35		
	(0.6)	(0.6)	(0.6)	(0.6)		
L1PriWRelated(HF)			7.16**	6.06*		
			(3.2)	(3.3)		
L1PriWUnRelated(HF)			1.31	-0.35	1.35	-0.33
			(2.9)	(2.9)	(2.9)	(2.9)
L1PriWUnknown(HF)			4.93	4.66		
			(3.0)	(3.0)		
L1PriWRelated(PEVC)			3.52	3.57		
			(8.7)	(8.9)		
L1PriWUnRelated(PEVC)			2.46	0.21	2.40	0.19
			(2.6)	(2.7)	(2.6)	(2.7)
L1PriWUnknown(PEVC)			4.76***	3.53**		
			(1.4)	(1.4)		
L1PriWRelBroad					-0.70	0.33
					(0.6)	(0.6)
L1PriWRelBroad(HF)					5.70**	5.16*
					(2.7)	(2.7)
L1PriWRelBroad(PEVC)					4.58**	3.49*
					(1.8)	(1.7)
L1PriWFAM	0.60	0.57	0.59	0.59	0.65	0.62
	(0.7)	(0.7)	(0.7)	(0.7)	(0.7)	(0.7)
L1PriWFAM(HF)			-3.86*	-4.26*	-4.60**	-4.74**
			(2.1)	(2.1)	(1.8)	(1.8)
L1PriWFAM(PEVC)			-5.16	1.31	-5.04	1.33
			(14.4)	(14.5)	(14.4)	(14.4)
Constant	0.64		0.65		0.65	
	(0.5)		(0.5)		(0.5)	
Ex. HF/PE/VC	Yes	Yes	No	No	No	No
Year FE	No	Yes	No	Yes	No	Yes
R-squared	0.56	0.58	0.56	0.57	0.56	0.57
N	17711	17711	19179	19179	19179	19179

Standard errors in parentheses
Standard errors clustered by calendar year. Fixed effects when indicated.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Private Tilt and Performance After Founder Retirement

This table reports results from a panel regression following the baseline specification in Equation (1), estimated on a matched sample of fund-insider and NONFIN foundations. Odd-numbered models use the difference in private weights between matched pairs as the dependent variable; even-numbered models use the difference in CAPM alpha. *Retire* equals one if the observation is in the post-retirement period. *pre0*, *pre1*, *pre2*, and *pre3* are indicator variables for the retirement year (t_0) and the three preceding years ($t_0 - 1$, $t_0 - 2$, $t_0 - 3$). Panel A uses one nearest neighbor propensity score matching (logit specification); Panel B implements Mahalanobis distance matching.

	Panel A: Logit				Panel B: Mahalanobis			
	(1) PrivW HF	(2) Alpha HF	(3) PrivW PEVC	(4) Alpha PEVC	(5) PrivW HF	(6) Alpha HF	(7) PrivW PEVC	(8) Alpha PEVC
retire	-0.02	0.24	0.05	0.26	-0.08	0.12	-0.00	0.08
	(0.0)	(2.6)	(0.0)	(3.4)	(0.0)	(1.6)	(0.0)	(2.7)
pre0	-0.00	10.11	0.05	-0.26	-0.05	9.97	0.08*	1.29
	(0.0)	(7.5)	(0.1)	(4.7)	(0.0)	(7.9)	(0.0)	(3.8)
pre1	-0.05	4.99	0.06*	10.42*	-0.08	3.07	0.06	9.83*
	(0.1)	(3.3)	(0.0)	(5.1)	(0.1)	(2.7)	(0.1)	(5.7)
pre2	0.06	3.44	-0.05	0.89	0.00	3.94	-0.09	1.37
	(0.1)	(2.8)	(0.1)	(6.6)	(0.0)	(2.9)	(0.1)	(6.0)
pre3	-0.05	-0.63	0.03	4.11	-0.01	0.02	0.05	5.07
	(0.0)	(3.8)	(0.0)	(4.8)	(0.0)	(2.9)	(0.0)	(4.3)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Foundation FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.38	0.41	0.37	0.54	0.23	0.43	0.34	0.58
N	878	965	464	510	878	965	464	510
N (EINs)	77	77	38	39	77	77	38	39

Standard errors in parentheses
Standard errors clustered by calendar year in Models (1), (3), (5) and (7).
Standard errors clustered by calendar year and EIN in Models (2), (4), (6) and (8).
Fixed effects when indicated.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendix A

Figure A1: Total Assets (1995–2023)
This figure plots the total assets across all foundations over the 1995-2023 period. Total assets controlled by private foundations increased from $100 billion in 1995 to more than $1 trillion in 2023.

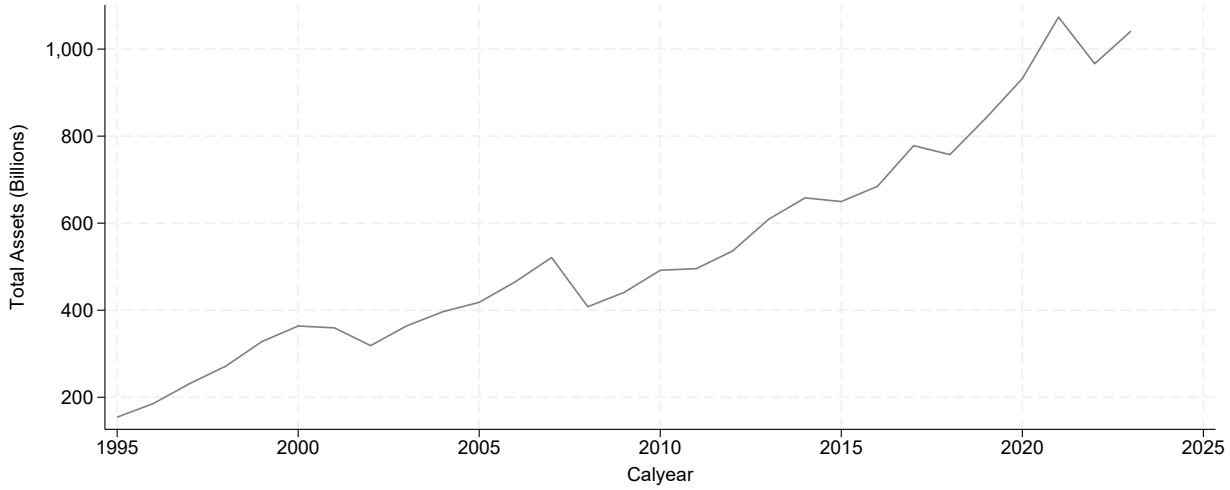


Figure A2: Share of Alternative Investments (1995–2023)
This figure plots the average alternative-asset weight for the full sample. LLM-based weights are complemented with SOI data when LLM weights are missing.

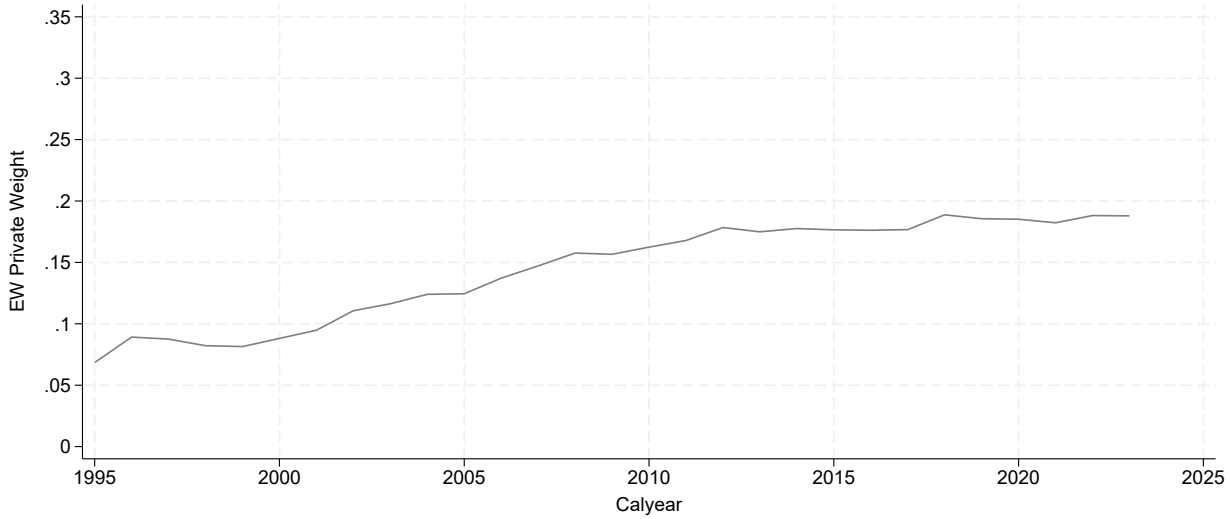
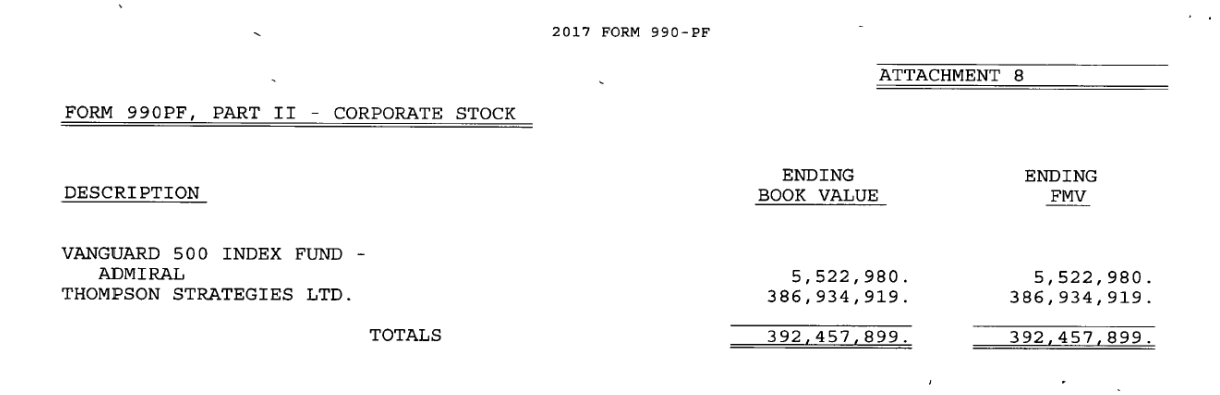


Figure A3: Misreporting - Overdeck Foundation
This figure shows a snapshot of the corporate stock schedule for the Overdeck Foundation for the fiscal year ending in 2017.



2017 FORM 990-PF

ATTACHMENT 8

FORM 990PF, PART II - CORPORATE STOCK

DESCRIPTION	ENDING BOOK VALUE	ENDING FMV
VANGUARD 500 INDEX FUND - ADMIRAL	5,522,980.	5,522,980.
THOMPSON STRATEGIES LTD.	386,934,919.	386,934,919.
TOTALS	392,457,899.	392,457,899.

Figure A4: Average Private Weight
This figure plots the cross-sectional average allocation to alternative investments for foundations in the $100M sample. Both LLM and SOI weights are required to be non-missing.



Table A1: LLM Classification Mapping

This table reports the asset classification assigned to each combination of *Investment Type* (rows) and *Broad Asset Type* (columns) in the LLM-based classification pipeline. CF = Non-Investment Assets; CH = Cash/Cash-like; EQ = Equity; FI = Fixed Income; RA = Real Assets; TBD = To Be Determined.

Investment Type	**CF**	**CH**	**EQ**	**FI**	**Private**	**RA**	**TBD**
Broker	CF	WealthMgmt	WealthMgmt	WealthMgmt	PrivateFund	WealthMgmt	WealthMgmt
WealthMgmt	CF	WealthMgmt	WealthMgmt	WealthMgmt	PrivateFund	WealthMgmt	WealthMgmt
Bank	CF	WealthMgmt	WealthMgmt	WealthMgmt	PrivateFund	WealthMgmt	WealthMgmt
MF	CF	CH	MF	MF	MF	MF	MF
ETF	CF	CH	MF	MF	MF	MF	MF
Commodities	CF	CH	EQ	RA	PrivateFund	RA	RA
Fixed Income	CF	CH	EQ	FI	PrivateFund	FI	FI
CorporateEquity	CF	CH	EQ	EQ	PrivateFund	EQ	EQ
Derivatives	CF	CH	EQ	Derivatives	PrivateFund	Derivatives	Derivatives
PE	CF	CH	PrivateFund	PrivateFund	PrivateFund	PrivateFund	PrivateFund
HF	CF	CH	PrivateFund	PrivateFund	PrivateFund	PrivateFund	PrivateFund
VC	CF	CH	PrivateFund	PrivateFund	PrivateFund	PrivateFund	PrivateFund
InvestmentManager	CF	CH	Alternative	Alternative	PrivateFund	Alternative	Alternative
AlternativeInvestment	CF	CH	Alternative	Alternative	PrivateFund	Alternative	Alternative
FamilyOffice	CF	CH	FamilyOffice	FamilyOffice	PrivateFund	FamilyOffice	FamilyOffice
RealEstate	CF	RA	RA	RA	PrivateFund	RA	RA
Impossible	CF	CH	EQ	FI	PrivateFund	Impossible	Impossible
Art	CF	CF	CF	CF	CF	CF	CF
Other non-investment	CF	CF	CF	CF	CF	CF	CF
Program-related Investment	CF	CF	CF	CF	CF	CF	CF

Table A2: LLM Classification Mapping

This table describes the mapping used to classify assets included in the LLM private weight. We define three specifications—strict, main (our preferred specification), and loose—that differ in how restrictively they assign the alternative asset label across asset classes. Private funds, which are always intermediated and private, are classified as alternative assets under all three definitions. Wealth management accounts are excluded under the strict definition but included under the main definition when the LLM classifies them as "intermediated/private." Assets flagged as "Impossible" are excluded under the strict definition; under the main definition they are included only when the underlying holding appears in the Other Investments Schedule, and under the loose definition when it appears in either the Other Investments Schedule or Other Assets.

Asset Class	**Strict Def.**	**Main Def.**	**Loose Def.**
WealthMgmt	No	Intermediated/Private	NonCorp/NonPublic
MF	No	Intermediated/Private	NonCorp/NonPublic
Fixed Income	No	Intermediated/Private	NonCorp/NonPublic
Equity	No	Intermediated/Private	NonCorp/NonPublic
Derivatives	No	Intermediated/Private	NonCorp/NonPublic
PrivateFund	Intermediated/Private	Intermediated/Private	Intermediated/Private
Alternative	Intermediated/Private	NonCorp/NonPublic	NonCorp/NonPublic
FamilyOffice	No	NonCorp/NonPublic	All
RA	No	Intermediated/Private	NonCorp/NonPublic
CF	No	Intermediated/Private	NonCorp/NonPublic
Impossible	No	OtherInv	NonCorp

Table A3: Equal-Weighted Size Sorted Portfolios of All Foundations (1995–2023)

This table reports performance of equal-weighted size-sorted portfolios, based on CPI-adjusted one-year-lagged total assets. $\alpha_i^{unadjusted}$ is estimated from the time-series regression:

$$R_{i,t} - R_{f,t} = \alpha_i^{unadjusted} + \beta_i MKT_t + \epsilon_{i,t}$$

The Dimson-adjusted α_i is estimated as:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_{i,0} MKT_t + \beta_{i,1} MKT_{t-1} + \varepsilon_{i,t}$$

where $R_{i,t}$ is portfolio i's annual return, $R_{f,t}$ is the risk-free rate, and MKT represents the market factor. Size groups are: Below 50M, 50M to 100M, Above 100M (Excludes top 100), and Top 100. *Mean Obs* represents the average number of foundations in a given portfolio-year; *Min Obs* represents the minimum number of foundations in a given portfolio-year.

	Below 50M	50M to 100M	Above 100M (not Top100)	Top100
Avg Excess Ret (in %)	6.19	6.36	6.62	7.06
$\alpha^{unadjusted}$ (in %)	0.91	1.11	1.05	1.35
α (in %)	1.09	1.09	1.03	0.89
β	0.53	0.53	0.56	0.57
β_0	0.53	0.53	0.56	0.57
β_1	-0.02	0.00	0.00	0.05
Avg Private Weight (in %)	12.80	18.54	20.94	33.04
Mean Obs	7288.28	571.03	520.52	75.03
Min Obs	1588	213	153	66

Table A4: Four-Factor Model and Alpha Adjustment

This table replicates the specification from Binfare and Zimmerschied (2026) using a four-factor model that includes the U.S. equity market factor, ex-U.S. equity market factor, U.S. bond market factor, and a hedge fund return factor, following their size bucket definitions. The four-factor model is also estimated augmented with the lagged market factor to adjust for return smoothing. We report equal-weighted (EW) and value-weighted (VW) results for both the four-factor alpha (4F) and the smoothing-adjusted version (DIM). EW_{Priv} and VW_{Priv} represent the equal-weighted and value-weighted private weight, respectively.

	Below 10M	10M to 50M	50M to 250M	250M to 500M	Above 500M	H–L
$\alpha_{4F,EW}$ (in %)	0.41	0.63	0.66	0.88	0.73	0.32
$\text{t}\alpha_{4F,EW}$	0.85	1.31	1.42	1.47	1.07	0.58
$\alpha_{4F,VW}$ (in %)	0.43	0.72	0.57	0.75	1.85	1.42
$\text{t}\alpha_{4F,VW}$	0.96	1.54	1.24	1.23	2.17	1.98
$\alpha_{DIM,EW}$ (in %)	0.28	0.42	0.30	0.23	-0.03	-0.31
$\text{t}\alpha_{DIM,EW}$	0.50	0.75	0.60	0.38	-0.04	-0.60
$\alpha_{DIM,VW}$ (in %)	0.28	0.49	0.25	0.08	1.08	0.80
$\text{t}\alpha_{DIM,VW}$	0.55	0.92	0.49	0.14	1.20	1.07
EW_{Priv} (in %)	11.40	14.55	18.68	25.13	30.12	17.51
VW_{Priv} (in %)	5.00	12.43	17.35	24.23	29.63	24.64