Private Equity for Pension Plans? Evaluating Private Equity Performance from an Investor's Perspective*

Arthur Korteweg† Stavros Panageas‡ Anand Systla§

November 2022

Abstract

We propose a methodology to evaluate private equity investments by using investor-specific stochastic discount factors. The methodology allows a direct way of decomposing an investor's private-equity return into a risk-compensation and an "alpha". It also helps determine whether a given investor could benefit from investing (more) in private equity. Applying our metrics to U.S. public pension plans, our key results are that: a) during our sample period, pension plan allocations to private equity funds were optimal overall, although the average plan was underexposed to buyout; b) plans invest in PE funds that have higher risk-adjusted performance, but this is because of some pension plans' superior access to successful private equity funds, c) the higher returns obtained by some pension plans in their private equity investments appear to be the result of a higher willingness to take risk rather than a manifestation of timing or selection: Differences in governance structure and funding ratios tend to correlate with the risk-compensation component of the private-equity return rather than the alpha component.

*We thank the Pacific Center for Asset Management (PCAM) for financial support, and seminar participants at EDHEC, Georgetown University, Lehigh University, University of Michigan, Michigan State, PCAM, and the University of Virginia for helpful comments. We are especially grateful to Aleksandar Andonov, Yael Hochberg, and Joshua Rauh for generously sharing their pension plan board membership data.

†Marshall School of Business, University of Southern California; korteweg@marshall.usc.edu

‡Anderson School of Management, UCLA and National Bureau of Economic Research (NBER); stavros.panageas@anderson.ucla.edu

§Anderson School of Management, UCLA; venkat.anand.systla.phd@anderson.ucla.edu.

The performance evaluation of investments is one of the most actively researched topics in finance. The typical approach is to compare historical rates of return against the return of similarly risky investments as predicted by various asset pricing models. In illiquid asset classes like private equity (PE), returns are not observed at regular intervals, and valuations are stale and potentially biased, complicating this task. The most commonly used performance measures in PE, such as the internal rate of return (IRR) or cash multiples, therefore rely primarily on fund cash flows (and do so exclusively in the case of fully liquidated funds). However, a high IRR or cash multiple could simply indicate a very risky investment, rather than an investment that raises a portfolio's Sharpe ratio upon inclusion. These limitations led to the development of performance metrics that use a stochastic discount factor (SDF) to discount cash flows. The key feature is the use of the cumulative rate of return of some portfolio (or a "levered" version of that portfolio) for discounting. For instance, the Kaplan and Schoar (2005) public market equivalent (PME) metric uses the public stock market return to discount cash flows back to the fund's inception date.

The SDF approach to performance evaluation has several theoretical advantages, as we discuss in detail in section 1.1.2. For example, if the cash flows of the PE fund can be replicated by some dynamic trading strategy in publicly traded assets, then the present value of the private equity fund's cash flows has a value of zero for all SDFs that price these same public assets. However, in the more realistic case where such a replication is not possible, there could be multiple SDFs that price all the publicly traded assets but assign different values to the "unspanned" risks of private equity. This problem becomes particularly important if investors hold different optimal portfolios (e.g., due to differences in risk-aversion, or to non-participation in certain markets), as they are likely to assign very different values to these unspanned risks.

In this paper we propose a pragmatic approach to determining the SDF that is to be used for discounting. The key idea is to use a given investor's own portfolio return to form the stochastic discount factor. We provide several theoretical arguments why using the investor's own return to form the SDF has some appealing properties, even if the financial market is incomplete. Specifically, we propose two variants of our measure. The first measure, the "investor portfolio equivalent" (IPE), is essentially the same measure as the Kaplan-Schoar PME, except that we use the investor's own portfolio return rather than the return of the market portfolio when discounting the cash flows of the private equity fund. We show that when this measure has a positive value, it indicates that an investor could raise the (logarithmic) growth rate of her investment portfolio by allocating a marginal dollar towards private equity. While simple to compute and easy to interpret, the IPE measure has the disadvantage that it could depend on different risk attitudes, which are reflected in an investor's optimal mix of stocks and bonds. For this reason, we also consider a generalized version (GIPE), which takes into account different investor risk aversions. In effect, the GIPE uses an investor-specific, levered version of the investor's portfolio to form the SDF. An attractive property of the GIPE is that it is zero, if a private equity investment is just a levered version of the return on public equity. Both the IPE and GIPE measure allow the computation of an "alpha". This alpha can be interpreted as the component of the internal rate of return of the investment that is not due to risk, but rather due to a meaningful expansion of the investment opportunity set for a specific investor.

We also develop some simple diagnostics to determine whether simple, *long-only* public market strategies would be able to produce the same gains as a given private equity investment (e.g., a value strategy for buyout funds). The comparison to long-only investment alternatives is important, because many large investors in private equity are constrained or altogether prohibited from shorting.

To illustrate our approach, we compute the IPE and GIPE values of several private equity strategies from the perspective of U.S. public pension plans.[1] We focus on public pension plans for several reasons. First, as Figure 1 illustrates, public pension plans have been very active investors in private equity, and many plans now have double-digit percentage allocations to this asset class. Second, underfunding and corporate governance concerns about pension plans make it especially important to risk-adjust their investments. Third, there is large heterogeneity in portfolios. For example, in 2018, the Public School Employees Retirement System of Pennsylvania had an allocation to public equities of 20.9%, compared to 70.5% for the Employee Retirement System of Alabama. The cross-sectional standard deviation for the year was 10.9 percentage points (the average was 47.6%). Dispersion in portfolio holdings was similar in other years. Fourth, data on pension plan investment returns are readily available, which helps in forming our SDFs.

Our results can be summarized as follows.

[1] For disambiguation, we use the nomenclature pension *plans* (rather than pension funds) to distinguish them from private equity *funds*. Also, when we say private equity we mean all forms of PE, including but not limited to VC, buyout, and real estate funds.

Figure 1: Average Public Pension Plan Allocation to Private Equity

Average portfolio allocation to private equity (PE) by 179 U.S. public pension plans from 1995 to 2018. Source: Comprehensive Annual Financial Report (CAFR) data from the Center for Retirement Research at Boston College, the Center for State and Local Government Excellence, and the National Association of State Retirement Administrators (available at `https://publicplansdata.org`), and authors' own work.

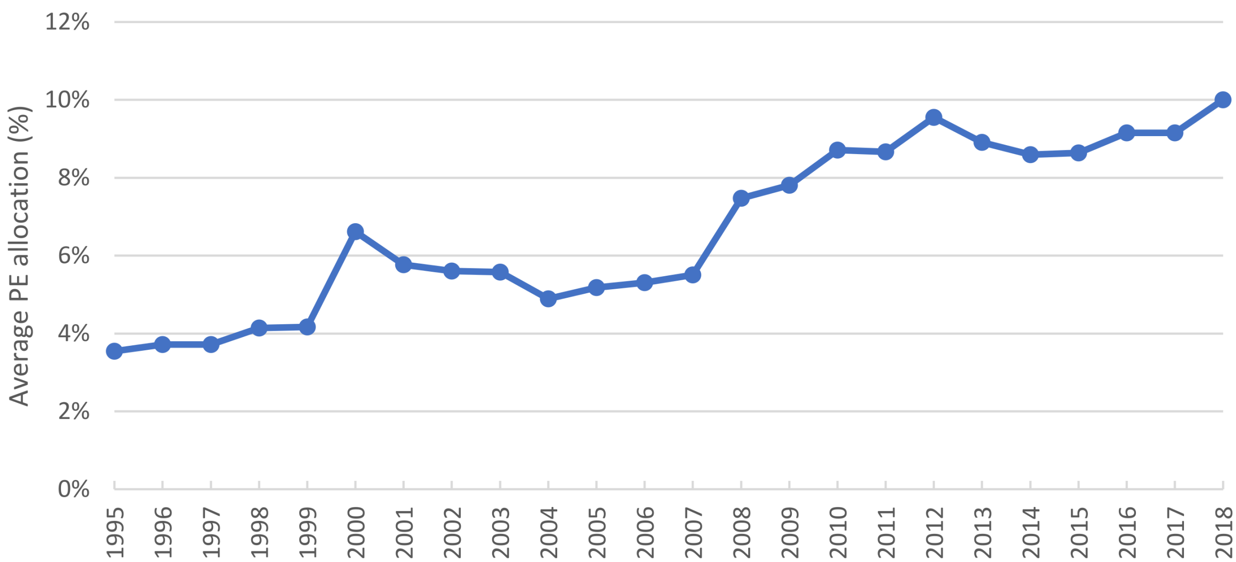

First, for our sample period that spans from 1995 to 2018, pension plans appear to have allocated their portfolios to long-only private equity optimally, in the sense that there was no benefit of investing more or less in a representative PE fund. Phrased more succinctly, the GIPE measure (averaged across pension plans) is approximately zero. The main exception to this statement is the performance of buyout funds, which has a positive GIPE. Additionally, the positive GIPE values of buyout funds do not appear a consequence of buyout having a "value exposure", nor do they seem to be the result of under-allocation to (publicly traded) value strategies by pension plans. While the average VC fund underperformed, this is driven by the poor performance of small-growth firms overall (including publicly traded firms).

Second, we do not find evidence that pension plans had market timing skill. There is some evidence of a "selectivity" skill in the sense that the PE funds that were selected by public pension plans outperform the average PE fund of the same vintage year. However, it appears that this is not the result of a genuine ability to select the better-performing PE investments. Rather it appears that the private equity funds select their investors and give them preferential access to the better PE investments. Indeed, when we consider first time funds (which are likely to be less selective about their clients), or when we confine attention to the universe of private equity investments that are continuations of ongoing relationships with a given pension plan, the "selectivity" skill disappears.

Third, we decompose the IRR of a pension plan's private equity investments into a risk-compensation and an alpha component and examine whether plan characteristics that correlate with a high IRR do so because of the riskiness of these PE investments, or because of a meaningful expansion of the investment opportunity set (alpha). To begin, we document a large degree of dispersion in (G)IPE across pension plans, implying that different pension plans obtain very different returns on their private equity portfolios. This contrasts with investments in publicly traded equities, which tend to be quite correlated across pension plans. We find that underfunded pension plans, and plans that have a larger fraction of state officials who serve ex-officio, and members of the public that were appointed by a government official, take more risk but earn lower risk-adjusted returns in their PE investments. Investments in PE funds that are located in the pension plan's state do not differ in risk from out-of-state funds, but they earn lower risk-adjusted returns. These results are suggestive of agency problems, such as "gambling for resurrection" and political influence, playing an important role in the selection of PE investments by pension plans.

The paper is related to several strands of literature. First, our SDF-based metrics are related to the public market equivalent measure of Kaplan and Schoar (2005) and, especially, the generalized PME of Korteweg and Nagel (2016). The key difference is that we use investor-specific SDFs, which allows for the possibility of heterogeneous investors with different risk assessments of the same investment opportunity.

The second related strand is the literature on limited partner performance in private equity. Prior work has shown that different types of limited partners experience different performance (e.g., Lerner et al., 2007; Sensoy et al., 2014; Cavagnaro et al., 2019). Korteweg and Westerfield (2022) survey this literature. For pension plans specifically, several papers consider the gambling for resurrection (also known as risk-shifting) concern for underfunded plans (e.g., Rauh, 2009; Pennacchi and Rastad, 2011; Mohan and Zhang, 2014; Bradley et al., 2016; Andonov et al., 2017; Myers, 2022). This literature studies whether underfunding changes the share of their portfolios that pension plans allocate to risky assets, and how it affects overall plan performance and the total return on their private equity investments. Our contribution is to separate risk and excess return within PE investments, which allows us to examine whether the difference in PE performance is simply due to a difference in risk-taking. Similarly, the literature that considers home bias (e.g., Lerner et al., 2007; Hochberg and Rauh, 2013; Bradley et al., 2016; Andonov et al., 2018) and board structure (Andonov et al., 2017, 2018) in public pension investing only considers broad plan performance (not specific to PE), or only total, not risk-adjusted PE performance.

The paper is organized as follows. Section 1 develops the theory behind our performance measures. Section 2 describes the pension plan and private equity fund data. Section 3 presents the empirical results, and Section 4 concludes.

1 Theoretical framework

1.1 Model setup

1.1.1 Definitions

Private equity funds are structured as limited partnerships. Investors, such as pension funds, are the limited partners (LPs) of the fund. The LPs are pure capital providers; they have no control over which deals are invested or exited. Capital is committed at fundraising but not immediately transferred to the fund. Instead, the fund manager (general partner, or GP) searches for deals, and "calls" capital from the LPs when they have identified an investment. Money from the sale of investments is distributed to the LPs, after fees to the GP. The fund has a limited life time to invest its committed capital and realize exits (typically 10 years, with limited extension options in case of unexited investments). When all portfolio investments have been sold, the fund is liquidated.[2]

From the perspective of an LP, committing to a PE fund may be viewed as producing a sequence of future net-of-fee fund cash flows $C = \{C_{t_1}, .., C_{t_K}\}$. Both the magnitude and timing of cash flows are random. We place no restriction on their sign, though typically, the first cash flow(s) are capital calls, which are negative flows for the LP, followed by positive cash distribution(s) later in the fund's life.

The fund's "net asset position", A_t, is the economic value of the assets that help finance the

[2] For an in-depth description of the PE fund industry, see, for example, Korteweg and Westerfield (2022).

cash flow C. It is defined as a continuous-time stochastic process with the properties[3]

$$A_{t_k^+} = A_{t_k} - C_{t_k} \text{ and } A_{t_0} = A_{t_K^+} = 0. \tag{1}$$

In other words, the net asset position starts at zero at fund inception, and increases every time there is a capital call ($C_{t_k} < 0$) and decreases whenever there is a distribution ($C_{t_k} > 0$) . After the final cash flow (i.e., at fund liquidation), the asset value is zero again. Between capital calls and distributions, the net asset position evolves randomly according to some diffusion process

$$\frac{dA_t}{A_t} = \mu_A dt + \sigma_A dB_t,$$

where μ_A is its expected return, B_t is a $d-$dimensional Brownian motion and σ_A is a ($d-$dimensional) row vector of exposures, governing the risk profile of $\frac{dA_t}{A_t}$.

We assume throughout that while a true process A_t exists, an econometrician cannot observe it. This is a realistic assumption in PE. While GPs do report quarterly fund net asset values (NAVs), these are subject to staleness and manipulation.[4] Therefore, all our performance measures rely exclusively on the observed cash flows, as is typical in the literature and in practice.

In addition to the cash flow sequence C, investors have access to N other risky investments (public equities, long term bonds, international equity, etc.) and a risk-free security yielding the instantaneous interest rate r_t. Letting S_{it} denote the price of the risky security i and $D_{i,t}$ its dividend, we assume that the total return $dR_{i,t}$ of security i follows the dynamics

$$dR_{i,t} \equiv \frac{dS_{i,t} + D_{i,t}dt}{S_{i,t}} = \mu_{i,t}dt + \sigma_{i,t}dB_t, \text{ for } i = 1, 2, .., N \tag{2}$$

where $\mu_{i,t}$ is the expected return process, $\sigma_{i,t}$ is a $d-$dimensional row vector of exposures to the $d-$dimensional Brownian motion. We assume that $d \geq N$ to capture the possibility that the market is incomplete.

[3]We set up the model in continuous time for tractability and because cash flows can happen at any time within a year.

[4]As of 2007, Accounting Standards Code (ASC) Topic 820 (formerly known as FAS 157) requires the disclosure of fair values, but there is no market to mark PE investments to. GPs and auditors usually rely on the pricing of recently traded comparable assets, but this is a subjective process. For empirical evidence of staleness and manipulation, see Phalippou and Gottschalg (2009); Jenkinson et al. (2013); Barber and Yasuda (2017); Chakraborty and Ewens (2018); Brown et al. (2019); Jenkinson et al. (2020).

Stacking the N equations (2) allows us to express the dynamics of all investible assets compactly as

$$\underbrace{dR_t}_{N\times 1} = \underbrace{\mu_t}_{N\times 1} dt + \underbrace{\sigma_t}_{N\times d} \underbrace{dB_t}_{d\times 1}.$$

1.1.2 Motivation: SDF-based performance evaluation

A common approach to evaluating whether an investment in the cash-flow stream C is an attractive investment opportunity is to use a stochastic discount factor (SDF) to discount the sequence of payments C and test whether the SDF-adjusted present value of C is zero or not.

A stochastic discount factor (SDF) H_t is defined as a process such that $H_t S_{i,t} + \int^t H_s D_{i,s} ds$ is a martingale for all assets i, so that the price of any security $S_{i,t}$ can be written in "present-value" form as

$$S_{i,t} = E_t \left(\int_t^T \left(\frac{H_s}{H_t} \right) D_{i,s} ds \right) + E_t \left(\frac{H_T}{H_t} S_{i,T} \right) \tag{3}$$

for any time $T > t$.

The justification for the SDF-approach is a replication argument. Specifically, suppose that there exists a portfolio ϕ_t of the N risky assets such that the cash flow sequence $C_{t_0}, C_{t_1}, ..., C_{t_K}$ is the result of a self-financing strategy in publicly traded assets. Specifically, suppose that the net asset position evolves as

$$\frac{dA_t}{A_t} = r_t dt + \phi_t' \left(dR_t - r_t 1_{N\times 1} dt \right), \tag{4}$$

where $1_{N\times 1}$ is a column vector of ones. Implicitly we defined the position in the risk-free asset as $1 - 1'_{N\times 1}\phi_t$ to ensure that the portfolio is self-financing. An implication of (4) is that the cash flows $C_{t_0}, C_{t_1}, ..., C_{t_K}$ can be exactly replicated with some dynamic, self-financing strategy ϕ_t in existing assets.

Letting H_t denote any SDF, a standard argument implies that[5]

$$0 = A_{t_0} = E_0 \sum_{k=0,..,K} \left(\frac{H_{t_k}}{H_{t_0}} \right) C_{t_k}. \tag{5}$$

[5] Since $H_t S_{i,t} + \int^t H_s D_{s,i} ds$ is a martingale, an application of Ito's Lemma together with equation (4) shows that $H_t A_t$ is a a martingale. Using (1) implies that

$$H_{t_k^+} A_{t_k^+} = E_{t_k} \left(H_{t_{k+1}} A_{t_{k+1}} \right) = E_{t_k} \left(H_{t_{k+1}} A_{t_{k+1}^+} \right) + E_{t_k} \left(H_{t_{k+1}} C_{t_{k+1}} \right).$$

Iterating forward, using the law of the iterated expectation and noting that $A_{t_0} = A_{t_K^+} = 0$ implies (5).

Equation (5) is a net present value (NPV) criterion. It states that if the return of the alternative investment can be replicated by trading in existing assets, then –by construction– the SDF-adjusted present value of C_{t_k} should be zero. Otherwise, there would be an arbitrage opportunity. Importantly, equation (5) holds for any possible stochastic discount factor H_t.

Matters become more complicated in the more realistic situation where the cash flow sequence sequence C_{t_k} cannot be replicated by some dynamic trading strategy in the publicly traded assets. In this case, a regression of $\frac{dA_t}{A_t}$ on the returns of all publicly traded assets allows us to decompose $\frac{dA_t}{A_t}$ as

$$\frac{dA_t}{A_t} = r_t dt + \phi_t' \left(dR_t - r_t 1_{N\times 1} dt \right) + \left(d\widetilde{R}_t - r_t dt \right), \tag{6}$$

where $d\widetilde{R}_t - r_t dt$ can be interpreted as the excess return of a fictitious asset, whose returns are orthogonal to all publicly traded asset returns.[6]

Now choosing an SDF H_t that satisfies equation (3) for all assets $i = 1..N$ no longer implies that the net present value $E_0 \sum_{k=0,,K} \left(\frac{H_{t_k}}{H_{t_0}} \right) C_{t_k}$ equals zero. The sign and magnitude of this net present value will depend on the choice of the SDF. The issue is that the usual "no-arbitrage" replication arguments cannot provide a unique and unambiguous way of pricing the unreplicable return component $d\widetilde{R}_t$.

One could argue that there may be some good reasons to simply assume that the unspanned components on the right-hand side of (6) have zero (or approximately zero) risk compensation under any stochastic discount factor H_t. Indeed, if the un-spanned risk , $d\widetilde{R}_t$, is independent across alternative investments, then an APT-like asymptotic arbitrage argument would imply that a well-diversified portfolio of such investments would be spanned by traded assets. Leaving aside the issue of whether there is a common factor in the un-spanned components, $d\widetilde{R}_t$, there is an additional problem with the idea of asymptotic arbitrage for alternative investments: The asymptotic arguments behind the APT were developed with easily tradable and divisible assets in mind, making diversification relatively easy. Alternative assets are not publicly traded, require minimum commitments, negotiations between general and limited partners, etc, so that it appears stretched to

[6] The evolution of A_t can be decomposed as

$$\frac{dA_t}{A_t} = a_t dt + \phi_t' \left(dR_t - 1_N \times r_t dt \right) + d\eta_t,$$

where $d\eta_t$ is a residual process orthogonal to the excess returns of all traded assets. Letting $d\widetilde{R}_t = d\eta_t + (a_t - r_t)\, dt$ yields (6).

apply such asymptotic arguments. More importantly, if it was simple to create such a diversified portfolio of these investments, we should observe that investors' returns in their alternative-assets positions should be highly correlated across investors, since investors should only retain the factor risk, but not idiosyncratic risk. As we show in the empirical section, this is a far cry from reality, especially for alternative investments.

1.1.3 Investor-specific SDFs

In light of these difficulties, in this paper we propose a relatively simple modification to the NPV criterion. The idea is to use different stochastic discount factors for different investors. The motivation is that in an incomplete market different investors may value the unspanned return components differently. Such incompleteness may arise, for instance if different investors invest in different assets, so that the returns on the investors' portfolios may differ.

To fix ideas, in this section we envisage an investor and let W_t denote the cumulative, continuously-compounded return of her portfolio, so that the dynamics of her investments evolve according to

$$\frac{dW_t}{W_t} = r_t dt + w_t' \left(dR_t - r_t 1_{N_j} dt \right) \tag{7}$$

where r_t is the interest rate and w_t is an $N_j \times 1$ vector of portfolio weights. We allow N_j to be investor specific and reflect the N_j assets that investor j is actively investing in.

We next present the investor with the opportunity to invest some small fraction ε of her portfolio in the cash flow $C = \{C_{t_0}, C_{t_1}...\}$. We fix some objective for the investor and ask whether this investor could improve her objective function from undertaking the investment in C.

For expositional reasons, and in order to relate our performance measures to the ones conventionally used in the literature, we start by considering an investor who is interested in maximizing the expected logarithmic growth rate of her investments:

$$V \equiv E_0 \log W_T - \log W_{t_0}, \tag{8}$$

where T is some distant time, larger than t_K. Specifically, we assume that the investor maximizes (8) over $w_t \in \mathcal{W}$, where the set $\mathcal{W}$ captures any constraints placed on the investor (portfolio limits, shorting constraints, borrowing constraints, etc.). The following proposition is an implication of the envelope theorem (the proof is in the appendix).

Proposition 1 *Suppose that an investor maximizes (8) over $w_t \in \mathcal{W}$. Let εC denote an investment of $\varepsilon > 0$ in the cash flow process C, and define*

$$IPE \equiv E_0 \sum_{k=1...K} \left(\frac{W_0}{W_{t_k}} \right) C_{t_k}, \tag{9}$$

where $\frac{W_0}{W_{t_k}}$ is the inverse of the cumulative return on the investor's portfolio. Then we have that $\frac{dV}{d\varepsilon} = V_W \left(W_{t_0} \right) \times IPE$, where $V_W \left(W_{t_0} \right)$ is the marginal valuation of wealth at time t_0.

Corollary 2 *Suppose that the investor does not maximize (8) but instead assume that the addition of the cash flow C does not change the investor's portfolio w_t. Then the conclusion of 1 remains valid, that is $\frac{dV}{d\varepsilon} = IPE$.*

Proposition 1 says that the change in the investor's objective per $\varepsilon > 0$ dollars of investment is given by the product of the IPE times the investor's marginal valuation of wealth at time t_0, $V_W \left(W_{t_0} \right)$. This means that the IPE is measured in time-t_0 dollars. It is the amount that keeps the investor indifferent between committing to the cash flow C or receiving IPE dollars in a lum sum fashion at time t_0. Obviously, an investor finds it meaningful to commit funds to the project C, when and only when $IPE > 0$.

Corollary 2 states that this conclusion does not depend on wheteher the investor chooses her portfolio to maximize (8) or not, as long as the introduction of the cash flow C does not change the investor's portfolio w_t.

Unlike other approaches to computing NPV-type criteria, the IPE is investor-specific. The IPE resembles the Korteweg-Nagel PME, except that the return on the investor's portfolio is used in place of the stock market return.

The theoretical advantage of the IPE is that it is no longer necessary to make assumptions about the pricing of the unspanned return components of the alternative investment. Rather than assuming that there is a "one-size-fits-all" pricing of the un-spanned return components, the IPE uses each investor's own investment return to price the unspanned components.

Under the additional assumption that (a) the investor's objective function (8) is correctly specified and (b) portfolio constraints are not binding, then $\left(\frac{W_t}{W_{t_0}}\right)^{-1}$ is also a valid stochastic discount factor for all the assets that comprise the investor's portfolio.[7] In light of the discussion of the previous section, this means that if there exists a trading strategy that can replicate the cash flow C using the N_j assets that the investor is already investing in, then the IPE will be zero and the cash flow C does not present a meaningful investment opportunity.

In section 1.1.4 we relax assumption (a) by allowing investors' risk aversion coefficients to differ. In section 1.1.5 we discuss the implications of portfolio constraints (assumption (b)).

1.1.4 Generalized IPE

In this section, we consider a simple generalization of our IPE, which applies in situations where the investor maximizes a CRRA objective function such as

$$V^{(\gamma)} = E_{t_0} \frac{(W_T)^{1-\gamma}}{1-\gamma}, \tag{10}$$

for $\gamma > 0$. We assume that γ may differ by investor.

With a straightforward adaptation of the arguments of proposition (1), we can show the following result

Proposition 3 *Consider the same assumptions as in proposition 1, except that the investor maximizes (10). Assume that $\mu_t = \mu, \sigma_t = \sigma, r_t = r$ and constant across time and define*

$$GIPE^{(\gamma)} \equiv E_{t_0} \sum_{k=t_0,t_1...t_K} e^{-r(t_k-t_0)} \frac{\left(\frac{W_{t_k}}{W_{t_0}}\right)^{-\gamma}}{E_0\left(\frac{W_{t_k}}{W_{t_0}}\right)^{-\gamma}} C_{t_k}. \tag{11}$$

Then $\frac{dV^{(\gamma)}}{d\varepsilon} = GIPE.$

[7] The fact that the return of the expected-logarithmic-growth-maximizing portfolio is a valid SDF is a well-known result. We provide a sketch of the argument for completeness. The optimal portfolio for the investor who maximizes (8) is $w_t = (\sigma_t \sigma_t')^{-1}\left(\mu_t - r_t 1_{\{N_j \times 1\}}\right)$. Substituting into (7) and solving the resulting linear stochastic differential equation for W_t gives

$$\left(\frac{W_t}{W_{t_0}}\right)^{-1} = \exp\left\{-\int r_u du - \frac{1}{2} w_t' (\sigma_t \sigma_t') w_t - w_t' \sigma_t dB_t\right\}$$

Letting $\widehat{H}_t = \left(\frac{W_t}{W_{t_0}}\right)^{-1}$ and using Ito's lemma to compute $d\left(\widehat{H}_t S_{i,t} + \int^t \widehat{H}_s D_{i,s} ds\right)$ shows that $\widehat{H}_t S_{i,t} + \int^t \widehat{H}_s D_{i,s} ds$ has zero drift. Assuming σ_t, μ_t, r_t are bounded (or more generally integrable) implies that $\widehat{H}_t S_{i,t} + \int^t \widehat{H}_s D_{i,s} ds$ is a martingale.

To operationalize the concept of the GIPE, one has to estimate a separate coefficient γ for each investor. In the empirical implementation, we assume an interior optimum between stocks and the risk free assets and infer γ from the "Euler" equation

$$E_t\left(\left(\frac{W_{t+1}}{W_t}\right)^{-\gamma}\left(R_{t+1}^S - R_{t+1}^f\right)\right) = 0, \tag{12}$$

where $R_{t+1}^S - R_t^f$ is the annual excess return on a broad stock market index.

The following proposition provides an equivalence between GIPE and the familiar "beta-based" approach to performance evaluation.

Proposition 4 *Assume that the investor is maximizing (10) and that no portfolio constraint is binding. Assume that an econometrician could observe* A_t *and let* β_t *denote the regression coefficient from regressing* $\frac{dA_t}{A_t} = \mu_A dt + \sigma_A dB_t$ *on the return of the investor's portfolio, where* μ_A *is the drift (expected return of* $\frac{dA_t}{A_t}$*) and* σ_A *is a row vector of exposures to the d−dimensional Brownian motion. Also, let* $\mu^W - r \equiv w'(\mu - r)$ *denote the expected excess return on the investor's portfolio and* μ_A *the drift (the expected return) of* $\frac{dA_t}{A_t}$*. Then if*

$$\mu_A - r = \beta_t\left(\mu^W - r\right), \tag{13}$$

the GIPE is zero.

Equation (13) shows that a zero GIPE follows from a zero "Jensen's alpha" in a regression of $\frac{dA_t}{A_t}$ on the investor's portfolio return. The advantage of the GIPE is that it does not require knowledge of the true economic value of the alternative investment. Instead it can be computed only by observing the cash flows C.

An additional feature of the GIPE is that different risk aversion coefficients, which would result in different degrees of leverage across investors, do not affect the evaluation of a private equity investment. The easiest way to see this is to note that neither the left-hand side nor the right-hand side of (13) depend on γ.[8]

While the GIPE is not affected by differences in the risk aversion coefficients across investors, it does depend on which assets an investor is actively participating in. (stocks, long term bonds,

[8] The fact that the left-hand side of (13) does not depend on γ follows by inspection. As for the right-hand side,

international investments, other alternative investments, etc.). For instance, adding the cash flow C to the portfolio of an investor who is already heavily invested in alternative-investment-type assets will generally result in a higher value of required risk compensation, $\beta_t\left(\mu^W - r\right)$, when compared to an investor who abstains from alternative investments. Intuitively, the same risk is more "diversifiable" in the second investor's portfolio.

Assuming that an investor's portfolio choice is unconstrained, then $H_t = e^{-rt}\frac{(W_t)^{-\gamma}}{E_0(W_t)^{-\gamma}}$ is a valid stochastic discount factor for the assets that the investor is already participating in. Therefore, if it were possible to replicate the cash flow C with some self-financing trading strategy in the assets that the investor is already participating in, then the GIPE would be zero.

We conclude this section with the following generalization of proposition 3, which applies even if μ_t, r_t, σ_t are time-varying, as long as changes in μ_t, r_t, σ_t are unrelated to the return to A_t:

Proposition 5 *Assume that* μ_t, r_t, σ_t *are functions of some vector of state variables* X_t *that follow some diffusion* $dX_t = \mu_X dt + \sigma_X dB_t$. *Assume also that* $\sigma_X \sigma'_A = 0$, *i.e., innovations to* dX_t *and* $\frac{dA_t}{A_t}$ *are independent. Assume that* $\mu_{A,t} - r_t + \frac{\langle dW_t^{-\gamma}; dA_t \rangle}{W_t^{-\gamma} A_t}$ *has the same sign for all* t *and define*

$$GIPE^{(\gamma)} \equiv E_{t_0} \sum_{k=t_0,t_1 \ldots t_K} e^{-\int_{t_0}^{t_k} r_u du} \frac{\left(W_{t_k}\right)^{-\gamma}}{E_0\left(W_{t_k}\right)^{-\gamma}} C_{t_k}. \tag{14}$$

Then $sign\left(\frac{dV^\gamma}{d\varepsilon}\right) = sign\left(GIPE^{(\gamma)}\right) = sign\left(\mu_{A,t} - r_t + \frac{\langle dW_t^{-\gamma}; dA_t \rangle}{W_t^{-\gamma} A_t}\right)$.

1.1.5 Portfolio constraints and a quasi-replicating portfolio

If investors are subject to portfolio constraints, then an investor's marginal valuation of wealth may fail to price the assets with constrained portfolio weights. We have already implicitly allowed a special case of such constraints in the previous sections by allowing for different investors to actively invest in different assets (this can be viewed as a mandate to maintain a zero weight on

note that an investor's optimal portfolio is proportional to $\frac{1}{\gamma}$,

$$w_t = \frac{1}{\gamma}\left(\sigma\sigma'\right)^{-1}(\mu - r).$$

Accordingly,

$$\beta_t\left(\mu^W - r\right) = \frac{\left(w'\sigma\sigma_A\right)\left(w'(\mu - r)\right)}{w'\sigma\sigma' w},$$

is independent of γ.

the excluded assets).

In the case of constrained portfolio choice, even if the cash flow C could be replicated by some portfolio of existing assets with constrained portfolio weights, the GIPE could be positive. In effect the cash flow C would be expanding the investor's constraint set.

It would seem useful to diagnose whether the attractiveness of the investment C is because of the attractive excess return in its unreplicable return component, or whether it is simply reflecting an investor's constrained portfolio allocation. As a practical matter, it would be particularly interesting to know whether simple, long-only positions in certain assets (that the investor may simply not be participating in) would be able to replicate the cash flows C, since such strategies provide investors with a realistic alternative.

With the goal of developing such a diagnostic, suppose that we consider a null hypothesis whereby $\frac{dA_t}{A_t}$ could be replicated by some benchmark portfolio of the assets $1, ..., N$. Let dR_t^b denote the (instantaneous) return of this portfolio and G_t the cumulative return on investing a dollar in the benchmark portfolio, so that $G_{t_0} = 1$, $\frac{dG_t}{G_t} = dR_t^b$.

Next we design a mimicking cash flow sequence $\widehat{C}$, with cash flows occurring at the same times t_k as the cash flow C_{t_k} and capital calls equal and identically timed to the capital calls of C, that is

$$\widehat{C}_{t_k} \equiv C_{t_k}, \text{ for } C_{t_k} < 0. \tag{15}$$

These capital calls are placed in a mimicking fund and invested at the rate of return dR_t^b.

To construct the distributions of the mimicking fund we start by letting ω_{t_k} denote the fraction of the present values of all distributions that occur at time k, discounted at the benchmark return:

$$\omega_{t_k} \equiv \frac{\frac{C_{t_k}}{G_{t_k}} 1_{\{C_{t_k}>0\}}}{\sum_{k=0..K} \frac{C_{t_k} 1_{\{C_{t_k}>0\}}}{G_{t_k}}} \text{ for } C_{t_k} > 0.$$

Define the distributions of the mimicking fund as

$$\widehat{C}_{t_k} = \widehat{A}_0 \omega_{t_k} G_{t_k}, \text{ where } \widehat{A}_0 \equiv \sum_{k=0..K} \frac{|C_{t_k}|}{G_{t_k}} 1_{\{C_{t_k}<0\}}. \tag{16}$$

We obtain the following proposition

Proposition 6 *a) The distributions (16) are non-negative. b) The cash flows $\widehat{C}_{t_k}$ given by (15) and (16) can be financed by the trading strategy of investing the capital calls at the benchmark return R_t^b and making distributions equal to (16). c) Under the null hypothesis*

$$\frac{dA_t}{A_t} = dR_t^b. \tag{17}$$

that the net asset position associated with the cash flow C grows at the benchmark return, the cash flows of the mimicking fund $\widehat{C}$ are identically equal to the cash flows of the fund C.

Parts a) and b) of Proposition 6 ensure that the distributions from the mimicking fund are positive and that the cash flows of the mimicking fund, $\widehat{C}$, can be financed by investing the capital calls in the benchmark portfolio. The more interesting property is property b), which ensures that under the null hypothesis that the return on the net asset position of the alternative investment, $\frac{dA_t}{A_t}$, is the same as the return of the benchmark portfolio (equation (17)), then the cash flows C and $\widehat{C}$ coincide.

The mimicking cash flow $\widehat{C}$ allows us to perform an exercise similar in spirit to the popular "style" analysis that Sharpe (1988, 1992) introduced for mutual funds. Specifically, we can choose a benchmark portfolio (say a value portfolio for private equity, or a small growth portfolio for venture capital). We can then construct the mimicking cash-flow sequence $\widehat{C}_{t_k}$, and examine its IPE and GIPE.

If the (G)IPE of C and $\widehat{C}$ are both positive but close to each other, this would indicate that the source of the positive (G)IPE is the investor's inability to obtain the benchmark return R^b due to some constraint (or sheer non-participation) in some of the assets that comprise the benchmark portfolio. If the (G)IPE of C is positive, but the (G)IPE of $\widehat{C}$ is zero, then this would indicate that the cash flow C helps raise the investor's objective function beyond what an investment in the benchmark portfolio could.

We conclude this section with two remarks on our construction of the mimcking cash flow $\widehat{C}$.

First, we note that our construction of the mimicking cash flow $\widehat{C}$ differs from the modified PME proposed by Cambridge Associates. Both approaches consider a mimicking fund that can be financed by investing the capital calls in the benchmark portfolio. In both approaches the distributions are positive. But our approach enforces that under the null hypothesis (17) the sequences C and $\widehat{C}$ coincide. More importantly, our approach does not rely on fund-provided NAVs, which may be problematic per the discussion above.

Second, in our construction the (G)IPE of the mimicking fund $\widehat{C}$ could be positive either a) because the investor is constrained in her ability to invest in assets that could replicate C or b) because the manager of the alternative investment possesses superior information. To see the latter possibility consider the following stylized example: Suppose that the manager of the alternative investment simply invests in the broad stock market, which the investor could do on their own. Suppose also that the expected return on the stock market, μ_t ,can have one of two values $\mu^H > \mu^L$ that switch according to some regime-switching process. The manager of the alternative fund has the ability to collect a capital call at the beginning of regime H, invest it in the stock market and liquidate it immediately before the regime is about to switch to L. By contrast the investor does not know which regime the economy is in and holds a constant portfolio in the stock market. In this example, once an econometrician constructs the cash flow $\widehat{C}$ according to our algorithm, she will notice that the cash flow C has a positive GIPE , but there is no difference between the cash flow C and $\widehat{C}$, since assumption (17) holds by construction in our example. In that sense observing a positive (G)IPE difference between the cash flows C and $\widehat{C}$ isolates the manager's skill to access investments that the investor is precluded from. Thus, it provides a diagnostic for "selectivity" but not "timing ability".

The above discussion also shows that the sign of the GIPE of the mimicking cash flow $\widehat{C}$ is of interest in its own right. If the GIPE of the mimicking cash flow $\widehat{C}$ is positive even for a benchmark return that the investor can easily obtain (say a long-only position in the market index), then this would indicate that the manager has timing ability.

1.2 Summary of IPE, GIPE, mimicking fund

Before proceeding with the empirical analysis it is useful to summarize the various information measures we have developed and their appropriate interpretation.

The IPE is perhaps the easiest measure to compute and interpret. A positive IPE shows that a given cash flow C helps increase the expected logarithmic growth rate of the investor's wealth. Remarkably, this interpretation continues to hold no matter how the rest of the investor's portfolio is chosen (Corollary 2).

The main limitation of the IPE is that unless the portfolio of the investor is chosen to maximize the expected logarithmic growth rate of wealth and is interior, the IPE may be positive even if it is possible to replicate the cash flow C with some trading strategy in the assets that the investor is already investing in.

The GIPE nests the IPE as a special case and allows investors to have different risk aversions (and hence investment objectives), which may result in different leverage choices across the investors. This benefit comes at the cost of having to estimate a risk aversion parameter.

Another possibility to diagnose whether an investment could be replicated by some simple trading strategy – be it in assets that the investor is already participating in, or assets that the investor abstains from – is to construct a mimicking cash flow $\widehat{C}$ and compute the IPE and GIPE of $C - \widehat{C}$. A zero (positive) IPE of $C - \widehat{C}$ indicates that the cash flow C does not (does) help raise the expected logarithmic growth rate of wealth beyond what investing in the benchmark portfolio would. Simiarly, a zero GIPE indicates that the same conclusion holds even after taking into account the leverage differences that result from heterogeneous investor risk preferences. A positive (G)IPE of $\widehat{C}$ using a benchmark portfolio that the investor already has access to allows one to diagnose timing ability. Similarly, a positive (G)IPE of $C - \widehat{C}$ using a benchmark portfolio that the investor could have acess to, (but may be constrained from participating) indicates that the cash flow C gives the investor access to assets that are not part of the benchmark portfolio, and the benefit of the cash flow C is not merely a reflection of constrained portfolio choice.

2 Data

We use data on pension plans and private equity funds to compute (G)IPE metrics. We describe each data source in turn.

2.1 Pension plans

We use pension plan data collected from Comprehensive Annual Financial Reports (CAFRs) of U.S. defined benefit public pension plans. These reports contain balance sheet and income statement information, returns, valuations, actuarial data, and other key statistics, audited to conform to Government Accounting Standards Board (GASB) reporting requirements.

We start with a CAFR data set developed and maintained by a collaboration of the Center for Retirement Research at Boston College, the Center for State and Local Government Excellence, and the National Association of State Retirement Administrators.[9] The data include 179 state and local pension plans for the years 2001 to 2018, and cover 95% of U.S. public pension membership and assets. We extend coverage by hand-collecting additional CAFRs going back to 1995 (few CAFRs are available online before 1995).

Next, we consolidate pension plans whose assets are jointly owned and managed, since their reported returns are virtually identical. For example, Maine's local employee plan that covers participating districts and its state employee and teacher plan are both managed by the Maine Public Employees' Retirement System (PERS), even though they are treated separately for accounting purposes and publish separate CAFRs. Maine PERS is also the investor listed in the private equity fund commitment data described below. After consolidation, this leaves 142 plans.

Our final filter is to drop four pension plans because they are missing a return observation for one or more years. We need a continuous time series of pension returns to discount private equity fund cash flows and calculate the (G)IPE performance measures. This leaves us with a final sample of 138 pension plans.

Panel A of Table 1 reports descriptive statistics of the pension plans. The first four columns show characteristics across all 138 plans. The most common plan covers state and/or local employees (96 plans, or 70% of the sample), followed by teachers (41 plans, 30%) and police or fire personnel (29 plans, 21%). Note that these categories are not mutually exclusive due to the merging of some jointly managed plans that cover multiple employee types. Most plans are administered at the state level (78 plans, 57%), with the remainder administered locally (either by a county, city, or school district). About two thirds cover multiple employers, the vast majority with a cost-sharing agreement that pools the employers' pension assets and obligations.

The size distribution of pension plans is highly skewed. Across all plan-years, the median assets under management (AUM) is $7.81 billion. The average, $20.45 billion, is pulled upwards by a few very large plans, such as CalPERS, which had $354 billion in AUM in 2018. Most plans are underfunded on an actuarial basis, with the average (median) funding rate across plan-years equal to 79% (80%).

[9] This data set is freely downloadable from: `https://publicplansdata.org`.

Table 1: Descriptive Statistics

Panel A shows descriptive statistics for public pension plans between 1995 and 2018. The first four columns show statistics for the full sample of 138 pension plans, the next four columns show the statistics for the sample of 103 plans for which we observe at least one private equity (PE) fund commitment in the Preqin data, and the column labeled "Diff. test p" shows the p-value for the test of equal proportions or means across the 103-plan subsample with commitment data and the 35-plan subsample without commitment data (for panel variables, *Assets*, *Funded ratio*, and *Annual return*, the tests include year fixed effects and double-cluster residuals by year and pension plan state). Numbers in the "%" column refer to the percentage of plans that belong to the category described in that row. The covered employee categories (*State/local employees*, *Teachers*, *Police/fire*) are not mutually exclusive. *State-administered plan* refers to a plan administered at the state level (rather than at the county, city, or school level). A *Multiple employer plan* covers employees for more than one employer, and *Cost sharing across employers* is the percentage of multiple employer plans in which employers are jointly responsible for each others' pension liabilities. *Assets* is the market value of pension assets in billions of U.S. dollars. *Funded ratio* is the ratio of actuarial assets to liabilities under GASB 25 accounting standards. *Annual return* is the one-year return on the pension plan (in %). Statistics for these three panel variables are calculated across all plan-years. *PE fund commitments per plan* is the plan's number of observed commitments to PE funds in the Preqin data over the sample. Panel B shows descriptive statistics for 1,303 North American private equity funds raised between 1995 and 2013, by strategy (venture capital (VC), buyout, and real estate). *Number of GPs* is the number of unique fund managers. *Fund size* is the fund's committed capital in millions of U.S. dollars. *Percentage of funds liquidated* is the percentage of funds that are fully liquidated (*100% liquidated*) or that have less than 5% of committed capital left in residual net asset value (*95% liquidated*). *Fund effective years* is the number of years between the first and the last observed cash flow for the fund, where cash flows are observed until the end of June 2018. *IRR* is the internal rate of return, *TVPI* the total value to paid-in capital, and *PME(KS)* and *PME(KN)* are the Public Market Equivalent of Kaplan and Schoar (2005) and Korteweg and Nagel (2016). *Size-weighted* is the NAV-weighted average performance metric. *Funds with matched LP data* is the number of funds for which we observe at least one public pension plan commitment in our data set. *Number of matched LPs / fund* is the number of observed pension plans commitments in a given fund in our data.

Panel A: Pension plans.

	Full Sample (N = 138)				Commitment Sample (N = 103)				Diff.
	%	Mean	Median	St.Dev.	%	Mean	Median	St.Dev.	Test p
Plans covering									
State/local employees	69.57				69.90				0.882
Teachers	29.71				29.13				0.797
Police/fire	21.01				18.45				0.204
State-administered plan	56.52				61.17				0.059
Inception year		1943.79	1944	18.35		1942.70	1942	18.29	0.223
Multiple employer plan, of which:	64.49				68.93				0.062
Cost sharing across employers	93.26				91.55				0.133
Assets (market value, $b)		20.45	7.81	34.48		19.59	8.90	33.19	0.537
Funded ratio (actuarial)		0.79	0.80	0.17		0.78	0.79	0.17	0.113
Annual return (%)		8.11	10.60	10.54		8.11	10.70	10.41	0.081
PE fund commitments per plan						52.51	27.00	60.38	N/A

Panel B: Private equity funds.

	VC funds (N = 527)			Buyout funds (N = 527)			Real Estate funds (N = 249)		
	Mean	Median	St.Dev.	Mean	Median	St.Dev.	Mean	Median	St.Dev.
Number of GPs	261			256			134		
Funds per GP	2.02	2.00	1.37	2.06	2.00	1.22	1.86	1.00	1.33
Fund size ($m)	348.19	250.00	356.60	1,588.85	700.00	2,634.58	929.05	535.00	1,429.91
Percentage of funds liquidated:									
100% liquidated	36.81			34.72			27.71		
95% liquidated	43.64			42.88			36.14		
Fund effective years	12.22	12.26	4.30	11.13	11.13	4.14	8.33	7.51	3.55
IRR (%)	8.27	4.48	36.74	13.66	13.39	14.89	12.46	12.26	11.55
Size-weighted	7.36		28.06	14.58		11.46	12.80		12.57
TVPI	1.54	1.22	1.85	1.67	1.63	0.65	1.45	1.44	0.40
Size-weighted	1.45		1.37	1.69		0.47	1.46		0.44
PME(KS)	1.00	0.79	1.22	1.18	1.13	0.49	0.98	1.01	0.31
Size-weighted	0.94		0.92	1.17		0.37	0.99		0.33
PME(KN)	-0.03	-0.17	1.01	0.16	0.11	0.42	-0.03	0.00	0.30
Size-weighted	-0.07		0.78	0.15		0.34	-0.02		0.33
Funds with matched LP data	465			488			211		
Number of matched LPs / fund	3.43	2.00	2.90	5.90	4.00	5.88	4.45	3.00	4.25

The average (median) annual reported return across plan-years is 8.11% (10.60%), with a standard deviation of 10.54%. Most of the variance is coming from the time-series, but there is also an economically meaningful degree of cross-sectional dispersion. This can be seen in Panel A of Figure 2, which shows the time series of the average return across plans, as well as the 10th and 90th percentiles of plan returns.

For a subset of plans we have data on their private equity fund commitments (described in detail in the next section). Panel A of Table 1 shows the descriptive statistics for the 103 pension plans that we can match to at least one investment in a PE fund (we call this the "commitment sample"). The commitment sample is statistically indistinguishable from the subsample of 35 plans without commitment data on most dimensions, as shown by the p-values of the difference in proportions and means tests in the final column of Panel A. The exceptions are that the commitment sample has a higher proportion of state-administered and multiple-employer plans, and its mean annual return is different (all three are significant at the 10% level). The latter result appears surprising, given that Table 1 shows the same average return for the full sample and the commitment sample, when computed across all plan-years. The explanation is that the test controls for year fixed effects, revealing a difference in mean returns when measured *within the same year*.[10] However, the difference is economically small: the commitment sample has only a 30 basis point lower average return (regression not reported for brevity). The small magnitude of the difference is confirmed visually in Panel B of Figure 2, which plots the average return for both the commitment sample and the subsample of plans without commitment data.

Over the period from 1995 to 2013, the commitment sample plans made an average (median) of 52.51 (27) commitments to PE funds in our data set (we drop PE funds with vintages post-2013 for reasons explained below). Similar to the size distribution, the number of commitments is highly skewed with a long right tail.

2.2 Private equity funds

Our private equity data is sourced from Preqin, and contains fund capital calls and distributions net of fees paid to the GPs, as well as quarterly net asset values (NAVs) for a large number of

[10] We run the same test (with year fixed effects) for all panel variables in Table 1: AUM, funded ratio, and annual returns. All other variables are measured only in the cross-section.

Figure 2: Pension Plan Returns

Panel A shows the time series of the average one-year return across pension plans, for the full sample of 138 plans described in Table 1. The shaded area represents the region between the 10th and the 90th percentile of plan returns. Panel B graphs the time series of average one-year plan returns for the subsample of 103 plans with at least one observed commitment to a PE fund (the solid line; this is the "commitment sample" described in Table 1) and the subsample of 35 plans without observed commitment data (the dashed line, labeled "No commitment sample").

Panel A: Pension plan returns for the full sample.

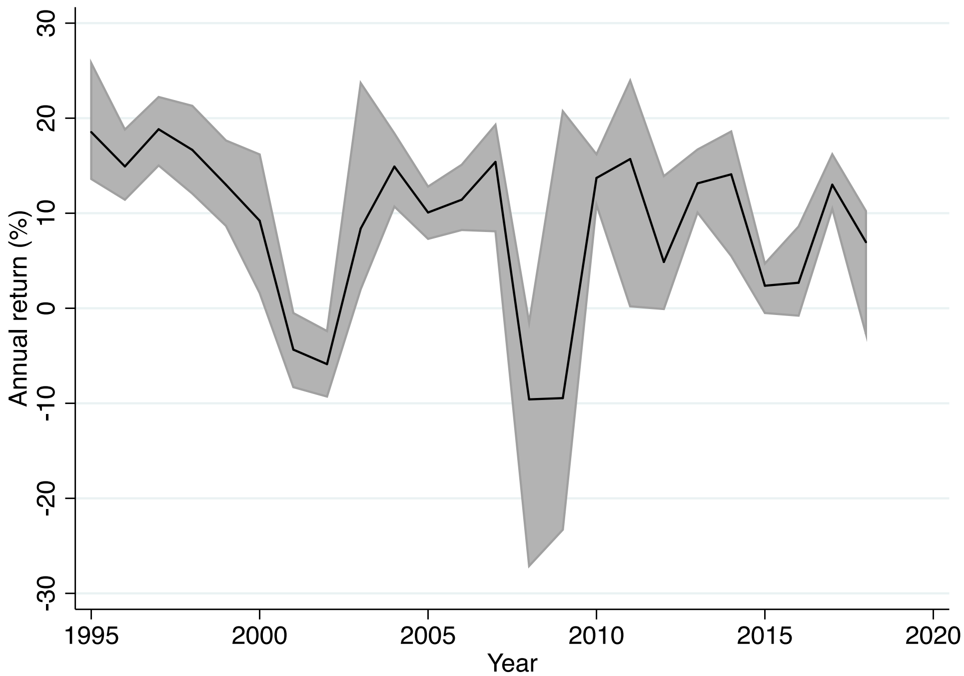

Panel B: Average pension plan returns for subsamples with and without PE fund commitment data.

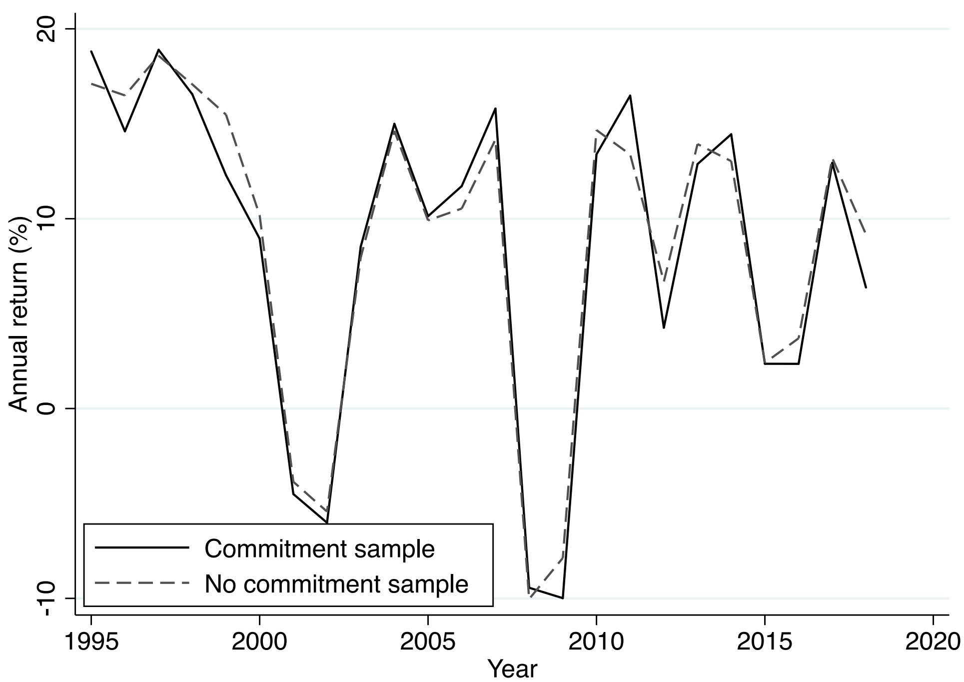

PE funds. Following the literature, we limit the sample to North American funds with at least $5 million in committed capital. We focus on the three main strategies in PE; venture capital (VC), buyout, and real estate. Given the data limitations for pension plans described above, we only include funds raised since 1995, and we use all cash flows until the end of June 2018.[11] Our final filter drops fund vintages after 2013, so we observe at least 5 years of cash flow data for each fund. The final sample contains 1,303 funds.

Panel B of Table 1 reports descriptive statistics by strategy. We observe 527 VC funds managed by 261 unique GPs, with the median GP raising two funds during 1995 to 2013 period. The median VC fund has $250 million in committed capital, whereas the average is higher at $348 million due to a number of very large funds. For buyout, the number of funds and GPs are similar to VC (527 funds by 256 GPs), but funds are substantially larger, at an average (median) size of $1,589 million ($700 million). There are fewer real estate funds (249 funds by 134 GPs, with the median GP raising just one fund). The average (median) real estate fund size of $929 million ($535 million) is between VC and buyout fund sizes.

For VC and buyout, just over a third of funds have been fully liquidated by the end of the sample period (June 2018). The liquidation rate increases to roughly 43% for both strategies if we include funds that have less than 5% of committed capital in remaining NAV, as it is not uncommon for funds to be extended after their original 10-year life if any un-exited portfolio companies remain. As Table 1 shows, the time between the first and final observed cash flow for the median VC (buyout) fund is 12.2 (11.1) years. The proportion of liquidated real estate funds is lower, with only 28% fully liquidated, and the time between first and final cash flow is shorter (8.3 years on average), in large part because PE real estate is a younger strategy with a higher proportion of funds raised in more recent times.

With respect to performance, we compute the standard metrics in the literature; total value to paid-in capital (TVPI), which is a cash multiple of total fund distributions to date divided by total capital calls, internal rate of return (IRR), and public market equivalent (PME). We compute two versions of the PME. The first is the Kaplan and Schoar (2005) defined as the sum of fund distributions discounted to fund inception at the public equity market rate of return, divided by

[11] We have pension plan returns through at least June of 2018 for all but two plans. The two exceptions report their financial results in March and April 2018. Dropping these two plans does not change our results.

the similarly discounted sum of capital calls. The second version is defined by Korteweg and Nagel (2016) as the sum (not the ratio) of discounted net cash flows (distributions minus capital calls, normalized by the fund's committed capital). Interpreting the PME as a benchmark against the public equity market, a PE fund has outperformed public equities if the Kaplan-Schoar (Korteweg-Nagel) PME is above one (zero). For funds that are not yet liquidated by the end of June 2018, we follow standard practice and include their final reported net asset value (NAV) as a pseudo-distribution in all return metrics. Panel B shows that VC funds had the worst performance during the sample period by most measures, and buyout funds experienced the best performance. VC also has by far the highest variance in fund outcomes (and real estate is the least variable), and its performance is the most skewed, as indicated by the difference between mean and median metrics.

Finally, our Preqin data includes commitments to PE funds by LPs, which we match to our pension plan data. For roughly 9 out of 10 PE funds we observe at least one investment by a pension plan in our data, depending on the strategy (465 out of 527 VC funds, 488 out of 527 buyout funds, and 211 out of 249 real estate funds). For the median VC fund we see 2 pension plan investments, 4 for the median buyout fund, and 3 for the median real estate fund. The averages are higher (3.43, 5.90, and 4.45, respectively) due to a number of large funds for which we see many commitments.

3 Results

3.1 Private equity fund performance

In this section we present performance results by fixing a given private equity fund and computing its IPE and $GIPE$ using the returns of the various pension plans in our sample. Specifically, for $i = 1...I$ an index denoting a private equity fund, and $j = 1..J$ denoting a pension plan, we compute $IPE_{i,j}$ and $GIPE_{i,j}$ using equations (9) and (14), including any final reported NAV at the end of the sample period as a pseudo-distribution. Because each observation is a combination of a particular PE fund and a particular pension plan, each fund and each plan can be represented in multiple observations (put differently, a single PE fund can be evaluated differently by multiple pension plans, and a single pension plan can price multiple funds, and all these combinations are in the data here). For now we consider all plan-fund combinations, whether or not pension plan j

actually invested in PE fund i.

The computation of $IPE_{i,j}$ is straightforward. We use the observed cash flows of private equity fund i and discount the cash flows with the compounded cumulative returns of pension plan j. Since the returns of pension plans are only available at annual frequency, we discount a cash-flow occurring on day d of year τ using the discount factor $\left(\prod_{t=1..\tau-1} R_{j,t}\right)^{-1} e^{-(\log R_{j,\tau})\frac{d}{365}}$, where $R_{j,t}$ is the annual gross return of pension plan j.

The computation of $GIPE_{i,j}$ also requires that we estimate a separate coefficient γ_j for each pension plan j to take into account that different pension plans may be choosing different degrees of leverage in their investments. We estimate γ_j by using all available annual returns of pension plan j, and determining γ_j so as to satisfy the Euler equation

$$\frac{1}{n_T - n_0 + 1} \sum_{t=n_0...n_T} (R_{j,t})^{-\gamma_j} \left(R_t^m - \left(1 + r_{t-1}^f\right)\right) = 0, \tag{18}$$

where $R_t^m - \left(1 + r_{t-1}^f\right)$ is the excess return on the CRSP value-weighted market portfolio. As such, each pension plan's gross return (raised to the power γ_j) correctly prices the (excess) market portfolio. Figure 3 shows the histogram of the estimated γ coefficients. The mass of the distribution is centered around a coefficient of 5, and all but one pension plan have coefficients below 10.

We compute the term $\frac{W_{j,t}^{-\gamma}}{E_0 W_{j,t}^{-\gamma}}$ in equation (14) as $\prod_{n=1..t} \frac{R_{j,n}^{-\gamma_j}}{\widehat{E}\left(R_{n,t}^{-\gamma_j}\right)}$, where $\widehat{E}\left(R_{n,t}^{-\gamma_j}\right)$ is the average value of $R_{n,t}^{-\gamma_j}$ for pension plan j across all available observations for plan j. For a cash flow occurring on day d of year τ we perform a similar within-year adjustment as in the case of the $IPE_{i,j}$.[12]

Table 2 reports the performance estimates. The average $IPE_{i,j}$ (in the first row of Panel A) across all plan-fund combinations is 0.127, which is statistically significant at the 1% level. A fund with this IPE generates 12.7 cents in risk-adjusted net present value terms over the life of the fund, per one dollar of commitment. The second to fourth columns in the table report estimates for individual PE strategies. Buyout funds have an average IPE of 0.222, which is also statistically significant at the 1% level. The average IPE for VC and real estate funds is 0.052 and 0.087. While positive, neither is statistically significant.

[12] Specifically, letting $\widehat{H}_t = \prod_{n=1..t} \left(1 + r_n^f\right)^{-1} \times \left(\frac{R_{j,n}^{-\gamma_j}}{\widehat{E}\left(R_{n,t}^{-\gamma_j}\right)}\right)$, we discount the cash flow on day d of year τ using the discount factor $\widehat{H}_{\tau-1} e^{\log\left(\frac{\widehat{H}_\tau}{\widehat{H}_{\tau-1}}\right)\frac{d}{365}}$.

Table 2: Private Equity Fund Performance Metrics

This table reports performance results for private equity funds. The first column shows the average performance across all possible pairs of pension plans and PE funds. The second column shows the standard error (s.e.) for the test of the null hypothesis that the mean is zero, double-clustered by vintage year and pension plan. The second, third, and fourth sets of two columns shows the average performance and its standard error separated by PE strategy: buyout, venture capital, and real estate funds. Panel A shows IPE-type metrics. The first row is the *IPE* of equation (9). The second row, *PME(KN)*, shows the PME of Korteweg and Nagel (2016), for comparison. The *IPE(mkt-repl.)* is the IPE of PE-mimicking funds with the same capital calls as the PE fund, and the same timing of distributions, but invested in the CRSP value-weighted market portfolio. Similarly, *IPE(value-repl.)* and *IPE(growth-repl.)* refer to mimicking funds that invest in the top quintile of value stocks, and those in the intersection of the lowest size and book-to-market quintiles. Panel B shows similar results for GIPE-type metrics, where *GIPE* is as defined in equation (14), and *GPME* is the Generalized PME of Korteweg and Nagel (2016). The row labeled α shows the annualized excess return that makes the (G)IPE equal to zero, computed as described in the text. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

	All funds		Buyout		VC		Real Estate	
	Mean	s.e.	Mean	s.e.	Mean	s.e.	Mean	s.e.
Panel A: IPE-type metrics.								
IPE	0.127	0.036***	0.222	0.025***	0.052	0.084	0.087	0.054
PME(KN)	0.037	0.028	0.157	0.037***	-0.049	0.069	-0.035	0.053
IPE(mkt-repl.)	0.086	0.025***	0.072	0.024***	0.078	0.031**	0.132	0.016***
IPE(value-repl.)	0.187	0.027***	0.183	0.033***	0.211	0.030***	0.148	0.017***
IPE(growth-repl.)	-0.107	0.027***	-0.101	0.027***	-0.128	0.032***	-0.075	0.023***
α(IPE)	0.035	0.011***	0.062	0.009***	0.004	0.021	0.043	0.018**
IPE - PME(KN)	0.091	0.024***	0.065	0.031**	0.100	0.024***	0.122	0.016***
IPE - IPE(mkt-repl.)	0.041	0.032	0.150	0.030***	-0.026	0.07	-0.045	0.062
IPE - IPE(value-repl.)	-0.06	0.05	0.039	0.046	-0.16	0.1	-0.061	0.062
IPE - IPE(growth-repl.)	0.234	0.026***	0.324	0.032***	0.18	0.064***	0.162	0.037***
IRR - α(IPE)	0.076	0.003***	0.079	0.003***	0.069	0.005***	0.082	0.004***
Panel B: GIPE-type metrics.								
GIPE	-0.036	0.032	0.106	0.037***	-0.14	0.063**	-0.113	0.075
GPME	-0.111	0.135	0.185	0.213	-0.283	0.110***	-0.365	0.148**
GIPE(mkt-repl.)	-0.049	0.021**	-0.046	0.024**	-0.055	0.022**	-0.045	0.021**
GIPE(value-repl.)	0.068	0.058	0.068	0.054	0.112	0.083	-0.022	0.024
GIPE(growth-repl.)	-0.242	0.041***	-0.219	0.041***	-0.265	0.040***	-0.241	0.051***
α(GIPE)	-0.017	0.010*	0.017	0.010	-0.049	0.015***	-0.025	0.016
GIPE - GPME	0.075	0.128	-0.079	0.187	0.143	0.104	0.252	0.085***
GIPE - GIPE (mkt-repl.)	0.014	0.027	0.152	0.034***	-0.085	0.067	-0.067	0.071
GIPE - GIPE(value-repl.)	-0.103	0.064	0.037	0.047	-0.251	0.129*	-0.09	0.070
GIPE - GIPE(growth-repl.)	0.206	0.024***	0.325	0.047***	0.125	0.057**	0.128	0.030***
IRR - α(GIPE)	0.128	0.012***	0.124	0.013***	0.122	0.016***	0.150	0.012***
N	157,025		63,182		62,758		31,085	

Figure 3: Histogram of Pension Plan Risk Aversion Coefficients

This figure shows the distribution of estimated risk aversion coefficients (γ) for U.S. public pension plans. Each pension plan has one risk aversion coefficient that is estimated such that the plan prices the excess return on the public stock market (i.e., it is the γ-coefficient that satisfies the Euler equation (18)), using all available annual returns for the plan.

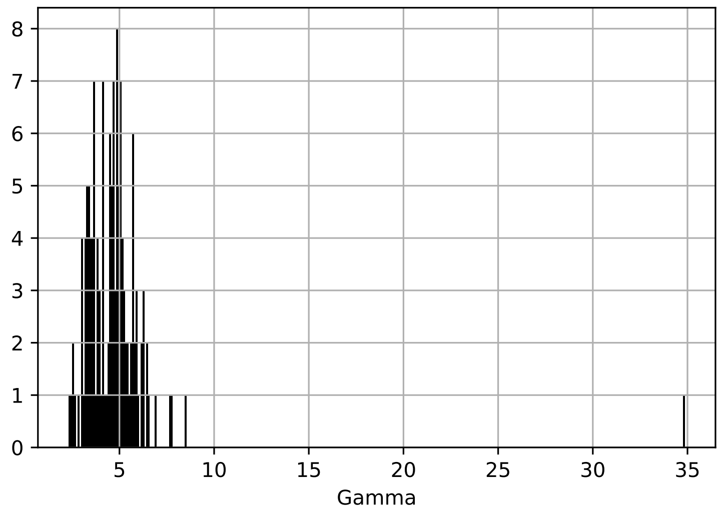

To show how performance has evolved over time, the left column of plots in Panel A of Figure 4 show the average $IPE_{i,j}$ by PE fund vintage year. The grey shaded area represents the 10th-90th percentile range across plan-fund combinations. The IPE across all plan-fund combinations (the top plot in the Figure) is high in the initial years of the sample (around 1995) but quickly drops to hover around zero from the late 1990s vintages until the mid 2000s, after which it is mildly positive until the end of our sample. The pattern varies by PE strategy: The average buyout IPE is positive throughout the sample period but does not exhibit the high initial IPEs. These high early numbers are driven by venture capital, which generally follows the overall IPE pattern but in a more exaggerated manner. The average real estate IPE is negative for the mid-to-late 2000s vintages, whose investments span the global financial crisis.[13]

To facilitate comparison with IRR and traditional performance metrics in public equities, Table 2 and Panel B of Figure 4 also show performance expressed in "alpha", the annualized excess return (not due to discounting or a risk premium) that is required to make the $(G)IPE$ equal to zero. Mathematically, we define α as the number that makes $\sum e^{-\alpha t} H_{j,t} C_{i,t} = 0$ where $H_{j,t}$ is the SDF

[13]We do not have any real estate funds for the 2001 vintage, hence the gap in Figure 4.

Figure 4: Time Series of IPE and GIPE

Panel A shows the average IPE (left column of graphs) and GIPE (right column of graphs) across pension plan-PE fund pairs, by fund vintage year. The first row depicts data for all funds, the second row for buyout funds only, the third row for venture capital funds and the fourth row for real estate funds. The horizontal axis corresponds to the vintage year of the funds. The solid line is the mean, the shaded area corresponds to the 10th-90th percentile range of the observations. The mean and standard deviation of all observations across all vintages is displayed on the top right corner of each graph. Panel B expresses the IPE and GIPE in units of alpha, computed as described in the text.

Panel A: IPE and GIPE by vintage year and PE strategy.

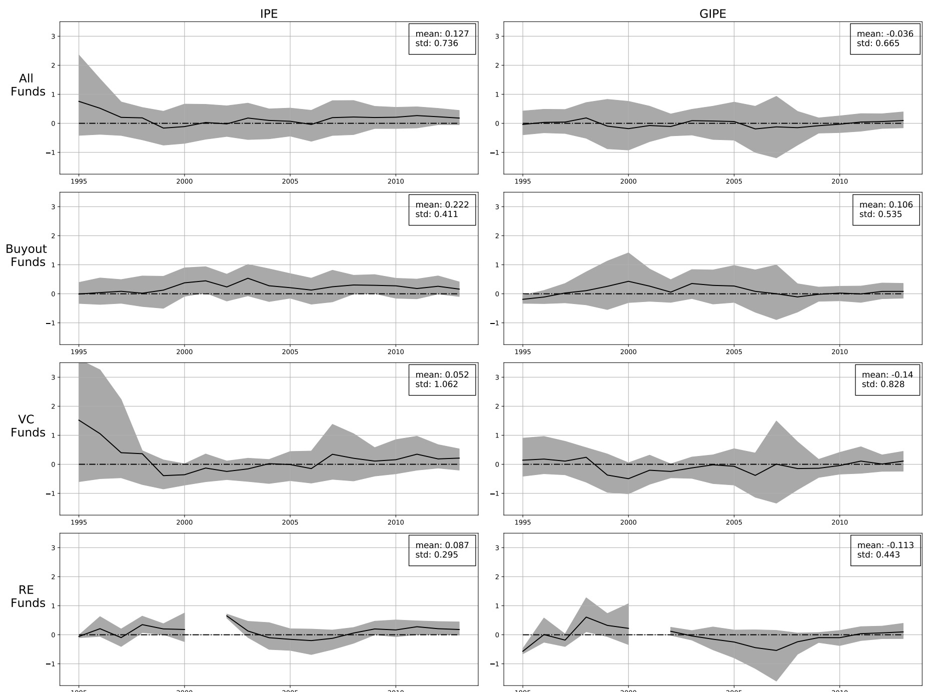

Panel B: IPE and GIPE alpha by vintage year and PE strategy.

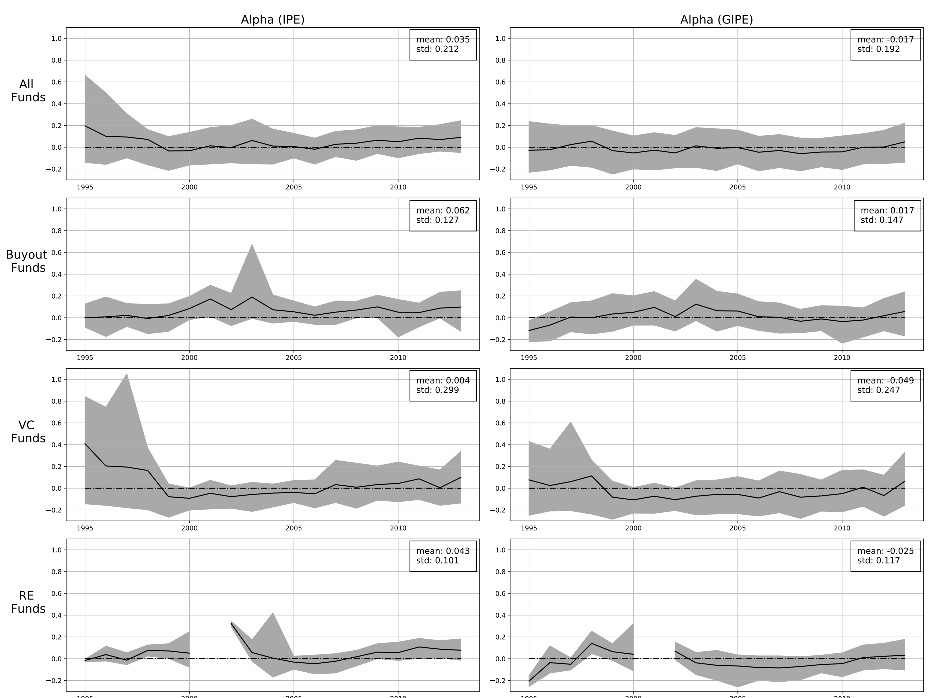

used in $(G)IPE$. An alpha of 0.05 represents an excess return of 5% per year.[14] Note that IPE and $GIPE$ use different SDFs, so that they result in different alphas. The overall IPE alpha across all strategies is 3.5% per year, which, like its IPE counterpart, is significant at the 1% level. Buyout has the highest α of 6.2% and venture capital has the lowest, at 40 basis points. The 4.3% alpha of real estate is statistically significant at the 5% level, unlike its IPE. The time series pattern of alpha across vintages is very similar to that of IPE.

Results are quite different for $GIPE$. The average $GIPE$ across all strategies is both economically small (-0.036) in magnitude and statistically insignificant (see Panel B of Table 2). The corresponding alpha of -1.7% per year is only marginally significant at the 10% level. The only strategy that has a positive and statistically significant $GIPE$ is buyout with an average $GIPE$ of 0.106. However, its α of 1.7%. per year is not significant. Venture capital has a negative $GIPE$ and alpha of -0.14 and -4.9%, both statistically significant. While real estate also has a negative $GIPE$ and alpha (of -0.112 and -2.5%), neither are significant. The right column of plots in Figure 4 show the time series of $GIPE$ (Panel A) and the corresponding alpha (Panel B). Generally, the $GIPE$ numbers are closer to zero than their IPE counterparts. Most striking is the disappearance of the strong performance of VC in the mid-1990s.

The generally positive IPE values imply that, on average, a typical pension fund would have increased the logarithmic growth rate of its assets by investing in private equity (especially buyout). However, the material differences in results between IPE and $GIPE$ suggest that this is a reflection of the higher exposure of private equity funds to market fluctuations, when compared to the portfolio of the typical public pension plan. Moreover, if an additional investment in public stocks were to have a higher $(G)IPE$, then pension plans would have been even better off investing more in the stock market rather than in private equity. To explore these points, the rows labeled "IPE(mkt-repl.)" and "GIPE(mkt-repl.)" in Table 2 report the means and standard errors for the average $(G)IPE$ of a mimicking cash flow process for the cash flows of each private equity fund. The mimicking funds make capital calls that are identical to the PE funds they are intended to replicate, per equation (15), but invest them in the CRSP-market weighted portfolio instead. Distributions are computed using equation (16). The average IPE of the mimicking funds is positive and

[14] The relation between our alpha definition and $(G)IPE$ is analogous to the relation between direct alpha and the PME performance metric (see Gredil et al., 2022)

statistically significant for all strategies, while the $GIPE$ values are negative, at around -0.050, and significant at the 5% level. The negative $GIPE$ values imply that there is little evidence that PE fund managers have superior market-timing ability, since the cash flows of the mimicking funds, $\widehat{C}_{i,t}$, exhibit (by construction) the same timing patterns as the actual cash flows, $C_{i,t}$.[15]

The difference in the $(G)IPE$ of a PE fund and that of its market-replicating portfolio is informative in determining whether pension plans would have been better off investing more in public or private equity.[16] The rows "$(G)IPE - (G)IPE(mkt - repl)$" in Table 2 shows that this difference is large, positive and significant for buyout funds, and statistically insignificant for all other strategies. Both the IPE and $GIPE$ of $C_{i,t} - \widehat{C}_{i,t}$ for buyout are equal to 0.150 with a standard error around 0.034. This shows that –at least historically– buyout funds gave pension funds access to investments that on average dominated the alternative of investing the funds in the stock market and making similar distributions to the buyout fund.

Figure 5 provides a visual impression of the distribution of $(G)IPE$ for the actual PE funds (in the first column of plots) and the mimicking funds (in the second column). A notable feature of this figure is that the dispersion of $(G)IPE$ for the mimicking funds is substantially smaller than for the actual funds. Strategies that simply invest in public stocks and mimic the capital calls and distributions of private equity funds produce $(G)IPEs$ that are very close to zero for the vast majority of observations. As a result, the distribution of the $(G)IPE$ of the difference in the cash flows between the actual and mimicking funds (in the third column if plots) looks much like the distribution of the actual PE fund $(G)IPEs$ from the first column. The much larger dispersion of the $(G)IPEs$ associated with the actual cash flows shows that, at the level of individual private equity funds, there are a substantial amount of shocks that appear unrelated to the benchmark portfolio (these could be idiosyncratic shocks or other risk factors).

The differences in the distributions of the $(G)IPE$ of the actual and mimicking funds may create the impression that their cash flow streams are not very correlated. However, this is not true. The average time-series correlation between the two cash flow processes across the different

[15]One might think that the market-replicating funds should have zero $GIPE$ by construction (since the pension plans price the overall stock market correctly), but this is not true. For example, the market-replicating fund $GIPE$ could be positive if GPs have market timing abilities that pension plans do not have. Conversely, as we discuss below, pension plans' shorting constraints could result in negative $GIPEs$.

[16]Note that the difference in $(G)IPE$ of the two cash flow streams is identical to the $(G)IPE$ of the difference in the two cash flow streams (i.e., the $(G)IPE$ of $C_{i,t} - \widehat{C}_{i,t}$), since these performance metrics are linear in cash flows.

Figure 5: IPE and GIPE of Actual and Mimicking Funds.
The first column of Panel A shows the IPE of private equity funds. The first row considers all combinations of PE funds and pension plans, and the next three rows show results for buyout, venture capital, and real estate funds separately. The second column shows the IPE of the respective public stock-market mimicking funds (IPE(mkt-repl.), as described in Table 2). The third column plots the difference IPE - IPE(mkt-repl.), which is mechanically the same as the IPE of the difference in cash flows. The last column plots the cross-sectional correlation between IPE and IPE(mkt-repl.), by vintage year. Panel B replicates the plots from Panel A for the GIPE performance metric.

Panel A: IPE comparisons.

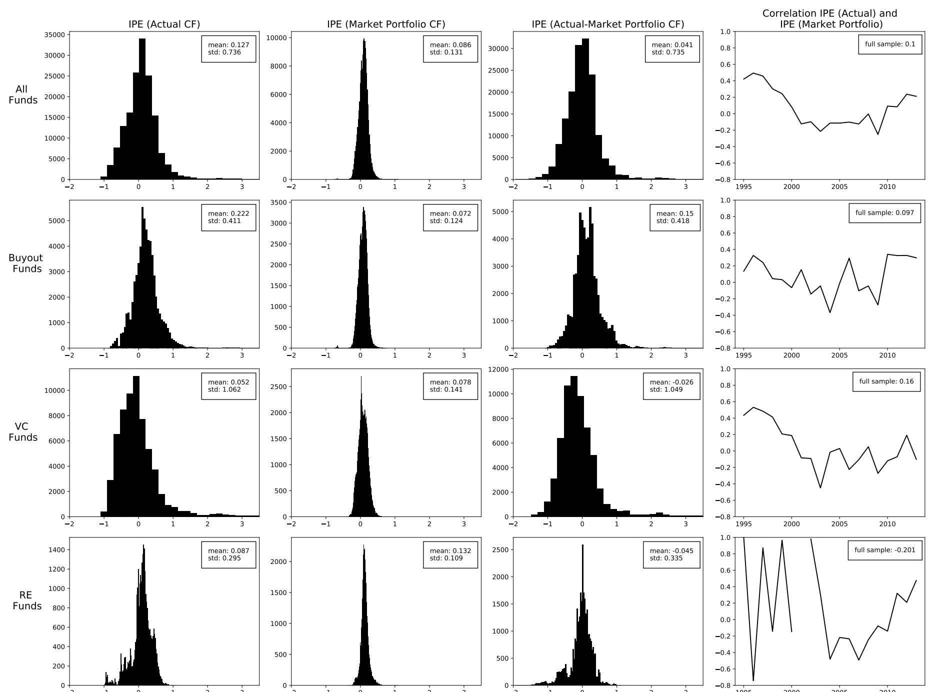

Panel B: GIPE comparisons.

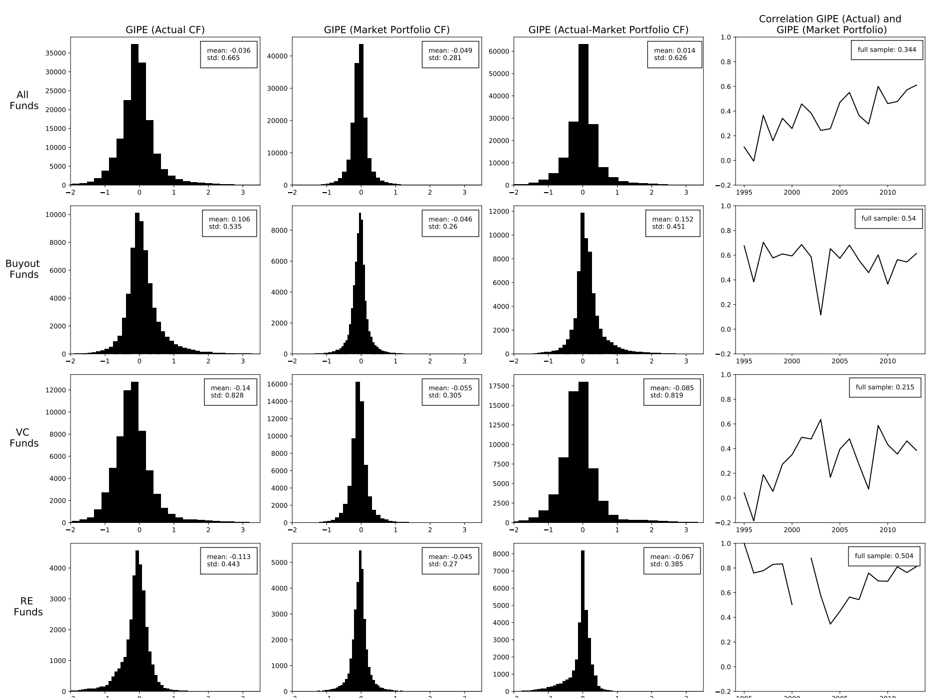

private equity funds is approximately 0.97. But this average correlation does not translate into a high correlation of their $(G)IPEs$, as seen in the fourth column of Figure 5. Intuitively, the reason is that the distributions of the mimicking funds are designed to be in similar in proportion to the distributions of their corresponding PE funds, whereas their capital calls are identical. As a result, when viewed as a time series, $C_{i,t}$ appears highly correlated with $\widehat{C}_{i,t}$. However, if the average *level* of the positive cash flows associated with $\widehat{C}_{i,t}$ is lower than for $C_{i,t}$, this significantly lowers the $(G)IPE$ of the mimicking fund relative to the actual PE fund, but without affecting the high correlation of the two cash flow series.

Thus far we have only considered mimicking funds that invests in the broad stock market. Since buyout funds historically tended to invest in value stocks (Stafford, 2022), and venture capital in small growth firms, we also construct mimicking funds that invest in publicly traded portfolios of value and small-growth stocks, downloaded from Kenneth French's website.[17] Specifically, the value portfolio return is the equal-weighted average return of the top quintile of book-to-market stocks, and the small-growth portfolio is the return of the stocks in the lowest size and book-to-market quintiles. Given the strong performance of buyout, our main goal is to examine whether buyout funds just represent "covert" value strategies. The row labeled "$GIPE(value - repl.)$" in Table 2 shows that the value-mimicking buyout fund $GIPE$ equals 0.068. While positive, this number is statistically not different from zero. In the language of asset pricing, one cannot reject the null hypothesis that the stochastic discount factor of an average pension plan can "price" the long-only value strategy of investing in a publicly traded value portfolio with cash flows that mimic those of a given buyout fund. The results do suggest, however, that while the $GIPE$ of buyout funds is the only positive and statistically significant $GIPE$, it owes its success partly due to the fact that it contains some value exposure. The difference between the $GIPE$ of the actual cash flows of buyout funds and their mimicking cash flows using the value return as a benchmark is positive at 0.042, but not statistically significant (see the row labeled $GIPE - GIPE(value - repl.)$). In that sense, the historical outperformance of buyout funds is due to both: a) their value exposure and; b) their ability to select better investments. Combining a) and b) leads to a positive and significant $GIPE$, even though individually the two components are not significant.

It is useful to contrast the buyout value results with the $GIPE$ of the mimicking funds that invest

[17] `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`

in small-growth stocks (see the row labeled "$GIPE(growth - repl.)$" in Table 2). The $GIPEs$ are significantly negative for all PE strategies, ranging from -0.218 to -0.263. We also found negative performance for the market-replicating funds, but the numbers here are a factor four to five larger. This implies that the SDF of a given pension plan cannot "price" these strategies. The existence of these large and negative $GIPEs$ may simply be due to shorting constraints, which prevent pension plans from capitalizing on strategies with negative long-only excess returns. Interestingly, the poor performance of VC is explained away by the poor performance of small-growth companies. Relative to the public small-growth mimicking funds, VC performance was positive and significant.

To summarize and conclude the replication exercises, the key takeaway is that some mimicking funds exhibit negative and statistically significant $GIPEs$, but none of the positive $GIPEs$ are significant. Thus, while the plan's portfolio decisions may be constrained with respect to shorting, they do not appear constrained or inconsistent with our proposed SDF with respect to long-only strategies.

We conclude this section with a comparison of the IPE and $GIPE$ measures to existing, related performance measures in the private equity literature. The measure closest to IPE is the $PME(KN)$ measure from Korteweg and Nagel (2016), described in section 2.2, while the measure closest to $GIPE$ is the $GPME$ from the same paper. The average $PME(KN)$ across all funds is 0.035, and the average $GPME$ is -0.114, neither of which is statistically different from zero (see Table 2, which also shows the numbers for the various PE strategies). The average difference between IPE and $PME(KN)$ across all plan-fund combinations is 0.091, which is statistically significant at the 1% level. In fact, the average difference is positive and statistically significant for all PE strategies, ranging from 0.067 for buyout to 0.123 for real estate. This is to be expected, since $PME(KN)$ uses the cumulative stock market return to discount the fund's cash flows, while IPE uses the cumulative returns of different pension plans. Since stock market investments are only a fraction of a pension plan's portfolio, the IPE in effect uses an SDF with a smaller exposure to the broad stock market, which results in a somewhat smaller risk premium when discounting cash flows.

The difference between a fund's $GIPE$ and $GPME$ is a more interesting object, because both measures are constructed with the aim of assigning a zero value to a fund whose cash flows are replicable with some (possibly) levered trading strategy in stocks and bonds. A non-zero $GIPE -$

$GPME$ reflects the fact that a typical pension plan's portfolio of *risky* assets (domestic but also international stocks, long term bonds, other private equity investments, etc.) differs from the CRSP-value weighted return. Table 2 shows that, across all plan-fund combinations, the average $GIPE$ is 0.077 higher than $GPME$, which is economically meaningful but statistically insignificant. The only PE strategy with a significant difference is real estate, where $GIPE$ is 0.254 higher than $GPME$.

Table 2 also reveals that the average $GPME$ has a substantially larger standard error than $GIPE$. This is due to the fact that the cross-sectional standard deviation of $GPME$ within a given vintage is much higher than for $GIPE$ (this result is not separately reported in Table 2). More, the cross-sectional dispersion of $GPME$ can change dramatically from year to year. It is reassuring that the distribution of $GIPE$ is more stable, as one would expect that the cross-sectional distribution of skill of managers does not change dramatically from one year to the next.

Table 2 also compares the IRR of each private equity fund with the $IPE-$implied α (Panel A) and the $GIPE-$implied α (Panel B), in the row labeled "$IRR - \alpha$". The average difference between these two numbers can be interpreted as the investment's risk premium. The IPE-implied and $GIPE$-implied risk premia are 7.6% and 12.9%, respectively, indicating that compensation for risk is a large component of a private equity fund's IRR. We employ this measure below, when we consider risk-taking behavior by pension plans.

3.2 Pension plan commitments and plan-level performance

In this section we turn our focus to the performance of the private equity portfolio that pension plans invested in. Specifically, we use the data set on pension plan commitments to identify the set of private equity funds $I_{j,t}$ that pension plan j invested in within vintage year t. We then compute an equal-weighted average $(G)IPE(\text{inv. EW})_{j,t} \equiv \frac{1}{N_{I_{j,t}}} \sum_{i \in I_{j,t}} (G)IPE_{i,j}$ of this portfolio, as well as a value-weighted average $(G)IPE(\text{inv. VW})_{j,t} \equiv \sum_{i \in I_{j,t}} w_{i,j} (G)IPE_{i,j}$, where the weight $w_{i,j}$ is based on the relative size of the commitments that plan j made in the private equity funds $I_{j,t}$ in vintage year t. For comparison, we also compute a "benchmark" $(G)IPE$, defined as the equal-weighted average of all the private equity funds I_t in vintage t, irrespective of whether plan j invested or not: $(G)IPE_{j,t} \equiv \frac{1}{N_{I_t}} \sum_{i \in I_t} (G)IPE_{i,j}$.

Table 3 shows that the average values of $IPE(\text{inv. EW})$ and $IPE(\text{inv. VW})$ across vintages and pension plans are very close to each other, at 0.193 and 0.199. Both numbers are statistically significant. They are also positive for the individual PE strategies, though not statistically significant in the case of VC. The left column of plots in Figure 7, Panel A, shows the time series of the average $IPE(\text{inv. EW})_{j,t}$. The average IPE values are positive in most vintage years for most strategies, with the exception of VC vintages in the early 2000s and real estate funds raised around 2005. Panel B shows the corresponding annualized alpha numbers.

Table 3: Pension Plan Portfolio-level Performance Metrics

This table reports performance metrics at the pension plan level. Panel A shows IPE-type metrics. In the *IPE* row, each observation is the IPE for a pension plan-year, equal-weighted across all possible PE funds of that vintage year, regardless of whether they were invested in by the pension plan. The first column shows the average across plan-years, and the second column shows the standard error (s.e.) for the test of the null hypothesis that the mean is zero, double-clustered by vintage year and pension plan. The second, third, and fourth sets of two columns repeat the calculation but using only PE funds of a given strategy (buyout, venture capital, real estate). In the *IPE(inv. EW)* row, an observation is the IPE for a plan-year, equal-weighted across the PE funds that the pension plan actually invested in in that year. *IPE(inv. VW)* performs the same calculation as IPE(inv. EW) but weighs each PE fund by the commitment size of the pension plan (see the text for a detailed description). We include only those plan-years in which the plan invested in at least one fund of the relevant PE strategy in that year (e.g., the "IPE" row in the "Buyout" column is computed only over plan-years in which the pension plan made at least one buyout investment in the year), so that the different rows have equal numbers of observations. Panel B shows the performance results using GIPE-type metrics. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

	All funds		Buyout		VC		Real estate	
	Mean	s.e.	Mean	s.e.	Mean	s.e.	Mean	s.e.
Panel A: IPE-type metrics.								
IPE	0.144	0.035***	0.230	0.028***	0.108	0.082	0.065	0.052
IPE(inv. EW)	0.193	0.034***	0.261	0.028***	0.151	0.101	0.081	0.045*
IPE(inv. VW)	0.199	0.029***	0.228	0.027***	0.118	0.078	0.070	0.030**
IPE(inv. EW) - IPE	0.049	0.018***	0.030	0.016*	0.043	0.042	0.016	0.023
IPE(inv. VW) - IPE	0.055	0.022**	-0.002	0.013	0.010	0.046	0.005	0.032
IPE(inv. VW) - IPE(inv. EW)	0.006	0.009	-0.032	0.007***	-0.033	0.032	-0.011	0.016
Panel B: GIPE-type metrics.								
GIPE	-0.013	0.027	0.123	0.037***	-0.084	0.057	-0.113	0.069
GIPE(inv. EW)	0.031	0.024	0.155	0.041***	-0.081	0.062	-0.080	0.055
GIPE(inv. VW)	0.047	0.026*	0.135	0.037***	-0.037	0.043	-0.034	0.032
GIPE(inv. EW) - GIPE	0.044	0.018**	0.032	0.015**	0.003	0.030	0.033	0.032
GIPE(inv. VW) - GIPE	0.060	0.024**	0.012	0.015	0.047	0.039	0.078	0.051
GIPE(inv. VW) - GIPE(inv. EW)	0.016	0.008**	-0.020	0.007***	0.044	0.027	0.046	0.029

Turning to $GIPE$, both the average equal and value-weighted invested portfolio performance of 0.031 and 0.047 are close to zero, with only the value-weighted number marginally significant (at the 10% level). The right column of plots in Figure 7 shows that the equal-weighted invested $GIPE$ hovers around zero across vintages, without any large outliers. These close-to-zero $GIPE$ values suggest that a typical pension plan would not have experienced a significant benefit by further increasing (or reducing) its allocation to its chosen portfolio of private equity investments.

The exception is in buyout, where both portfolio $GIPEs$ are positive and statistically significant, primarily due to the strong performance in the late 1990s and early 2000s vintages. In contrast, the invested portfolio performance for both VC and real estate is negative but statistically no different from zero. Importantly, based on their actual investments, pension plans were not overallocated to VC, contrary to what the negative GIPE for an average VC fund that we documented in the previous section, might suggest.

A comparison of the benchmark $(G)IPE$ with the equal and value-weighted invested portfolios in Table 3 suggests that the average pension plan invests in better-performing private equity investments compared to an uninformed investor who simply allocates an equal amount to each PE fund that is raised in a given vintage year. To provide a visual impression, The left column of graphs in Figure 7 show the time series of the average $IPE(\text{inv. EW})_{j,t}$, $IPE(\text{inv. VW})_{j,t}$, and the benchmark $IPE_{j,t}$, with the corresponding $GIPEs$ in the right column. Casual inspection suggests that the pension plan portfolios outperform the benchmark more often than not, albeit with some variation across strategies and vintages. Table 3 provides formal confirmation that the differences $(G)IPE(\text{inv. EW})_{j,t} - (G)IPE_{j,t}$ and $(G)IPE(\text{inv. VW})_{j,t} - (G)IPE_{j,t}$ are positive and significant (at the 5% level of higher) for the portfolio of all PE funds. The difference is also significant for the equal-weighted buyout portfolio, but insignificant for the value-weighted portfolio. For VC and real estate, the differences are positive but insignificant. We also test the difference between the value-weighted and equal-weighted invested portfolio, which informs whether pension plans invest more in better-performing funds (if the VW portfolio outperforms the EW portfolio). The VW-EW difference is positive but small for the overall PE portfolios (though significant for $GIPE$). For buyout, the value-weighted portfolio underperforms significantly (both in IPE and $GIPE$), whereas in VC and real estate, the differences are insignificant. The next section will dive deeper into the role of pension plan characteristics in PE fund selection.

A final noteworthy aspect of Figure 7 is the remarkably large cross-sectional dispersion of the portfolio-level $(G)IPEs$ across pension plans. While the mean value of portfolio $GIPE$ is approximately zero for almost all vintages, the unconditional dispersion is 0.335. This number is only about half of the 0.659 cross-sectional standard deviation of individual PE fund $GIPEs$ (as reported in the top right plot of Figure 4). A back-of-the envelope calculation suggests that, on average, a public pension plan invests in only about 4 private equity funds per vintage, resulting

Figure 6: Dispersion in Alternative Asset Portfolio Returns

The left plot shows the histogram of the 15-year geometric average return on the alternative assets and public equities portfolios reported by pension plans for the period from 2002 to 2017. The right plot shows the histogram of the 5-year geometric average return on the same portfolios, over the 2012 to 2017 period.

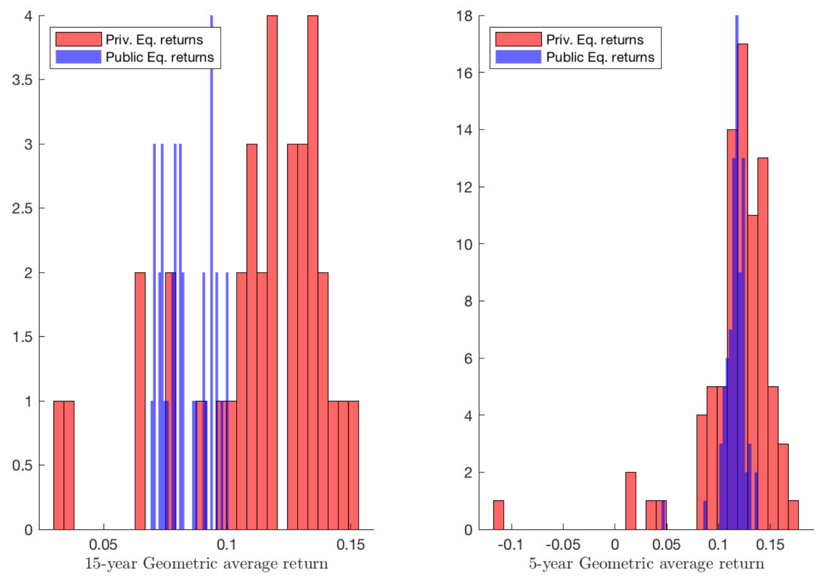

in a reduction of the standard deviation of $\frac{1}{\sqrt{4}} = 0.5$. Indeed, the average number of commitments per pension plan per vintage in our data is about 4.7 (not elsewhere reported).

This large cross-sectional dispersion is reflected in the rather large cross-sectional dispersion of the returns that pension plans report for their alternative assets portfolio more generally. For example, Figure 6 shows histograms of the 15-year (left plot) and 5-year (right plot) geometric average returns on the alternative-assets and public-equities portfolios, as reported by pension plans in our CAFR data for the period ending in 2017. The cross-sectional dispersion of public equity investment returns across pension plans is rather small, suggesting that the returns of their public equity portfolios are highly correlated. By contrast, the cross sectional standard deviation of long-run private equity returns is approximately 3 to 4 times larger than for public equities.

To appreciate the economic significance of this observation, assuming an annual stock market volatility of 16%, the volatility of a 15-year geometric average of investing in stocks is approximately $\frac{1}{\sqrt{15}}16\% \approx 4\%$. This is approximately the same order of magnitude as the *cross-sectional* standard deviation of alternative investment returns in figure 6. This means that even if there was no aggregate risk in private equity returns, just the cross-sectional standard deviation of private-equity portfolio returns would be of the same order of magnitude as the time-series uncertainty from investing in the stock market over a 15-year period.

3.3 Pension plan heterogeneity and PE performance

The observation that pension plan portfolios tend to have higher risk-adjusted returns than a "naive" portfolio, and the large cross-sectional dispersion in PE returns across plans, beg the question to what degree PE performance can be explained by pension plan characteristics? To dig deeper into the drivers of performance across pension plans, we switch back to plan-fund combinations as the unit of observation. We focus our analysis on the IRR and $GIPE$-implied alpha performance measures. These two metrics are useful to compare, since the former measures a fund's total annualized return and the latter its risk-adjusted annualized return. For brevity we do not report results based on IPE-implied alphas. All regressions in this section double-cluster standard errors by vintage and by pension plan.

First, we test whether risk-adjusted returns at the plan-fund level are indeed higher for funds that pension plans actually invested in. This augments the plan-level results from the previous section, as we consider not only the pension plan's individual PE fund investments, but also the funds it did not invest in, and we allow for additional controls such as vintage year fixed effects. Table 4 shows estimates from regressions with the performance of all plan-fund combinations as the dependent variable, measured in percentage points per year (for example, an IRR of 1.0 means a return of 1 percent per year). Panel A shows specifications with vintage fixed effects, and Panel B adds pension plan fixed effects. We first consider the coefficient on the indicator variable *Active*, which measures if a pension plan was an active investor in PE in a given year: It equals one if the plan made a commitment to any PE fund for the year, and zero otherwise. For IRR, the estimated coefficient is economically very small across all specifications and PE strategies, and statistically insignificant for most specifications. This means that pension plans do not tend to enter PE when IRRs are expected to be higher, or to get out of PE when future returns are lower. Results are different for alpha. Without plan fixed effects, the average alpha is about 1.5 percentage point higher for plans that invest in PE than alphas as evaluated by plans that do not (and this is fairly stable across PE strategies). This could indicate that some plans make a rational decision to stay

Figure 7: Time Series of Pension Plan Portfolio-level Performance.

Panel A shows the time series of the equal-weighted IPE (in the left column of graphs) and equal-weighted GIPE (right column of graphs) of the portfolio of PE investments by pension plans in a given vintage year. These correspond to IPE(inv. EW) and GIPE(inv. EW) in Table 3. The dashed line marked “EW Average” corresponds to the IPE row in Table 3, and considers all PE funds, regardless of whether they were invested in by the pension plan. The solid “EW Invested” line corresponds to IPE(inv. EW) in Table 3, and equal-weights across the PE funds invested in by the pension plan in the given vintage year. The dash-dotted “VW Invested” line corresponds to IPE(inv. VW), the value-weighted invested portfolio performance. The first row depicts data for all PE funds, the second row for buyout funds, the third row for venture capital funds and the fourth row for real estate funds. Panel B expresses the IPE and GIPE in units of alpha.

Panel A: IPE and GIPE by vintage year and type of strategy.

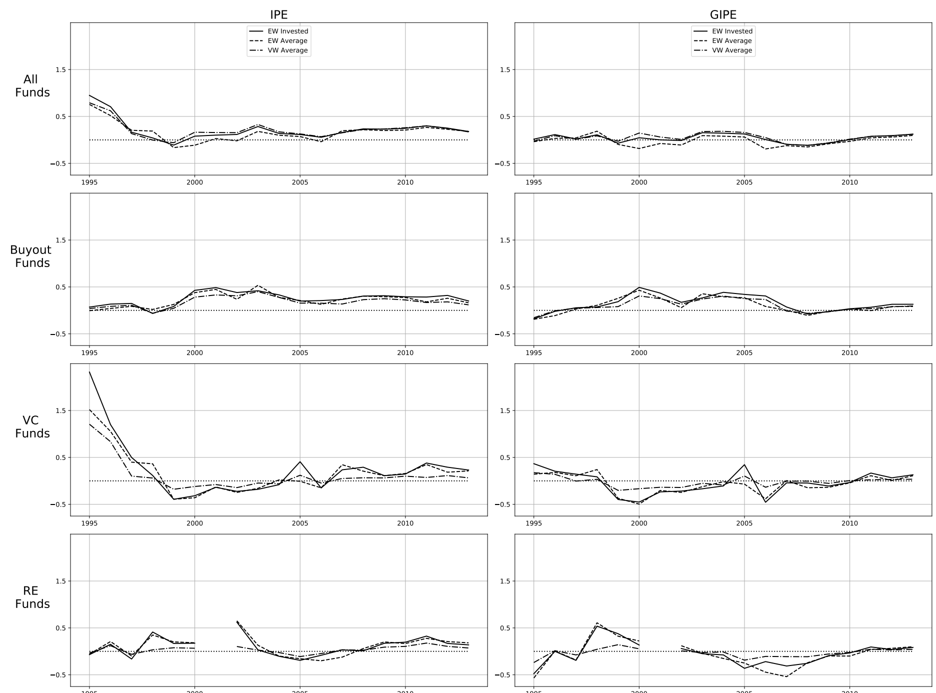

Panel B: IPE and GIPE alpha by vintage year and type of strategy.

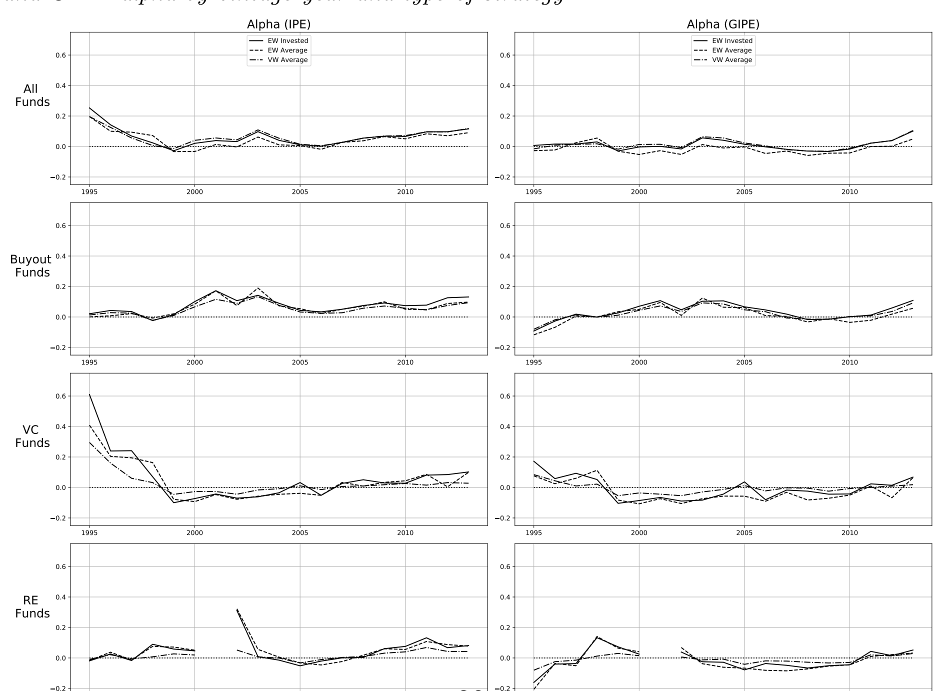

out of PE investments, given the low alphas that they expect to earn. Indeed, the result is due to heterogeneity across plans: with plan fixed effects, the *Active* coefficient in alpha regressions is close to zero and insignificant across the board, suggesting that pension plans do not have skill in timing their entry or exit from PE investing altogether.

Plans do experience better performance in the PE funds that they actually invest in, as evidenced by the positive and economically large coefficient on the indicator *Commit*, which equals one only if, in a given plan-fund combination, the pension plan invested in that specific PE fund. The coefficient is of similar magnitude for both IRR and alpha as the dependent variable, and with and without plan fixed effects. There is some heterogeneity across PE strategies, but all estimates are positive. This improved performance could come from pension plans having selection skill or differential access to better PE funds. To disentangle these two channels, we consider the performance of funds of GPs with which the pension plan has invested before, measured by the indicator *Relationship*. This indicator equals one only if the pension plan has made a commitment to a prior fund of the same GP. Prior investors are usually given the option to reinvest in follow-on funds (Lerner et al., 2007), such that access is not a concern in the performance of reinvestment decisions. The large, positive coefficient on *Relationship* shows that such funds tend to perform well, both in terms of IRR and alpha, likely due to the fact that these managers have survived to raise follow-on funds. What is more important for our purpose, is the interaction between *Relationship* and *Commit*. The coefficient is negative and almost of equal magnitude of *Commit*. This means that, conditional on a pre-existing relationship with a GP, a commitment is not associated with higher performance. This suggests that the unconditional higher performance from committed funds is driven by access, not skill. We find similar results in a subsample of first-time funds, which typically accept investments from any LP that is interested in investing (Lerner et al., 2011). For brevity, these results are reported in appendix B.

The literature also considers reinvestments and first-time funds to distinguish access and skill (Lerner et al., 2007; Sensoy et al., 2014; Andonov et al., 2018; Cavagnaro et al., 2019). The empirical evidence is mixed. For public pension plans, Lerner et al. (2007) and Cavagnaro et al. (2019) conclude that performance is due to more than just access (but not as much as for endowments), whereas Sensoy et al. (2014) find no evidence of skill after controlling for access, especially post-1999. Given that competition in PE has increased, and performance persistence has weakened for some strategies (e.g., Harris et al. (2022)), the fact that our sample period is more recent than these prior studies may help explain why we find little evidence of skill.

Next, we introduce additional pension plan-specific covariates to further explore performance differences between plans. We are especially interested in agency problems that may arise depending on the funding status of a plan, its governance structure, and the location of the PE fund that the plan invests in. We focus on the subsample of plan-fund combinations for which the plan actually made an investment in the PE fund. Furthermore, since many of the explanatory variables, such as the structure of the board, do not vary much (or at all) over time, we do not include pension plan fixed effects. Instead, we use plan-state fixed effects (in addition to vintage fixed effects), so that we are effectively comparing the outcome of investments made by different pension plans within the same state in the same year.

Table 5 reports the estimated coefficients of regressions of IRR and alpha on plan characteristics. To determine the effect of geography we include an indicator variable that equals one if the PE fund has the same state of domicile as the pension plan, and zero otherwise. A negative coefficient means that local investments underperform relative to out-of-state investments. This could come about, for example, as a result of political pressure to invest in the home state (Hochberg and Rauh, 2013). Conversely, a positive coefficient suggests that pension plans may have superior information about funds in their home state. Across all plan-fund investments, we find a negative coefficient on IRR, consistent with Hochberg and Rauh (2013). The coefficient on home-state investments is also negative (and marginally significant) when we use alpha as the dependent variable. This result complements the literature by showing that local underperformance is due to poorer investments in a risk-adjusted sense, and not only due to investments being potentially less risky (we consider risk-taking below). The coefficients are also negative for VC and real estate strategies, and significant in the case of real estate IRRs. For buyout, however, the coefficients are positive, and significant (at the 10% level) for IRR, suggesting that for this strategy, pension funds tend to have an information advantage when it comes to investing in local funds.

Turning to the funding status of the plan, we find that more highly funded plans, measured as the ratio of actuarial assets to actuarial liabilities, have lower IRR but higher alpha when we consider all plan-fund investments. The coefficients are insignificant, though economically meaningful: a one standard deviation increase in funded ratio lowers the expected annual IRR by 0.79 percentage

Table 4: **Pension Plan Commitments and Performance**

Each column shows the coefficients of a regression of private equity fund performance on two indicator variables plus controls. The observations are all possible combinations of pension plans and PE funds. The indicator *Active* equals one if pension plan made a commitment to any PE fund in that vintage year. and zero otherwise. *Commit* equals one when the pension plan committed to this particular PE fund, and zero otherwise. *Relationship* equals one for all plan-fund combinations where a prior commitment has been made between the pension plan and PE fund manager. The first two columns consider all PE funds. Columns (3) and (4) only include buyout funds, columns (5) and (6) only venture capital funds, and columns (7) and (8) only real estate funds. The odd-numbered columns use IRR as the performance metric, and the even-numbered columns use the GPME-implied α (both measured in percentage points per year, i.e., an IRR of 1 means one percent per year). All regressions include vintage year fixed effects. The regressions in Panel B also include pension plan fixed effects. Standard errors are double clustered by vintage year and pension plan, and are shown in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level

Panel A: Without pension plan fixed effects

	All Funds		Buyout Funds		Venture Capital Funds		Real Estate Funds	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$
Active=1	-0.166**	1.465**	-0.104	1.665**	-0.111	1.250*	-0.0557	1.785**
	(0.0662)	(0.686)	(0.0697)	(0.704)	(0.113)	(0.635)	(0.0880)	(0.828)
Commit=1	2.333***	2.237***	0.842	1.077*	3.295	2.268*	0.271	0.441
	(0.739)	(0.491)	(0.624)	(0.528)	(2.290)	(1.178)	(0.697)	(0.625)
Relationship=1	2.876***	3.527***	2.288	2.794**	4.469**	5.503***	0.927	1.486
	(0.867)	(0.887)	(1.339)	(1.218)	(1.731)	(1.523)	(1.315)	(1.453)
Relationship x Commit	-2.800*	-2.695*	-1.885	-1.711	-5.604	-5.066	-0.318	-1.204
	(1.521)	(1.345)	(1.159)	(1.288)	(4.361)	(3.010)	(1.097)	(1.117)
N	156190	155022	62790	62265	62315	61921	31085	30836
adj. R^2	0.059	0.034	0.137	0.093	0.119	0.060	0.319	0.188
Vintage FE	Y	Y	Y	Y	Y	Y	Y	Y
Pension Plan FE	N	N	N	N	N	N	N	N

Panel A: Without pension plan fixed effects

	All Funds		Buyout Funds		Venture Capital Funds		Real Estate Funds	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$
Active=1	-0.0173	-0.164	-0.0213	-0.0748	0.0425	-0.145	0.0448	-0.0678
	(0.0516)	(0.321)	(0.0354)	(0.343)	(0.110)	(0.331)	(0.0390)	(0.207)
Commit=1	2.450***	1.981***	0.912	0.867	3.602	1.779	0.295	0.504
	(0.726)	(0.519)	(0.648)	(0.524)	(2.308)	(1.371)	(0.723)	(0.670)
Relationship=1	3.013***	2.658***	2.375	1.949	4.785**	4.703***	0.969	0.711
	(0.884)	(0.828)	(1.384)	(1.226)	(1.787)	(1.594)	(1.358)	(1.359)
Relationship x Commit	-2.908*	-2.157	-1.944	-1.359	-5.918	-4.456	-0.335	-0.423
	(1.524)	(1.357)	(1.183)	(1.275)	(4.404)	(3.167)	(1.115)	(0.959)
N	156190	155022	62790	62265	62315	61921	31085	30836
adj. R^2	0.058	0.100	0.135	0.210	0.117	0.095	0.316	0.411
Vintage FE	Y	Y	Y	Y	Y	Y	Y	Y
Pension Plan FE	Y	Y	Y	Y	Y	Y	Y	Y

Table 5: **Pension Plan Characteristics and Performance**

Each column shows the coefficients of a regression of private equity fund performance on pension plan characteristics plus controls. The observations are all plan-fund combinations for which the pension plan invested in the PE fund. The first two columns consider all PE funds. Columns (3) and (4) only include buyout funds, columns (5) and (6) only venture capital funds, and columns (7) and (8) only real estate funds. The odd-numbered columns use IRR as the performance metric, and the even-numbered columns use the GPME-implied α. *Home state = 1* is an indicator variable that equals one if the PE fund and pension plan are located in the same state, and zero otherwise. *Funded ratio* is the ratio of actuarial assets to actuarial liabilities. *Public equity weight* is the fraction of the plan's portfolio allocated to public equity. *State appointed*, *State ex-officio*, *Member elected*, and *Public appointed* are the fraction of the plan's board members who are state officials appointed by government official, state officials serving ex-officio, plan members elected by their peers, and members of the public appointed by a government official, respectively. *Other trustees* is the fraction of the board who were installed by other means, with the omitted category being plan members who were appointed by an official. *Board size* is the number of pension plan board members. *Log(AUM)* is the natural logaritm of the pension plan's assets under management, and *Log(commitment %)* is the natural logarithm of the commitment of the plan to the fund, as a percentage of the plan's assets under management. All regressions include vintage year and plan-state fixed effects. Standard errors are double clustered by vintage year and pension plan, and are shown in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level

	All Funds		Buyout Funds		Venture Capital Funds		Real Estate Funds	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$
Home state = 1	-1.641*	-1.398*	1.596*	1.358	-1.182	-1.207	-3.133**	-1.012
	(0.896)	(0.726)	(0.873)	(1.146)	(1.869)	(1.628)	(1.053)	(1.133)
Funded ratio	-4.636	3.472	-2.226	4.164	-11.25	-0.425	0.333	8.935*
	(4.054)	(3.027)	(3.767)	(2.987)	(11.38)	(6.554)	(2.850)	(4.537)
Public equity weight	-15.26**	-23.08***	-5.103	-13.73***	-19.46	-21.49*	-9.846**	-20.39**
	(6.701)	(5.191)	(4.334)	(3.213)	(15.21)	(10.84)	(3.732)	(7.401)
State appointed	-3.959	-6.860	-5.184	-7.683	9.916	6.280	0.606	0.333
	(4.613)	(4.944)	(5.886)	(5.969)	(10.31)	(8.100)	(4.964)	(5.065)
State ex-officio	-5.269	-9.563**	-5.781	-13.24***	-5.278	-6.823	5.867	2.698
	(3.272)	(3.686)	(3.416)	(3.647)	(7.162)	(4.863)	(4.005)	(5.018)
Member elected	-0.946	-1.721	-0.810	-2.518	-0.846	-4.108	3.383	3.955
	(2.136)	(2.299)	(1.782)	(2.213)	(3.650)	(2.909)	(3.408)	(6.124)
Public appointed	-0.341	-5.535*	-2.409	-7.091*	0.105	-7.914**	0.226	0.778
	(1.946)	(2.972)	(1.936)	(3.404)	(4.457)	(3.653)	(3.300)	(3.541)
Other trustees	7.193**	1.336	4.484	-1.635	9.889	2.546	1.879	-0.872
	(3.077)	(2.944)	(3.922)	(4.427)	(6.604)	(3.868)	(3.105)	(4.332)
Board Size	-0.00811	0.202***	-0.0791	0.140**	0.378*	0.447***	-0.0489	0.198
	(0.0540)	(0.0469)	(0.0521)	(0.0574)	(0.195)	(0.113)	(0.0328)	(0.114)
Log(AUM)	0.597	1.030*	0.870*	1.539***	-2.224	-1.743	-0.246	-0.595
	(0.603)	(0.577)	(0.416)	(0.458)	(1.572)	(1.111)	(0.550)	(0.516)
Log(commitment %)	1.256*	1.267**	1.013**	0.970*	-1.199	-1.024	-0.693	0.190
	(0.705)	(0.545)	(0.456)	(0.492)	(0.956)	(0.652)	(0.860)	(0.832)
N	3933	3933	2139	2139	1072	1072	712	712
adj. R^2	0.101	0.096	0.235	0.208	0.134	0.075	0.364	0.268
Vintage FE	Y	Y	Y	Y	Y	Y	Y	Y
Plan State FE	Y	Y	Y	Y	Y	Y	Y	Y

points, and raises the expected alpha by 0.59 percentage points. Across all strategies, the coefficient estimates imply that IRR and alpha are closer together for higher funding ratios, an observation that we will return to below when we consider risk-taking.

To analyze the impact of governance structure on performance, we follow Andonov et al. (2018) and use the fraction of trustees on the pension plan's board, depending on whether they are state officials, plan participants, or members of the public, and whether they were elected by plan members, appointed by a government official, or serving ex-officio. Across these 9 categories, the most common ones are elected plan participants (27% of the average board), state officials, such as the state treasurer, serving ex-officio (25%), appointed members of the public (25%), appointed plan participants (12%), and appointed state officials (8%). The other four categories are less than 2% of the average board each, and we lump them together as *Other trustees*. Following Andonov et al. (2018), the omitted category is appointed plan participants. Despite covering a different sample period, and a different regression specification, our IRR results are qualitatively in line with Andonov et al.: plans with a higher fraction of state-appointed and ex-officio trustees, participant-elected, or appointed members of the public (relative to appointed plan participants), experience lower IRRs on their PE investments, and the effect is largest for the two categories of state official trustees. Unlike Andonov et al., we can also consider the effect on risk-adjusted returns. With alpha as the dependent variable, we find even larger (i.e., more negative) coefficients. However, the only statistically significant results are for state ex-officio and appointed members of the public. Economically, a 10 percentage point increase in the proportion of state ex-officio (appointed public) trustees lowers the annual alpha by 0.96 (0.55) percentage points. For the individual PE strategies, buyout looks similar to PE overall. VC looks similar with the exception of state-appointed trustees, who have positive (though insignificant) coefficient. The coefficient signs are all reversed in real estate, though none are statistically significant, so it's difficult to draw strong conclusions for this strategy.

With respect to size, we find that larger pension plans (measured by the natural logarithm of assets under management) have better performance. The relation is strongest, and marginally significant, for alpha. This result is driven by buyout, as the signs reverse in VC and real estate investments (but the coefficients are insignificant). This result is consistent with Dyck and Pomorski (2015), who find that plans that have more dollars invested in PE, realize better plan returns.

Possible explanations are that larger LPs have access to better GPs (Lerner et al., 2007), and a wider scope of due diligence, monitoring, and other related activities (Da Rin and Phalippou, 2017). Interestingly, we find a positive and significant effect of board size on alpha, even after conditioning on plan size. Adding another member to the board increase the annual alpha by 0.20 percentage points.

Finally, we control for the natural logarithm of the size of the PE fund commitment by the pension plan, as a percentage of the plan's assets under management, and the fraction of the plan's portfolio that is allocated to public equity. Larger commitments tend to have better performance, although this effect reverses and becomes insignificant for VC and real estate. Plans with a higher allocation to public equity, indicative of lower risk aversion, have worse performance in private equity, both in terms of IRR and alpha, for all strategies. This echoes the results in Andonov et al. (2017), who find that pension plans with a higher allocation to risky assets have lower benchmark-adjusted plan returns.

The relation between performance and the funded ratio, as well as the equity share, suggest that risk-taking by pension plans may play an important role. A unique advantage of our approach is that we can measure the degree of risk-taking by pension plans more directly, at the individual investment level, as measured by the difference between PE fund IRR and alpha. Table 6 reports regression results with $IRR - \alpha$ as the dependent variable.[18]

A key result is that better-funded pension plans take less risk in their private equity investments, consistent with gambling for resurrection by underfunded plans (Rauh, 2009; Pennacchi and Rastad, 2011; Mohan and Zhang, 2014; Bradley et al., 2016; Myers, 2022). At the same time, plans with a higher public equity allocation take more risk in PE funds, consistent with these plans having lower risk aversion. As to governance and risk-taking, plans with a higher fraction of trustees who are appointed members of the public take more risk, especially in buyout and VC. Boards with a higher fraction of state officials who serve ex-officio also take more risk, as do boards with more appointed state officials, though in the latter case the coefficients are insignificant. Plan size, measured by assets under management, is not significantly related to risk-taking, whereas having more board members has a small, negative (but statistically significant) relation with risk-taking. Regarding

[18]Note that the coefficients in the last four columns of Table 6 are equal to the difference in the corresponding coefficients of the IRR and alpha regressions in Table 5.

Table 6: **Pension Plan Risk-taking**

This table shows results for regressions of the difference between IRR and the GIPE-implied α, a measure of compensation for risk, on pension plan characteristics plus controls. The observations are all plan-fund combinations for which the pension plan invested in the PE fund. Columns (1) and (5) include all PE funds. Columns (2) and (6) only include buyout funds, columns (3) and (7) only venture capital (VC) funds, and columns (4) and (8) only real estate (RE) funds. All explanatory variables are as described in Table 5. All regressions include vintage year and plan-state fixed effects. Standard errors are double clustered by vintage year and pension plan, and are shown in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level.

Dependent variable: $IRR - \alpha(GIPE)$

	(1) All	(2) Buyout	(3) VC	(4) RE	(5) All	(6) Buyout	(7) VC	(8) RE
Home state = 1	-0.285	0.344	-0.152	-1.988*	-0.243	0.238	0.0250	-2.121*
	(0.611)	(0.754)	(0.463)	(1.076)	(0.574)	(0.746)	(0.528)	(0.990)
Funded ratio	-7.818***	-7.765***	-9.111**	-7.607**	-8.109***	-6.390**	-10.83*	-8.602*
	(2.292)	(2.597)	(3.531)	(2.946)	(2.392)	(2.831)	(5.456)	(4.210)
Public equity weight	8.477**	9.570***	2.465	10.93*	7.820**	8.623**	2.030	10.54
	(3.714)	(3.253)	(7.300)	(5.671)	(3.603)	(3.505)	(7.493)	(5.962)
Log(AUM)	-0.695	-0.774**	-1.001**	0.601	-0.433	-0.669	-0.481	0.349
	(0.419)	(0.348)	(0.377)	(0.470)	(0.447)	(0.526)	(0.732)	(0.444)
State appointed					2.901	2.500	3.636	0.273
					(4.368)	(5.200)	(6.302)	(4.695)
State ex-officio					4.294*	7.454**	1.546	3.169
					(2.299)	(2.756)	(4.422)	(5.003)
Member elected					0.775	1.707	3.262*	-0.572
					(2.001)	(1.919)	(1.709)	(3.547)
Public appointed					5.194**	4.682**	8.019***	-0.552
					(2.363)	(1.866)	(2.722)	(3.615)
Board size					-0.210***	-0.220***	-0.0694	-0.246*
					(0.0580)	(0.0591)	(0.117)	(0.135)
Log(commitment %)					-0.0112	0.0430	-0.175	-0.883
					(0.332)	(0.388)	(0.399)	(0.682)
Other trustees					5.857**	6.119*	7.343	2.752
					(2.752)	(3.144)	(4.364)	(3.390)
N	3933	2139	1072	712	3933	2139	1072	712
adj. R^2	0.341	0.441	0.337	0.299	0.345	0.446	0.338	0.301
Vintage FE	Y	Y	Y	Y	Y	Y	Y	Y
Plan State FE	Y	Y	Y	Y	Y	Y	Y	Y

location, home-state investments are not related to risk-taking, except for real estate, where local investments tend to be less risky.

To summarize, the main results in this section are as follows. Pension plans appear to have some skill in selecting the PE funds they invest in, but some plans do better than others. Across all PE investments, underfunded plans take more risk, which translates into higher IRRs but lower alphas. Pension plan boards composed of a higher fraction of state officials and appointed members of the public tend to invest in more risky PE funds, but earn lower alphas on these investments. Finally, private equity funds in the home-state of a pension plan tend to underperform, and this is not due to these investments being less risky.

4 Conclusion

Evaluating performance and the distinction between risk-taking and informed trading in observed returns are fundamental issues in the theory and practice of investments. A key step in addressing such questions is to separate the compensation for bearing risk from excess returns. With heterogeneous agents, the assessed riskiness of the same investment opportunity can vary across investors. The measures developed in this paper extend existing metrics to allow for the estimation of investor-specific risk-adjusted performance in long-dated investments, without relying on stale or potentially biased valuations.

We apply our method to U.S. public pension plan investments in private equity funds. We find that for the 1995 to 2018 sample period, pension plans appear to have pursued an optimal allocation to private equity investments overall, in the sense that there was no significant benefit in changing the investment in a representative PE fund. Looking at individual PE strategies, the average plan could have benefited from a higher allocation to buyout funds. While there is no evidence of market timing skill, pension plans did realize higher risk-adjusted returns in the funds they chose to invest in, compared to an average PE fund of the same vintage. However, this appears to be due to differences in access rather than skill in picking outperforming funds.

We find systematic differences across plans: Underfunded plans take more risk, which yields higher total returns, but lower risk-adjusted returns. Similarly, pension plan boards that have a higher fraction of state officials and appointed members of the public tend to invest in riskier funds, but earn lower alpha. Home-state PE investments also earn lower alpha, but do not differ in risk compensation, so that total returns are lower. These results are broadly consistent with agency problems within pension plans, such as gambling for resurrection and political influence, playing an important role in investment decisions. Our findings augment the prior literature that considers total, but not risk-adjusted, plan returns, and papers that only consider pension returns at a highly aggregated level.

The paper does not explore the mechanisms by which pension plans identify attractive PE funds, or the underlying reasons why certain types of board members have a preference for riskier investments. We also note a large degree of cross-pension plan dispersion in alternative asset returns, but we do not pursue pension plans' diversification strategies, and whether they price idiosyncratic risk. These and other questions we leave for future research.

References

Andonov, A., R. M. Bauer, and K. M. Cremers (2017). Pension Fund Asset Allocation and Liability Discount Rates. *The Review of Financial Studies 30*(8), 2555–2595.

Andonov, A., Y. V. Hochberg, and J. D. Rauh (2018). Political Representation and Governance: Evidence from the Investment Decisions of Public Pension Funds. *The Journal of Finance 73*(5), 2041–2086.

Barber, B. M. and A. Yasuda (2017). Interim fund performance and fundraising in private equity. *Journal of Financial Economics 124*(1), 172 – 194.

Bradley, D., C. Pantzalis, and X. Yuan (2016). The influence of political bias in state pension funds. *Journal of Financial Economics 119*, 69–91.

Brown, G. W., O. R. Gredil, and S. N. Kaplan (2019). Do private equity funds manipulate reported returns? *Journal of Financial Economics 132*(2), 267 – 297.

Cavagnaro, D. R., B. A. Sensoy, Y. Wang, and M. S. Weisbach (2019). Measuring Institutional Investors' Skill at Making Private Equity Investments. *The Journal of Finance 74*(6), 3089–3134.

Chakraborty, I. and M. Ewens (2018). Managing Performance Signals Through Delay: Evidence from Venture Capital. *Management Science 64*(6), 2875–2900.

Da Rin, M. and L. Phalippou (2017). The importance of size in private equity: Evidence from a survey of limited partners. *Journal of Financial Intermediation 31*, 64 – 76.

Dyck, A. and L. Pomorski (2015, 07). Investor Scale and Performance in Private Equity Investments. *Review of Finance 20*(3), 1081–1106.

Gredil, O. R., B. Griffiths, and R. Stucke (2022). Benchmarking Private Equity: The Direct Alpha Method. Unpublished working paper. Tulane University, Ares Management Corporation and Warburg Pincus LLC.

Harris, R. S., T. Jenkinson, S. N. Kaplan, and R. Stucke (2022). Has Persistence Persisted in Private Equity? Evidence From Buyout and Venture Capital Funds. Unpublished working paper. Darden School of Business.

Hochberg, Y. V. and J. D. Rauh (2013). Local overweighting and underperformance: Evidence from limited partner private equity investments. *Review of Financial Studies 26*(2), 403–451.

Jenkinson, T., W. R. Landsman, B. Rountree, and K. Soonawalla (2020). Private Equity Net Asset Values and Future Cash Flows. *The Accounting Review 95*(1), 191–210.

Jenkinson, T., M. Sousa, and R. Stucke (2013). How Fair are the Valuations of Private Equity Funds? Unpublished working paper. University of Oxford.

Kaplan, S. N. and A. Schoar (2005). Private Equity Performance: Returns, Persistence, and Capital Flows. *The Journal of Finance 60*(4), 1791–1823.

Korteweg, A. and S. Nagel (2016). Risk-Adjusting the Returns to Venture Capital. *The Journal of Finance 71*(3), 1437–1470.

Korteweg, A. and M. M. Westerfield (2022). Asset Allocation with Private Equity. *Foundations and Trends in Finance 13*(2), 95–204.

Lerner, J., F. Hardymon, and A. Leamon (2011). Note on the private equity fundraising process. Harvard Business School Case 9-201-042.

Lerner, J., A. Schoar, and W. Wongsunwai (2007). Smart Institutions, Foolish Choices: The Limited Partner Performance Puzzle. *The Journal of Finance 62*(2), 731–764.

Mohan, N. and T. Zhang (2014). An analysis of risk-taking behavior for public defined benefit pension plans. *Journal of Banking and Finance 40*, 403–419.

Myers, S. (2022). Public Employee Pensions and Municipal Insolvency. Unpublished working paper. University of Pennsylvania.

Pennacchi, G. and M. Rastad (2011). Portfolio allocation for public pension funds. *Journal of Pension Economics and Finance 10*(2), 221–245.

Phalippou, L. and O. Gottschalg (2009, 03). The Performance of Private Equity Funds. *The Review of Financial Studies 22*(4), 1747–1776.

Rauh, J. D. (2009). Risk Shifting versus Risk Management: Investment Policy in Corporate Pension Plans. *Review of Financial Studies 22*(7), 2687–2733.

Sensoy, B. A., Y. Wang, and M. S. Weisbach (2014). Limited partner performance and the maturing of the private equity industry. *Journal of Financial Economics 112*(3), 320 – 343.

Sharpe, W. F. (1988, December). Determining a Fund's Effective Asset Mix. *Investment Management Review 2*(6), pp. 59–69.

Sharpe, W. F. (1992). Asset allocation. *The Journal of Portfolio Management 18*(2), 7–19.

Stafford, E. (2022). Replicating Private Equity with Value Investing, Homemade Leverage, and Hold-to-Maturity Accounting. *Review of Financial Studies 35*(1), 299–342.

A Proofs

Proof. [Proof of Proposition 1]Applying Ito's Lemma to compute the dynamics of $d\log W_t$ and taking expectations, implies that for any t_k the value function $V\left(W_{t_k}, t_k\right)$ satisfies

$$
\begin{aligned}
V\left(W_{t_k}, t_k\right) &= E_{t_k} \log W_T = \\
& \log\left(W_{t_k}\right) + E_{t_k} \int_{t_k}^{T} \max_{w_u} \left(r_u + w_u' \left(\mu_u - r_u 1_{N_j}\right) - \frac{1}{2} w_u' \sigma_u \sigma_u' w_u \right) du.
\end{aligned} \tag{19}
$$

Accordingly, an implication of the envelope theorem is

$$
\frac{d\left(E_{t_k} \log W_T\right)}{d\varepsilon} = \sum_{k=0..K} V_W\left(W_{t_k}, t_k\right) C_{t_k} = \sum_{k=0..K} \frac{C_{t_k}}{W_{t_k}}. \tag{20}
$$

Therefore

$$
\frac{d\left(E_{t_k} \log W_T\right)}{d\varepsilon} = \frac{1}{W_{t_0}} \times IPE = V_W\left(t_0, W_{t_0}\right) \times IPE.
$$

■

Proof. [Proof of Corollary 2]All the steps of the proof of proposition 1 are unchanged, except that the maximization in equation (19) is no longer needed, and hence equation (20) follows without invoking the Envelope theorem. ■

Proof. [Proof of Proposition 3] For any t_k we have that that the value function $V\left(W_{t_k}, t_k\right)$ satisfies

$$
V\left(W_{t_k}, t_k\right) = \frac{E_{t_k} W_T^{1-\gamma}}{1-\gamma} = \frac{W_{t_k}^{1-\gamma}}{1-\gamma} f\left(t_k\right), \tag{21}
$$

where

$$
f\left(t_k\right) = E_{t_k} \exp \left\{ \begin{array}{c} (1-\gamma)\left(r + w'\left(\mu - r1_{N_j}\right) - \frac{1}{2} w'\sigma\sigma' w\right)(T - t_k) \\ + (1-\gamma) w'\sigma\left(B_T - B_{t_k}\right) \end{array} \right\}
$$

and the optimal portfolio is $w = \frac{1}{\gamma}\left(\sigma\sigma'\right)^{-1}\left(\mu - r1_{N_j}\right)$. Accordingly, an implication of the envelope theorem is

$$
\frac{d\left(\frac{E_{t_0} W_T^{1-\gamma}}{1-\gamma}\right)}{d\varepsilon} = E_{t_0} \sum_{k=0..K} V_W\left(W_{t_k}, t_k\right) C_{t_k} = E_{t_0} \sum_{k=0..K} W_{t_k}^{-\gamma} f\left(t_k\right) C_{t_k}. \tag{22}
$$

The Euler equation for bonds states that

$$e^{r(t_k-t_0)}E_{t_0}\frac{V_W\left(W_{t_k},t_k\right)}{V_W\left(W_{t_0},t_0\right)}=1. \tag{23}$$

Using (21) inside (23) leads to

$$e^{r(t_k-t_0)}\frac{f\left(t_k\right)}{f\left(t_0\right)}E_{t_0}\left(\frac{W_{t_k}}{W_{t_0}}\right)^{-\gamma}=1, \tag{24}$$

and using (24) inside (22) gives

$$\frac{d\left(\frac{E_{t_0}W_T^{1-\gamma}}{1-\gamma}\right)}{d\varepsilon}=f\left(t_0\right)W_{t_0}^{-\gamma}\times GIPE=V_W\left(t_0,W_{t_0}\right)\times GIPE.$$

■

Proof of Proposition 4. Let $H_t=e^{-rt}\frac{W_t^{-\gamma}}{E_{t_0}W_t^{-\gamma}}$. Noting that $\frac{W_t^{-\gamma}}{E_{t_0}W_t^{-\gamma}}$ is a martingale, an application of Ito's Lemma yields

$$\frac{dH_t}{H_t}=-rdt-\gamma w'\sigma dB_t$$

where the optimal portfolio $w=\frac{1}{\gamma}\left(\sigma\sigma'\right)^{-1}\left(\mu-r1_{N_j}\right)$. Applying Ito's Lemma again shows that

$$\frac{d\left(H_tA_t\right)}{H_tA_t}=\left(\mu_A-r-\gamma w'\sigma\sigma_A'\right)dt-\left(\gamma w'\sigma-\sigma_A\right)dB_t \tag{25}$$

Next note that

$$\begin{aligned}\beta_t\left(\mu^W-r\right) &= \beta_t w'\left(\mu-r1_{N_j}\right)=\frac{w'\sigma\sigma_A'}{w'\sigma\sigma'w}w'\left(\mu-r1_{N_j}\right)\\ &= \frac{\gamma w'\sigma\sigma_A'}{\left(\mu-r1_{N_j}\right)'\left(\sigma\sigma'\right)^{-1}\left(\mu-r1_{N_j}\right)}\left(\mu-r1_{N_j}\right)'\left(\sigma\sigma'\right)^{-1}\left(\mu-r1_{N_j}\right)\\ &= \gamma w'\sigma\sigma_A'.\end{aligned}$$

Accordingly, assumption (13) implies that $\mu_A-r=\beta_t\left(\mu^W-r\right)$ is equivalent to $\mu_A-r-\gamma w'\sigma\sigma_A'=0$. Accordingly, (25) implies that H_tA_t is a martingale. Using (1) implies that

$$H_{t_k^+}A_{t_k^+}=E_{t_k}\left(H_{t_{k+1}}A_{t_{k+1}}\right)=E_{t_k}\left(H_{t_{k+1}}A_{t_{k+1}^+}\right)+E_{t_k}\left(H_{t_{k+1}}C_{t_{k+1}}\right).$$

Iterating forward, using the law of the iterated expetation and noting that $A_{t_0} = A_{t_K^+} = 0$ implies that $GIPE = 0$. ■

Proof. [Proof of Proposition 5] With a (Markovian) time-varying opportunity set, the value function is multiplicatively separable in $W_{t_k}^{1-\gamma}$ and X_{t_k} :

$$V\left(W_{t_k}, X_{t_k}, t_k\right) = \frac{E_{t_k} W_T^{1-\gamma}}{1-\gamma} = \frac{W_{t_k}^{1-\gamma}}{1-\gamma} f\left(X_{t_k}, t_k\right), \tag{26}$$

The marginal valuation of C is now

$$\begin{aligned}
\frac{d\left(\frac{E_{t_0} W_T^{1-\gamma}}{1-\gamma}\right)}{d\varepsilon} &= E_{t_0} \sum_{k=0..K} V_W\left(W_{t_k}, t_k\right) C_{t_k} \qquad (27)\\
&= E_{t_0} \sum_{k=0..K} V_W\left(W_{t_k}, t_k\right) \left(A_{t_k} - A_{t_k^+}\right)\\
&= E_{t_0} \sum_{k=0..K-1} \left(V_W\left(W_{t_{k+1}}, t_{k+1}\right) A_{t_{k+1}} - V_W\left(W_{t_k}, t_k\right) A_{t_k^+}\right)\\
&= E_{t_0} \sum_{k=0..K-1} \int_{t_k^+}^{t_{k+1}} d(V_W\left(W_t, t\right) A_t)
\end{aligned}$$

Using Ito's Lemma gives

$$\begin{aligned}
\frac{d(V_W\left(W_t, t_k\right) A_t)}{V_W\left(W_t, t_k\right) A_t} &= \frac{d(V_W\left(W_t, t_k\right))}{V_W\left(W_t, t_k\right)} + \frac{dA_t}{A_t} + \frac{\left\langle dV_W\left(W_t, t_k\right); dA_t\right\rangle}{V_W\left(W_t, t_k\right) A_t}\\
&= -r_t dt + \frac{dA_t}{A_t} + \frac{\left\langle dW_t^{-\gamma}; dA_t\right\rangle}{W_t^{-\gamma} A_t}
\end{aligned}$$

Since $V_W\left(W_t, t_k\right) A_t$ is positive, then $E_{t_0} \sum_{k=0..K-1} \int_{t_k^+}^{t_{k+1}} d(V_W\left(W_t, t_k\right) A_t)$ has the same sign as $\mu_{A,t} - r_t + \frac{\left\langle dW_t^{-\gamma}; dA_t\right\rangle}{W_t^{-\gamma} A_t}$. Using $H_t = e^{-\int_{t_0}^{t} r_u du} \frac{W_t^{-\gamma}}{E_{t_0}\left(W_t^{-\gamma}\right)}$ as a stochastic discount factor and repeating the same steps as in (27) leads to

$$GIPE = E_{t_0} \sum_{k=0..K} H_{t_k} C_{t_k} = E_{t_0} \sum_{k=0..K-1} \int_{t_k^+}^{t_{k+1}} d(H_t A_t).$$

Applying Ito's Lemma to $d(H_t A_t)$, noting that $\frac{W_t^{-\gamma}}{E_{t_0}\left(W_t^{-\gamma}\right)}$ is a martingale and taking expectations shows that $GIPE$ has the same sign as $\mu_{A,t} - r_t + \frac{\left\langle dW_t^{-\gamma}; dA_t\right\rangle}{W_t^{-\gamma} A_t}$. ■

Proof of proposition 6. a) All terms in (16) are non-negative for all t_k. b) It suffices to show that the present value of the cash flows $\widehat{C}_{t_k}$ discounted at the benchmark rate of return is zero, $\sum_{k=t_0,t_1...t_K} \frac{\widehat{C}_{t_k}}{G_{t_k}} = 0$. We have that

$$\sum_{k=0..K} \frac{\widehat{C}_{t_k}}{G_{t_k}} 1_{\{\widehat{C}_{t_k}>0\}} = \widehat{A}_0 \times \sum_{k=0..K} \omega_{t_k} = \widehat{A}_0. \tag{28}$$

Equation (28) implies $\sum_{k=0..K} \frac{\widehat{C}_{t_k}}{G_{t_k}} = 0$, which shows that the strategy can be financed by investing the funds in the benchmark portfolio. c) By construction, $C_{t_k} = \widehat{C}_{t_k}$ whenever $\widehat{C}_{t_k} < 0$. So we focus on $C_{t_k} > 0$ to obtain

$$\begin{aligned} \widehat{C}_{t_k} &= \widehat{A}_0 \omega_{t_k} G_{t_k} \\ &= \frac{\widehat{A}_0}{\sum_{k=0..K} \frac{C_{t_k} 1_{\{C_{t_k}>0\}}}{G_{t_k}}} C_{t_k} \\ &= C_{t_k}, \end{aligned} \tag{29}$$

where the last equality follow from assumption (17) and (28). ■

B Robustness

Table 7: **Pension Plan Commitments and Performance: First-time Funds**

Each column shows the coefficients of a regression of private equity fund performance on two indicator variables plus controls. The observations only include the first-time funds by a PE manager in our sample. The indicator *Active* equals one if pension plan made a commitment to any PE fund in that vintage year. and zero otherwise. *Commit* equals one when the pension plan committed to this particular PE fund, and zero otherwise. The first two columns consider all PE funds. Columns (3) and (4) only include buyout funds, columns (5) and (6) only venture capital funds, and columns (7) and (8) only real estate funds. The odd-numbered columns use IRR as the performance metric, and the even-numbered columns use the GPME-implied α. All regressions include vintage year fixed effects. The regressions in Panel B also include pension plan fixed effects. Standard errors are double clustered by vintage year and pension plan, and are shown in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level

Panel A: Without pension plan fixed effects

	All Funds		Buyout Funds		Venture Capital Funds		Real Estate Funds	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$
Active=1	-0.0453	1.376**	0.0179	1.479**	0.0941	1.197*	0.0231	1.792**
	(0.0767)	(0.651)	(0.0907)	(0.647)	(0.168)	(0.677)	(0.0197)	(0.823)
Commit=1	0.282	0.896	-1.796	0.138	-1.561	-0.966	-1.467*	-1.420
	(1.265)	(1.194)	(1.950)	(1.432)	(3.877)	(3.383)	(0.737)	(1.119)
N	25465	25255	9020	8944	9763	9682	6682	6629
adj. R^2	0.084	0.111	0.277	0.203	0.218	0.201	0.524	0.285
Vintage FE	Y	Y	Y	Y	Y	Y	Y	Y
Pension Plan FE	N	N	N	N	N	N	N	N

Panel B: With pension plan fixed effects

	All Funds		Buyout Funds		Venture Capital Funds		Real Estate Funds	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$	IRR	$\alpha(GIPE)$
Active=1	0.0198	-0.167	-0.000805	0.0274	0.157	-0.114	0.0317	-0.127
	(0.0780)	(0.333)	(0.0689)	(0.301)	(0.164)	(0.380)	(0.0228)	(0.264)
Commit=1	0.284	0.328	-1.903	-0.334	-1.856	-2.596	-1.584*	-1.356
	(1.341)	(1.238)	(2.075)	(1.631)	(4.505)	(3.497)	(0.781)	(0.907)
N	25465	25254	9020	8944	9763	9681	6682	6629
adj. R^2	0.079	0.170	0.266	0.269	0.207	0.232	0.514	0.549
Vintage FE	Y	Y	Y	Y	Y	Y	Y	Y
Pension Plan FE	Y	Y	Y	Y	Y	Y	Y	Y