Valuing Private Equity Investments Strip by Strip

ARPIT GUPTA and STIJN VAN NIEUWERBURGH*

ABSTRACT

We propose a new valuation method for private equity investments. It constructs a replicating portfolio using cash flows on listed equity and fixed-income instruments (strips). It then values the strips using an asset pricing model that captures the risk in the cross-section of bonds and equity factors. The method delivers a risk-adjusted profit on each PE investment and a time series for the expected return on each fund category. We find negative risk-adjusted profits for the average PE fund, with substantial heterogeneity and some persistence in performance. Expected returns and risk-adjusted profit decline in the later part of the sample.

JEL codes: G24, G12

*Gupta (Corresponding Author): Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; arpit.gupta@stern.nyu.edu; http://arpitgupta.info. Van Nieuwerburgh: Department of Finance, Columbia Business School, Columbia 3022 Broadway, Uris Hall 809, New York, NY 10027; svnieuwe@gsb.columbia.edu; Tel: (212) 854-02289; https://www0.gsb.columbia.edu/faculty/svannieuwerburgh/. The authors would like to thank Stefan Nagel (the Editor), an Associate Editor, and two anonymous referees; Thomas Gilbert, Arthur Korteweg, Stavros Panageas, Antoinette Schoar, and Morten Sorensen (discussants); as well as Ralph Koijen, Clive Lipshitz, Ingo Walter, Neng Wang, and seminar and conference participants at the the 2018 UNC Real Estate Research Symposium, the 2019 AFA annual meetings, the Southern California Private Equity Conference, NYU Stern, 2019 NBER Long-Term Asset Management, the University of Michigan, SITE 2019: Asset Pricing Theory and Computation, the Western Finance Association, Arizona State University, the MIT Junior Finance Conference, the Advances in Financial Research Conference, Columbia GSB, Carnegie Mellon University, the Five-Star Conference, and the 11th Annual PERC Symposium. We thank Burgiss and the Private Equity Research Consortium for assistance with data. For underlying dividend strip data, visit https://github.com/arpitrage/Dividend_Strip. This work was supported by the NYU Stern Infrastructure Initiative. The authors have no conflicts of interest to declare.

Private equity investments have risen in importance over the past 25 years. In contrast, the number of publicly listed firms has been falling since 1997, especially among smaller firms. Private equity funds account for $5.8 trillion in assets under management, and raised nearly $800 billion in new capital in 2018 alone (Bökberg et al. (2019)). Large institutional investors such as pension and sovereign wealth funds allocate substantial fractions of their portfolios to such private investments. For example, the celebrated Yale University endowment invests more than 50% of its portfolio in alternative assets. As the fraction of overall wealth held in the form of private investments grows, so does the importance of developing appropriate valuation methods. The nontraded nature of these assets and their irregular cash flows make this a challenge.

As with any investment, the value of a private equity (PE) investment equals the present discounted value (PDV) of its cash flows. The general partner (GP, fund manager) deploys the capital committed by the limited partners (LPs, investors) by investing in a portfolio of risky projects. The risky projects pay some interim cash flows that are distributed back to the LPs. The bulk of the cash flows accrue when the GP sells the projects and distributes the proceeds, net of fees, to the LPs.

The main challenge in evaluating a PE investment is how to adjust the distributions the LPs receive for the systematic risk inherent in the cash flows. Industry practice is to report the ratio of distributions to capital contributions (TVPI) and the internal rate of return (IRR), both of which ignore the risk. Standard risk-adjustment procedures, namely, the public market equivalent (PME) approach of Kaplan and Schoar (2005) and the generalized PME approach of Korteweg and Nagel (2016), only consider aggregate stock market risk.

We propose a novel, two-step methodology that broadens the nature and refines the timing of the cash-flow risk for PE investments. In a first step, we estimate the exposure of PE funds' cash flows to the cash flows of a set of publicly listed securities. We consider a much richer cross-section of risks than prior PE literature. To capture the temporal nature of risk, we estimate the exposure of PE cash flows to dividends and capital gains on the listed factors at each cash-flow horizon. Intuitively, exposure to dividend strips captures how PE cash flows from operations covary with stocks and bonds. Exposure to gain strips captures the covariance of PE cash flows that arise from asset dispositions. Exposures may depend on the market environment at the time of fund origination. For identification, we assume that all PE funds within a given category and vintage have the same exposures to the public market strips. Estimating the many potential exposures across horizons and factors is greatly facilitated by using an elastic net approach. The first step results in a replicating portfolio of strips that has the same amount of systematic risk as the PE investment.

With the exception of zero-coupon bonds and a short history of dividend strips on the aggregate stock market, data on strip prices are not available. Therefore, the second step of the approach sets up and estimates a flexible asset pricing model to obtain dividend and capital gain strip prices for all the listed securities that feature in the replicating portfolio. This estimation is disciplined by the observed prices on nominal and real Treasury bonds of various maturities as well as by the prices and dividends on the various equity factors. The asset pricing model also delivers a time series of the expected return for all strips. The strip prices and expected returns are new and of independent interest to the asset pricing literature.

Combining the replicating portfolio of strips obtained from the first step with the strip prices from the second step, we obtain the fair price for the PE-replicating portfolio for each PE category and vintage. We define the risk-adjusted profit (RAP) of a fund as the difference between the net present value of the PE cash flows and the net present value of the replicating portfolio. A fund has a positive RAP because it delivers positive idiosyncratic cash flows or because it delivers systematic cash-flow exposure at lower cost than that available in public markets. Under the joint null hypothesis of no outperformance (after GP fees) and the correct asset pricing model, RAP should be zero.

We also obtain a time series of the expected return on each PE category-vintage pair and decompose the expected return into its various horizon components (strips) and cross-sectional exposures. This decomposition provides deeper insight into what risk sources PE investors are exposed to and compensated for. By providing expected returns of PE funds and their covariances with traded securities, our approach allows for standard portfolio analysis despite the absence of a time series for realized PE returns.

We apply our method to the universe of all PE fund categories and vintages. Our main sample contains 4,474 funds with $4.3 trillion of assets under management (AUM) across eight PE categories. We follow funds started between 1981 and 2018, and we use cash flow data through the third quarter of 2019. Buyout is the largest category with 1,145 funds and $1.9 trillion in AUM, followed by Real Estate ($592 billion), Fund of Funds ($528 billion), and Venture Capital ($489 billion). Infrastructure, Restructuring, Debt Fund, and Natural Resources make up the remaining PE categories. Our main data source is Preqin, but all results replicate on a different data set provided by Burgiss.

Our first main finding is that PE funds display substantial exposure to risk factors beyond the traditional Treasury bond and aggregate stock market factors. The nature of this factor exposure varies in ways related to the nature of the underlying assets the fund invests in. Real estate funds, for instance, take on listed real estate exposure; infrastructure fund cash flows have listed infrastructure factor exposure, and venture capital (VC) funds have distribution payoffs best proxied by growth gain strips, corresponding to a strategy of selling growth stocks. The replicating portfolio for Buyout funds includes substantial amounts of small, growth, and value dividend and capital gain strips. This accords well with a Buyout fund's strategy to buy small companies, harvest some dividends early in the life of the fund, and then gradually sell the companies near the end of the life of the fund. Ignoring these cross-sectional exposures may lead investors and researchers to conclude that VC funds have large aggregate market betas rather than average small-growth exposures, for example. This not only affects our understanding of what risks VC funds expose their investors to, but may also result in distorted inference on the expected return.

Our second main finding is that accounting for a richer factor exposure reduces average RAP across PE categories. A substantial component of the return to PE investment, which previous research has considered outperformance (abnormal return after fees), can instead be attributed to missing factor exposure. We find that the average PE fund creates little value for its LPs after accounting for a broader spectrum of risk. We estimate average RAP of -6 cents per $1 of committed capital for Buyout, -9 cents for VC, -16 cents for Real Estate, -19 cents for Fund of Funds, -0.1 cents for Restructuring, -13 cents for Debt funds, -6 cents for Infrastructure, and -6 cents for Natural Resources. The corresponding PMEs (after subtracting the initial investment) are 36 cents for Buyout, 22 cents for VC, -4 cents for Real Estate, 17 cents for Fund of Funds, 20 cents for Restructuring, 12 cents for Debt funds, 17 cents for Infrastructure, and 26 cents for Natural Resources. Hence, the richer risk adjustment turns substantial outperformance into substantial underperformance.

Third, we find that there is large cross-sectional dispersion in performance. A non-trivial fraction of funds in each category delivers a substantially positive RAP. There is persistence in the identities of GPs that outperform, and this persistence is about as large as for simpler risk-adjustment methods like PME.

Fourth, both average RAPs and expected returns have been trending downward, and are especially low in recent periods. The decline in expected returns for PE funds reflects a broad-based decline in expected returns in public markets. The decline in RAP indicates that the PE industry has not been able to repeat the early successes in Buyout and VC for funds started in the 1980s and 1990s.

Our paper makes four contributions to the literature. First, we contribute a rich stochastic discount factor (SDF) model to the asset pricing literature that provides prices of dividend and capital gain strips for cross-sectional risk factors. A literature started by Lettau and Wachter (2011), van Binsbergen, Brandt, and Koijen (2012), and van Binsbergen et al. (2013) studies claims that pay a single dividend on the aggregate stock market. Our model provides dividend strip prices for a much longer sample period than is available from futures data and for every maturity. It prices dividend strips for size, value, growth, real estate, infrastructure, and natural resource factors, shedding new light on the temporal composition of risk in the cross-section of equities. We also introduce the concept of gain strips, that is, assets that pay the realized stock price at a future date. The term structures of expected returns on dividend strips in these various factors display a range of levels and shapes that provide new targets for asset pricing models. Recent work by Weber (2018) and Giglio, Kelly, and Kozak (2020) has a similar goal.

The SDF model follows in a long tradition of combining a vector autoregression model for the state variables (Campbell (1991, 1993, 1996)) with a no-arbitrage model for the SDF (Duffie and Kan (1996), Dai and Singleton (2000), Ang and Piazzesi (2003), Lustig, Van Nieuwerburgh, and Verdelhan (2013)). Our model's state vector includes a broader cross-section of equity factors and the market prices of risk feature richer dynamics, in light of new evidence in the literature (Haddad, Kozak, and Santosh (2020)). The estimation of the market prices of risk matches a broad set of asset pricing moments. It also imposes good-deal bounds that limit the maximum Sharpe ratio (Cochrane and Saa-Requejo (2000)).

Our finding that PE funds are exposed to the cross-sectional risk factors underscores the importance of including more than just the aggregate stock market risk factor. Adjusting for this exposure results in substantially lower estimates of performance. Our study presents the most comprehensive risk adjustment to date.

In this context, it is often argued that the long-term capital lock-up feature inherent in PE fund structures should entitle GPs and/or LPs to an illiquidity premium. Like any other performance metric in the PE literature, RAP does not distinguish skill from an illiquidity premium. The fact that we find low average RAPs suggests that true skill would be lower still if there were an illiquidity premium.[1]

[1] To the best of our knowledge, there is no hard evidence for the existence of an illiquidity premium in PE. In contrast, many institutional investors such as pension funds seem to value the fact that they do

Our second contribution is a reassessment of valuation approaches that use the realized SDF. Rather than discounting PE cash flows by the realized SDF, as is done, for example, under the PME and GPME approaches, we propose using strip prices to compute the PDV of cash flows. Strip prices are *expectations* of the SDF and avoid the use of *realizations* of the SDF. Monte Carlo simulations show that using the realized SDF results in large cross-fund dispersion in value-added where there should be none. The issue is that realizations of long-horizon SDFs, used to discount long-horizon PE cash flows, tend to be far below their unconditional mean.[2] The Monte Carlo exercise shows that our proposed statistic, the RAP, is tightly estimated around the truth in samples of the size of the data: it reliably recovers GP skill (outperformance after fees) if there is skill, and recovers no skill if there is none.

Third, the insight that strip prices can be productively used to value nontraded cash flows is new to the PE literature. It applies more broadly to any valuation problem with a stream of private cash flows, such as valuing an individual private firm, building, or project. Applying the method to the PE context, we contribute to a large empirical literature on the performance evaluation of PE funds, such as Kaplan and Schoar (2005), Phalippou and Gottschalg (2009), Cochrane (2005), Harris, Jenkinson, and Kaplan (2014a), Korteweg and Sorensen (2017), and Robinson and Sensoy (2016), among many others. Most of this literature focuses on Buyout and Venture Capital funds. Recent work in valuing privately held real estate assets includes Peng (2016) and Sagi (2017). Ammar and Eling (2015) and Andonov, Kräussl, and Rauh (2020) study infrastructure investments.

not have to mark-to-market PE investments. Given that public pensions are the largest investors in PE, the equilibrium illiquidity premium may well be negative to reflect the "convenience" of the illiquidity. Recently, Asness (2019) and Riddiough (2020) express a similar view. Further exploration of this possibility represents an interesting direction for future work.

[2]The issue is similar to that in Martin (2012). Let M be the SDF and X be the cash flow. Martin (2012) shows that the sample mean of $M_1X_1, \cdots, M_tX_t$ converges to zero almost surely even though $\mathbb{E}[MX] = 1$. While this result applies only asymptotically, we find that it already has bite for the cash-flow horizons relevant to PE analysis.

The literature reports mixed results regarding PE outperformance and its persistence, depending on the data set and period in question; see Korteweg (2019) for a recent review. Our analysis spans the full sample from 1981 to 2019 and covers all PE investment categories. Our approach results in substantially lower estimates of average risk-adjusted profits for PE funds across all categories, albeit with large cross-sectional and time-series variation and some evidence of persistent outperformance for a small group of funds.

Another difference between our approach and extant methods like PME and GPME is that our approach provides exposures to risk factors.[3] The replicating-portfolio step estimates exposures of PE funds with respect to a range of cross-sectional risk factors that vary by (i) PE category, (ii) vintage, and (iii) horizon of the cash flow. This decomposition provides new insight into the nature of risk that PE investors are exposed to.

In complementary work, Ang et al. (2018) filter a time series of realized private equity returns using Bayesian methods. They then decompose that time series into a systematic component, which reflects compensation for factor risk exposure, and an idiosyncratic component (alpha). While our approach does not recover a time series of realized PE returns, it does recover a time series of *expected* PE returns. At each point in time, the asset pricing model can be used to revalue the replicating portfolio for the PE fund. Since it does not require a Bayesian estimation step, our approach is easier to implement, and hence more flexible in terms of the number of factors as well as the factor risk premium dynamics. Other important methodological contributions to PE valuation in-

[3]The GPME approach estimates market price of risk parameters, while the PME approach does not estimate any parameters. Like GPME, our approach estimates market prices of risk, but it considers a larger set of risks and allows the prices of risk to vary over time. Sorensen and Jagannathan (2015) assess the PME approach from a SDF perspective. Like ours, the PME and GPME approaches avoid making strong assumptions on the return-generating process of the PE fund because they work directly with the cash flows. Cochrane (2005) and Korteweg and Sorensen (2010) discuss this distinction. In contrast, much of the earlier literature assumes linear beta-pricing relationships, for example, Ljungqvist and Richardson (2003) and Driessen, Lin, and Phalippou (2012).

clude Driessen, Lin, and Phalippou (2012), Sorensen et al. (2014), and Metrick and Yasuda (2010).

Our fourth contribution is to use the elastic net approach to estimate fund exposures to a large set of risk factors. While the use of machine learning tools in asset pricing has gained substantial traction recently (Kozak et al. (2017), Gu, Kelly, and Xiu (2020), Karolyi and Van Nieuwerburgh (2020)), the tools are equally relevant in the PE context because (i) considering exposures to a broader range of risk factors is indispensable when valuing alternative asset categories, and (ii) the amount of PE fund data available is not that large relative to the number of exposures to be estimated. The combination of limited data and a large number of factors necessitates the use of dimension-reduction techniques.

The rest of the paper is organized as follows. Section I describes our methodology. Section II sets up and solves the asset pricing model. Section III presents our main results on the risk-adjusted profits and expected returns of PE funds. Section IV concludes. The Internet Appendix provides additional derivations (Section II), additional details on the VAR estimation (Section III), estimates on shock-exposure elasticities of our estimates (Section IV), results on additional fund categories (Section V), a validation exercise on public equities (Section VI), robustness of our estimates across different choices of hyperparameters in the elastic net estimation (Section VII), and estimates on the Burgiss data set (Section VIII).[4]

I. Methodology

PE investments are finite-horizon strategies. Upon inception of the fund, the investor (LP) commits capital to the fund manager (GP). The GP deploys that capital at her discretion, but typically within the first two to four years. Intermediate cash flows may accrue from the operation of the assets, such as net operating income from renting out an office building. Towards the end of the fund's life (typically in years 5 to 12), the GP "har-

[4]The Internet Appendix is available in the online version of this article on the *Journal of Finance* website.

vests" the assets and distributes the proceeds to the LPs after subtracting fees (including performance fees called the carry or promote). These distribution cash flows are risky. Understanding the nature of the risk in these cash flows is the key question in this paper.

Denote the sequence of net-of-fee cash-flow distributions for fund i by $\{X^i_{t+h}\}^T_{h=0}$. Time t is the inception quarter of the fund, the vintage, defined here as the quarter in which the fund GP makes the first capital call to the LPs. The horizon h indicates the number of quarters since inception; we also refer to it as the age of the fund. The maximum horizon H is set to 64 quarters to allow for "zombie" funds that continue past their expected life span of approximately 10 years. Any cash flows observed in the data after quarter 64 are discounted and allocated evenly across quarters 61 to 64. Monthly fund cash flows are aggregated to the quarterly frequency.

All PE cash flows in our data are reported for a \$1 investor commitment. In practice, GPs do not always call the full \$1, maybe because they lack profitable investment opportunities, and/or they call in the capital over multiple years. Our baseline results assume that the LP commits the present value of actual calls made, which we label $C_t \leq 1$, where the discounting uses the nominal term structure of Treasury yields. This is equivalent to assuming that the LP invests the committed but yet-to-be-called capital in a portfolio of Treasuries, with a maturity profile that matches that of the actual calls. This assumption is conservative in that it results in a higher RAP than the alternative assumption that the LP sets aside \$1 in cash, earning a zero return, while she waits for the GP to call in the committed funds. This alternative assumption penalizes the GP for time lost in deploying the money, but rewards for skillful delay. Under this alternative assumption, the LP's counterfactual investment strategy (to be compared to the actual fund investment) invests the full \$1 in the replicating portfolio at time t.[5] We return to the role of calls in Section III.G,

[5] We truncate call amounts above \$1. While infrequent, such instances occur more frequently in more recent periods. They arise because of recycling provisions that allow GPs to reinvest capital proceeds from early exits back into the fund to be drawn down later for new investments. From the LP's perspective, the call amount never exceeds \$1.

where we present results for the alternative treatment of calls. Having addressed capital calls and their timing, our method focuses on valuing the distribution of cash flows.

A. Two-Step Approach

In a first step, we use our asset pricing model to price the time series and cross-section of zero-coupon bond and equity strips. Let $\boldsymbol{F}_{t,t+h}$ be the $K \times 1$ vector of cash-flow realizations on the public securities in the replicating portfolio, where K is the number of factors. The first element of $\boldsymbol{F}_{t,t+h}$ is the payoff on a nominal zero-coupon bond that is bought at time t and matures at time $t+h$, namely, \$1. The second element of $\boldsymbol{F}_{t,t+h}$ is the payoff on a dividend strip on the aggregate stock market of maturity h, $\frac{D^m_{t+h}}{D^m_t}$. The realized dividend at time $t+h$ is scaled by the dividend at vintage origination t to create a cash flow that is comparable in magnitude to the payoff on the zero-coupon bond of \$1. The third element of $\boldsymbol{F}_{t,t+h}$ is the payoff on a capital gain strip on the aggregate stock market of maturity h, $\frac{P^m_{t+h}}{P^m_t}$. The gain strip is an asset that pays off the realized stock price at time $t+h$. This payoff is scaled by the stock price at time t. For each additional cross-sectional equity factor, we include both dividend and gain strips. For example, the payoff on a value dividend strip of horizon h is $\frac{D^{value}_{t+h}}{D^{value}_t}$ while the payoff on a value gain strip is $\frac{P^{value}_{t+h}}{P^{value}_t}$. In the full model, we have zero-coupon bonds, seven dividend strips, and seven gain strips for a total of $K = 15$ risk factors. Since there are $H = 64$ horizons, we have $KH = 960$ strips in total.

Intuitively, PE cash flows that result from operating the assets in its portfolio are akin to dividends earned on listed securities. PE cash flows that result from asset dispositions are akin to realizing capital gains on listed securities. We expect overall PE cash flows to be more strongly correlated with dividend strips early in the life cycle (small h), when few portfolio assets have been sold, while late-in-life cash flows should have greater exposure to the listed gain strips.

Denote the $K \times 1$ vector of strip prices by $\boldsymbol{P}_{t,h}$. The first element is the time-t price of a nominal zero-coupon bond of maturity h, which we also denote by $P^{\$}_{t,h}$. The second element denotes the time-t price of the h-period market dividend strip. The third element is the price of the market gain strip, and so on. The time-t price of the gain strip of maturity h is less than \$1 since the stock price at time $t+h$ reflects only the dividends after period $t+h$ while the stock price at time t reflects all dividends after period t.

Let the one-period nominal SDF be M_{t+1}. Then the h-period cumulative SDF is

$$M_{t,t+h} = \prod_{j=1}^{h} M_{t+j}.$$

The (vector of) strip prices satisfy the (system of) Euler equation

$$\boldsymbol{P}_{t,h} = \mathbb{E}_t[M_{t,t+h}\boldsymbol{F}_{t,t+h}] = \mathbb{E}_t[M_{t,t+h}]\mathbb{E}_t[\boldsymbol{F}_{t,t+h}] + Cov_t[M_{t,t+h}, \boldsymbol{F}_{t,t+h}].$$

Strip prices reflect expectations of the SDF, expectations of cash flows, and their covariance. Using strip prices to value PE cash flows avoids using the realized SDF.

In the second step of our approach we obtain the replicating portfolio, which consists of positions in dividend and gain strips, for the PE cash-flow distributions. Denote the vector of exposures of PE fund i's cash flow at time $t+h$ to the cash flow in the replicating portfolio by $\boldsymbol{\beta}^i_{t,h}$. The exposure vector describes how many units of each strip the replicating portfolio contains. We estimate the exposures from a projection of realized PE cash flows on the cash flows of the listed strips,

$$X^i_{t+h} = \boldsymbol{\beta}^i_{t,h}\boldsymbol{F}_{t,t+h} + e^i_{t+h}, \tag{1}$$

where e^i_{t+h} denotes the idiosyncratic cash-flow component, which is orthogonal to $\boldsymbol{F}_{t,t+h}$. Below explain the cross-equation restrictions imposed on the estimation of (1).

Expected Returns The expected return on PE fund i, measured over the life of the fund, is given by

$$\mathbb{E}_t\left[R^i\right] = \sum_{h=1}^{H}\sum_{k=1}^{K} w^i_{t,h}(k)\mathbb{E}_t\left[R_{t+h}(k)\right], \tag{2}$$

where $\boldsymbol{w}^i$ is a $1 \times HK$ vector of replicating portfolio weights with generic element $w^i_{t,h}(k) = \beta^i_{t,h}(k)P_{t,h}(k)$. The $HK \times 1$ vector $\mathbb{E}_t[\boldsymbol{R}]$ denotes the expected returns on the K traded asset strips at each horizon h, obtained from the asset pricing model. Fund expected returns vary over time for two reasons. First, expected returns on the listed strips vary (because the market prices of risk in the SDF model vary over time). Second, the fund exposures $\boldsymbol{\beta}^i_{t,h}$ also vary over time due to vintage effects, as explained below. Equation (2) decomposes the risk premium into the sum of compensation earned for exposure to each of the listed risk factors at each horizon, that is, "strip by strip."

The expected return in (2) is measured over the life of the fund. For comparison with IRRs, for example, it is useful to annualize it. Akin to a Macauley duration in fixed income, we define the maturity of the fund, expressed in years (rather than quarters), as

$$\delta^i_t = \frac{1}{4}\sum_{h=1}^{H}\sum_{k=1}^{K} \tilde{w}^i_{t,h}(k)h, \tag{3}$$

where the weights $\tilde{w}^i_{t,h}(k)$ are the original weights $w^i_{t,h}(k)$ rescaled to sum to one. The annualized expected PE fund return is then

$$\mathbb{E}_t\left[R^i_{ann}\right] = \left(1 + \mathbb{E}_t\left[R^i\right]\right)^{1/\delta^i_t} - 1. \tag{4}$$

This is the first main object of interest.[6]

[6]For the annualized return to correctly reflect how a dollar grows in the fund, one must first calculate the total return over the fund life by weighting life-time strip returns, and then annualize that life-time return. First annualizing strip returns and then calculating their weighted average does not result in the correct future value of the fund. This alternative approach assumes that the fund earns the average annualized strip return for the average number of years. The resulting annualized return tends to understate the

Risk-Adjusted Profit Performance evaluation of PE funds requires quantifying the LP's profit after taking into account the riskiness of the PE investment. This ex-post realized, risk-adjusted profit is the second main object of interest. Under the maintained assumption that all of the relevant sources of systematic risk are captured by the payoffs of the assets in the replicating portfolio, PE cash flows consist of one component that reflects compensation for risk and a second component that reflects a RAP.

We define the RAP for fund i in vintage t as

$$\begin{aligned} RAP_t^i &= \left(\sum_{h=1}^{H} X_{t+h}^i P_{t,h}^{\$} - C_t \right) - \left(\sum_{h=1}^{H} \sum_{k=1}^{K} \beta_{t,h}^i(k) F_{t,t+h}(k) P_{t,h}^{\$} - \beta_{t,h}^i(k) P_{t,h}(k) \right) \\ &= \sum_{h=1}^{H} e_{t+h}^i P_{t,h}^{\$} + \left(\sum_{h=1}^{H} \sum_{k=1}^{K} \beta_{t,h}^i(k) P_{t,h}(k) - C_t \right). \end{aligned} \tag{5}$$

The risk-adjusted profit is the difference between the net present value (NPV) of the PE fund and the NPV of the replicating portfolio. The NPV of the PE fund equals the future cash flows of the PE fund, discounted at the risk-free term structure of interest rates (recall that nominal bond prices are $P_{t,h}^{\$}$), minus the $\$C_t \leq 1$ of capital committed to the fund. Apart from discounting, the first term would be the traditional TVPI measure that captures cash received out of the investment relative to cash put in. It is also like the PME measure, except that Treasury yields are used for discounting rather than the aggregate stock market return. The second term measures the NPV of the replicating portfolio: the discounted value of all realized cash flows minus the cost of purchasing the replicating portfolio.

Rewriting, we can express the RAP as the sum of two components. The first is the discounted sum of the idiosyncratic fund cash flows e^i. Since the idiosyncratic cash flows are orthogonal to the priced cash flows, they are discounted at the risk-free interest rate. The second component is the difference between the purchase price of the replicating

expected return if the term structure of expected strip returns is flat or upward-sloping and overstate the expected return if the term structure is downward-sloping.

portfolio of strips and the purchase price of the PE fund, C_t.

The first term in (5) attributes out-performance to funds that deliver a stream of high idiosyncratic cash flows, by selecting the right portfolio of assets (asset selection skill). The second term in RAP credits PE funds with outperformance to the extent that they are able to deliver a set of factor exposures at an (after-fee) cost C_t that is lower than the cost of that portfolio in public asset markets. An outperforming fund is one that generates cash flows without taking commensurate risk. This can manifest as a replicating portfolio that contains a large quantity of risk-free bonds. Since risk-free bonds are valuable, the second term is positive.

A PE fund with market timing skills, which buys assets at the right time (within the investment period) and sells at the right time (within the harvesting period) will have a positive RAP.[7] The RAP measure does not credit the GP for lucky realizations of the risk factors.

The null hypothesis of no outperformance is $\mathbb{E}[RAP_t^i] = 0$, where the expectation is taken across funds. Under the null, the idiosyncratic cash flows average to zero across funds, and purchasing a portfolio of strips that has the same systematic risk as the PE fund has the same cost as the PE fund itself.

This null is a joint null of also having a correctly specified SDF. All relevant risk factors for the evaluation of PE cash flows are included. Much of the modern asset pricing literature finds that a fairly low-dimensional factor structure spans the cross-section of stock returns (e.g., Kozak, Nagel, and Santosh (2017), Gu, Kelly, and Xiu (2020)). The 15 cross-sectional factors that we include should go a long way towards capturing this factor structure. Our method can easily accommodate extra factors.[8]

[7] The fund's horizon is endogenous because it is correlated with the success of the fund. As noted by Korteweg and Nagel (2016), this endogeneity does not pose a problem as long as cash flows are observed. They write, "Even if there is an endogenous state-dependence among cash-flows, the appropriate valuation of a payoff in a certain state is still the product of the state's probability and the SDF in that state."

[8] Omitted risk factors may bias RAP estimates. Which way the bias goes depends on the covariance of the

The RAP measure can be computed for each fund. To assess the performance of PE funds, we report the distribution of RAP across all funds in the sample as well as the equal-weighted average RAP by vintage. When calculating our RAP measure (and only then), we exclude vintages after 2010 Q4 for which we are still missing a substantial fraction of the cash flows as of 2020.

B. Identifying and Estimating Cash-Flow Exposures

The replicating portfolio must be rich enough that it spans all priced sources of risk, yet it must be parsimonious enough that its exposures can be estimated with sufficient precision. Allowing every fund in every category and vintage to have its own unrestricted cash-flow exposure profile for each risk factor leads to parameter proliferation and lack of identification. We impose cross-equation restrictions to aid identification.

Identifying Assumptions Identification is achieved both from the cross-section and from the time series. We make four assumptions. First, the cash flows $X^{i\in c}_{t+h}$ of all funds i in category c (category superscripts are omitted below for ease of notation) and vintage t have the same risk factor exposures at horizon h, $\beta^i_{t,h}(k) = \beta^c_{t,h}(k), \forall i \in c$. We drop the category superscript in what follows but note that exposures are estimated separately for every fund category. Second, the risk exposures $\beta_{t,h}(k)$ are the sum of a *vintage effect* a^k_t and an

omitted factor with the included factors and on the magnitudes of omitted and included strip prices. If the PE cash flows have a strong positive loading on the omitted factor's cash flow, the direct effect is to reduce the cost of the replicating portfolio since the portfolio does not buy the omitted strips by definition. The indirect effect is that the missing exposure will be partially picked up by higher exposures to the included factors. More of the included strips will be bought, increasing the cost of the replicating portfolio. The net effect is ambiguous, and depends on the net change in the cost of the replicating portfolio under the misspecified model. As an example, we consider the setting of Section I.A of the Internet Appendix, which assumes that the true model is a three-factor model. We estimate RAP for a misspecified model that omits the third factor. We find that the direct effect dominates because the included factors are not that highly correlated with the omitted factor. The result is a downward bias for RAP.

age effect b_h^k for each factor k. Third, horizon effects are constant for the four quarters in a calendar year. This reduces the number of horizon effects that need to be estimated for each factor from $H = 64$ to $H/4 = 16$.[9] Fourth, the vintage effects depend on the price-dividend ratio of the aggregate stock market in the quarter of fund inception: $a_t^k = a_{pd(t)}^k$. The vintage effects thus capture dependence on the overall investment climate at the time of PE fund origination. Haddad, Loualiche, and Plosser (2017) emphasize the importance of price-divident ratios and aggregate equity premia in explaining Buyout activity. The choice of the pd_t^m ratio is also motivated by the asset pricing model of Section II, where the pd_t^m ratio is one of the key state variables driving time-variation in risk premia.[10] To simplify the time dimension, vintages are separated into four groups by the quartile of the pd_t^m ratio distribution at the time of fund inception. Quartile breakpoints are based on the full 1974 to 2019 sample. Only three of the four vintage effects are identified so we normalize the vintage effects to zero on average across quartiles.

To summarize, we estimate $(H/4 + 3) \times K$ risk exposures for each fund category, rather than $H \times T \times K$ exposures in an unrestricted model. We use $N_f \times T \times H$ fund cash-flow observations to do so, where T reflects the number of different vintage quarters in the sample, N_f the average number of funds in a category per vintage quarter, and H the life span of a fund in quarters ($H = 64$). We include all available vintages in estimating the exposures, including very recent ones, because the early cash flows from recent vintages still aid in the estimation of the first few elements of b_h and the vintage effects $a_{pd(t)}$.

Two-Factor Model We start with a model in which all PE cash flows are only exposed to bonds and aggregate stock market capital gain strips. We refer to this as the two-factor

[9] We find no systematic evidence of seasonality in PE fund cash flows.

[10] A natural alternative would be to consider total capital raised by vintage. One challenge with this metric is that it is nonstationary since the PE industry has grown. Even scaling by GDP does not remove this trend. Nevertheless, we find similar results under this alternative.

model ($K = 2$). Fund cash flows for the two-factor model can be expressed as

$$\begin{aligned} X_{t+h}^{i \in c} &= \beta_{t,h}^{bond} + \beta_{t,h}^{equity} F_{t,t+h}^{equity} + e_{t+h}^{i} \\ &= a_{pd(t)}^{bond} + b_{h}^{bond} + \left(a_{pd(t)}^{equity} + b_{h}^{equity} \right) F_{t,t+h}^{equity} + e_{t+h}^{i}. \end{aligned} \tag{6}$$

We estimate equation (6) by OLS. This model extends the PME and GPME approaches in that it (i) estimates a richer SDF model with time-varying prices of bond and aggregate stock market risk, and (ii) results in exposures of fund cash flows that differ by vintage and cash-flow horizon.

K-Factor Model Our main model is a K-factor model in which we add cross-sectional equity market factors beyond the two factors from the previous model to better capture the systematic risk in PE fund cash flows. PE fund cash flows are modeled as

$$\begin{aligned} X_{t+h}^{i \in c} &= \beta_{t,h}^{bond} + \sum_{k=2}^{K} \beta_{t,h}(k) F_{t,t+h}(k) + e_{t+h}^{i} \\ &= a_{pd(t)}^{bond} + b_{h}^{bond} + \sum_{k=2}^{K} \left(a_{pd(t)}^{k} + b_{h}^{k} \right) F_{t,t+h}^{k} + e_{t+h}^{i}. \end{aligned} \tag{7}$$

In the empirical implementation, $K = 15$. The 15 factors are bond strips plus both dividend strips and capital gain strips on seven equity factors: the aggregate stock market, small stocks, growth stocks, value stocks, REITs, infrastructure stocks, and natural resource stocks. Under our identifying assumptions, we estimate $3K = 45$ vintage effects (the a's) and $KH/4 = 240$ age effects (the b's) for a total of 285 coefficients. The replicating portfolio for PE funds takes time-varying positions in $KH = 960$ strips that are obtained from these 285 coefficients.

Because of the large number of exposure coefficients to be estimated and the relative data scarcity, it is paramount to use a dimension-reduction technique. We use the well-known elastic net approach, which selects only some of the 285 potential exposure coefficients and shrinks others to zero. Using dimension reduction avoids having to take a stance on the identity of a small number of factors that drive PE cash flows, a problem with the OLS approach. Furthermore, we impose a non-negativity constraint on all estimated positions in the replicating portfolio. This avoids spurious long-short positions that arise due to the high correlation among some of the listed factors, as well as difficulties and costs related to taking on short positions (especially in cross-sectional risk factors) that investors face in reality. For example, some pension funds may be prohibited from taking short positions.

The elastic net estimation of equation (1) can be written as

$$\hat{\beta}_{EN} = \arg\min_{\beta \in \mathbf{R}^{KH}} \|X^i_{t+h} - \boldsymbol{\beta}_{t,h}\boldsymbol{F}_{t,t+h}\|_2^2 + \lambda_0 \mathbf{1}\{\beta < 0\} + \lambda \left[(1-\alpha)\|\beta\|_2^2/2 + \alpha\|\beta\|_1\right]. \quad (8)$$

We set the hyper-parameter $\lambda_0 = \infty$, which ensures only positive coefficients. The parameter α governs the lasso component, zeroing out a subset of coefficients (factor selection), and λ is the ridge regression penalty, which shrinks the magnitude of coefficient estimates closer to zero. Setting $\alpha = 1$ reduces the problem to the case of a lasso specification only, while $\alpha = 0$ corresponds to the ridge regression. The λ parameter determines the total penalty amount. We use cross-validation to tune the hyper-parameters α and λ for each fund category separately. Section VII of the Internet Appendix details the hyper-parameter choices and shows robustness of the results to these choices.

C. *Approaches that Use Realized SDF*

If one has estimated a rich SDF that fits the listed asset data well, why not directly use the realized SDF time series and the fund cash flows to measure firms' value-added? Define the value-added of a PE fund using the realized SDF as

$$VAsdft^i_t = \sum_{h=1}^{H} M_{t,t+h} X^i_{t,t+h} - C_t. \quad (9)$$

While in expectation this measure is identical to the RAP measure, sample averages show wide divergence. A Monte Carlo study in Section I of the Internet Appendix shows that the distribution of the *VAsdf* across funds and vintages shows left-skewness and substantial dispersion. Martin (2012) shows that there is a generic problem with discounting a long-run stream of cash flows using the realized SDF. Even though his results are asymptotic in nature, the problem materializes at horizons relevant for PE valuation. To compensate for a few high realizations of $M_{t,t+h}$ (aggregate disasters), most SDF realizations need to be vanishingly small so as to enforce the Euler equation. The first term in (9) converges to zero almost surely as h is very large, but is already close to zero for $h < 64$ in rich models of the SDF. For the SDF that we estimate in the next section, the estimated mean of *VAsdf* is indeed close to $-C_t$. According to the $VAsdf$ estimation, LPs lose their entire investment on a risk-adjusted basis when, in truth, funds have zero value-added. This problem arises not only for our rich SDF model but also in much simpler SDF models such as the CAPM, as we show in Section I of the Internet Appendix.

Our RAP approach in (5) is fundamentally different because it uses strip prices, which are *expectations* of the SDF. This avoids having to use *realizations* of the SDF (times the cash flow).

Our approach requires estimation of exposures or PE fund cash flows to strip cash flows, but this has the advantage of generating useful information about the underlying risk exposures of PE funds. Section I of the Internet Appendix shows that our strip-based approach reliably recovers the true risk factor exposures, and results in a RAP distribution that is tightly centered around zero if in truth RAP is zero for each fund: it recovers skill (heterogeneity) where there truly is skill (heterogeneity), and finds no evidence for skill when there is none.

II. Asset Pricing Model

The second main step is to obtain prices for the dividend and gain strips. If the only source of risk were fluctuations in the term structure of interest rates, this step would be straightforward. After all, we directly observe zero-coupon bond prices of all maturities at each date. However, fluctuations in interest rates are not the only, and indeed not even the main, source of risk in the cash flows of PE funds. If fluctuations in the aggregate stock market were the only other source of aggregate risk, then we could use price information from dividend strips. Those prices can be observed directly from dividend strip futures markets (van Binsbergen et al. (2013)) or inferred from options and stock markets (van Binsbergen, Brandt, and Koijen (2012)). However, the available time series is too short for our purposes – strips are not available for horizons beyond seven years and do not come in one-quarter-horizon increments. Moreover, the only dividend strip data correspond to the aggregate stock market – there are no strip data for the additional traded factors we wish to include in our analysis such as publicly listed real estate or infrastructure assets or a small stock, value stock, or growth stock index. Finally, we do not observe expected returns on the available strips, only realized returns, expected returns are difficult to infer from short time series of realized returns. We therefore need an asset pricing model to generate the time series of strip prices, $\boldsymbol{P}_{t,h}$, and the corresponding expected returns for each strip.

We propose a reduced-form SDF model that prices publicly traded assets well, including the available dividend strip data. A virtue of the reduced-form model is that it can accommodate a substantial number of risk factors. We argue that it is important to go beyond the aggregate stock and bond markets to capture the risk embedded in PE fund cash flows. As Korteweg and Nagel (2016) note, the objective is not to test the asset pricing model itself but rather to investigate whether a potential PE investment adds value to an investor who already has access to securities whose sources of risk are captured by the SDF. In complementary work, Gredil, Sorensen, and Waller (2020) investigate consumption-based asset pricing models' ability to price PE cash flows.

A. Setup

A.1. State Variable Dynamics

Time is denoted in quarters. We assume that the $N \times 1$ vector of state variables follows a Gaussian first-order VAR,

$$z_t = \boldsymbol{\Psi} z_{t-1} + \boldsymbol{\Sigma}^{\frac{1}{2}} \boldsymbol{\varepsilon}_t, \tag{10}$$

with shocks $\boldsymbol{\varepsilon}_t \sim i.i.d.\mathcal{N}(0, I)$ whose variance is the identity matrix. The companion matrix Ψ is a $N \times N$ matrix. The vector z is demeaned. The covariance matrix of the innovations to the state variables is Σ; the model is homoskedastic. We use a Cholesky decomposition of the covariance matrix, $\Sigma = \Sigma^{\frac{1}{2}} \Sigma^{\frac{1}{2}\prime}$, which has nonzero elements only on and below the diagonal. The Cholesky decomposition of the residual covariance matrix allows us to interpret the shock to each state variable as the shock that is orthogonal to the shocks of all state variables that precede it in the VAR. We discuss the elements of the state vector and their ordering below. The (demeaned) one-quarter bond nominal yield is one of the elements of the state vector: $y^{\$}_{t,1} = y^{\$}_{0,1} + e'_{yn} z_t$, where $y^{\$}_{0,1}$ is the unconditional average one-quarter nominal bond yield and e_{yn} is a vector that selects the element of the state vector corresponding to the one-quarter yield. Similarly, the (demeaned) inflation rate is part of the state vector: $\pi_t = \pi_0 + e'_{\pi} z_t$ is the (log) inflation rate between $t-1$ and t. Lowercase letters denote logs.

A.2. Stochastic Discount Factor

The nominal SDF $M^{\$}_{t+1} = \exp(m^{\$}_{t+1})$ is conditionally log-normal,

$$m^{\$}_{t+1} = -y^{\$}_{t,1} - \frac{1}{2} \boldsymbol{\Lambda}'_t \boldsymbol{\Lambda}_t - \boldsymbol{\Lambda}'_t \boldsymbol{\varepsilon}_{t+1}. \tag{11}$$

Note that $y^{\$}_{t,1} = -\mathbb{E}_t[m^{\$}_{t+1}] - 0.5\mathbb{V}_t[m^{\$}_{t+1}]$. The real log SDF $m_{t+1} = m^{\$}_{t+1} + \pi_{t+1}$ is also conditionally Gaussian. The innovations in the vector ε_{t+1} are associated with a $N \times 1$ market price of risk vector $\mathbf{\Lambda}_t$ of the affine form,

$$\mathbf{\Lambda}_t = \mathbf{\Lambda}_0 + \mathbf{\Lambda}_1 z_t. \tag{12}$$

The $N \times 1$ vector $\mathbf{\Lambda}_0$ collects the average prices of risk while the $N \times N$ matrix $\mathbf{\Lambda}_1$ governs the time-variation in risk premia. Asset pricing amounts to estimating the market prices of risk $(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1)$.

A.3. Bond Pricing

Proposition 1 in Section II of the Internet Appendix shows that nominal bond yields of maturity τ are affine in the state variables,

$$y^{\$}_{t,\tau} = -\frac{1}{\tau} A^{\$}_{\tau} - \frac{1}{\tau}\left(B^{\$}_{\tau}\right)' z_t.$$

The scalar $A^{\$}(\tau)$ and the vector $B^{\$}_{\tau}$ follow ordinary difference equations that depend on the properties of the state vector and on the market prices of risk. The appendix also calculates the real term structure of interest rates, the real bond risk premium, and the inflation risk premium on bonds of various maturities. We include the cross-section of nominal and real bond yields (price levels) in the set of moments used to estimate the market price of risk coefficients. We put more weight on matching the time series of one- and 20-quarter nominal bond yields since those yields are part of the state vector z_t. We also fit the dynamics of 20-quarter nominal bond risk premia (price changes).

A.4. Equity Pricing

The VAR contains both the log price-dividend ratio and log dividend growth for each equity risk factor. Together these two time series imply a time series for log stock returns.

The VAR implies linear dynamics for the expected excess stock return, or equity risk premium, for each equity risk factor. We choose market prices of risk to match these dynamics (price changes).

The price of a stock equals the PDV of its future cash flows. By value-additivity, the price of the aggregate stock index, P_t^m, is the sum of the prices to each of its future cash flows D_t^m. These future cash-flow claims are the so-called market dividend strips or zero-coupon equity (Wachter (2005)). Dividing by the current dividend D_t^m yields

$$\frac{P_t^m}{D_t^m} = \sum_{\tau=1}^{\infty} P_{t,\tau}^d \tag{13}$$

$$\exp\left(\overline{pd} + e'_{pd^m} z_t\right) = \sum_{\tau=0}^{\infty} \exp\left(A_\tau^m + B_\tau^{m\prime} z_t\right), \tag{14}$$

where $P_{t,\tau}^d$ denotes the price of a τ-period dividend strip divided by the current dividend. Proposition 2 in Section II of the Internet Appendix shows that the log price-dividend ratio on each dividend strip, $p_{t,\tau}^d = \log\left(P_{t,\tau}^d\right)$, is affine in the state vector and provides recursions for the coefficients $\left(A_\tau^m, \boldsymbol{B}_\tau^m\right)$. Since the log price-dividend ratio on the stock market is an element of the state vector, it is affine in the state vector by assumption. Equation (14) restates the present value relationship from equation (13). It articulates a nonlinear restriction on the coefficients $\{(A_\tau^m, B_\tau^m)\}_{\tau=1}^{\infty}$ at each date (for each state z_t), which we impose in the estimation (price levels). Analogous present value restrictions are imposed for each of the six other traded equity factors, whose price-dividend ratios and dividend growth rates are also included in the state vector.

If dividend growth were unpredictable and its innovations carried a zero price of risk, then dividend strips would be priced like real zero-coupon bonds. The strips' dividend-price ratios would equal yields on real bonds with the coupon adjusted for deterministic dividend growth. All variation in the price-dividend ratio would reflect variation in the real yield curve. In reality, the dynamics of real bond yields account for only a small fraction of the variation in the price-dividend ratio, implying large prices of risk associated with shocks to dividend growth that are orthogonal to shocks to bond yields.

A.5. Dividend Futures

The model readily implies the price of a futures contract that receives the single realized nominal dividend at some future date, $D^{\$}_{t+k}$. That futures price, $F^{d}_{t,\tau}$, scaled by the current nominal dividend $D^{\$}_{t}$, is

$$\frac{F^{d}_{t,\tau}}{D^{\$}_{t}} = P^{d}_{t,\tau} \exp\left(\tau y^{\$}_{t,\tau}\right).$$

The one-period realized return on this futures contract for $k > 1$ is

$$R^{fut,d}_{t+1,\tau} = \frac{F^{d}_{t+1,\tau-1}}{F^{d}_{t,\tau}} - 1.$$

Section II of the Internet Appendix shows that $\log(1 + R^{fut,d}_{t+1,\tau})$ is affine in the state vector z_t and in the shocks ε_{t+1}. It is straightforward to compute average realized returns over any subsample and for any portfolio of futures contracts.

B. Estimation

B.1. State Vector Elements

The state vector contains $N = 18$ variables. The first six variables, in order of appearance, are (1) GDP price inflation, (2) real GDP growth, (3) the nominal short rate (three-month nominal Treasury bill rate), (4) the spread between the yield on a five-year Treasury note and a three-month Treasury bill, (5) the log price-dividend ratio on the CRSP value-weighted stock market, (6) the log real dividend growth rate on the CRSP stock market. Variables 7, 9, 11, 13, 15, and 17 are the log price-dividend ratios on the REIT index of publicly listed real estate companies, a listed infrastructure index (infra), the first size quintile of stocks (small), the first book-to-market quintile of stocks (growth), natural resource stocks (nr), and the fifth book-to-market quintile of stocks (value), while variables 8, 10, 12, 14, 16, and 18 are the corresponding log real dividend growth rates:[11]

$$
\begin{aligned}
z_t = \quad & \Big[\pi_t, x_t, y^{\$}_{t,1}, y^{\$}_{t,20} - y^{\$}_{t,1}, pd^{m}_t, \Delta d^{m}_t, pd^{reit}_t, \Delta d^{reit}_t, pd^{infra}_t, \Delta d^{infra}_t, \\
& pd^{small}_t, \Delta d^{small}_t, pd^{growth}_t, \Delta d^{growth}_t, pd^{nr}_t, \Delta d^{nr}_t, pd^{value}_t, \Delta d^{value}_t\Big]'.
\end{aligned} \tag{15}
$$

This state vector is observed at the quarterly frequency from 1974 Q1 until 2019 Q4 (184 observations). This is the longest time series for which all variables are available.[12] Our PE cash-flow data start shortly thereafter in the early 1980s. While the bulk of PE cash flows occur after 1990, we use the longest possible sample to more reliably estimate the VAR dynamics and especially the market prices of risk. All state variables are demeaned with the observed full-sample mean.[13]

The VAR is estimated using OLS in the first stage of the estimation. We recursively

[11]The ordering of the state variables is not that important for our purposes since we are not interested in structurally interpreting the risk prices, but rather in finding a good fit for the asset pricing moments.

[12]We use the average of daily constant-maturity Treasury yields within the quarter. The REIT index is the NAREIT All Equity index, which excludes mortgage REITs. The first observation for REIT dividend growth is in 1974 Q1. All dividend series are deseasonalized by summing dividends across the current month and past 11 months. This means we lose the first eight quarters of data in 1972 and 1973 when computing dividend growth rates. The infrastructure stock index is measured as the value-weighted average of the eight relevant Fama-French industries (Aero, Ships, Mines, Coal, Oil, Util, Telcm, Trans). The natural resource index is measured from the Alerian Master Limited Partnership from 1996 Q1 thereafter and as the Fama-French Oil industry index beforehand.

[13]The VAR literature finds that results can be sensitive to the choice of state variables. Campbell and Voulteenaho (2004) emphasize the role of the small value spread as a predictor of aggregate market returns, while Liu and Zhang (2008) suggest using the market-to-book spread of value-minus-growth and the book-to-market spread of value-minus-growth as separate predictors. Haddad, Kozak, and Santosh (2020) argue that each risk factor's expected return is driven by its own dividend-price ratio. Our model allows each factor's *pd* ratio to affect the expected return, but also allows for cross-predictability and for the level and slope of the yield curve to predict stock returns, for example. With $N = 18$ state variables, our VAR model is large and includes the candidate state variables highlighted in the literature.

zero out all elements of the companion matrix Ψ whose t-statistic is below 1.96. Section III of the Internet Appendix contains the resulting point estimates for Ψ and $\Sigma^{\frac{1}{2}}$.

B.2. Market Prices of Risk

The state vector contains both priced sources of risk as well as predictors of bond and stock returns. We estimate 12 nonzero parameters in the constant market price of risk (MPR) vector Λ_0 and 92 nonzero elements of the matrix Λ_1, which governs the dynamics of the risk prices. The point estimates are reported in Section III.B of the Internet Appendix. We employ the following target moments to estimate the MPR parameters.

First, we target the average 20-quarter bond yield and its dynamics. This delivers one restriction on Λ_0 and $N = 18$ restrictions on Λ_1:

$$-A^{\$}_{20}/20 = y^{\$}_{0,20} \quad \text{and} \qquad -B^{\$}_{20}/20 = [0,0,1,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0].$$

Because the demeaned five-year bond yield is the sum of the third and fourth elements in the state vector, the MPR must be such that $-B^{\$}_{20}/20$ has a value of one in the third and fourth places and zeroes everywhere else.

Second, we match the time series of nominal bond yields for maturities of one quarter, one year, two years, five years, 10 years, 20 years, and 30 years. This leads to about $7 \times T$ moments, where $T = 184$ quarters.[14]

Third, we match the time series of real bond yields for maturities of five, seven, 10, 20, and 30 years. They constitute about $5 \times T_2$ moments, where $T_2 = 68$ quarters.[15] Having both nominal and real bonds helps disentangle the respective roles of growth and inflation risks.

[14] The 20-year bond yield is missing prior to 1993 Q4 while 30-year bond yield data are missing from 2002 Q1 to 2005 Q4. In total 107 observations are missing, so that we have 1,232-107=1,125 bond yields to match.

[15] FRED data on Treasury Inflation Indexed bond yields start in 2003 Q1. Real yields for the 20-year and 30-year bonds are available only for 61 and 39 quarters, respectively.

Fourth, we require that the time series of risk premia for the aggregate stock market, real estate stocks, infrastructure stocks, small stocks, growth stocks, natural resource stocks, and value stocks match the expected excess returns implied by the VAR, that is, from the data. The expected excess return in logs, including a Jensen adjustment, equals minus the conditional covariance between the log SDF and the log return. For example, for the aggregate stock market we have

$$\begin{aligned}
E_t\left[r_{t+1}^{m,\$}\right] - y_{t,1}^{\$} + \frac{1}{2}V_t\left[r_{t+1}^{m,\$}\right] &= -Cov_t\left[m_{t+1,}^{\$} r_{t+1}^{m,\$}\right] \\
r_0^m + \pi_0 - y_0^{\$}(1) + \left[(e_{divm} + \kappa_1^m e_{pd} + e_\pi)'\Psi - e_{pd}' - e_{yn}'\right] z_t & \\
+\frac{1}{2}\left(e_{divm} + \kappa_1^m e_{pd} + e_\pi\right)' \Sigma \left(e_{divm} + \kappa_1^m e_{pd} + e_\pi\right) &= \left(e_{divm} + \kappa_1^m e_{pd} + e_\pi\right)' \Sigma^{\frac{1}{2}} \Lambda_t.
\end{aligned}$$

The left-hand side is given by the VAR (data), while the right-hand side is determined by the market prices of risk Λ_0 and Λ_1 (model). This provides $(N+1) \times 7$=133 additional restrictions. These moments identify the 6th, 8th, 10th, 12th, 14th, 16th, and 18th elements of Λ_0 and the corresponding rows of Λ_1.

Fifth, we match the time series of log price-dividend ratios (price levels) on the seven stock indices. The model's price-dividend ratios come from 3,600 quarterly dividend strips according to equation (13). We impose these present value relationships in each quarter, which results in $7 \times T$ moments.

Sixth, we price a claim that pays the next eight quarters of realized nominal dividends on the aggregate stock market. The value of this claim is the sum of the prices on the nearest eight dividend strips. Data on the price-dividend ratio of this claim and the fraction it represents of the overall stock market value (S&P500) for the period 1996 Q1 to 2009 Q3 (55 quarters) come from van Binsbergen, Brandt, and Koijen (2012). This procedure delivers 2×55 moments. We also ensure that the model is consistent with the high average realized returns on short-horizon dividend futures documented by van Binsbergen et al. (2013). Table 1 in van Binsbergen and Koijen (2017) reports that the average return on an equally weighted portfolio of one- through seven-year U.S. SPX dividend futures over the period November 2002 to July 2014 is 8.71% per year. We construct an average return for the same short-maturity futures portfolio (paying dividends 2 to 29 quarters from now) in the model:

$$R_{t+1}^{fut,portf} = \frac{1}{28} \sum_{\tau=2}^{29} R_{t+1,\tau}^{fut,d}.$$

We evaluate the realized return on this dividend futures portfolio using the state variables observed between 2003 Q1 and 2014 Q2, average it, and annualize it. This procedure results in one additional restriction. These dividend strip moments identify (some of) the MPR parameters associated with the market price-dividend ratio shock (fifth element of Λ_0 and first six elements of the fifth row of Λ_1).

Seventh, we impose a good deal bound on the standard deviation of the log SDF, the maximum Sharpe ratio, in the spirit of Cochrane and Saa-Requejo (2000). We also impose a penalty on choosing excessively large values for Λ_t. These $1 + T$ constraints help reduce the entropy of the SDF.

Eighth, we impose regularity conditions on bond yields. We require that very long-term real bond yields have average yields that weakly exceed average long-run real GDP growth, which is 2.63% per year in our sample. Long-run nominal yields must exceed long-run real yields by 2%, an estimate of long-run average inflation. Nominal and real bond yields must flatten out as the maturity grows. These regularity conditions are satisfied at the final solution.

Not counting the regularity conditions, we have $5,245$ moments to estimate 116 MPR parameters. Section III.C of the Internet Appendix contains a detailed discussion of the estimation algorithm and argues that the parameters are identified.

Still, concerns about estimation error around the many MPR point estimates are natural.[16] In the PE analysis of the next section we take these MPR estimates as given. For

[16]We have calculated standard errors on the MPR parameters using GMM and find them to be modest. Intuitively,by matching the time series of bond yields of various maturities, the stock price dividend ratios,

the purposes of PE performance measurement this is fine. An analogy to the CAPM may be useful here. If the β^i of fund i were known, we could estimate abnormal returns as $avg(R^i) - \beta^i avg(R^m) = \alpha + avg(e^i)$. The "MPR estimate" inside $avg(R^m)$ drops out. It is efficient to subtract $\beta\, avg(R_m)$, not $\beta\, \mathbb{E}[R^m]$.[17]

B.3. Model Fit

Figure 1 plots the nominal bond yields on bonds of maturities one quarter, one year, five years, and 10 years. Those are the most relevant horizons for the PE cash flows. The model matches the time series of bond yields in the data closely. It matches nearly perfectly the one-quarter and five-year bond yields, which are part of the state space. Figure 2 shows that the model also does a good job matching real bond yields. The top panels of Figure 3 show the model's implications for the average nominal (left panel) and real (right panel) yield curves at longer maturities. These long-term yields are well behaved. The bottom right panel shows a decomposition of the yield on a five-year nominal bond into the five-year real bond yield, annual expected inflation over the next five years, and the five-year inflation risk premium. The importance of these components fluctuates over time. The bottom left panel shows that the model matches the dynamics of the nominal bond risk premium, defined as the expected excess return on five-year nominal bonds. The compensation for interest rate risk varies substantially over time, both in data and in the model.

---

and the expected returns on each equity factor, as well as by imposing an adding-up constraint that the dividend strip prices add up to the total equity price for each equity index, there is little room for the estimation to arrive at sufficiently different strip prices that it would make a material difference for PE performance analysis.

[17]We thank the Editor for pointing this analogy out to us. We note that this argument requires estimating $avg(R^i)$ and $avg(R^m)$ on the sample sample. We estimate our MPR on a slightly longer sample (1974 to 2019) than we have for the fund data (1981 to 2019). In a robustness check, we have reestimate our MPR on the 1990 to 2019 subsample, over which most of the fund data are concentrated. The results are similar. The code for this subsample is provided as part of the code replication package.

— Figures 1–3 go about here —

Figures 4 and 5 show the equity risk premium and the expected excess return in the left panels and the price-dividend ratio in the right panels. The various rows cover the seven equity indices that we price. The dynamics of the risk premia in the data are dictated by the VAR. The model chooses the MPR to fit these risk premium dynamics as closely as possible. The price-dividend ratios in the model are based on the price-dividend ratios on the strips of maturities ranging from 1 to 3,600 quarters, as explained above. The figures show an excellent fit for price-dividend levels and a good fit for risk premium dynamics. Some of the VAR-implied risk premia have outliers which the model does not fully capture. This is due in part to the fact that the good deal bounds restrict the SDF from becoming too volatile and extreme. We note large level differences in valuation ratios across the various stock factors, as well as big differences in the dynamics of both risk premia and price levels, which the model is able to capture well.

— Figures 4–5 go about here —

C. *Temporal Pricing of Risk*

The first key inputs from the model into the PE valuation exercise are the prices of the various bond and stock strips. Figure 6 plots zero-coupon bond and dividend strip prices, the latter scaled by the current quarter is dividend. For readability, we plot only three maturities: one, five, and 10 years. The model implies substantial variation in strip prices over time, across maturities, as well as across risky assets.

— Figure 6 goes about here —

As part of the estimation, the model fits several features of traded dividend strips on the aggregate stock market. Figure 7 shows the observed time series of the price-dividend ratio on a claim to the first eight quarters of dividends (red line, left panel), as well as the share of the total stock market value that these first eight quarters of dividends represent (red line, right panel). The blue line represents the model. The model generates the right level for the price-dividend ratio on the short-horizon claim. For the same 55 quarters for which the data are available, the average is 7.85 in the model and 7.65 in the data. The first eight quarters of dividends represent 3.4% of the overall stock market value in the data and 3.3% in the model over the period in which data are available. The model captures the dynamics of this share reasonably well, as shown in the right panel, including the decline over 2000 Q4 to 2001 Q1 when the short-term dividend strip prices fell by more than the overall stock market. The market clearly perceived the 2001 recession to be short-lived. In contrast, the contribution of short-term strips to the overall stock market value increases in the Great Recession, both in the data and in the model, in recognition of the persistent nature of the crisis.

— Figure 7 goes about here —

The second key inputs from the model into the PE valuation exercise are the expected excess returns on the bond and stock strips of horizons of 1 to 64 quarters. After all, the expected return of the PE fund is a linear combination of these expected returns per equation (2). Figure 8 plots the average risk premium on nominal zero-coupon bond yields and on dividend strips. Risk premia on nominal bonds (top left panel) are increasing in maturity, from zero to 4.5%. The top right panel shows the (spot) risk premia on dividend strips on the aggregate stock market (solid blue line). It also plots the dividend futures risk premium (red line). The difference between the spot and futures risk premia is approximately equal to the nominal bond risk premium. The unconditional dividend futures risk premium is downward-sloping in maturity at the short end of the curve and then flattens out. The graph also plots the model-implied dividend futures risk premium, averaged over the period 2003 Q1 to 2014 Q2 (yellow line). If anything, this risk premium is less downward-sloping than that averaged over the entire 1974 to 2019 sample (red line). The model matches the *realized* portfolio return on dividend futures of maturities from one to seven years over the period 2003 Q1 to 2014 Q2, which is 8.7% in the data and 8.6% in the model.[18]

The remaining panels of Figure 8 show the unconditional dividend strip (spot and future) risk premia for the other cross-sectional factors. There are interesting differences in the levels of future risk premia especially at shorter horizons and in the shapes of the term structures. Average futures risk premia are generally declining to flat in maturity, but they are increasing for small and value firms beyond the five-year horizon. Heterogeneity in risk premia by asset class, by horizon, and over time will give rise to heterogeneity in the risk premia on the PE-replicating portfolios.

— Figure 8 goes about here —

Figure 9 plots the time series of expected returns on bonds and on both dividend and gain strips for the seven equity factors; the maturity of all plotted strips is 20 quarters. Expected returns are annualized. We observe rich cross-sectional heterogeneity in levels and dynamics across panels, a low-frequency decline over time in the level of expected returns common across most panels, and high pairwise correlation between dividend and capital gain strip expected returns in each panel.

— Figure 9 goes about here —

Section IV of the Internet Appendix provides further insight into how the model prices risk at each horizon using the tools developed by Hansen and Scheinkman (2009) and Borovička and Hansen (2014). Specifically, we show that the various equity factors have very different risk exposures from each other, and at various horizons.

[18]As an aside, the conditional risk premium, which is the *expected* return on the dividend futures portfolio over the 2003 Q1 to 2014 Q2 period, is 9.8% per year in the model. The risk premium on the dividend futures portfolio over the full sample is 5.7%.

III. Expected Returns and Risk-Adjusted Profits in PE Funds

In this section, we combine the cash-flow exposures from Section I with the asset prices from Section II to obtain expected returns and risk-adjusted profits on PE funds.

A. *PE Cash Flow Summary Statistics*

Our fund data cover the period January 1981 to June 2019. Our main data source is Preqin, but we find comparable results using the Burgiss data set as shown in Section VIII of the Internet Appendix.[19] We group PE funds into eight categories: Buyout (LBO), Venture Capital (VC), Real Estate (RE), Infrastructure (IN), Natural Resources (NR), Fund of Funds (FF), Debt Funds (DF), and Restructuring (RS). Our FF category contains the Preqin categories Fund of Funds, Hybrid Equity, and Secondaries. The Buyout category is commonly referred to as Private Equity, whereas we use the PE label to refer to the combination of all investment categories.

We include all funds with nonmissing cash-flow information. All cash flows are net of fees imposed by the GPs. Table I reports the number of funds and the aggregate AUM in each vintage-category pair. In total, we have 4,474 funds with $4.3 trillion in AUM. Buyout is the largest category by AUM ($1,888 billion), followed by RE ($592 billion), FF ($528 billion), and VC ($489 billion). We group funds by their vintage, defined as the quarter in which they make their first capital call. The last column of the table shows the quartile of the price-dividend ratio on the aggregate stock market, which we use to sort funds into vintage bins. The table reports the average price-dividend quartile across the four quarters in the calendar year. We observe clear business-cycle variation in the

[19] Preqin data are substantially sourced by FOIA requests made to public pensions, which may have differential pricing terms in side letters and "Most Favored Nation" clauses. However, Da Rin and Phalippou (2017) suggest that public pensions are not statistically different from other investors in their access to these clauses. Burgiss data come from a more representative set of institutional investors.

timing of fund starts as well as in their size (AUM). The last cash flow we include in the estimation corresponds to for June 2019, since cash flows in the second half of 2019 have not been fully reported as of the timing of writing. The last vintage we consider in the estimation of exposures is 2017 Q4. For the RAP analysis, which requires a full life cycle, the last vintage we consider is 2010 Q4.

— Table I goes about here —

Figure 10 shows the average cash-flow profile in each category for distribution events, pooling all funds and vintages together and equally weighting them. For this graph, we combine all monthly cash flows into one yearly cash flow for each fund and then average across funds within the category. The first 15 orange bars are for the first 15 years since the first capital call. The last bar (in green) represents the cash flows that occur in year 16 and the discounted sum of cash flows that occur after year 16.[20] The literature typically treats PE vehicles as lasting 10 years. While the majority of distribution cash flows occur between years 5 and 10, cash flows after year 10 still account for a substantial portion of the total cash received by LPs.[21]

— Figure 10 goes about here —

Figure 11 zooms in on the four largest investment categories: Buyout, VC, RE, and FF. The figure shows the average cash-flow profile for each vintage. Since there are few

[20]We discount cash flows after quarter 64 at the nominal term structure. For infrastructure, we have a smaller and more recent sample. We consider 12 regular cash-flow years and year 13 as the terminal year.

[21]Industry publications have noted the increasing lifespan of PE funds. For instance, a Preqin report from 2016 notes that "The average lifespan of funds across the whole private capital industry is increasing beyond the typical 10 years... older funds of vintages 2000 to 2005 still hold a substantial $204bn worth of investments, equating to 7.2% of total unrealized assets" (Preqin (2016)). In the sample of funds with vintages before 2011, 48.3% of funds distribute more than 10% of cash flows after year 10. In a robustness check, we reestimate our results on a subsample of funds that distribute 10% or less of their cash after 10 years and find comparable results.

Buyout and VC funds prior to 1990 and few RE and FF funds prior to 2000, we start the former two panels with vintage year 1990 and the latter two panels with vintage year 2000. The figure reveals substantial variation in cash flows across vintages, even within the same investment category. This variation allows us to identify vintage effects. The figure also shows that there is a lot of variation in cash flows across calendar years. VC funds started in the mid- to late-1990 vintages realized very high average cash flows around calendar year 2000 and a sharp drop thereafter. Since growth stocks had very high stock price realizations in 2000 and a sharp drop thereafter, this type of variation will lead the model to estimate a high exposure of VC funds to growth gain strips. Internet Appendix IA.F shows cash-flow profiles for the remaining four PE categories.

— Figure 11 goes about here —

B. *Factor Estimation in OLS and Elastic Net*

We compare the results of two estimation approaches, run separately for each fund category. The first is a two-factor model (bond and aggregate stock market gain strips) estimated using OLS; recall equation (6). The second is an elastic net model estimated on the full set of 15 factors; recall equations (7) and (8). The estimated parameters are the factor exposures across horizon, b_h^k, and how these exposures shift by vintage (price-dividend quartile), captured by $a_{pd(t)}^k$. Their sum, $a_{pd(t)}^k + b_h^k$, measures the number of units of strip k with maturity h that the PE-replicating portfolio buys. The elastic net approach finds a parsimonious replicating portfolio consisting of long-only positions in some of the strips. While we have not constrained the elastic net estimation to require that adjacent years have similar exposures, we frequently find that factors have some periodic tendencies, with rising and falling exposures over stretches of the fund's life cycle.[22]

[22] To minimize overfitting, we rely on a cross-validation exercise in which we use a leave-out sample to fit the α and λ hyper-parameters. Section VII of the Internet Appendix discusses the details and provides robustness checks on the benchmark hyper-parameter choices.

Figure 12 shows the estimated age effects $\widehat{b}_h$ for the two-factor model estimated using OLS (left panels) and the 15-factor model estimated using elastic net (right panel). The rows correspond to the four main PE categories. Internet Appendix Figure IA.8 contains these estimates for the other four PE categories. Appendix Figures IA.9 and IA.10 show the estimates for the price-dividend quartile effects $\widehat{a}_{pd(t)}$.

— Figure 12 goes about here —

Buyout For the two-factor model in the top left panel of Figure 12, Buyout displays substantial positive exposure to market gain strips throughout the life cycle, with peak exposure in years 3 to 6. A strategy that sells the aggregate stock market is correlated with the distribution cash flows made by Buyout funds. Years 5 to 12 show substantial bond exposure in the replicating portfolio. The large bond exposure hints at out-performance. Buyout managers produce cashflows in the peak harvesting period that appear to be risk-free, according to the simple two-factor model.

The 15-factor model, estimated using elastic net and plotted in the right panel, gives a very different account of the riskiness of Buyout funds' cash flows. The two factors in the OLS model receive much less weight in the elastic net model. Instead, a rich set of cross-sectional risk factors contributes to describe the systematic riskiness of Buyout funds. The positions in each of the individual strips is much smaller. Early cash flows are exposed to value dividend strips, consistent with the findings of Stafford (2017). Later cash flows tend to load more on gain strips such as small and NR gain strips. This overall pattern corresponds to Buyout fund activities that consist of purchasing a broad range of companies, restructuring the operations, harvesting some initial cash flows (for instance, through dividend recapitalization), and ultimately selling these assets.

The takeaway is that Buyout vehicles do not simply take on bond and equity exposure, as is commonly assumed. Our best estimate for fund cash flow paints a more complex picture of rich factor exposures across a range of cross-sectional equity factors and horizons. Portfolio management of PE within institutional investor portfolios should consider this rich pattern of risk exposure of Buyout funds.

Venture Capital We see further evidence of the importance of considering a broad cross-section of factor exposures in the second row of Figure 12, which depicts results for Venture Capital funds. Our OLS two-factor model in the left panel places a large, hump-shaped weight on stock market gain strips and takes a mirror-image short position in bonds. This pattern suggests that VC funds are levered bets on the aggregate stock market.

In contrast, the 15-factor model in the right panel describes VC funds as loading mostly on growth gain strips. VC distribution cash flows are like those obtained from initially investing and eventually selling growth stocks to capture the capital gain. Figure IA.9 suggests that this growth gain strip loading is higher for funds started when the price-dividend ratio is in the second and fourth quartiles. Vintages in the early 1990s are such second-quartile price-dividend vintages, which ended up with very high cash flows. Vintages in the fourth price-dividend quartile, such as the 1997 to 2004 vintages, also have higher growth gain exposure.

Our findings for VC funds accord with economic intuition. While Buyout funds acquire a range of companies that may differ in their underlying risk exposures, VC funds invest in early stage and rapidly growing entrepreneurial companies that distribute little cash prior to exiting the investments. Since the same cross-validation procedure is used for each fund category, the elastic net will pick up a small number of dominant risk factor exposures if their payoffs are strongly related to PE cash flows. VC funds, unlike Buyout (or Real Estate or Infrastructure) funds typically harvest few cash flows from operations prior to deal exit. Correspondingly, we find that the bulk of VC fund cash flow exposure can be accounted for by growth gains strips (rather than growth dividend strips).

Real Estate The third row of Figure 12 considers Real Estate funds. The two-factor model characterizes RE PE funds as levered positions in stock market gain strips, especially for the first seven years of cash flows. Cash flows in some of the later years have positive bond exposure, hinting at out-performance.

The 15-factor model assigns no weight to bonds nor to market gain strips, again highlighting the need to consider a broader cross-section. Instead, it retains substantial weight on REIT dividend strips and value dividend strips in early years, and on small and REIT gain strips in years 5 to 8. Reassuringly, REIT dividends and REIT gain strips are important components of the replicating portfolio of RE PE funds. The small and value exposures accord well with the fact that listed REITS' returns behave like those of small value stocks (Van Nieuwerburgh (2019)). In later years, RE PE fund cash flows are exposed to the same risk as infrastructure and natural resources dividend strips, real asset cash flows that bear a certain intuitive resemblance to real estate cash flows. These results suggest that Real Estate funds take on a distinct factor exposure profile from Buyout and VC funds, and an exposure not well described by the two-factor model.

Fund of Funds The fourth row reports on the Fund of Funds category. The two-factor model shows modest market gain strip exposure and rising bond strip exposure, which becomes substantial in later years. The 15-factor elastic net model estimates a rich set of factor exposures, including to small, value, and NR gain and dividend strips, which corresponds to the miscellaneous nature of Fund of Funds strategies.

Other Categories The remaining four categories, which have substantially fewer fund observations, are shown in Figure IA.8. Infrastructure and natural resources show substantial exposure to NR, Infrastructure, and REIT gain and dividend strips. These exposures again point to the role of underlying asset characteristics in driving the fund-level cash flow risk profile.

Take-Aways The risk loadings on PE funds cannot be assumed to be static either in the time series or across fund age (maturity). A simple bond-stock portfolio typically does not survive inclusion of cross-sectional risk factors. The relevant factor identities differ across PE categories. We conduct the first systematic analysis of the risk properties of some of the alternative fund categories (RE, IN, NR), and find that they carry important sector-specific asset exposures. These exposures are frequently concentrated in the first half of the fund's life. Our estimation approach allows us to translate these complex risk dynamics into the expected return for different fund categories and to revisit the question of performance evaluation. We turn to expected returns next.

C. *Expected Returns*

With the replicating portfolio of zero-coupon bonds and dividend strips in hand, we can calculate the expected return on PE funds in each investment category using equation (2). The expected return measures compensation to systematic sources of risk. It excludes any abnormal performance, which is contained in the RAP. Figure 13 plots the time series of the expected return for the four main PE categories. It aggregates over all of the horizon effects and annualizes the resulting expected return per equation (4). The left panels of this figure correspond to the two-factor OLS model and the right panels to the 15-factor elastic net model. Figure IA.11 shows the results for the four other PE categories.

— Figure 13 goes about here —

On average over time, the expected return on PE vehicles is 9.5% for Buyout, 8.4% for VC, 8.7% for RE, and 9.8% for FF. RS has a 7.5% average expected return, DF 7.2%, Infrastructure 5.9%, and NR 8.2%. These expected returns are in the vicinity of IRRs calculated under the same assumption on calls. Expected returns are higher under the 15-factor model than under the two-factor model. The additional risk factor exposures result in a higher required return.

Vintage effects in the exposures (the a's for each of the factors) generate time-variation in the exposures, and time-varying MPR generate time-variation in the expected returns on dividend and gain strips. Combined, they lead to time-variation in the expected return of PE funds. The annualized expected return that investors can anticipate on their PE investments as compensation for systematic risk has seen large variation over time, with a declining pattern at low frequencies. The low-frequency decline is inherited from a low-frequency decline in strip expected returns. For example, the bond, growth, REIT, and infrastructure risk premia in Figure 9 all show strong secular declines. The low risk premia for PE at the end of the sample reflects the elevated prices for all risky assets at that time.

At higher frequencies we note the low expected return around 2000, when the stock market peaked, and an increase in risk premia during the Great Recession for several of the PE categories. These dynamics are driven by vintage effects, which switch discretely between price-dividend quartiles and sometimes result in spikes, and by dynamics on the listed risk premia. For example, the aggregate MRP is very low around 2000, while small, value, and growth risk premia are elevated in the Great Recession.

D. *Risk-Adjusted Profit*

Next, we turn to performance evaluation, the main result in the paper. Figure 14 plots the histogram of RAP, computed from equation (5), for the two-factor OLS (gray) and 15-factor elastic net (yellow) models for all fund categories. A kernel density, estimated from the discrete histogram, is superimposed.

— Figure 14 goes about here —

In all eight fund categories, the RAP distribution is shifted down when accounting for the cross-sectional risk factors in the 15-factor model compared to the two-factor model. For VC and real estate, a substantial part of the right tail of the distribution is removed and shifted to the left.

Across categories, adjusting for risk removes all (and for some categories more than all) of the excess cash flows; the average TVPI and RAP are reported above each panel. As a result, an LP who uses traditional approaches (TVPI) or even a flexible two-factor model would attribute a fund's excess payouts to outperformance, while the 15-factor elastic net model instead attributes this profit to compensation for risk.

Average RAP under our benchmark (NVP Call) 15-factor model is -6 cents for the average Buyout fund, -9 cents for the average VC fund, -16 cents for the average RE fund, and -19 cents for the average FF. For the remaining four categories, we find -0.1 cents for RS, -13 cents for DF, -6 cents for IN, and -6 cents for NR funds. With the exception of VC, the risk-adjusted profit under the 15-factor model is far lower than under the two-factor model.

The average RAP masks substantial cross-sectional dispersion. The RAP histograms are wide. The Monte Carlo exercise in Section I.A of the Internet Appendix suggests that the bulk of the dispersion in RAP reflects skill heterogeneity, with the remainder attributable to idiosyncratic cash-flow risk (luck). As indicated above each panel of Figure 14, anywhere between 16% and 41% of funds have a RAP in excess of 10 cents. The fraction is highest for natural resources and lowest for FF. Large dispersion also means that a large fraction of PE funds destroy substantial value on a risk-adjusted basis, as indicated by the large mass of the RAP distribution that is well to the left of zero.

Table II compares average outperformance. Panel A reports results from our main categories, Panel B for the additional categories. The first three rows report standard measures to evaluate PE funds: TVPI, IRR, and PME.[23] The next two rows display results

[23]TVPI and PME subtract out the initial investment to express them as excess performance metrics. We calculate the IRR since Preqin computes IRRs only for funds that are fully liquidated, which leads to a bias. Like we do in the benchmark RAP calculation, we set the initial investment in the IRR calculation equal to the discounted sum of calls, discounted at the Treasury yield curve. The implicit assumption that the standard IRR calculation makes on calls—that any uncalled amounts are invested at the IRR—leads to an IRR that is about 3% points higher. In related work, Phalippou (2009) discusses how cash-flow timing

of our two-factor OLS and 15-factor elastic net models. The columns report the average RAP and the cross-sectional standard deviation of the RAP. The mean performance shows a downward shift for performance models that adjust for risk (PME and our models compared to TVPI). Richer risk adjustment (15-factor model) leads to lower performance than simpler risk adjustment (PME and two-factor models). The large standard deviation confirms the wide dispersion in fund performance. The 15-factor model tends to result in a tighter distribution than the two-factor model, which in turn has a smaller dispersion than TVPI. While there is wide dispersion in performance, at least some of that dispersion is accounted for by risk.

— Table II goes about here —

Figure 15 plots the average RAP by vintage for both the two-factor (left panels) and 15-factor elastic net (right panels) models. Figure IA.12 plots these estimates for alternative fund categories. Consistent with our earlier results, average RAP by vintage in the 15-factor model is shifted down from average RAP in the two-factor model. While the time-series for RAP show commonality across both models, there are also notable differences. For instance, in the Buyout category, we observe substantial positive profits continuing through the 2000s in the two-factor model. By contrast, the elastic net model estimates mildly negative average profits for most vintages in the 2000s.

— Figure 15 goes about here —

Similarly, while we observe extremely high profits for FF originated before 1994, these profits fall by a factor of two in the elastic net model. Recent FF vintages have generated positive RAP according to the two-factor model but consistently negative RAP according to the full model. RE funds launched before the Great Recession (vintages 2002 to 2006) have performed reasonably according to the two-factor model, but poorly according to

distorts IRR measurement in the PE context.

the full model. The only category that sees its performance improve in the 15-factor model is VC. That said, VC has generated negative average RAP in each of the last 14 vintage years.

E. Model Comparison

It is instructive to benchmark our results against other approaches used in the literature. Figure 16 graphically compares the results from the elastic net model for our main PE categories against two commonly used PE fund performance metrics: IRR and PME. Figure IA.13 repeats this analysis for the additional categories. The left panels plot fund-level IRR against our fund-level RAP measure; the right panels plot fund-level PME (subtracting the initial investment) against fund-level RAP. The key takeaway from this comparison is a broadly similar ranking of fund performance. Our measure of RAP correlates between 83% to 87% with our IRR measure and between 66% to 87% with the PME measures in the cross-section of funds. The broad similarity lends credibility to our measure of RAP. The measures are not identical, however, so there exist funds that conventional measures assess to be high-performing but that our estimates suggest offer only fair or even too little compensation given their factor risk exposure, and vice versa.

— Figure 16 goes about here —

F. Performance Persistence

While we find that about one-third of funds deliver meaningfully positive RAP, these may not be the same funds that consistently outperform. We therefore examine the persistence of the various performance metrics. Specifically, we look at the relationship between the performance of every pair of adjacent funds by the same GP in the same fund category. For example, we compare Blackstone's Real Estate Partners fund I with Blackstone's Real Estate Partners fund II, REP fund II with REP fund III, etc. We only study the four main fund categories since there are insufficiently many fund pairs in the remaining four categories for a reliable analysis.

We consider three performance metrics in Table III: (i) the pairwise correlation between the RAP of the two funds in each pair (labeled "Persistence"), (ii) the likelihood that the second fund in the pair is in the top quartile of the RAP distribution given that the first fund was in the top quartile ("Top Quart."), and (iii) the likelihood that the second fund in the pair is in the bottom quartile of the RAP distribution given that the first fund was in the bottom quartile ("Bottom Quart"). The last two measures are based on rankings and are inspired by Korteweg and Sorensen (2017), who also study performance at the top and bottom of the fund distribution. We compare the persistence of our two- and 15-factor models to the persistence of the traditional performance metrics TVPI, IRR, and PME.

— Table III goes about here —

Consistent with prior research (Harris et al. (2014b), Korteweg and Sorensen (2017)), we find modest persistence in performance across funds managed by the same GP. The persistence metric for the two-factor model is between 0.16 for RE and 0.35 for VC. For the 15-factor model, it is between 0.14 for VC and 0.31 for Buyout. For PME, persistence tends to be higher, for example, 0.47 for VC and 0.35 for RE. These results illustrate that the additional risk adjustment makes a difference for the assessment of fund persistence.

For Buyout and VC, we find substantial persistence both at the top and at the bottom of the fund distribution. Buyout funds that have a fund in the top quartile of the RAP distribution according to our 15-factor model have a 36% probability of the next fund also being in the top quartile of the distribution; a 25% transition probability denotes no persistence. The corresponding number is 39% for VC. Bottom performance is also persistent, especially among RE funds where the likelihood of the next fund in the series also being a bottom performer is 40%. For Buyout and VC, there is more persistence in the top than in the bottom of the distribution. Compared to risk-unadjusted performance

(TVPI), our benchmark 15-factor model tends to show less persistence. Thus, at least some of the persistence in the traditional metrics can be ascribed to persistent risk factor exposure.

G. *Importance of Calls*

In our baseline results, we subtract the NPV of all calls discounted at the Treasury yield curve, C_t, in the RAP measure (5). This implicitly assumes that the LP has perfect foresight over the call schedule and invests the capital committed but not called in the quarter of the first call in the right portfolio of Treasury bonds of various maturities. The LP sells the Treasuries when the GP calls in the remaining committed capital in the ensuing quarters and this liquidation does not incur transaction costs. This investment strategy requires a fairly sophisticated understanding of the future call schedule on the part of the LP. This treatment of calls also omits from consideration any committed capital that was never called. Implicitly, the LP is assumed to earn the same return on calls never made as on calls made. We make this assumption because it is conservative - it results in the lowest call amount C_t and hence the highest possible risk-adjusted performance. It does not penalize the GP for delays in calling some of the committed capital or for never calling some of the capital at all.

We now entertain two alternative and equally defensible assumptions on the calls. The first is that the LP earns a zero return on the capital committed but not called on the first call date. The corresponding call amount C_t in the RAP calculation (5) is the undiscounted sum of all calls. Since this call amount is larger than the benchmark one (unless all calls are made on the first call date), the RAP is lower. Intuitively, the replicating portfolio, which is the benchmark for the PE investment, invests an amount equal to the sum of all calls on the first date. The GP is penalized for delaying to call some of the capital until after the first call date. Table II reports this case with the label "Sum Call." These RAPs are about 4-9 cents lower per dollar of committed capital than our benchmark RAP results.

The second alternative is to assume that the benchmark for the GP is a replicating portfolio that invests the full \$1 of committed capital. Equivalently, the LP gives \$1 to the GP. If the GP delays to deploy some of the capital or decides not to invest some of the capital at all, this hurts the PE fund's performance relative to the benchmark. Any amount never invested is returned to the LP at the end of the fund's life. This is equivalent to an initial call amount of \$1 minus a return of the never-invested capital at the end of the life of the investment, that is, discounted at the 64-quarter Treasury bond price. This case has the highest call amount and therefore the lowest RAP. Table II reports this case with the label "Residual Call." The resulting RAPs are another 1-3 cents lower.

In sum, when we assume that the LP does not earn a return on the committed but never-called capital or on the capital that is not yet but eventually called, the RAP is about 10 cents lower. The difference with the benchmark RAP ranges from 4 to 11 cents across categories. Average RAP is now solidly in negative territory in every investment category for the main 15-factor model. RS remains the best PE category, but still loses 6.5 cents on a risk-adjusted basis for every \$1 of committed capital. Buyout loses 16 cents, VC loses 19 cents, and RE loses 22 cents. FF is the worst category with a 30 cent loss.

H. Discussion

Our pervasive finding of negative average RAP stands in contrast to previous literature on PE, which generally finds evidence of fund outperformance (Brown et al. (2015), Kaplan and Sensoy (2015), Harris, Jenkinson, and Kaplan (2014a)). The difference is made apparent when comparing standard performance metrics in the first three rows of Table II, which tend to show more favorable performance, to our risk adjusted metrics. Even the two-factor model generally shows positive risk-adjusted profits. Our more negative conclusion about outperformance in the PE industry is a natural consequence of richer risk-adjustment due to cross-sectional factor exposure. Similar risk adjustments in the mutual fund (Fama and French (2010)) or hedge fund (Fung et al. (2008)) literatures likewise point to lower outperformance.

We also find evidence for substantial cross-sectional dispersion in RAPs, as well as sizable persistence in risk-adjusted performance over time. This suggests that a subset of PE managers consistently deliver outsized returns. This finding is consistent with Kacperczyk, Van Nieuwerburgh, and Veldkamp (2014, 2016), who find that a small right tail of mutual fund managers consistently outperform, although the earlier mutual fund literature has been skeptical of skill in any part of the fund distribution (Fama and French (2010)). Decomposing PE performance into market timing and asset selection components is left for future study.

Similar findings of limited excess returns across categories of delegated asset managers are suggestive of similar economic forces at work. First, investors may find it difficult to replicate complex factor strategies on their own, and hence may be willing to pay PE managers to generate factor strategies. Second, superior performance tends to be associated with higher fund flows, leading to increased capital commitments that reduce returns in the presence of decreasing returns to scale, along the lines of Berk and Green (2004). As the PE industry has grown substantially, there may be insufficient economies of scale to adequately manage a growing asset base to generate the same outsize returns of previous decades. Our finding of declining RAPs over time is consistent with this conjecture. Finally, delegated asset managers charge sizable management and performance fees especially when financial products are more opaque and complex (Célérier and Vallée (2015)) as they are in PE. This reduces after-fee returns for investors further. We have a deeper analysis of before-fee performance for future work.

IV. Conclusion

We provide a novel valuation method for PE cash flows that decomposes the cash flow at each horizon into a systematic component that reflects exposure to multiple sources of aggregate risk, captured by the cross section of listed securities, and a component that reflects the risk-adjusted profit to the PE investor. The systematic component represents a portfolio of stock and bond strips paying safe or risky cash flows at horizons over which PE funds make cash-flow distributions. A state-of-the-art no-arbitrage asset pricing model estimates prices and expected returns for these strips by closely fitting the time series of bond yields and stock prices, including dividend strips. The asset pricing model provides the first estimates of the term structure of risk and return in the cross-section of equity factors.

Using a two-factor OLS and 15-factor elastic net approach, we estimate rich heterogeneity in PE fund risk exposures across horizons, in the cross-section, and in the time series. PE funds' risk exposure is best modeled using not only bonds and the aggregate stock market, but also sector-specific equity factor exposures. The estimated exposures are sensible given the underlying nature of PE assets, and indicate an important role for growth stocks in VC and for REITs in real estate funds. In the time series, we find that expected returns on PE funds have been declining substantially since the 1980s and especially since the Great Recession, reflecting declining risk premia in public markets.

On average, PE funds tend to underperform their replicating portfolio benchmark, suggesting that while PE funds offer investors access to complex risk exposures, they do so at a cost that is higher than that offered in public markets. While our resulting profit measures correlate well with existing measures of outperformance in the cross-section of funds, they imply substantially lower average performance than traditional measures. Performance deteriorates further if we penalize the fund manager for not calling all committed capital or for calling some of it with a delay. Under this alternative assumption, all PE fund categories underperform substantially on a risk-adjusted basis. One potential interpretation of this underperformance is that investors such as pension funds may be willing to pay an illiquidity premium for the convenience of not having to mark-to-market investments that display much of the same risk characteristics as a portfolio of stocks and bonds. Exploring this conjecture more thoroughly represents interesting ground for future research.

The negative average performance hides substantial dispersion across funds. A substantial fraction outperforms, and we find some evidence for persistence in outperformance. Exploring the characteristics of out-performing funds are also merits further inquiry.

Our analysis highlights the value of a methodological advance in the assessment of risk and return for unlisted assets, which are an increasing component of the total investable universe for many institutional investors. While PE, given its size, is an especially important application, our method can be applied more broadly to study the risk characteristics and risk-adjusted performance of any other cash-flowing asset. Individual private firms, real estate assets, or infrastructure investment projects are applications left for future work.

REFERENCES

Ammar, Semir Ben, and Martin Eling, 2015, Common risk factors of infrastructure investments, *Energy Economics* 49, 257–273.

Andonov, Aleksandar, Roman Kräussl, and Joshua Rauh, 2020, The subsidy to infrastructure as an asset class, SSRN Working Paper No. 3245543.

Ang, Andrew, Bingxu Chen, William N. Goetzmann, and Ludovic Phalippou, 2018, Estimating private equity returns from limited partner cash flows, *Journal of Finance* 73, 1751–1783.

Ang, Andrew, and Monika Piazzesi, 2003, A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables, *Journal of Monetary Economics* 50, 745–787.

Asness, Cliff, 2019, The illiquidity discount?, AQR Perspectives.

Berk, Jonathan B., and Richard C. Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112, 1269–1295.

Bökberg, Anna, Julia Carrellas, Andrea Chau, and Charlie Duane, 2019, Private markets come of age, Technical report, McKinsey Global Institute.

Borovička, Jaroslav, and Lars Peter Hansen, 2014, Examining macroeconomic models through the lens of asset pricing, *Journal of Econometrics* 183, 67–90.

Brown, Greg, Robert Harris, Tim Jenkinson, Steven N. Kaplan, and David Robinson, 2015, What do different commercial data sets tell us about private equity performance?, Working paper, University of North Carolina Kenan Institute.

Campbell, John Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.

Campbell, John Y., 1993, Intertemporal asset pricing without consumption data, *American Economic Review* 83, 487–512.

Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.

Campbell, John Y., and Tuomo Voulteenaho, 2004, Good beta, bad beta, *American Economic Review* 94, 1249–1275.

Célérier, Claire Myriam, and Boris Vallée, 2015, The motives for financial complexity: An empirical investigation, Working paper, University of Zurich.

Cochrane, John H., 2005, The risk and return of venture capital, *Journal of Financial Economics* 75, 3–52.

Cochrane, John H., and Jesus Saa-Requejo, 2000, Beyond arbitrage: ’good deal’ asset price bounds in incomplete markets, *Journal of Political Economy* 108, 79–119.

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

Dai, Qiang, and Kenneth J. Singleton, 2000, Specification analysis of affine term structure models, *Journal of Finance* 55, 1943–1978.

Driessen, Joost, Tse-Chun Lin, and Ludovic Phalippou, 2012, A new method to estimate risk and return of nontraded assets from cash flows: The case of private equity funds, *Journal of Financial and Quantitative Analysis* 47, 511–535.

Duffie, Darrell, and Raymond Kan, 1996, A yield factor model of interest rates, *Mathematical Finance* 6, 379–406.

Fama, Eugene F, and Kenneth R French, 2010, Luck versus skill in the cross-section of mutual fund returns, *Journal of Finance* 65, 1915–1947.

Fung, William, David A Hsieh, Narayan Y Naik, and Tarun Ramadorai, 2008, Hedge funds: Performance, risk, and capital formation, *Journal of Finance* 63, 1777–1803.

Giglio, Stefano, Bryan Kelly, and Serhiy Kozak, 2020, Equity term structures without dividend strips data, Working Paper Yale School of management.

Gredil, Oleg, Morten Sorensen, and William Waller, 2020, Evaluating private equity performance using stochastic discount factors.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, 2020, Empirical asset pricing via machine learning, *Review of Financial Studies* 33, 2223–2273.

Haddad, Valentin, Serhiy Kozak, and Shrihari Santosh, 2020, Factor timing, *Review of Financial Studies* 33, 1980–2018.

Haddad, Valentin, Erik Loualiche, and Matthew Plosser, 2017, Buyout activity: The impact of aggregate discount rates, *Journal of Finance* 72, 371–414.

Hansen, Lars Peter, and Jose Scheinkman, 2009, Long-term risk: An operator approach, *Econometrica* 77, 177–234.

Harris, Robert S., Tim Jenkinson, and Steven N. Kaplan, 2014a, Private equity performance: What do we know?, *Journal of Finance* 69, 1851–1882.

Harris, Robert S., Tim Jenkinson, Steven N. Kaplan, and Rüdiger Stucke, 2014b, Has persistence persisted in private equity? evidence from buyout and venture capital funds, SSRN Working Paper No. 2304808.

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, 2014, Time-varying fund manager skill, *Journal of Finance* 69, 1455–1484.

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, 2016, Rational attention allocation over the business cycle, *Econometrica* 84, 571–626.

Kaplan, Steven, and Berk Sensoy, 2015, Private equity performance: A survey, *Annual Review of Financial Economics* 7, 597–614.

Kaplan, Steven N., and Antoinette Schoar, 2005, Private equity performance: Returns, persistence, and capital flows, *Journal of Finance* 60, 1791–1823.

Karolyi, G. Andrew, and Stijn Van Nieuwerburgh, 2020, New methods for the cross-section of returns, *Review of Financial Studies* 33, 1879–1890.

Korteweg, Arthur, and Stefan Nagel, 2016, Risk-adjusting the returns to venture capital, *Journal of Finance* 71, 1437–1470.

Korteweg, Arthur, and Morten Sorensen, 2010, Risk and return characteristics of venture capital-backed entrepreneurial companies, *Review of Financial Studies* 23, 3738–3772.

Korteweg, Arthur, and Morten Sorensen, 2017, Skill and luck in private equity performance, *Journal of Financial Economics* 124, 535–562.

Korteweg, Arthur G., 2019, Risk adjustment in private equity returns, *Annual Review of Financial Economics* 11, 131–152.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2017, Shrinking the cross section, Working Paper No. 24070, National Bureau of Economic Research.

Lettau, Martin, and Jessica Wachter, 2011, The term structures of equity and interest rates, *Journal of Financial Economics* 101, 90–113.

Liu, Naiping, and Lu Zhang, 2008, Is the value spread a useful predictor of returns?, *Journal of Financial Markets* 11, 199–277.

Ljungqvist, Alexander, and Matthew Richardson, 2003, The cash flow, return and risk characteristics of private equity, Working paper, NYU Stern.

Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan, 2013, The wealth-consumption ratio, *Review of Asset Pricing Studies* 3, 38–94.

Martin, Ian, 2012, On the valuation of long-dated assets, *Journal of Political Economy* 120, 346–358.

Metrick, Andrew, and Ayako Yasuda, 2010, The economics of private equity funds, *Review of Financial Studies* 23, 2303–2341.

Peng, Liang, 2016, The risk and return of commercial real estate: A property level analysis, *Real Estate Economics* 44, 555–583.

Phalippou, Ludovic, 2009, The hazards of using irr to measure performance: The case of private equity, SSRN Working Paper No. 1111796.

Phalippou, Ludovic, and Oliver Gottschalg, 2009, The performance of private equity funds, *Review of Financial Studies* 22, 1747–1776.

Preqin, 2016, *The 2016 Preqin Global Private Equity and Venture Capital Report*, (Preqin, London, UK).

Riddiough, Timothy J., 2020, Pension funds and private equity real estate: History, performance, pathologies, risks, Working paper, University of Wisconsin.

Robinson, David T., and Berk A. Sensoy, 2016, Cyclicality, performance measurement, and cash flow liquidity in private equity, *Journal of Financial Economics* 122, 521–543.

Sagi, Jacob S., 2017, Asset-level risk and return in real estate investments, Working Paper, UNC Kenan-Flagler Business School.

Sorensen, Morten, and Ravi Jagannathan, 2015, The public market equivalent and private equity performance, *Financial Analysts Journal* 71, 43–50.

Sorensen, Morten, Neng Wang, and Jinqiang Yang, 2014, Valuing private equity, *Review of Financial Studies* 27, 1977–2021.

Stafford, Erik, 2017, Replicating private equity with value investing, homemade leverage, and hold-to-maturity accounting, SSRN Working Paper No. 2720479.

van Binsbergen, Jules, Michael Brandt, and Ralph Koijen, 2012, On the timing and pricing of dividends, *American Economic Review* 102, 1596–1618.

van Binsbergen, Jules H., Wouter H. Hueskes, Ralph Koijen, and Evert B. Vrugt, 2013, Equity yields, *Journal of Financial Economics* 110, 503–519.

van Binsbergen, Jules H., and Ralph Koijen, 2017, The term structure of returns: Facts and theory, *Journal of Financial Economics* 124, 1–21.

Van Nieuwerburgh, Stijn, 2019, Why are reits currently so expensive?, *Real Estate Economics* 47, 18–65.

Wachter, Jessica, 2005, Solving models with external habit, *Finance Research Letters* 2, 210–226.

Weber, Michael, 2018, Cash flow duration and the term structure of equity returns, *Journal of Financial Economics* 128, 486–503.

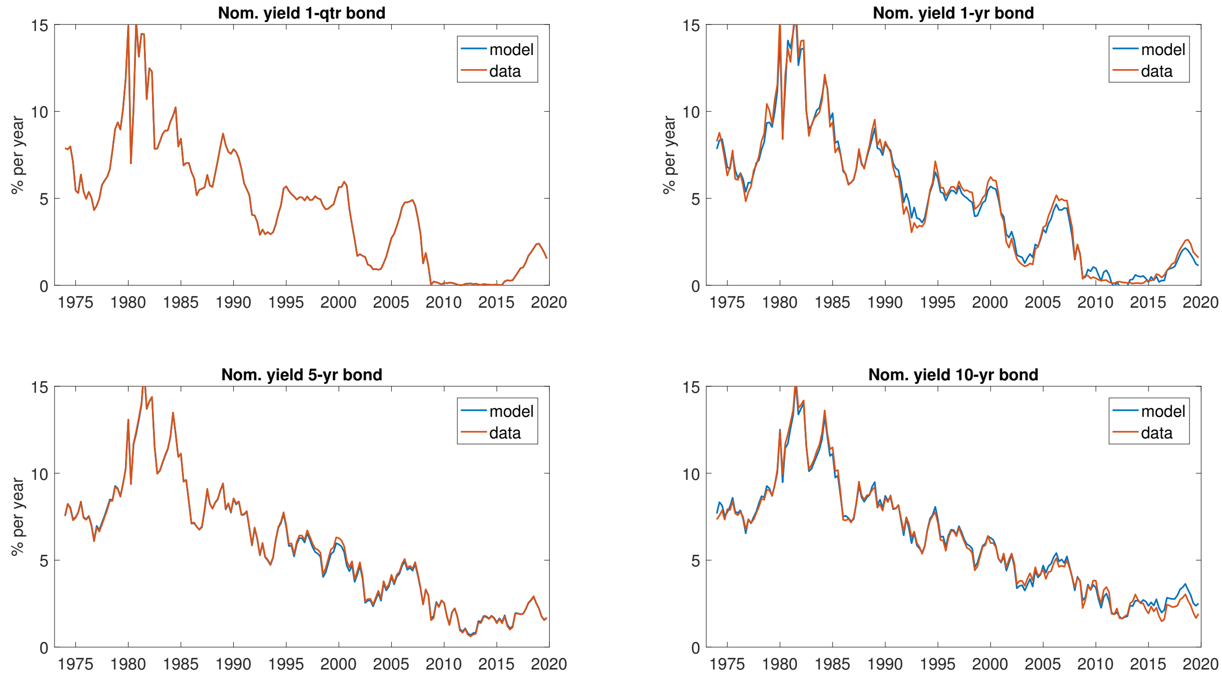

Figure 1. Dynamics of the nominal term structure of interest rates.

The figure plots the observed and model-implied one-, four-, 20-, 40-quarter nominal bond yields.

Figure 2. Dynamics of the real term structure of interest rates.

The figure plots the observed and model-implied 20-, 28-, 40-, and 80-quarter real bond yields.

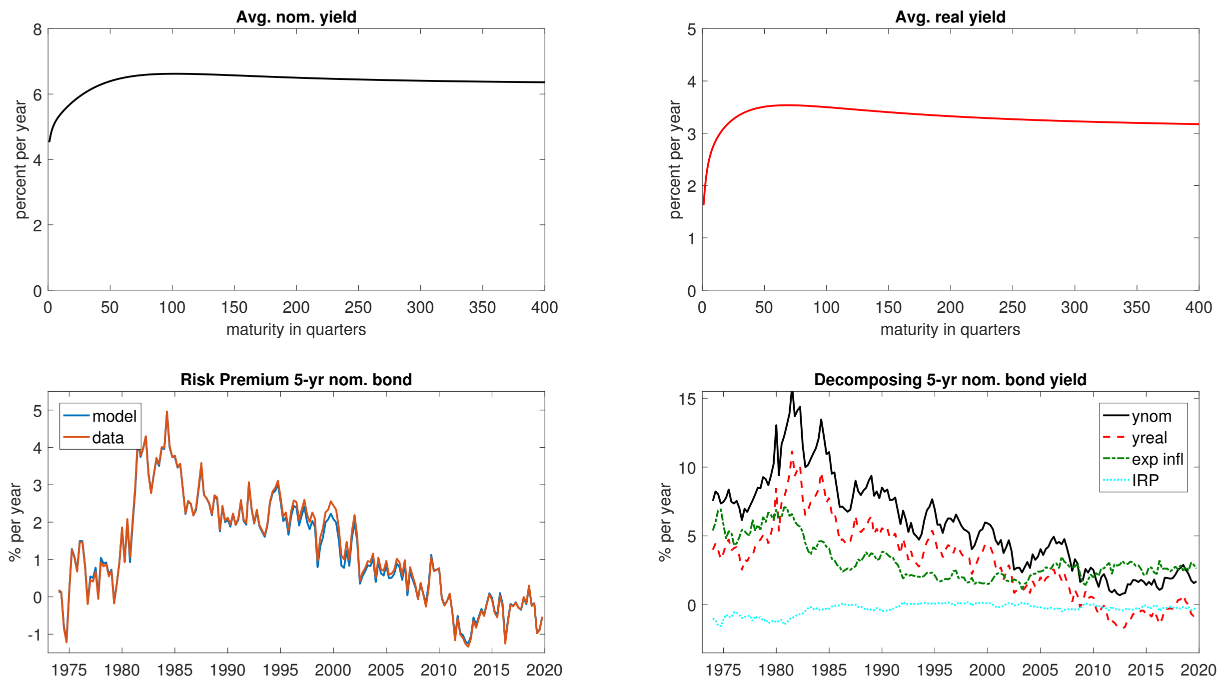

Figure 3. Long-term yields and bond risk premia.

The top panels plot the average bond yield on nominal (left panel) and real (right panel) bonds for maturities ranging from one quarter to 200 quarters. The bottom left panel plots the nominal bond risk premium in model and data. The bottom right panel decomposes the model's five-year nominal bond yield into the five-year real bond yield, the five-year inflation risk premium, and the five-year real risk premium.

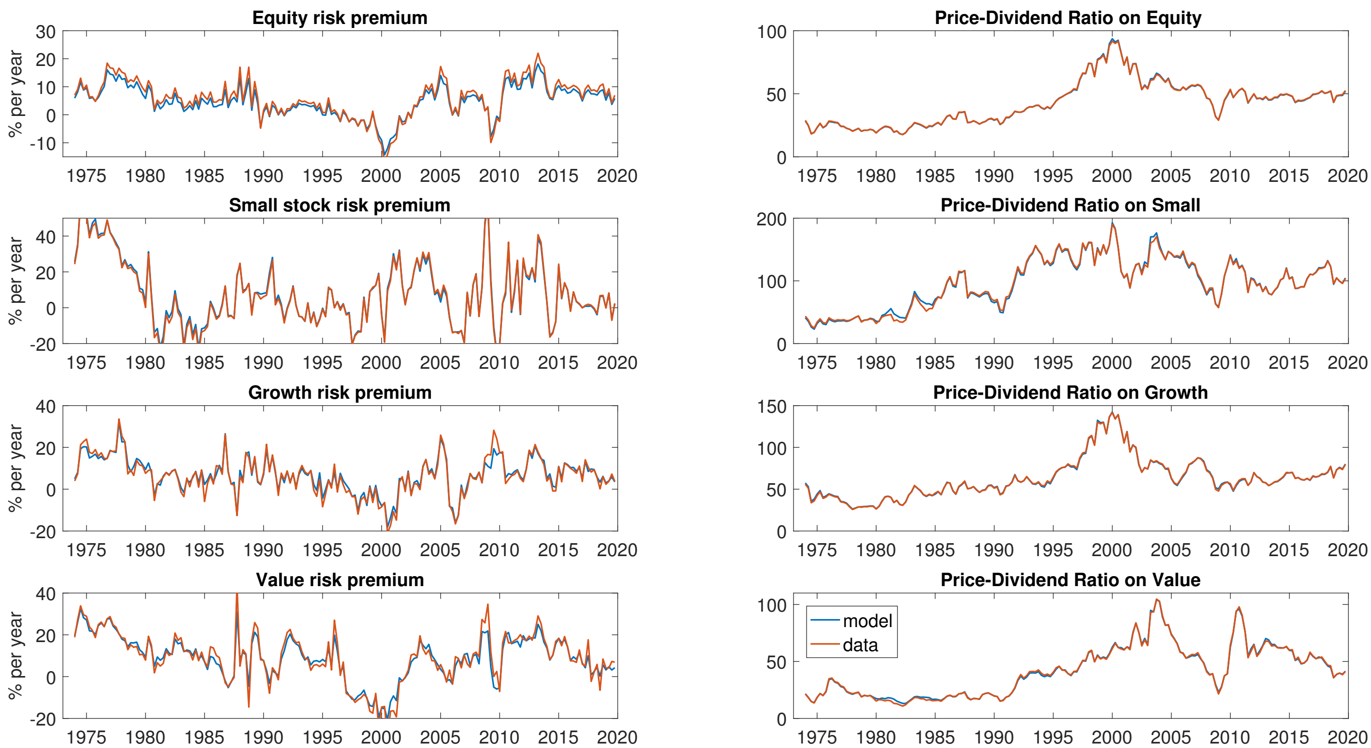

Figure 4. Equity risk premia and price-dividend ratios (part 1).

The figure plots the observed and model-implied equity risk premium on the overall stock market, small stocks, growth stocks, and value stocks in the left panels, as well as the corresponding price-dividend ratios in the right panels. The model is reported by the blue lines, while the data are depicted by the red lines.

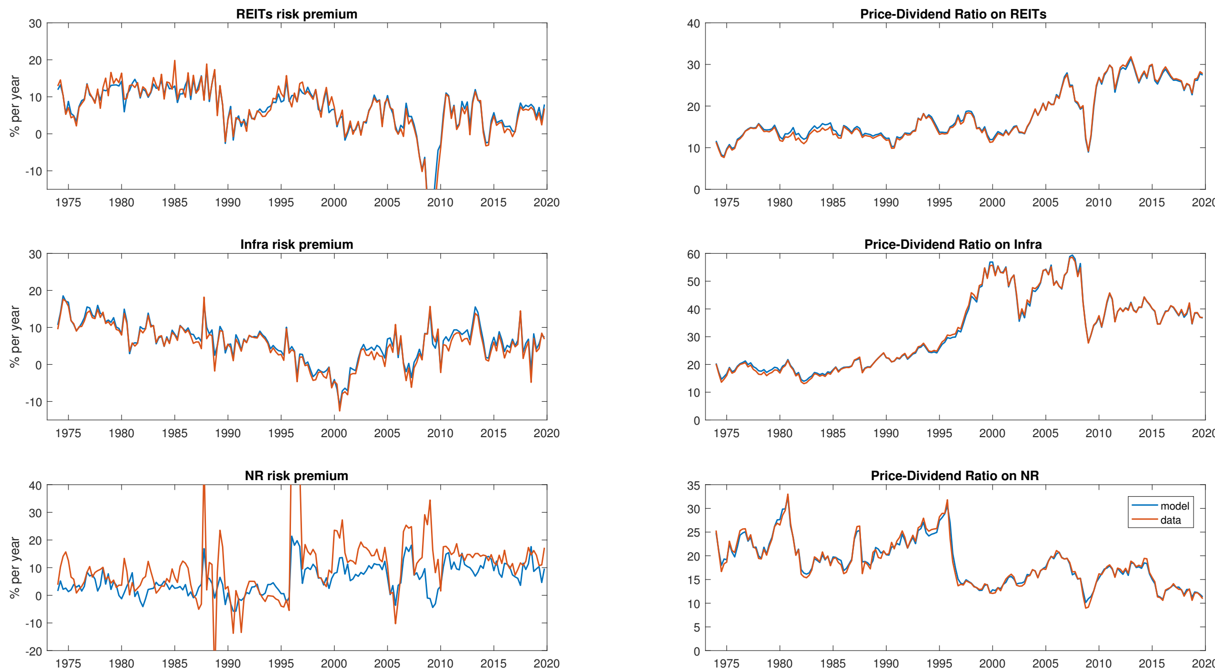

Figure 5. Equity risk premia and price-dividend ratios (part 2).

The figure plots the observed and model-implied equity risk premium on REIT stocks, infrastructure stocks, and natural resource stocks in the left panels, as well as the corresponding price-dividend ratios in the right panels. The model is reported by the blue lines, while the data are depicted by the red lines.

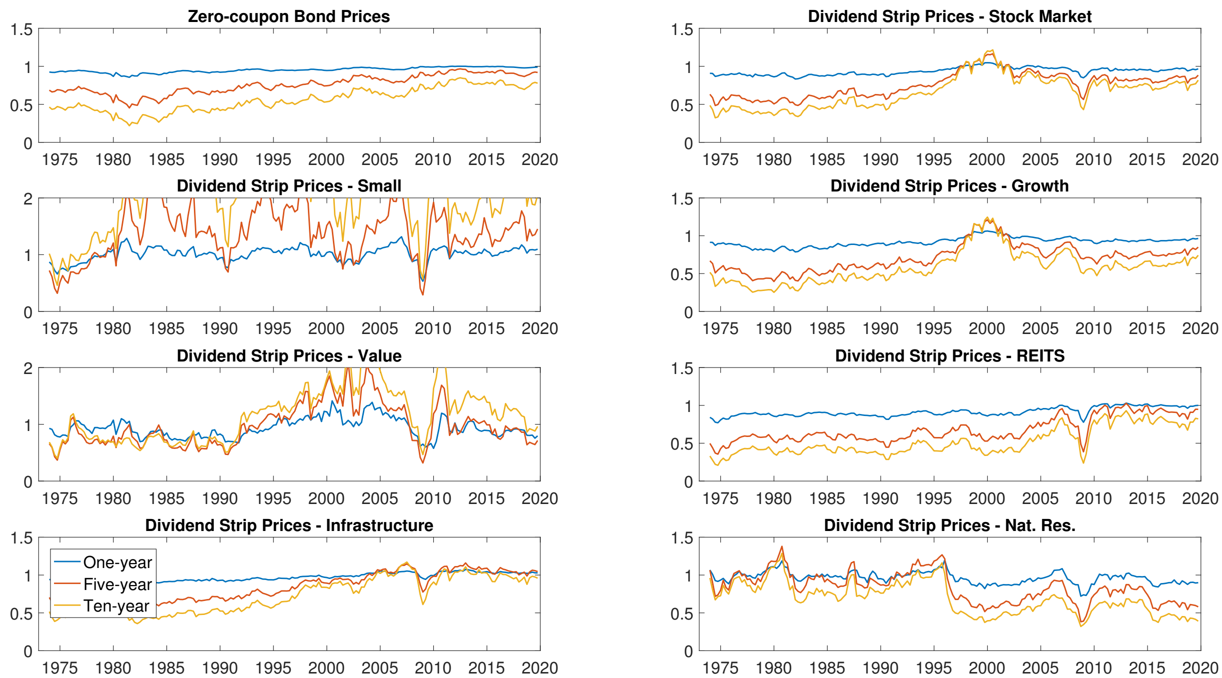

Figure 6. Zero-coupon bond prices and dividend strip prices.

The figure plots the model-implied prices on zero-coupon Treasury bonds in the first panel, and price-dividend ratios for dividend strips on the overall stock market, small stocks, growth stocks, value stocks, REIT market stocks, infrastructure stocks, and natural resources stocks in the next seven panels, for maturities of four, 20, and 40 quarters. The prices/price-dividend ratios are expressed in levels and each claim pays out a single cash flow.

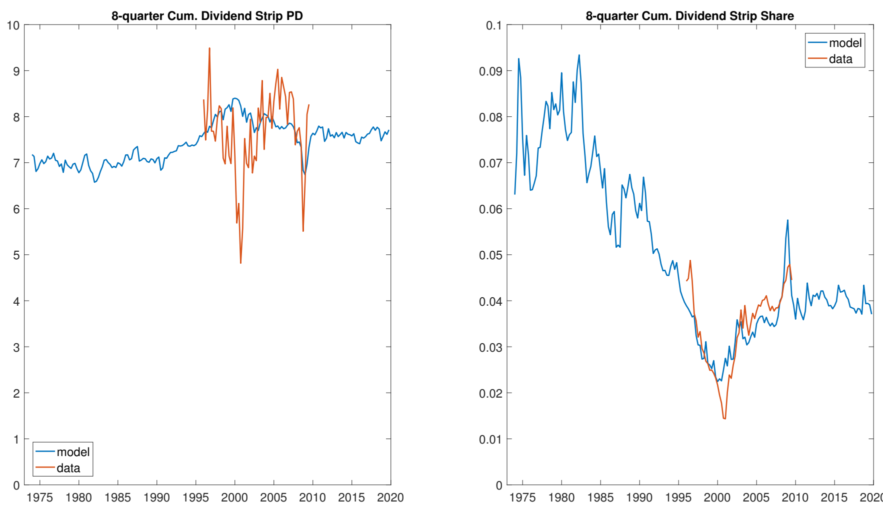

Figure 7. Short-run cumulative dividend strips.

The left panel plots the model-implied price-dividend ratio on a claim that pays the next eight quarters of dividends on the aggregate stock market. The right panel plots the share that this claim represents in the overall value of the stock market. The data are from van Binsbergen, Brandt, and Koijen (2012) and are available from 1996 Q1-2009 Q3.

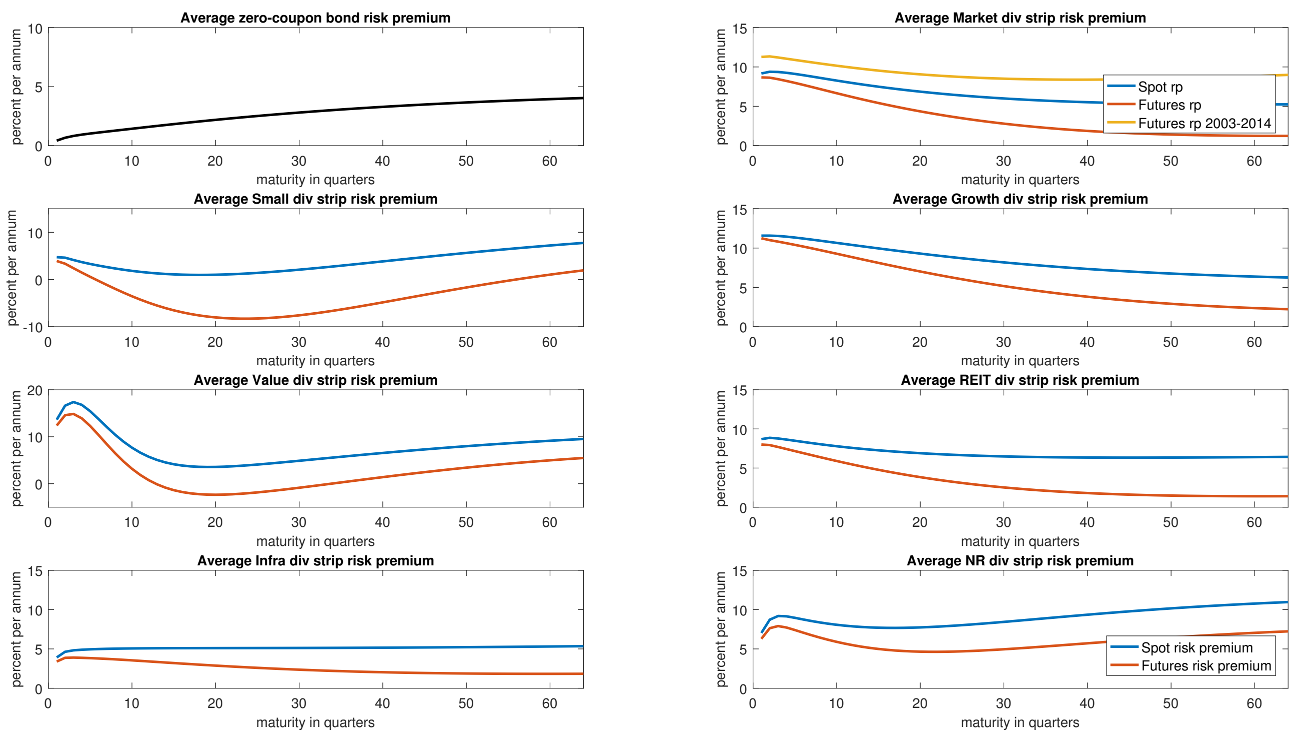

Figure 8. Strip expected returns by horizon.

The figure plots the model-implied average risk premia on nominal zero-coupon Treasury bonds in the first panel, and on dividend strips on the overall stock market, small stocks, growth stocks, value stocks, REITs, infrastructure stocks, and natural resource stocks in the next seven panels, for maturities ranging from one to 64 quarters.

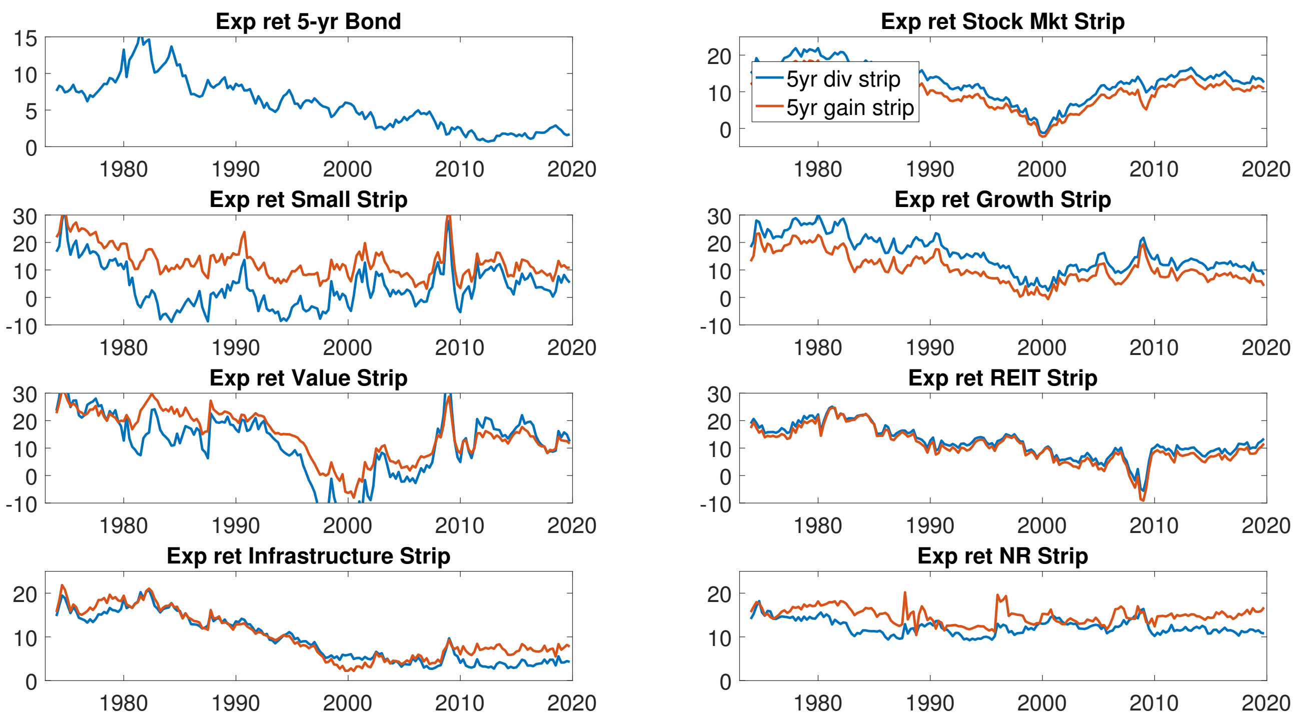

Figure 9. Strip expected returns in the time series.

The figure plots the model-implied time series of risk premia on nominal zero-coupon Treasury bonds in the first panel, and on dividend and capital gain strips on the overall stock market, small stocks, growth stocks, value stocks, REITs, infrastructure stocks, and natural resource stocks in the next seven panels. The maturity for each strip that is plotted is five years (20 quarters).

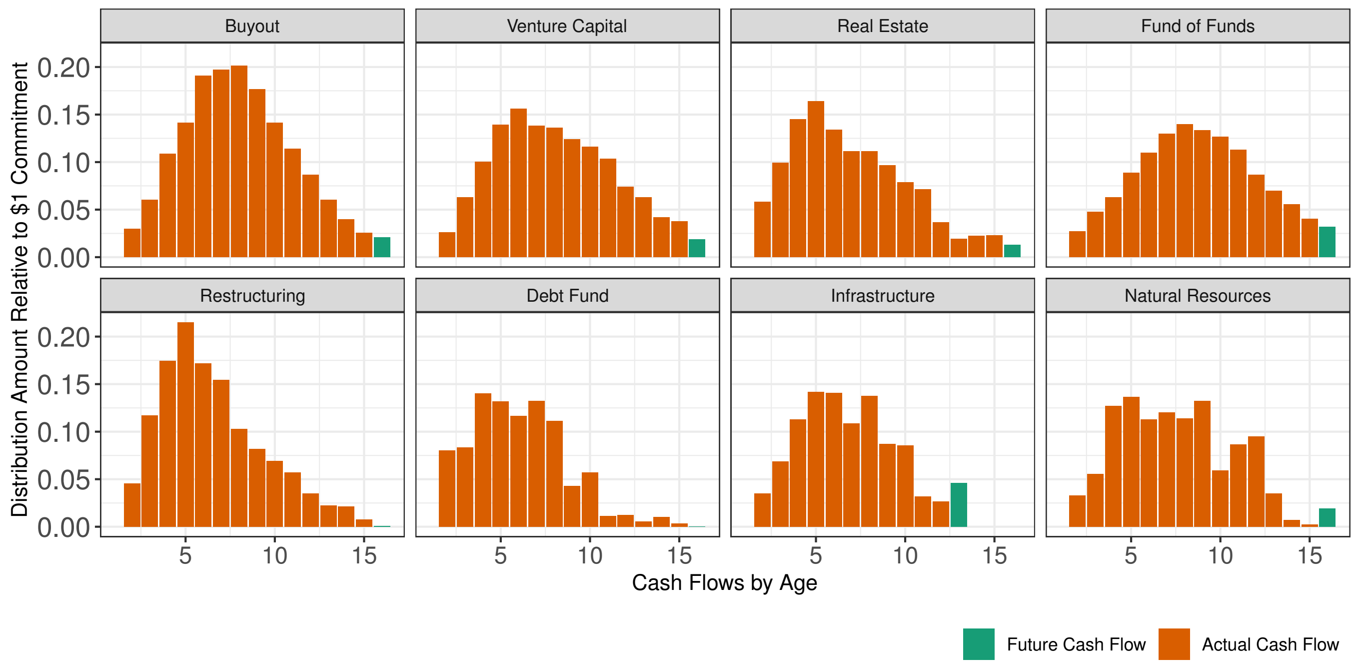

Figure 10. Distribution cash-flow profiles.

The figure plots distribution cash flows by PE fund category. Monthly cash flows are aggregated by year and then averaged across all funds in the category. Year zero is the year of fund inception, the year in which the first capital call is made. The last bar for year 16 contains not only the cash flows in year 16 but also the discounted cash flows after year 16, discounted back to year 16 at the Treasury yield curve.

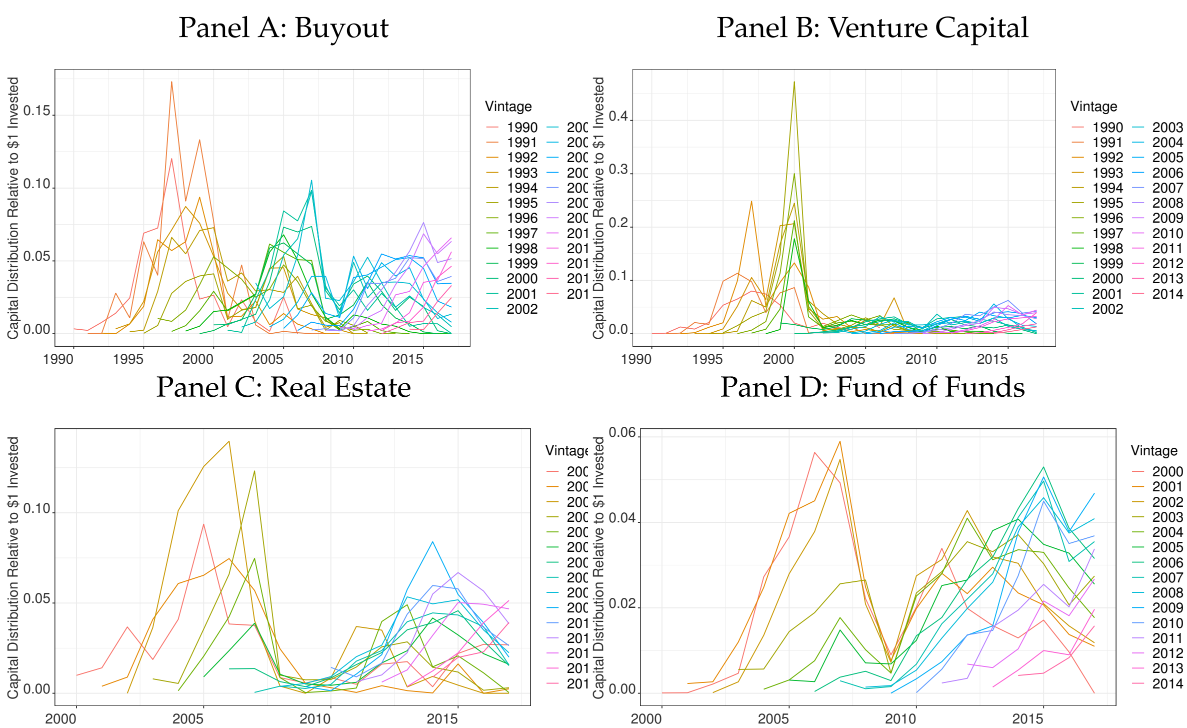

Figure 11. Cash flows by vintage.

This figure plots average cash-flow profiles by vintage. Only vintages from 1990 onwards are plotted for Buyout and Venture Capital, and only vintages from 2000 onwards are plotted for Real Estate and Fund of Fund categories.

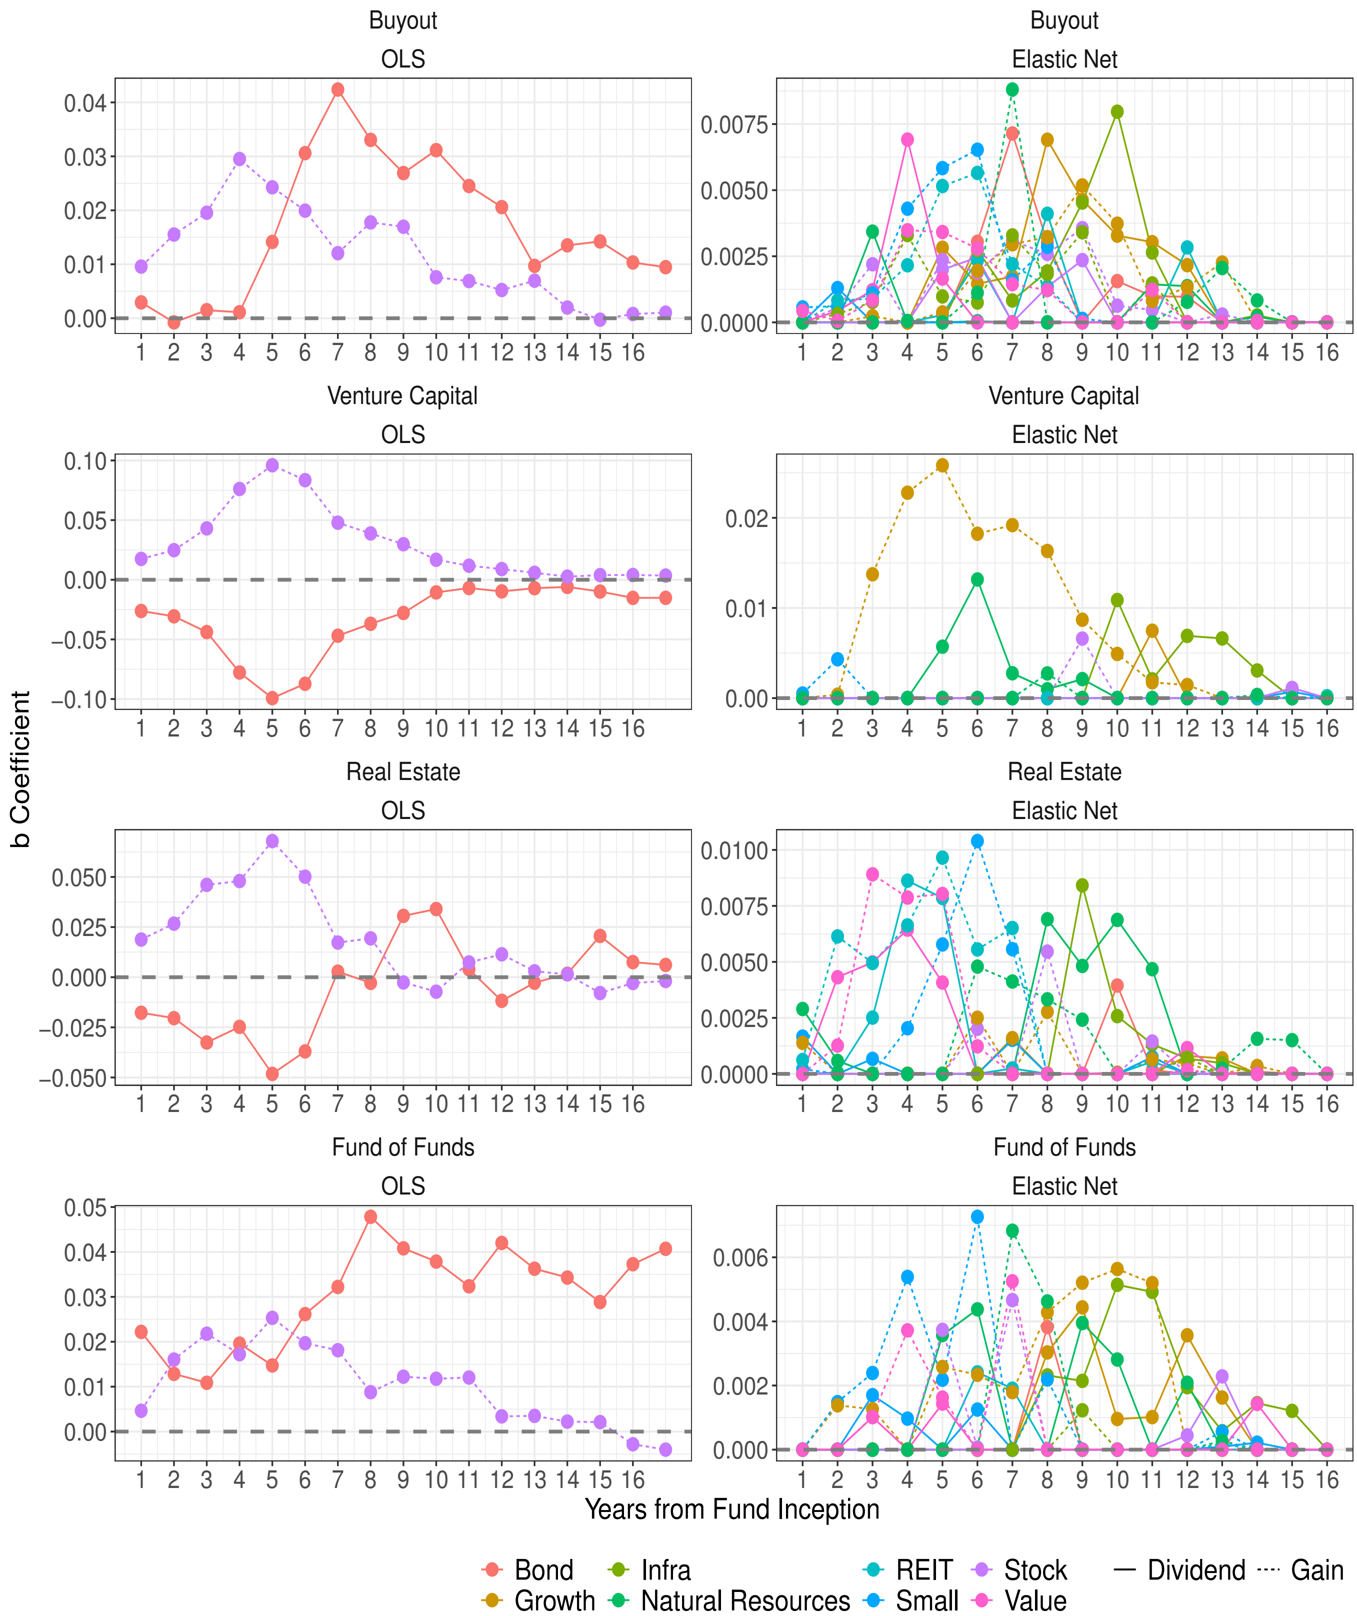

Figure 12. Factor exposure over fund horizon.

Two-Factor | Elastic Net

Panel A: Buyout

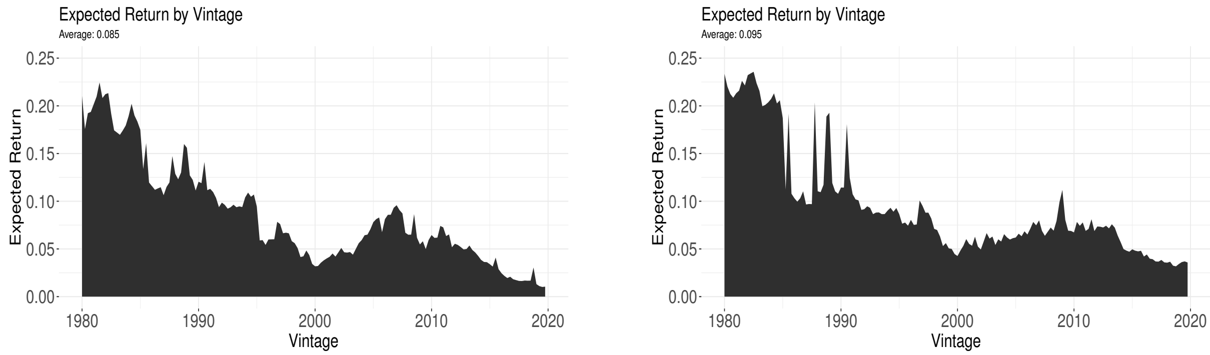

Panel B: Venture Capital

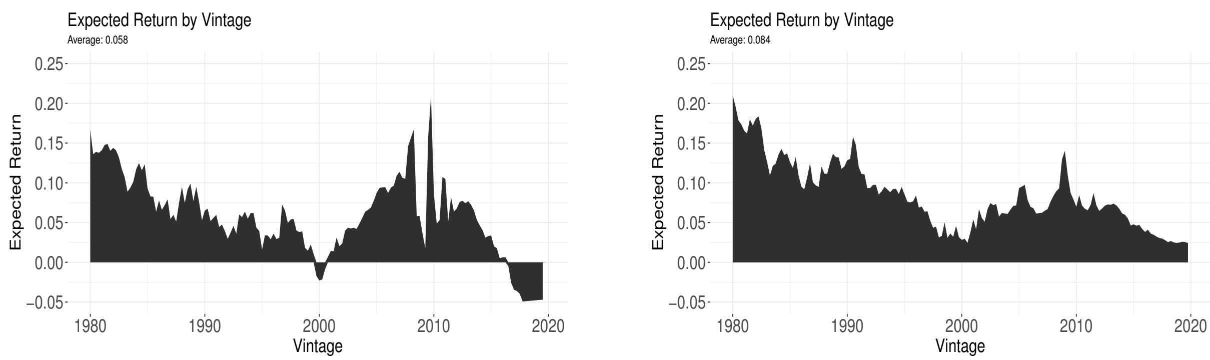

Panel C: Real Estate

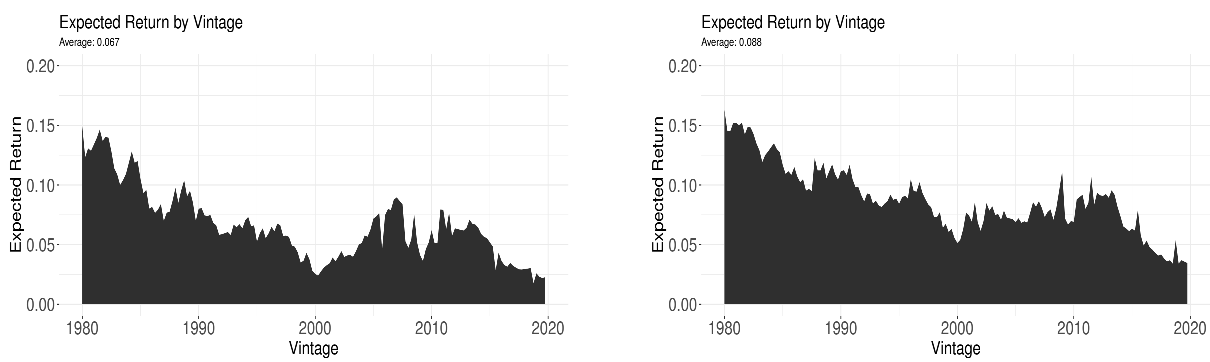

Panel D: Fund of Funds

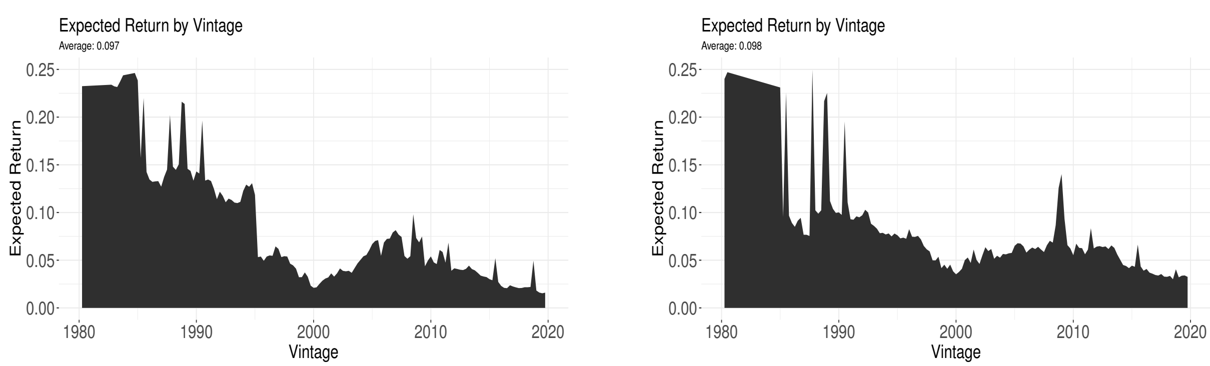

Figure 13. Expected returns by vintage.

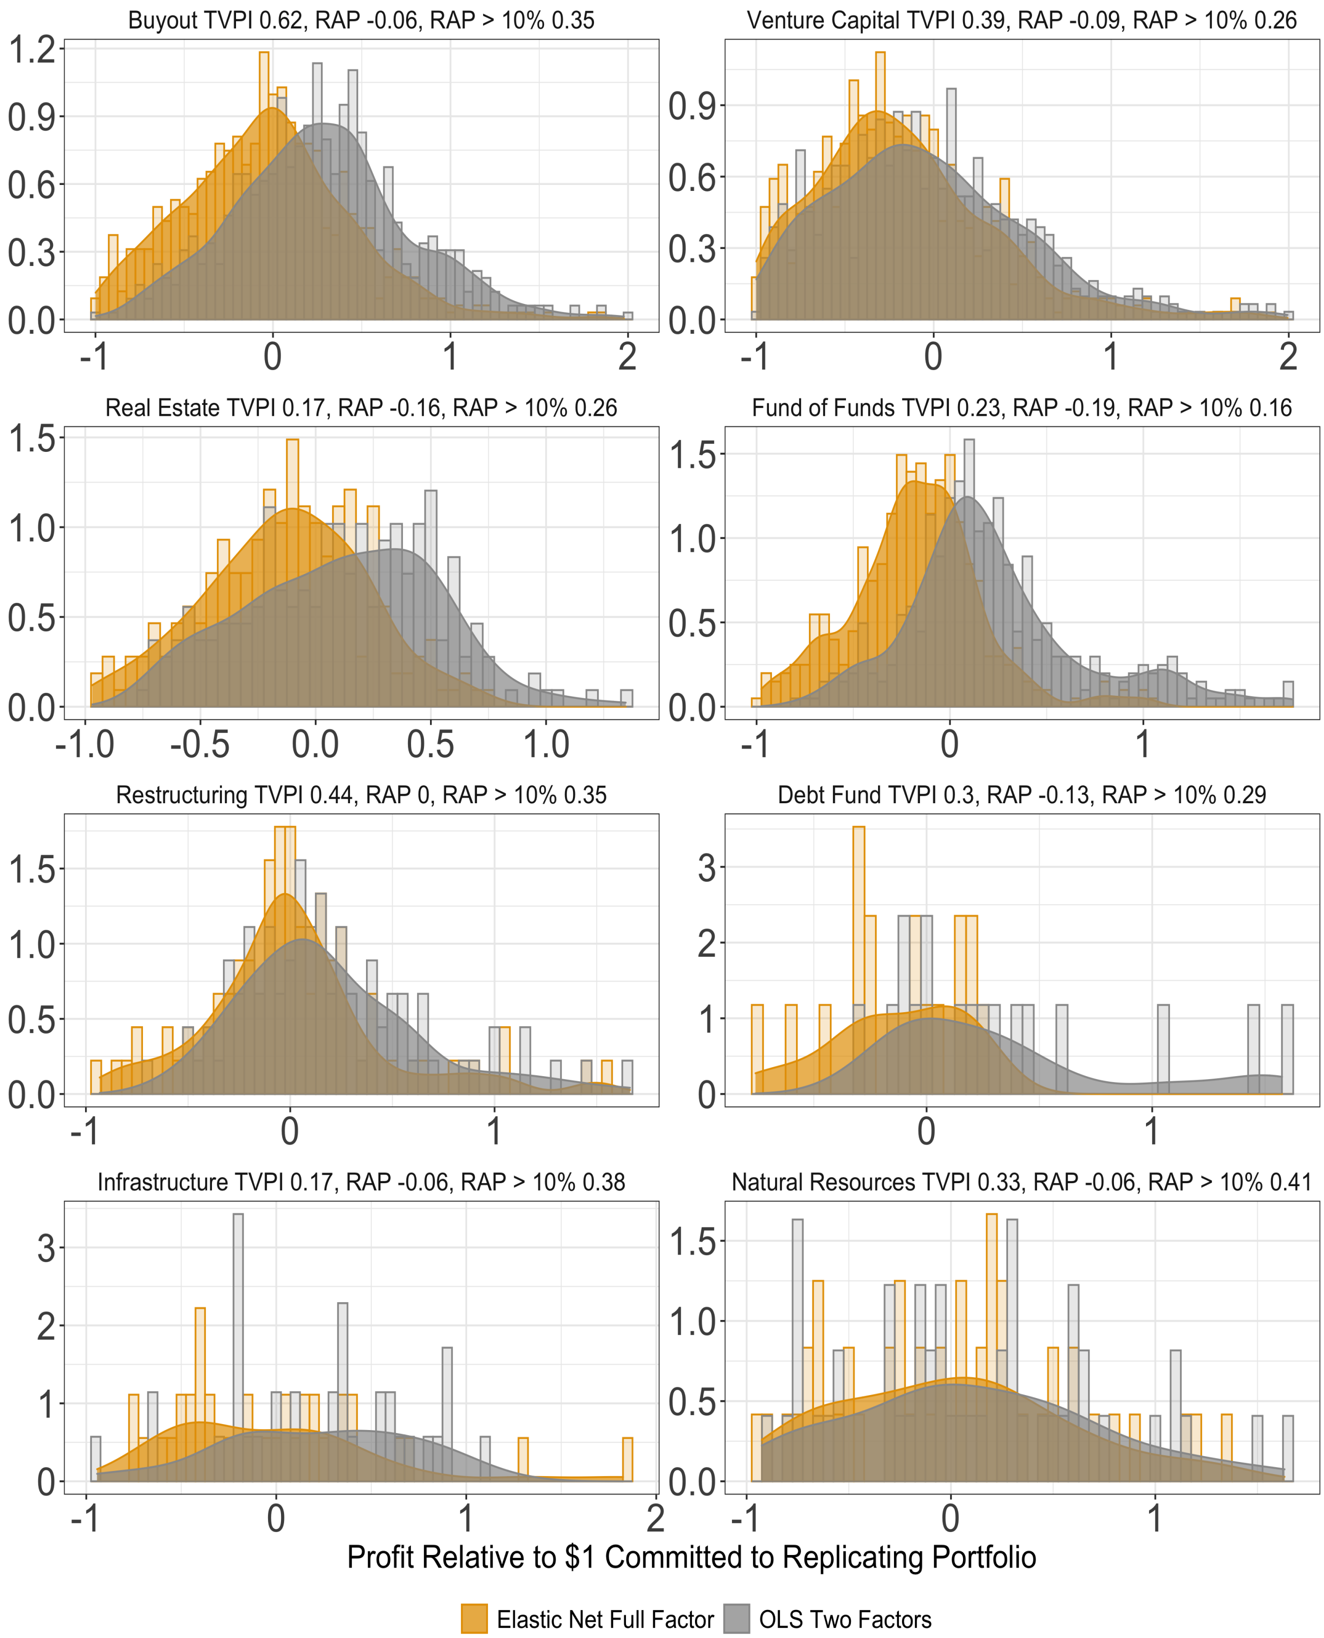

Figure 14. Histogram of fund-level profit relative to replicating portfolio.

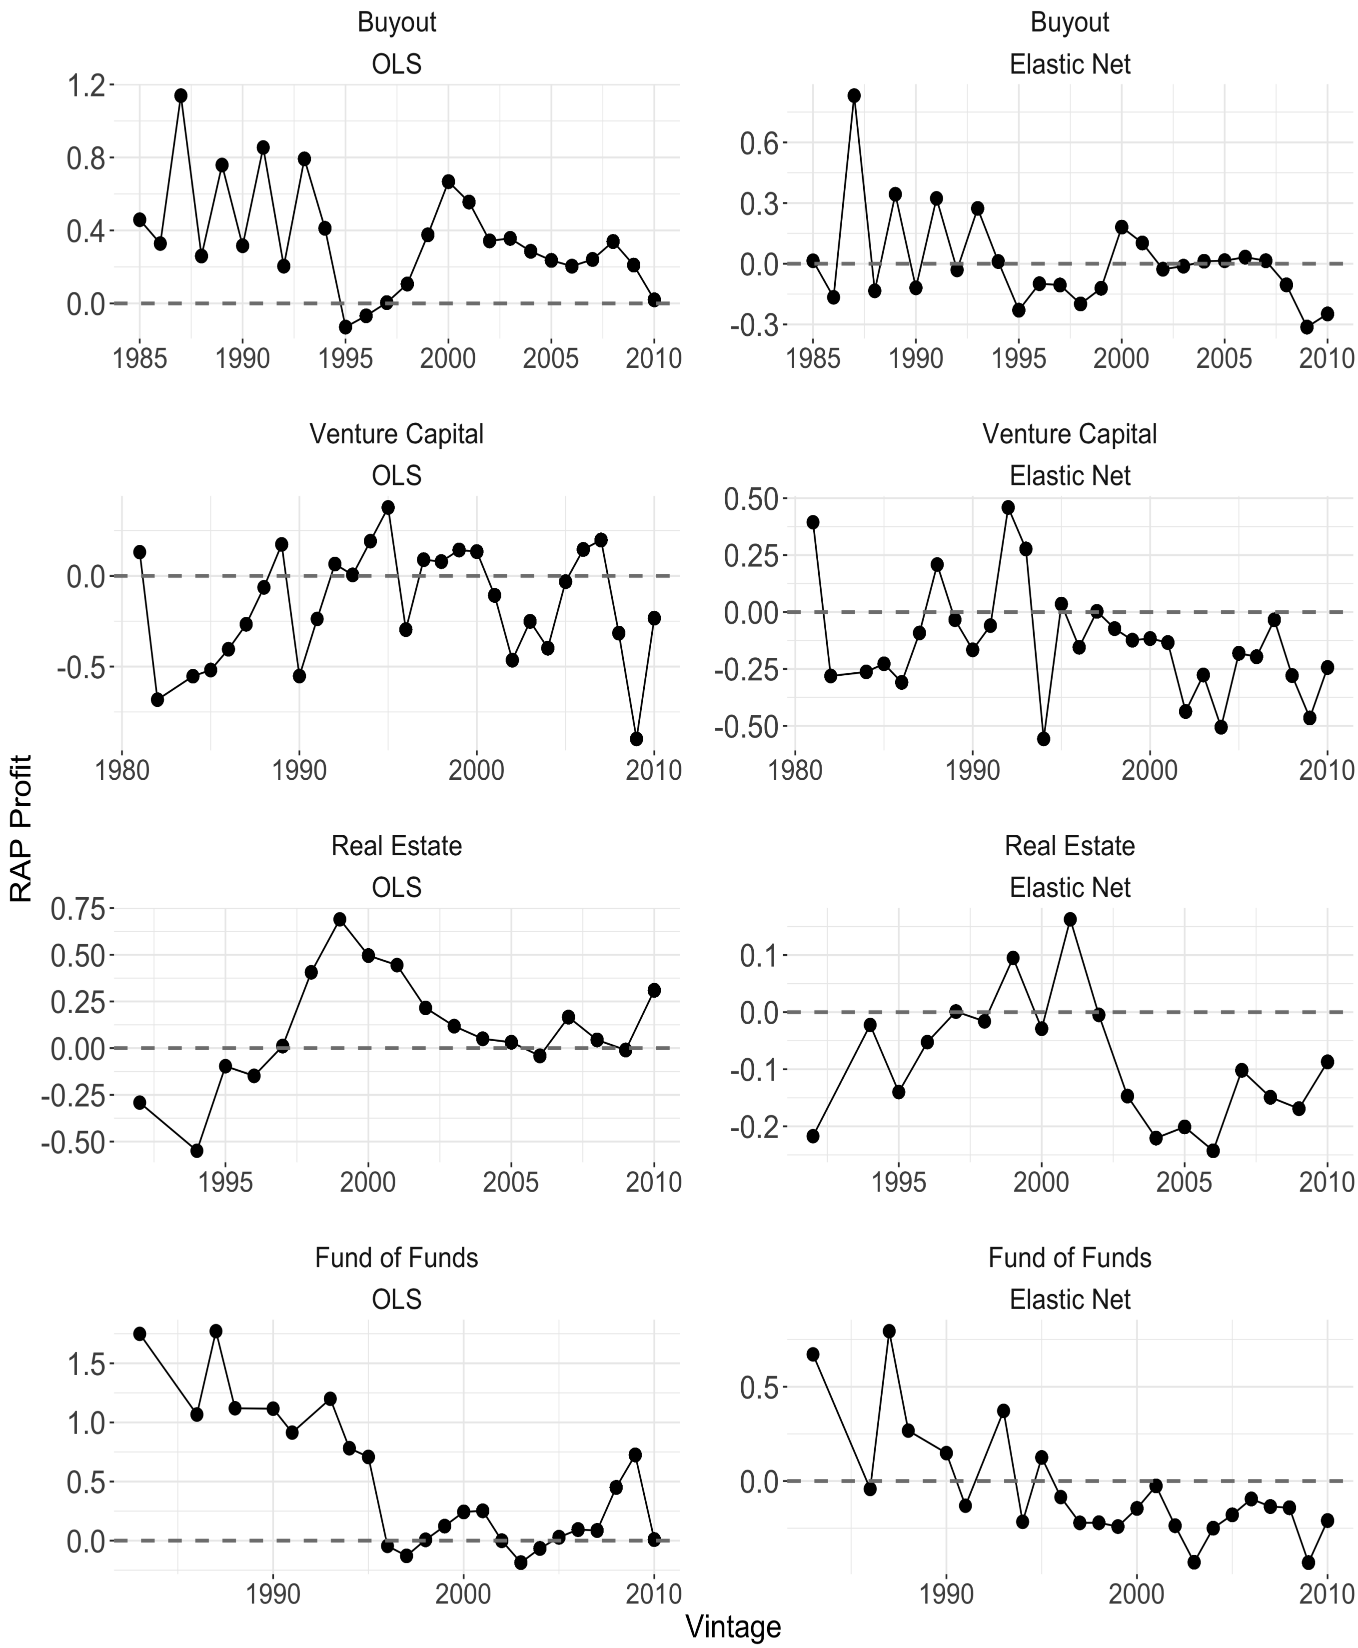

Figure 15. Average fund-profits over time.

Panel A: Buyout

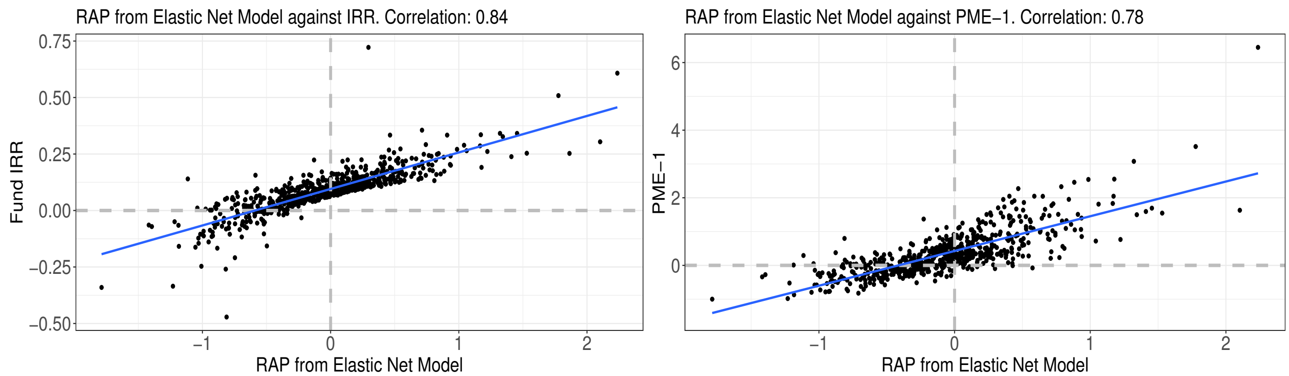

Panel B: Venture Capital

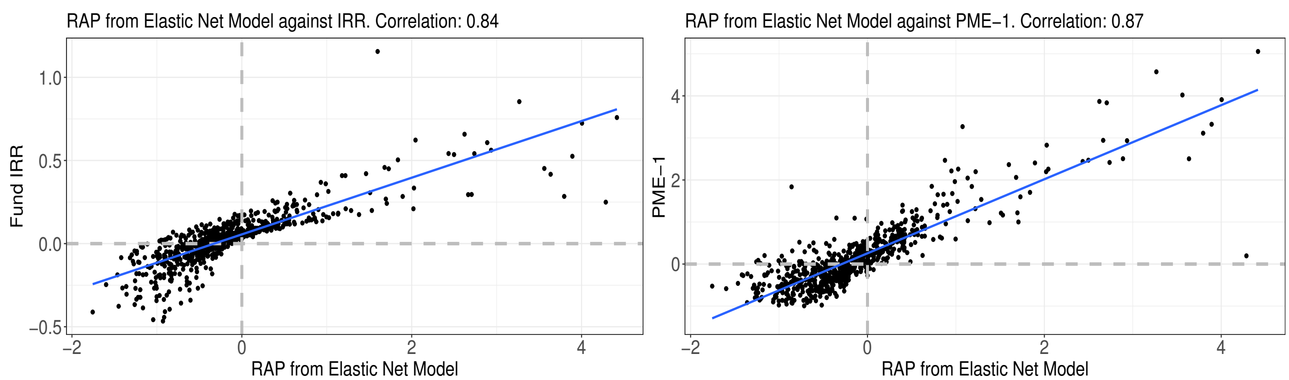

Panel C: Real Estate

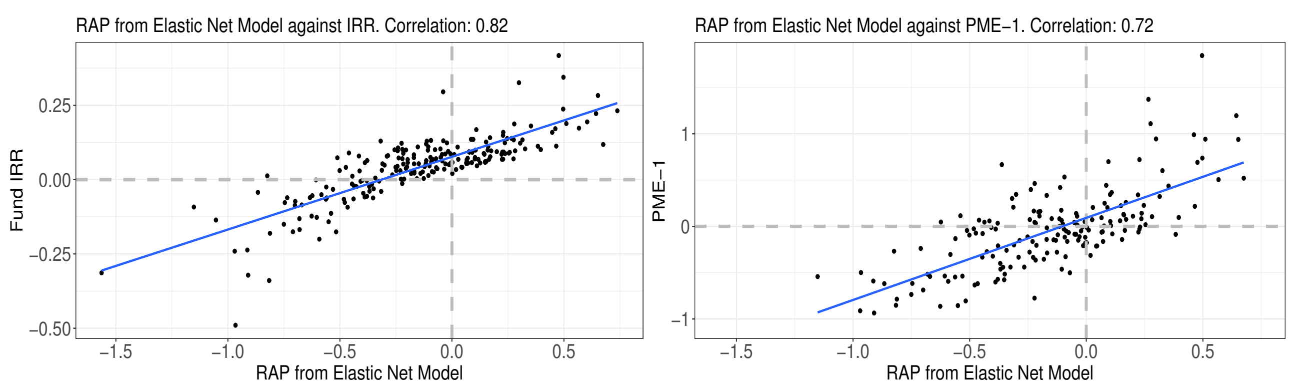

Panel D: Fund of Funds

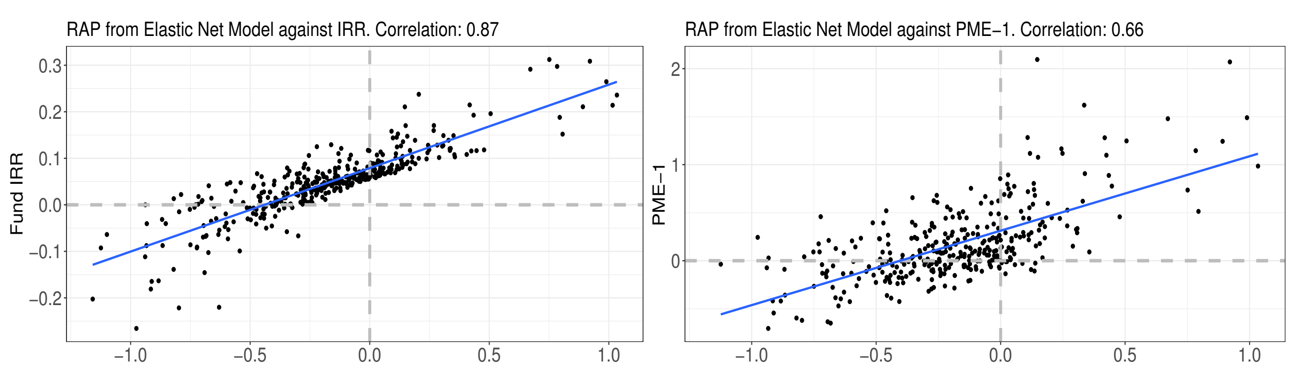

Figure 16. Comparing RAP to IRR and PME.

Table I
Summary Statistics

This table presents all funds in Preqin that have cash-flow data. Not all funds have AUM data, particularly in the early years of the sample.

Panel A: Fund Count

Vintage	Buyout	Venture Capital	Real Estate	Infrastructure	Restructuring	Fund of Funds	Debt Fund	Natural Resources	Total	PD Ratio
1981	0	1	0	0	0	0	0	0	1	1
1982	0	3	0	0	0	0	0	0	3	1
1983	0	1	0	0	0	1	0	0	2	1
1984	1	3	0	0	0	0	0	0	4	1.25
1985	4	6	0	0	0	1	0	0	11	1.64
1986	1	7	0	0	0	2	0	0	10	1.8
1987	5	5	0	0	0	0	0	0	10	1.7
1988	7	4	0	0	0	1	0	0	12	1.5
1989	3	5	0	0	0	1	0	2	11	2
1990	7	8	0	0	1	2	0	0	18	1.78
1991	3	4	0	0	2	0	0	0	9	2
1992	9	12	1	0	2	0	0	1	25	2
1993	9	11	0	0	0	2	0	0	23	2
1994	16	11	1	1	1	2	0	1	33	2.09
1995	15	17	2	0	0	4	0	1	39	2.82
1996	21	24	4	1	3	1	0	0	56	3.32
1997	23	23	5	0	2	5	0	1	60	4
1998	40	33	3	1	1	12	0	3	93	4
1999	31	50	2	0	3	9	1	1	97	4
2000	35	87	7	0	3	17	1	0	150	4
2001	20	52	2	0	5	19	0	1	99	4
2002	23	30	2	1	4	13	1	2	76	4
2003	18	20	7	1	4	13	1	1	65	3.99
2004	28	34	11	4	2	24	1	2	106	3.99
2005	56	48	20	0	6	35	2	5	172	3.74
2006	76	60	35	5	10	54	0	4	244	3.96
2007	74	69	37	6	14	49	1	7	257	3.55
2008	65	63	39	4	11	72	5	8	268	2.55
2009	31	31	14	6	8	36	2	4	132	2.49
2010	43	41	34	9	9	41	3	8	189	3.5
2011	55	52	50	9	10	70	2	9	260	2.94
2012	64	46	41	8	13	55	4	11	245	3
2013	69	50	60	11	18	70	14	8	302	3
2014	67	65	55	15	16	77	12	15	322	2.99
2015	77	79	83	14	20	75	18	8	375	2.78
2016	96	79	59	16	11	92	13	16	383	3
2017	53	76	65	14	14	49	29	11	312	3
Total:	1,145	1,210	639	126	193	904	110	130	4,474	-

Panel B: Fund AUM ($m)

Vintage	Buyout	Venture Capital	Real Estate	Infrastructure	Restructuring	Fund of Funds	Debt Fund	Natural Resources	Total
1981	0	0	0	0	0	0	0	0	0
1982	0	55	0	0	0	0	0	0	55
1983	0	0	0	0	0	75	0	0	75
1984	59	189	0	0	0	0	0	0	248
1985	1,580	74	0	0	0	200	0	0	1,854
1986	0	335	0	0	0	1,310	0	0	1,645
1987	1,608	1,061	0	0	0	0	0	0	2,669
1988	2,789	463	0	0	0	0	0	0	3,252
1989	805	305	0	0	0	1,775	0	210	3,095
1990	2,553	1,134	0	0	153	381	0	0	4,221
1991	1,068	450	0	0	329	0	0	0	1,847
1992	1,150	1,320	0	0	59	0	0	184	2,713
1993	3,192	1,433	0	0	0	597	0	0	5,276
1994	7,577	1,189	488	861	93	357	0	658	11,223
1995	10,648	2,645	523	0	0	1,042	0	205	15,063
1996	8,279	4,558	2,681	1,013	1,600	242	0	0	18,474
1997	23,534	5,245	2,812	0	1,700	1,852	0	480	35,623
1998	39,410	9,001	3,461	1,671	52	11,144	0	2,262	67,001
1999	34,418	17,850	2,293	0	3,133	9,323	109	42	67,168
2000	54,995	39,206	7,324	0	3,320	13,540	230	0	118,615
2001	26,870	23,441	3,225	0	7,461	11,607	0	1,375	73,979
2002	23,233	8,065	4,940	950	2,844	8,699	100	845	49,676
2003	32,629	6,670	3,085	734	5,105	9,059	366	150	57,798
2004	34,523	9,702	6,145	2,725	2,580	5,390	215	2,721	64,001
2005	97,698	15,226	25,036	0	5,830	25,103	412	6,353	175,658
2006	218,225	32,407	45,376	8,054	19,728	41,740	0	9,472	375,002
2007	186,374	23,970	44,794	9,773	40,995	42,708	400	11,795	360,809
2008	164,993	31,870	44,285	8,418	26,158	43,258	4,697	20,416	344,345
2009	40,037	11,716	10,396	9,480	11,235	18,282	195	3,450	104,791
2010	33,186	22,030	19,593	10,920	12,955	12,994	1,164	8,637	121,823
2011	104,281	23,608	55,532	8,325	13,453	29,670	1,720	10,152	247,901
2012	93,805	32,067	28,541	13,306	22,828	40,276	1,054	21,646	253,903
2013	94,105	22,975	61,832	21,509	28,020	22,372	14,211	13,177	279,471
2014	123,339	34,146	40,855	31,666	20,996	38,449	5,619	24,060	319,130
2015	126,648	31,464	73,480	14,233	32,970	65,562	17,113	14,584	376,304
2016	194,563	39,640	46,805	45,637	14,803	48,513	11,773	18,400	420,402
2017	99,981	33,041	58,128	12,051	12,741	22,090	33,231	16,692	288,343
Total:	1,888,155	488,551	591,630	201,326	291,141	527,610	92,609	187,966	4,273,453

Table II
Model Comparison

This table reports the cross-sectional mean and standard deviation of the risk-adjusted profit (RAP) for the main four fund categories in Panel A and the remaining four categories in Panel B. The benchmark RAP metric is labeled "NPV call." RAP metrics for two alternative assumptions on calls are labeled "Sum Call" and "Residual Call." The first row reports the TVPI (total distributions minus total calls). The second row reports the IRR, assuming an initial investment equal to all calls discounted at the Treasury yield curve. The third row reports the PME (public market equivalent), subtracting the initial investment of $1. All metrics are relative to a $1 capital commitment.

Panel A: Main Categories								
	Buyout		VC		Real Estate		Fund of Funds	
	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev
TVPI	0.62	(0.74)	0.39	(1.69)	0.17	(0.52)	0.23	(0.51)
IRR (%)	0.09	(0.10)	0.03	(0.20)	0.04	(0.11)	0.05	(0.07)
PME-1	0.36	(0.67)	0.22	(1.49)	-0.04	(0.44)	0.17	(0.40)
RAP 2-factor (NPV Call)	0.28	(0.53)	-0.15	(1.36)	0.09	(0.45)	0.24	(0.50)
RAP 15-factor (NPV Call)	-0.06	(0.51)	-0.09	(1.27)	-0.16	(0.38)	-0.19	(0.35)
RAP 2-factor (Sum Call)	0.20	(0.53)	-0.25	(1.36)	0.04	(0.45)	0.15	(0.51)
RAP 15-factor (Sum Call)	-0.14	(0.51)	-0.18	(1.27)	-0.20	(0.38)	-0.28	(0.36)
RAP 2-factor (Residual Call)	0.18	(0.53)	-0.26	(1.36)	0.02	(0.45)	0.13	(0.51)
RAP 15-factor (Residual Call)	-0.16	(0.51)	-0.19	(1.27)	-0.22	(0.38)	-0.30	(0.36)
Panel B: Additional Categories								
	Restructuring		Debt Fund		Infrastructure		Natural Resources	
	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev
TVPI	0.44	(0.57)	0.30	(0.27)	0.17	(0.65)	0.33	(0.91)
IRR (%)	0.09	(0.10)	0.07	(0.04)	0.03	(0.11)	0.02	(0.18)
PME-1	0.20	(0.56)	0.12	(0.17)	0.17	(0.57)	0.28	(0.86)
RAP 2-factor (NPV Call)	0.17	(0.47)	0.34	(0.55)	0.33	(0.65)	0.07	(0.66)
RAP 15-factor (NPV Call)	-0.001	(0.46)	-0.13	(0.30)	-0.06	(0.58)	-0.06	(0.62)
RAP 2-factor (Sum Call)	0.13	(0.47)	0.31	(0.56)	0.27	(0.66)	0.00	(0.65)
RAP 15-factor (Sum Call)	-0.04	(0.46)	-0.16	(0.31)	-0.12	(0.59)	-0.13	(0.60)
RAP 2-factor (Residual Call)	0.11	(0.47)	0.28	(0.57)	0.24	(0.66)	-0.03	(0.66)
RAP 15-factor (Residual Call)	-0.07	(0.47)	-0.18	(0.31)	-0.15	(0.58)	-0.16	(0.61)

Table III
Persistence in Performance

This table reports various fund persistence measures for TVPI, IRR, PME-1, the two-factor OLS model, and the 15-factor elastic net model under the baseline assumption on calls (NPV of calls). Persistence measures the average correlation of RAP between successive funds by the same firm in the same PE category. Top Quart is the probability that if the first fund in the pair is in the top quartile of the distribution of RAP of that vintage-category, the second fund in the pair is also in the top quartile of RAP in its vintage-category. Bottom Quart is the corresponding probability that if the first fund is in the bottom quartile, the second fund in the pair is also in the bottom quartile.

	Buyout			VC		
	Persistence	Top Quart	Bottom Quart	Persistence	Top Quart	Bottom Quart
TVPI	0.46	0.36	0.32	0.41	0.45	0.43
IRR	0.40	0.38	0.29	0.41	0.45	0.35
PME-1	0.44	0.29	0.27	0.43	0.41	0.47
RAP 2-factor (NPV Calls)	0.44	0.34	0.25	0.41	0.33	0.35
RAP 15-factor (NPV Calls)	0.42	0.36	0.31	0.39	0.35	0.14
	Real Estate			**Fund of Funds**		
	Persistence	Top Quart	Bottom Quart	Persistence	Top Quart	Bottom Quart
TVPI	0.33	0.46	0.20	0.34	0.31	0.21
IRR	0.25	0.41	0.20	0.28	0.27	0.22
PME-1	0.27	0.44	0.35	0.30	0.35	0.22
RAP 2-factor (NPV Calls)	0.26	0.44	0.16	0.27	0.24	0.20
RAP 15-factor (NPV Calls)	0.26	0.40	0.22	0.29	0.29	0.17