

Valuing Private Equity Investments Strip by Strip

Arpit Gupta and Stijn Van Nieuwerburgh*

January 28, 2021

Abstract

We propose a new valuation method for private equity investments. It constructs a replicating portfolio using cash-flows on various listed equity and fixed income instruments (strips). It then values the listed strips using an asset pricing model that accurately 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. Applying the method to the universe of PE funds, we find negative risk-adjusted profits for the average fund, 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: 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 twenty-five years. Indeed, 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, Carrellas, Chau, and Duane, 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 non-traded 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 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 LP receives 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, the public market equivalent 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 the PE literature has done hitherto. 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 (from asset dispositions) 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 the same 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 is 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 either because it delivers positive idiosyncratic cash flows or because it delivers systematic cash-flow exposure at a cost that is lower 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 break down that 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 enables 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 across eight PE categories. We follow funds started from 1981 until 2018, and 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 exposures 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 gains 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 risk-adjusted profits (RAP) across PE categories. A substantial component of the return to PE investment, which previous research has considered to be out-performance (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 PME (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 out-performance into substantial under-performance.

Third, 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 out-perform, 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 [Let-](#)

tau and Wachter (2011), van Binsbergen, Brandt, and Koijen (2012), and van Binsbergen, Hueskes, Koijen, and Vrugt (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 temporal composition of risk in the cross-section of equities. We also introduce the concept of gain strips, assets which 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 auto-regression 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

The second contribution is a re-assessment of valuation approaches that use the realized SDF. Rather than discounting PE cash flows by the realized SDF, as is done for example in 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 risk-adjusted profit (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 non-traded 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 performance evaluation of PE funds, such as [Kaplan and Schoar \(2005\)](#), [Phalippou and Gottschalg \(2009\)](#), [Cochrane \(2005\)](#), [Harris, Jenkinson, and Kaplan \(2014\)](#), [Korteweg and Sorensen \(2017\)](#), [Robinson and Sensoy \(2011\)](#), among many others. Most of this lit-

¹To 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 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 point of view. Further exploration of this possibility is an interesting direction for future work.

²The issue is similar to the one in [Martin \(2012\)](#). Let M be the SDF and X be the cash-flow. [Martin \(2012\)](#) shows that the sample mean of $M_1 X_1, \dots, M_t X_t$ 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.

erature 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. The literature has found 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 until 2019 and covers all PE investment categories. Our approach results in a 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 ones, like PME and GPME, is that our method 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 which vary by (i) PE category, (ii) vintage, and (iii) horizon of the cash flow. This decomposition provides new insight into the nature of risk PE investors are exposed to.

In complementary work, Ang, Chen, Goetzmann, and Phalippou (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

³The 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 considers a larger set of risks and allows the risk prices 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, e.g., Ljungqvist and Richardson (2003), Driessen, Lin, and Phalippou (2012).

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 number of factors as well as the factor risk premium dynamics. Other important methodological contributions to the valuation of private equity include [Driessen, Lin, and Phalippou \(2012\)](#), [Sorensen, Wang, and Yang \(2014\)](#), and [Metrick and Yasuda \(2010\)](#).

The 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, Nagel, and Santosh, 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 the main results on the risk-adjusted profits and expected returns of PE funds. Section [IV](#) concludes. Appendices [B](#) provide additional derivations, while Appendix [C](#) has additional detail on the VAR estimation. Appendix [D](#) provides estimates on shock-exposure elasticities of our estimates. Appendix [E](#) includes results on additional fund categories. Appendix [F](#) performs a validation exercise on public equities. Appendix [G](#) shows robustness of our estimates across different choices of hyper-parameters in the Elastic Net estimation, and Appendix [H](#) provides estimates on the Burgiss data set.

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 2-4 years. Intermediate cash-flows may accrue from the operations of the assets, such as net operating income from renting out an office building. Towards the end of the life of the fund (typically in years 5-12), the GP “harvests” 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, and 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 $\{X_{t+h}^i\}_{h=0}^T$. 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 ten years. Any cash flows observed in the data after quarter 64 are discounted and allocated evenly to quarters 61-64. 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 de-

ploying the money, but rewards for skillful delay. Under this alternative assumption, the LP's counter-factual investment strategy (to be compared to the actual fund investment) invests the full \$1 in the replicating portfolio at time t .⁴ We return to the role of calls in Section 4.7 and present results there for the alternative treatment of calls. Having dealt with the capital calls and their timing, our method focuses on valuing the distribution cash flows.

I.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 $F_{t,t+h}$ 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 $F_{t,t+h}$ is the pay-off on a nominal zero-coupon bond that is bought at time t and matures at time $t + h$, namely \$1. The second element of $F_{t,t+h}$ is the pay-off on a dividend strip on the aggregate stock market of maturity h , $\frac{D_{t+h}^m}{D_t^m}$. 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 $F_{t,t+h}$ is the pay-off on a capital gain strip on the aggregate stock market of maturity h , $\frac{P_{t+h}^m}{P_t^m}$. 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 $\frac{D_{t+h}^{value}}{D_t^{value}}$ while the payoff on a value gain strip is $\frac{P_{t+h}^{value}}{P_t^{value}}$. 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, there are $KH = 960$ strips in total.

⁴We truncate call amounts above \$1. While infrequent, such instances occur more frequently in more recent periods. They arise because of recycling provisions which 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.

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 $\mathbf{P}_{t,h}$. The first element is the time- t price of a nominal zero-coupon bond of maturities 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 stochastic discount factor (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:

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

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.

The second step of our approach is to obtain the replicating portfolio, consisting 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 $\beta_{t,h}^i$. 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:

$$X_{t+h}^i = \beta_{t,h}^i \mathbf{F}_{t,t+h} + e_{t+h}^i \quad (1)$$

where e_{t+h}^i denotes the idiosyncratic cash-flow component, orthogonal to $\mathbf{F}_{t,t+h}$. We explain below 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:

$$\mathbb{E}_t [R^i] = \sum_{h=1}^H \sum_{k=1}^K \tilde{w}_{t,h}^i(k) \mathbb{E}_t [R_{t+h}(k)] \quad (2)$$

where w^i is a $1 \times HK$ vector of replicating portfolio weights with generic element $w_{t,h}^i(k) = \beta_{t,h}^i(k) P_{t,h}(k)$. The $HK \times 1$ vector $\mathbb{E}_t[\mathbf{R}]$ 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 $\beta_{t,h}^i$ 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, i.e., “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:

$$\delta_t^i = \frac{1}{4} \sum_{h=1}^H \sum_{k=1}^K \tilde{w}_{t,h}^i(k) h, \quad (3)$$

where the weights $\tilde{w}_{t,h}^i(k)$ are the original weights $w_{t,h}^i(k)$ rescaled to sum to 1. The annualized expected PE fund return is then:

$$\mathbb{E}_t [R_{ann}^i] = \left(1 + \mathbb{E}_t [R^i] \right)^{1/\delta_t^i} - 1. \quad (4)$$

This is the first main object of interest.⁵

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 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 risk-adjusted profit (RAP).

We define the RAP for fund i in vintage t :

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

The risk-adjusted profit is the difference between the net present value 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 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. Or it is 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

⁵For 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 calculate 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 under-state the expected return if the term structure of expected strip returns is flat or upward sloping and over-state the expected return if the term structure is downward sloping.

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 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 out-performance 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.⁶ 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 liter-

⁶The 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."

ature 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 we include should go a long way towards capturing this factor structure. Our method can easily accommodate extra factors.⁷

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.

I.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.

⁷Omitted risk factors may bias RAP estimates. Which way the bias goes depends on the covariance of the 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 cost of the replicating portfolio under the misspecified model. As an example, we consider the setting of Appendix A.1, which assumes 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.

Identifying Assumptions Identification is achieved both from the cross-section and from the time series. We make four assumptions. First, the cash-flows $X_{t+h}^{i \in c}$ of all funds i in the same category c (category superscripts are omitted below for ease of notation) and vintage t have the same risk factor exposures at horizon h , $\beta_{t,h}^i(k) = \beta_{t,h}^c(k), \forall i \in c$. 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_t^k and an *age effect* b_h^k for each factor k . Third, horizon effects are constant for the four quarters within the same calendar year. This reduces the number of horizon effects that need to be estimated for each factor from $H = 64$ to $H/4 = 16$.⁸ 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 captures dependence on the overall investment climate at the time of PE fund origination. [Haddad, Loualiche, and Plosser \(2017\)](#) emphasizes the importance of P/D 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.⁹ To simplify the time dimension, vintages are grouped in 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 1974-2019 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 for pur-

⁸We find no systematic evidence of seasonality in PE fund cash flows.

⁹A natural alternative would be to consider total capital raised by vintage. One challenge with this metric is that it is non-stationary since the PE industry has grown. Even scaling by GDP does not remove this trend. Nevertheless, we find similar results under this alternative. Results are available upon request.

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

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

We estimate equation (6) by OLS. This model extends 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:

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

In the empirical implementation, $K = 15$. The 15 factors are bond strips, and 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:

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

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 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. Appendix G details the hyper-parameter choices and shows robustness of the results to these choices.

I.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 the value added of firms? Define the value-added of a PE fund using the realized SDF as:

$$VAsdf_t^i = \sum_{h=1}^H M_{t,t+h} X_{t,t+h}^i - C_t \quad (8)$$

While in expectation, this measure is identical to the RAP measure, sample averages show wide divergence. A Monte Carlo study in Appendix A 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 (8) 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 Appendix A.

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 necessitates the estimation of exposures or PE fund cash flows to strip cash flows, but this has the advantages of generating useful information about the underlying risk exposures of PE funds. Appendix A 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 was 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 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 either be observed directly from dividend strip futures markets (van Binsbergen, Hueskes, Kojien, and Vrugt, 2013) or inferred from options and stock markets (van Binsbergen, Brandt, and Kojien, 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. Most importantly, the only dividend strip data are for 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, 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. For all these reasons, we need an asset pricing model to generate the time series of strip prices, $P_{t,h}$, and the corresponding expected returns for each strip.

We propose a reduced-form stochastic discount factor (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 noted by [Korteweg and Nagel \(2016\)](#), the objective is not to test the asset pricing model itself but 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 \(2020\)](#) investigates the consumption-based asset pricing models' ability to price PE cash flows.

II.A Setup

II.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:

$$z_t = \Psi z_{t-1} + \Sigma^{\frac{1}{2}} \varepsilon_t, \quad (9)$$

with shocks $\varepsilon_t \sim i.i.d. \mathcal{N}(0, I)$ 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 homoscedastic. We use a Cholesky decomposition of the covariance matrix, $\Sigma = \Sigma^{\frac{1}{2}} \Sigma^{\frac{1}{2}'}$, which has non-zero 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 1-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.

II.A.2 Stochastic Discount Factor

The nominal SDF $M_{t+1}^{\$} = \exp(m_{t+1}^{\$})$ is conditionally log-normal:

$$m_{t+1}^{\$} = -y_{t,1}^{\$} - \frac{1}{2}\Lambda_t'\Lambda_t - \Lambda_t'\varepsilon_{t+1}. \quad (10)$$

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 Λ_t of the affine form:

$$\Lambda_t = \Lambda_0 + \Lambda_1 z_t. \quad (11)$$

The $N \times 1$ vector Λ_0 collects the average prices of risk while the $N \times N$ matrix Λ_1 governs the time variation in risk premia. Asset pricing amounts to estimating the market prices of risk (Λ_0, Λ_1) .

II.A.3 Bond Pricing

Proposition 1 in Appendix B 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 twenty-quarter nominal bond yields since those yields are part of the state vector z_t . We also fit the dynamics of twenty-quarter nominal bond risk premia (price changes).

II.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 present-discounted value 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 :

$$\frac{P_t^m}{D_t^m} = \sum_{\tau=1}^{\infty} P_{t,\tau}^d \quad (12)$$

$$\exp\left(\overline{pd} + e'_{pd^m} z_t\right) = \sum_{\tau=0}^{\infty} \exp\left(A_{\tau}^m + B_{\tau}^{m'} z_t\right), \quad (13)$$

where $P_{t,\tau}^d$ denotes the price of a τ -period dividend strip divided by the current dividend. Proposition 2 in Appendix B shows that the log price-dividend ratio on each dividend strip, $p_{t,\tau}^d = \log(P_{t,\tau}^d)$, is affine in the state vector and provides recursions for the coefficients $(A_{\tau}^m, B_{\tau}^{m'})$. 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 (13) restates the present-value relationship from equation (12). It articulates a non-linear 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 risk price, 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 only account for 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.

II.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, $F_{t,\tau}^d$, scaled by the current nominal dividend $D_t^{\$}$, is:

$$\frac{F_{t,\tau}^d}{D_t^{\$}} = P_{t,\tau}^d \exp\left(\tau y_{t,\tau}^{\$}\right),$$

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

$$R_{t+1,\tau}^{fut,d} = \frac{F_{t+1,\tau-1}^d}{F_{t,\tau}^d} - 1.$$

Appendix B shows that $\log(1 + R_{t+1,\tau}^{fut,d})$ 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.

II.B Estimation

II.B.1 State Vector Elements

The state vector contains the following $N = 18$ variables, in order of appearance: (1) GDP price inflation, (2) real GDP growth, (3) the nominal short rate (3-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. Elements 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). Elements 8, 10, 12, 14, 16, and 18 are the corresponding log real dividend growth rates:¹⁰

$$z_t = \left[\pi_t, x_t, y_{t,1}^{\$}, y_{t,20}^{\$} - y_{t,1}^{\$}, pd_t^m, \Delta d_t^m, pd_t^{reit}, \Delta d_t^{reit}, pd_t^{infra}, \Delta d_t^{infra}, \right. \\ \left. pd_t^{small}, \Delta d_t^{small}, pd_t^{growth}, \Delta d_t^{growth}, pd_t^{nr}, \Delta d_t^{nr}, pd_t^{value}, \Delta d_t^{value} \right]'. \quad (14)$$

This state vector is observed at quarterly frequency from 1974.Q1 until 2019.Q4 (184 observations). This is the longest available time series for which all variables are available.¹¹ Our PE cash flow data starts shortly thereafter in the early 1980s. While the bulk

¹⁰The 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 with finding a good fit for the asset pricing moments.

¹¹We 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 8 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 onwards and as the Fama-French Oil

of PE cash flows occur after 1990, we deem it advantageous to 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.¹²

The VAR is estimated by OLS in the first stage of the estimation. We recursively zero out all elements of the companion matrix Ψ whose t-statistic is below 1.96. Appendix C contains the resulting point estimates for Ψ and $\Sigma^{\frac{1}{2}}$.

II.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 estimated 12 non-zero parameters in the constant market price of risk vector Λ_0 and 92 non-zero elements of the matrix Λ_1 which governs the dynamics of the risk prices. The point estimates are listed in Appendix C.2. We use the following target moments to estimate the market price of risk parameters.

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

$$-A_{20}^{\$/20} = y_{0,20}^{\$/20} \quad \text{and} \quad -B_{20}^{\$/20} = [0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]$$

Because the demeaned five-year bond yield is the sum of the third and fourth element in the state vector, the market prices of risk must be such that $-B_{20}^{\$/20}$ has a one in the third industry index beforehand.

¹²The VAR literature has found that results can be sensitive to the choice of state variables. [Campbell and Voulteenahe \(2004\)](#) emphasize the role of the small value spread as a predictor of aggregate market returns, [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 for each factor's *pd* ratio to affect the expected return, but also allows for cross-predictability and for 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.

and fourth place 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, ten years, twenty years, and thirty years. They constitute about $7 \times T$ moments, where $T = 184$ quarters.¹³

Third, we match the time-series of real bond yields for maturities of five, seven, ten, twenty, and thirty years. They constitute about $5 \times T_2$ moments, where $T_2 = 68$ quarters.¹⁴ Having both nominal and real bonds helps disentangle the respective roles of growth and inflation risks.

Fourth, we impose 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, equal minus the conditional covariance between the log SDF and the log return. For example, for the aggregate stock market:

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

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 corresponding rows of Λ_1 .

¹³The 20-year bond yield is missing prior to 1993.Q4 while the 30-year bond yield data is missing from 2002.Q1-2005.Q4. In total 107 observations are missing, so that we have $1232 - 107 = 1125$ bond yields to match.

¹⁴FRED data on Treasury Inflation Indexed bond yields start in 2003.Q1. Real yields for the 20-year and 30-year bonds are only available for 61 and 39 quarters, respectively.

Fifth, we match the time-series of log price-dividend ratios (price levels) on the seven stock indices. The model's price-dividend ratios are built up from 3,600 quarterly dividend strips according to equation (12). We impose these present-value relationships in each quarter, delivering $7 \times T$ 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 to the nearest eight dividend strips. Data for the price-dividend ratio on this claim and the fraction it represents of the overall stock market value (S&P500) for the period 1996.Q1-2009.Q3 (55 quarters) are from [van Binsbergen, Brandt, and Koijen \(2012\)](#). This 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, Hueskes, Koijen, and Vrugt \(2013\)](#). 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 Nov 2002 - Jul 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 at the state variables observed between 2003.Q1 and 2014.Q2, average it, and annualize it. This results in one additional restriction. These dividend strip moments identify (some of) the market price of risk 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 impose 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. Appendix C.3 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.¹⁵ In the PE analysis of the next section we take these MPR estimates as given. For 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 “market price of risk estimate” inside $avg(R^m)$ drops out. It is efficient to subtract $\beta avg(R_m)$, not $\beta E[R^m]$.¹⁶

II.B.3 Model Fit

Figure 1 plots the nominal bond yields on bonds of maturities 1 quarter, 1 year, 5 years, and 10 years. Those are the most relevant horizons for the private equity cash-flows.

¹⁵We have calculated standard errors on the MPR parameters by GMM and find them to be modest. Intuitively, by matching the time series of bond yields of various maturities, and stock-price dividend ratios and expected returns on each equity factor plus and 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.

¹⁶We thank the Editor for pointing this analogy out to us. One detail here is 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 (1974-2019) than we have for the fund data (1981-2019). As a robustness check, we have re-estimated our MPR on the 1990-2019 subsample, the sample where most of the fund data is concentrated. The results are similar. These results are available upon request.

The model matches the time series of bond yields in the data closely. It matches nearly perfectly the 1-quarter and 5-year bond yield 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.

Figures 4 and 5 show the equity risk premium, 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 we price. The dynamics of the risk premia in the data are dictated by the VAR. The model chooses the market prices of risk to fit these risk premium dynamics as closely as possible. The price-dividend ratios in the model are formed from the price-dividend ratios on the strips of maturities ranging from 1 to 3600 quarters, as explained above. The figure shows 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 in part because 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.

II.C Temporal Pricing of Risk

The first key input from the model into the private equity 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 dividend. For readability, we plot only three maturities: one, five, and ten years. The model implies substantial variation in strip prices over time, across maturities, as well as across risky assets.

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 8 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 is the model. The model generates the right level for the price-dividend ratio for 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 8 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 there are data. The model captures the dynamics of this share reasonably well, as shown in the right panel, including the decline in 2000.Q4-2001.Q1 when the short-term dividend strip prices fall 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.

The second key input from the model into the private equity valuation exercise are the expected excess returns on the bond and stock strips of horizons of 1-64 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 with maturity from 0 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 premium 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.Q1-2014.Q2 (yellow line). It is, if anything, less downward sloping than the risk premium averaged over the entire 1974-2019 sample (red line). The model matches the *realized* portfolio return on dividend futures of maturities 1-7 years over the period 2003.Q1-2014.Q2, which is 8.7% in the data and 8.6% in the model.¹⁷

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 5-yr 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 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 note 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.

Appendix D provides further insight into how the model prices risk at each horizon

¹⁷As an aside, the conditional risk premium, which is the *expected* return on the dividend futures portfolio over the 2003.Q1-2014.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%.

using the tools developed by Hansen and Scheinkman (2009) and Borovička and Hansen (2014). It shows that the various equity factors have very different risk exposures from each other, and at various horizons.

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 private equity funds.

III.A PE Cash Flow Summary Statistics

Our fund data cover the period January 1981 until June 2019. Our main data source is Preqin, but we find comparable results on the Burgiss data set as shown in Appendix H.¹⁸ We group private equity 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 non-missing 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 assets under management. Buyout is the largest category by AUM (\$1,888 billion), followed by RE

¹⁸Preqin data is substantially sourced by FOIA requests made to public pensions, who may have differential pricing terms in side letters and “Most Favored Nation” clauses. However, Da Rin and Phalippou (2017) suggests that public pensions are not statistically different from other investors in their access to these clauses. Burgiss data is sourced from a more representative set of institutional investors.

(\$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 pd-quartile across of the four quarters in the calendar year. There is clear business cycle variation in when funds get started as well as in their size (AUM). The last cash flow we consider in the estimation is for June 2019, since cash flows in the second half of 2019 have not been fully reported yet. 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.

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 16.¹⁹ The literature typically treats PE vehicles as lasting ten 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.²⁰

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

¹⁹We 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.

²⁰Industry publications have noted the increasing lifespan of PE funds. For instance, a Preqin report from 2016 remarks: “The average lifespan of funds across the whole private capital industry is increasing beyond the typical 10 years... older funds of vintages 2000-2005 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. As a robustness check, we re-estimate our results on a subsample of funds which 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 shows that there is substantial variation in cash-flows across vintages, even within the same investment category. This variation will allow us to identify vintage effects. The figure also highlights that there is a lot of variation in cash-flows across calendar years. VC funds started in the mid- to late-1990s 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 the year 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. Appendix Figure E1 shows cash-flow profiles for the remaining four PE categories.

III.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 by OLS; recall equation (6). The second is an Elastic Net model estimated on the full set of fifteen factors; recall equations (7) and (34). The estimated parameters are the factor exposures across horizon, b_h^k , and how these exposures shift by vintage (pd quartile), captured by the $a_{pd(t)}^k$. Their sum, $a_{pd(t)}^k + b_h^k$, measures the number of units of strip k with maturity h 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.²¹

²¹To minimize over-fitting, we rely on a cross-validation exercise in which we use a leave-out sample to fit the α and λ hyper-parameters. Appendix G discusses the details and provides robustness checks on the benchmark hyper-parameter choices.

Figure 12 shows the estimated age effects \hat{b}_h for the 2-factor model estimated by OLS (left panels) and the 15-factor model estimated by Elastic Net (right panel). The rows correspond to the four main PE categories. Appendix Figure E2 contains these estimates for the other four PE categories. Appendix Figures E3 and E4 contain the estimates for the PD-quartile effects $\hat{a}_{pd(t)}$.

Buyout For the 2-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 3-6. A strategy that sells the aggregate stock market is correlated with the distribution cash flows made by Buyout funds. The later years 5-12 show substantial bond exposure in the replicating portfolio. The large bond exposure hint at out-performance. Buyout managers produce cash-flows in the peak harvesting period that appear to be risk-free, according to the simple 2-factor model.

The 15-factor model, estimated by 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 which 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 examines Venture Capital funds. Our OLS 2-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. It describes VC funds as 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 E3 suggests that this growth gain strip loading is higher for funds started when the pd ratio is in the second and fourth quartiles. Vintages in the early 1990s are such second-quartile pd vintages, which ended up with very high cash flows. Vintages in the fourth pd-quartile, such as the 1997-2004 vintages, also have higher growth gain exposure.

Our findings for VC funds accord with economic intuition. While Buyout funds acquire a range of companies which may differ in their underlying risk exposures, VC funds invest in early-stage and rapidly growing entrepreneurial companies which 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 pay-offs 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 REPE funds as levered positions in stock market gain strips, especially for the first 7 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 to REIT dividend strips and value dividend strips early on, and to small and REIT gain strips in years 5-8. Reassuringly, REIT dividends and REIT gains strips are important components of the replicating portfolio of REPE funds. The small and value exposures accord well with the fact that listed REITS have returns that behave like those of small value stocks (Van Nieuwerburgh, 2019). In the later years, REPE fund cash flows are exposed to the same risk as infrastructure and natural resources dividends 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 contains 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, that 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 E2. 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 provide the first systematic analysis of the risk properties of these 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.

III.C Expected Return

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 different horizon effects and annualizes the resulting expected return as per equation (4). The left panels of this figure are for the 2-factor OLS model; the right panels for the 15-factor Elastic Net model. Appendix Figure E5 reports the results for other four PE categories.

On average over time, the expected returns on PE vehicles has been 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 market prices of risk generate time variation in the expected returns on dividend and gain strips. Combined, they generate 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 the year 2000, when the stock market peaked, and an increase in risk premia during the Great Recession for several of the PE categories. Those dynamics are driven by vintage effects, which switch discretely between pd quartiles sometimes resulting in spikes, and by dynamics on the listed risk premia. For example, the aggregate market risk premium is very low around the year 2000, while small, value, and growth risk premia are elevated in the Great Recession.

III.D Risk-Adjusted Profit

Next, we turn to performance evaluation, the main result in the paper. Figure 14 plots the histogram of Risk-Adjusted Profits (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.

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 2-factor model. For VC and real estate, a substantial part of the right tail of the distribution is removed and shifted to the left side.

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 that under the two-factor model.

The average RAP masks substantial cross-sectional dispersion. The RAP histograms are wide. The Monte Carlo exercise in Appendix A.1 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. That 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 shows results from our main categories, Panel B for the additional categories. The first three rows report standard measures to evaluate PE funds: the TVPI, IRR, and PME.²² The next two rows display 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

²²TVPI and PME subtract out the initial investment to express them as excess performance metrics. The IRR is calculated by us since Preqin only computes IRRs 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 distorts IRR measurement in the PE context.

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.

Figure 15 plots the average RAP by vintage for both the two-factor (left panels) and 15-factor Elastic Net (right panels) models. Appendix Figure E6 plots these estimates for alternate 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 2-factor model. By contrast, the Elastic Net model estimates mildly negative average profits for most vintages in the 2000s.

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 2-factor model but consistently negative RAP according to the full model. RE funds launched before the Great Recession (vintages 2002-2006) have performed reasonably according to the two-factor model, but poorly according to the full model. The only category which 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.

III.E Model Comparison

It is instructive to benchmark our results across other approaches pursued 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. Appendix Figure E7 repeats the 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 83-87% with our IRR measure and between 66-87% 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 that there are funds which conventional measures assess to be high-performing but our estimates suggest only offer fair or even too little compensation given their factor risk exposure, and vice versa.

III.F Performance Persistence

While we found that about one-third of funds deliver meaningfully positive RAP, it may not be the same funds that consistently outperform. Therefore, we 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 the performance at the top and bottom of the fund distribution. We compare the persistence of our two- and 15-factor models to those of the traditional performance metrics TVPI, IRR, and

PME.

Consistent with prior research (Harris, Jenkinson, Kaplan, and Stucke, 2014; 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 there is 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 having the next fund also 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 remaining 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. At least some of the persistence in the traditional metrics can be ascribed to persistent risk factor exposure.

III.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 leaves out of 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 neither penalizes the GP for delays in calling some of the committed capital nor for never calling some of the capital at all.

We now entertain two alternative and equally defensible assumptions on the calls. The first one is that the LP earns a 0% 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 life of the fund. 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 as well as 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.

III.H Discussion

Our pervasive finding of negative average RAP stands in contrast to previous literature in private equity which has generally found evidence of fund outperformance (Brown, Harris, Jenkinson, Kaplan, and Robinson, 2015; Kaplan and Sensoy, 2015; Harris, Jenkinson, and Kaplan, 2014). 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 literatures (Fung, Hsieh, Naik, and Ramadorai, 2008) likewise lowered 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 there is a subset of PE managers that consistently delivers 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 outperforms, 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 different 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 so 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 which diminish 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. Finally, delegated asset managers charge sizable management and performance fees especially when financial products are more opaque and complex ([C  lerier and Vall  e, 2015](#)) as they are in PE. This lowers after-fee returns for investors further. A deeper analysis of before-fee performance is left for future work.

IV Conclusion

We provide a novel valuation method for private equity 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 which 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 not only using bonds and the aggregate stock market, but also with sector-specific equity factor exposures. The estimated exposures are sensible given the nature of underlying nature of PE assets, and indicate an important role for growth stocks in venture capital 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, and for calling some of it with a delay. Under this alternative assumption, all PE fund categories under-perform 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 is 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 out-performance. Exploring more deeply what 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 Private Equity, 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*.
- 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, Working paper, 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, Adi Sunderam, and Luis Viceira, 2017, Inflation Bets or Deflation Hedges? The Changing Risk of Nominal Bonds, *Critical Finance Review* 6, 263–301.

- 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, *The Journal of Political Economy* 104, 298–345.
- Campbell, John Y., and Robert J. Shiller, 1991, Yield Spreads and Interest Rate Movements: A Bird's Eye View, *Review of Economic Studies* 58, 495–514.
- Campbell, John Y., and Tuomo Vuolteenaho, 2004, Good Beta, Bad Beta, *American Economic Review* 94, 1249–1275.
- C  l  rier, Claire Myriam, and Boris Vall  e, 2015, The Motives for Financial Complexity: An Empirical Investigation, *SSRN*.
- Cochrane, John H., 2005, The risk and return of venture capital, *Journal of Financial Economics* 75, 3–52.
- Cochrane, John H., and Monika Piazzesi, 2006, Decomposing the Yield Curve, Working Paper, University of Chicago.
- 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.
- Cox, J., J. Ingersoll, and S. Ross, 1985, An Intertemporal General Equilibrium Model of Asset Pricing, *Econometrica* 53, 363–384.
- 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, *The Journal of Finance* 65, 1915–1947.
- Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.
- Fung, William, David A Hsieh, Narayan Y Naik, and Tarun Ramadorai, 2008, Hedge Funds: Performance, Risk, and Capital formation, *The 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, *Working Paper*.
- 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.
- Haddad, Valentin, Erik Loualiche, and Matthew Plosser, 2017, Buyout activity: The impact of aggregate discount rates, *The Journal of Finance* 72, 371–414.

- Hansen, Lars Peter, and Jose Scheinkman, 2009, Long-Term Risk: An Operator Approach, *Econometrica* 77 (1), 177–234.
- Harris, Robert S., Tim Jenkinson, and Steven N. Kaplan, 2014, Private Equity Performance: What Do We Know?, *Journal of Finance* 69.
- Harris, Robert S, Tim Jenkinson, Steven N Kaplan, and Rüdiger Stucke, 2014, 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, *The Review of Financial Studies* 33, 1879–1890.
- Koijen, Ralph S. J., Hanno Lustig, and Stijn Van Nieuwerburgh, 2017, The Cross-Section and Time-Series of Stock and Bond Returns, *Journal of Monetary Economics* 88, 50–69.
- 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, *National Bureau of Economic Research Working Paper No. 24070*.
- Lettau, Martin, and Jessica Wachter, 2011, The Term Structures of Equity and Interest Rates, *Journal of Financial Economics* 101 (1), 90–113.
- Liu, Naiping, and Lu Zhang, 2008, Is the value spread a useful predictor of returns?, *Journal of Financial Markets*.
- 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(1), 38–94. *Review of Asset Pricing Studies*.
- 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, .
- Robinson, David T., and Berk A. Sensoy, 2011, Cyclical, Performance Measurement, and Cash Flow Liquidity in Private Equity, *Working Paper*.
- 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.

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 1-, 4-, 20-, 40-quarter nominal bond yields.

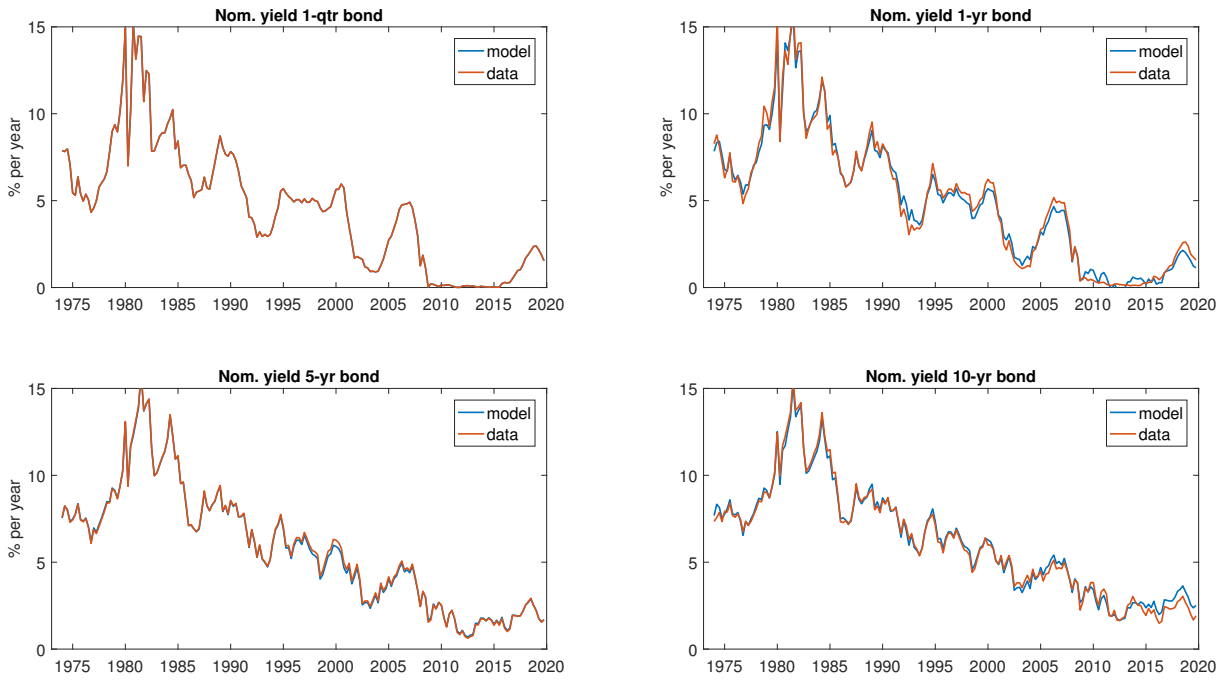


Figure 2. Dynamics of the Real Term Structure of Interest Rates

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

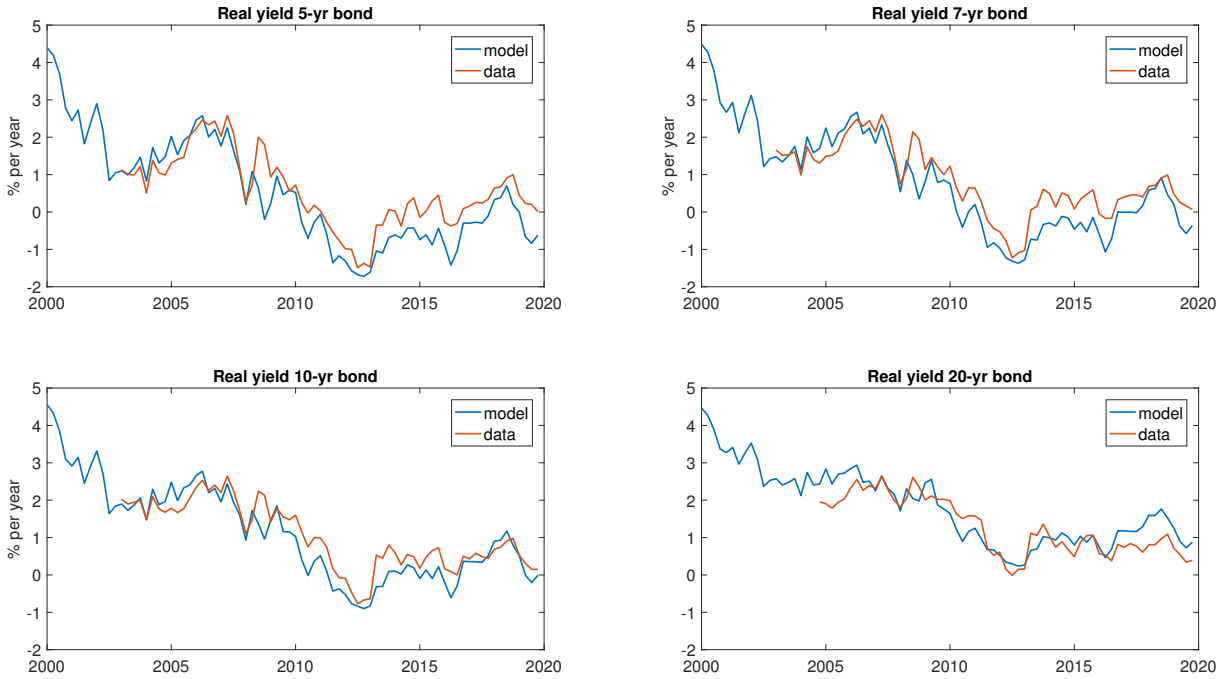


Figure 3. Long-term Yields and Bond Risk Premia

The top panels plot the average bond yield on nominal (left panel) and real (right panel) bonds for maturities ranging from 1 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.

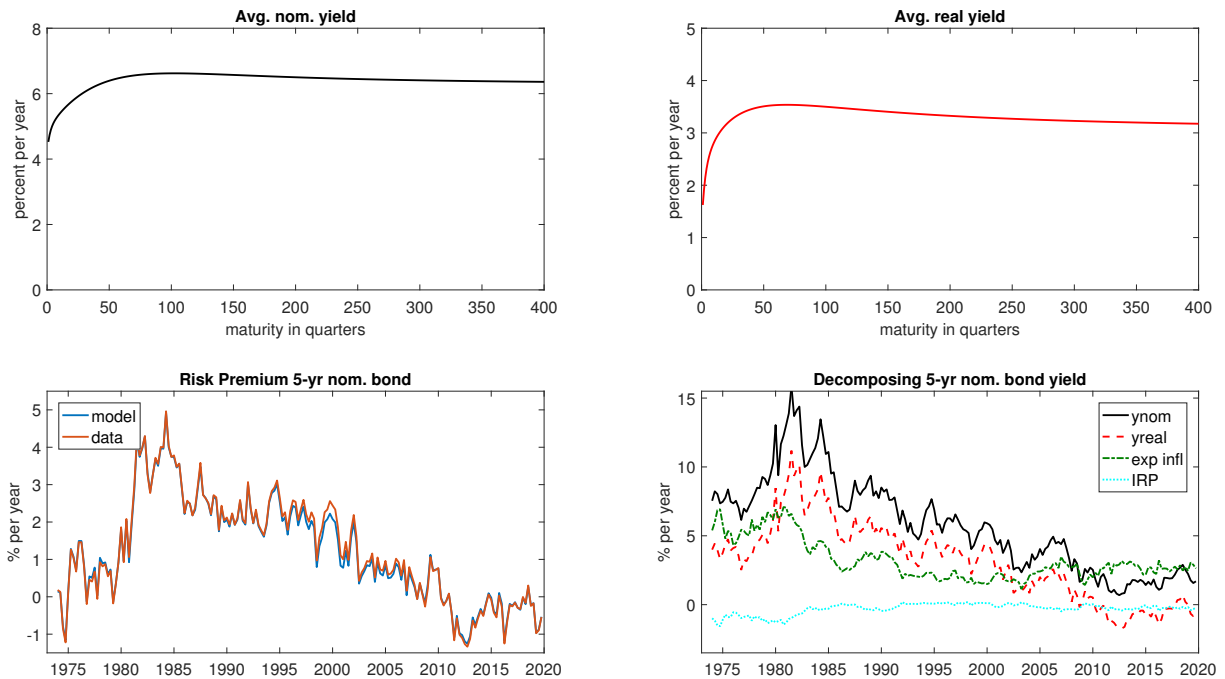


Figure 4. Equity Risk Premia and Price-Dividend Ratios (part 1)

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 ratio in the right panels. The model is the blue line, the data are the red line.

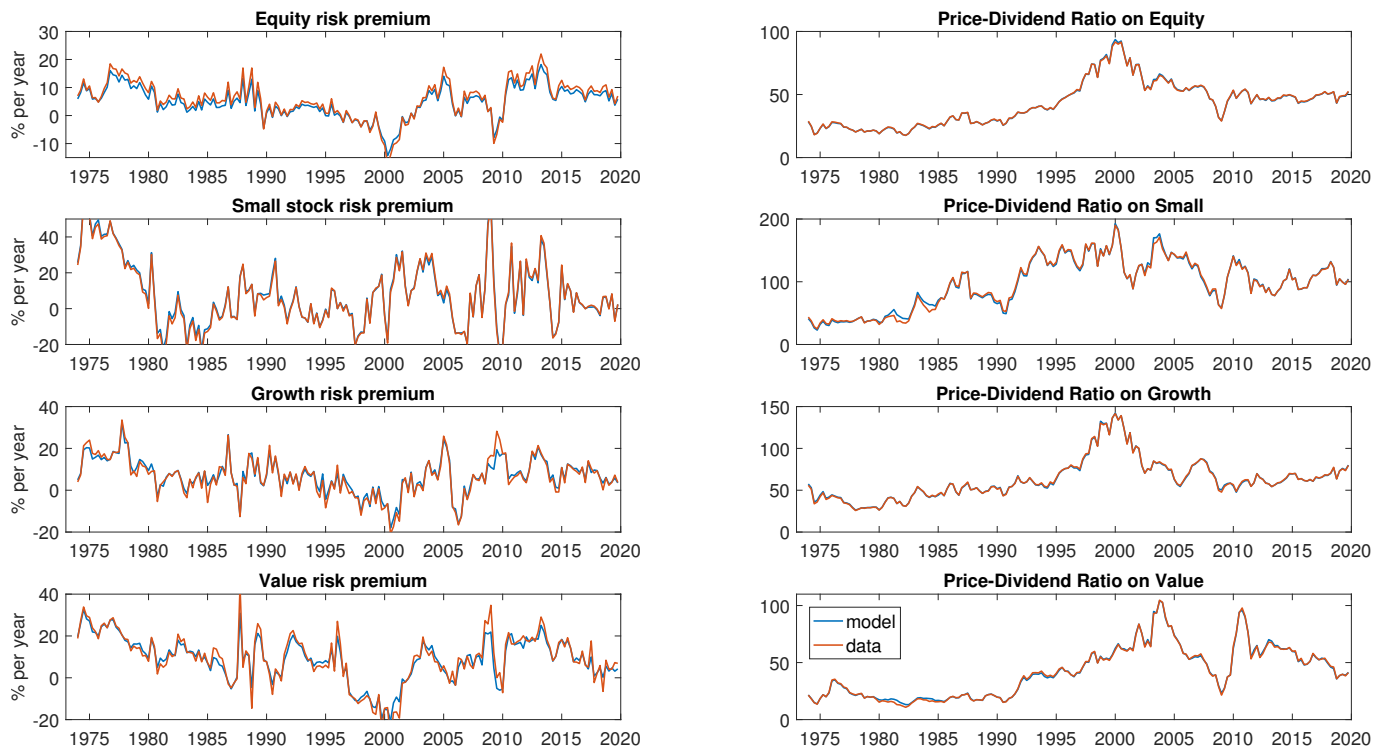


Figure 5. Equity Risk Premia and Price-Dividend Ratios (part 2)

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 ratio in the right panels. The model is the blue line, the data are the red line.

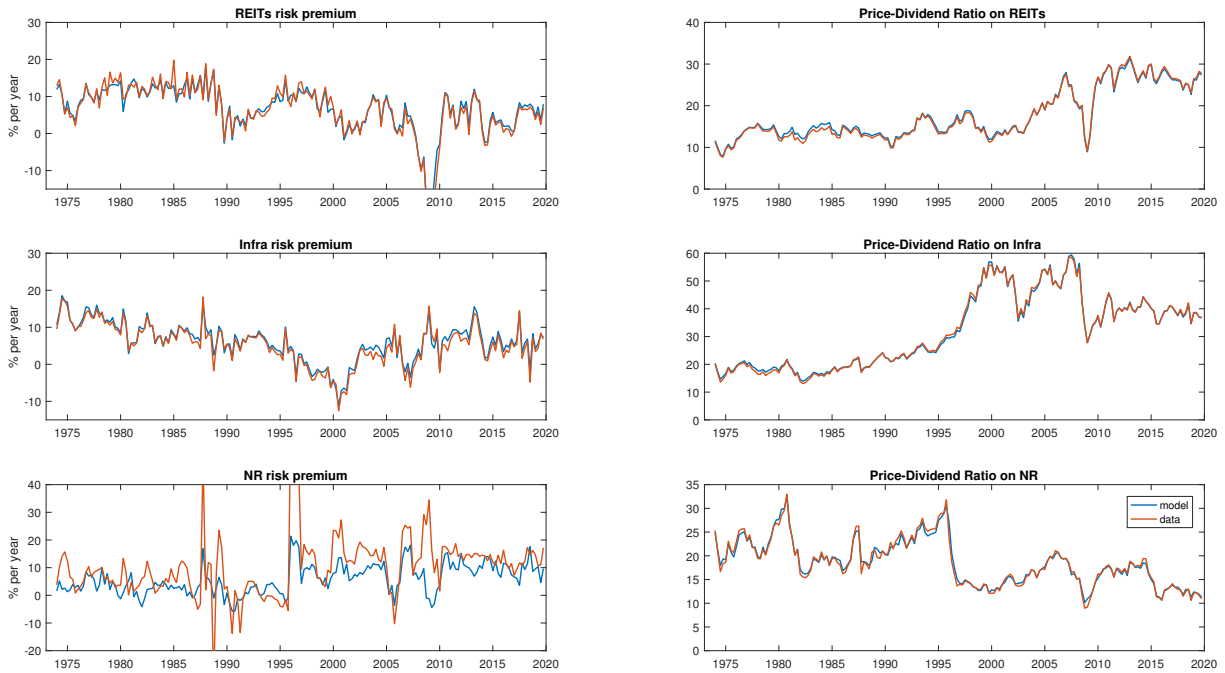


Figure 6. Zero Coupon Bond Prices and Dividend Strip Prices

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 4, 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 Kojen \(2012\)](#) and 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 1 to 64 quarters.

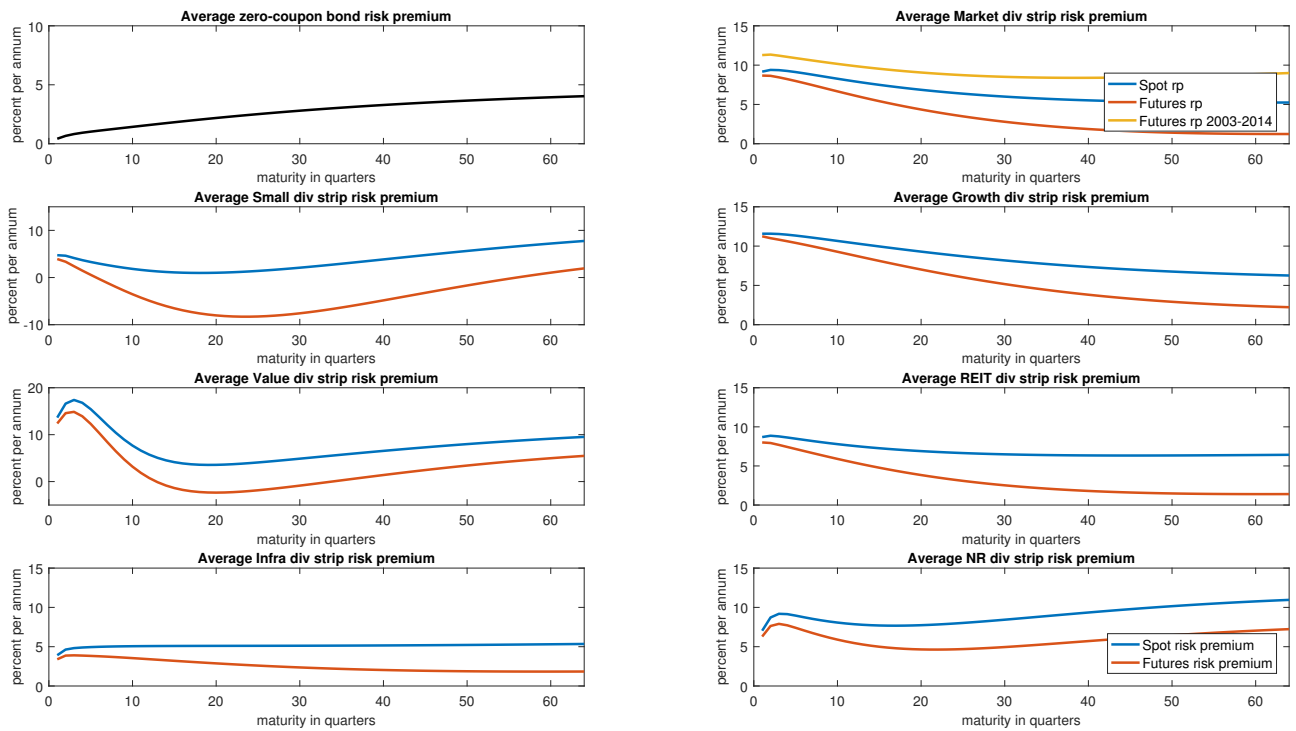


Figure 9. Strip Expected Returns in the Time Series

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 gains 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 5 years (20 quarters).

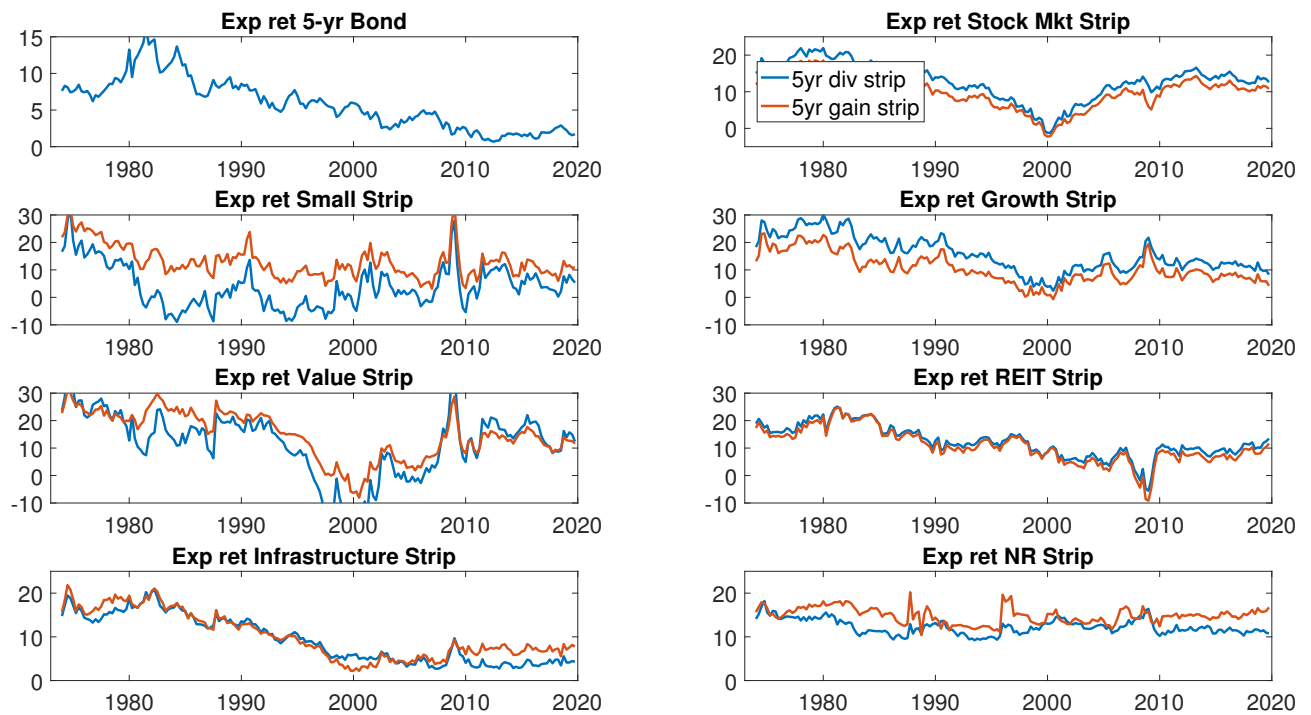


Figure 10. Distribution Cash-flow Profiles

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

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.

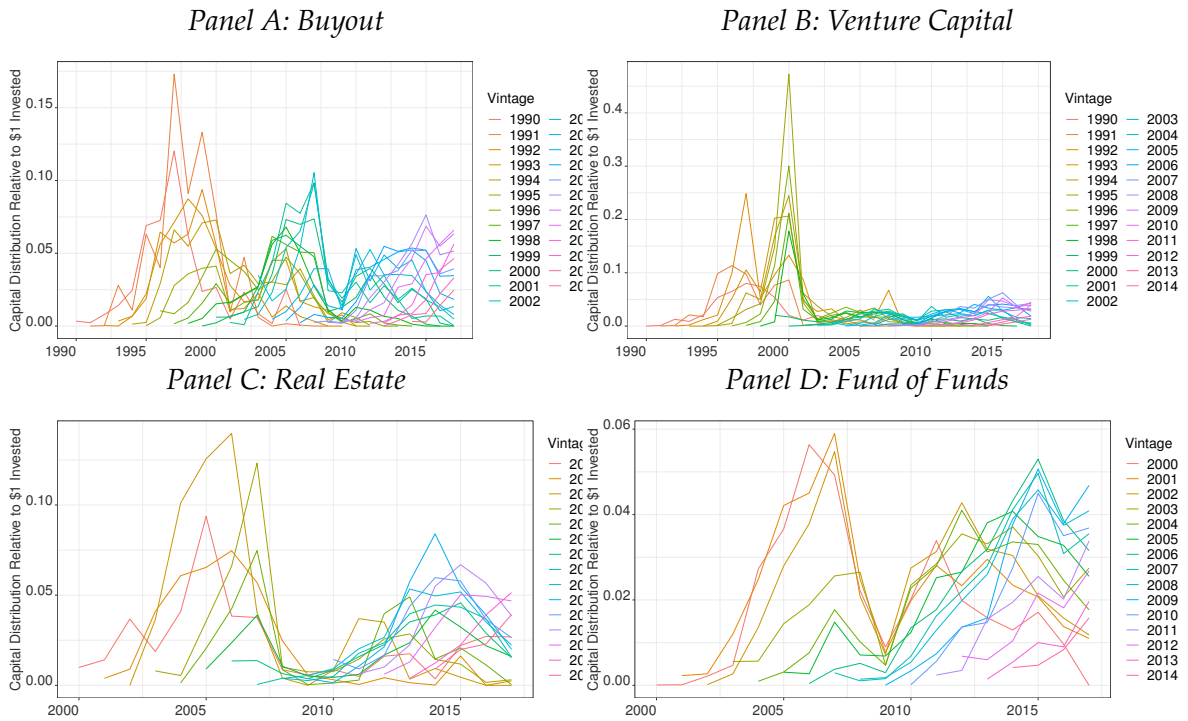


Figure 12. Factor Exposure over Fund Horizon

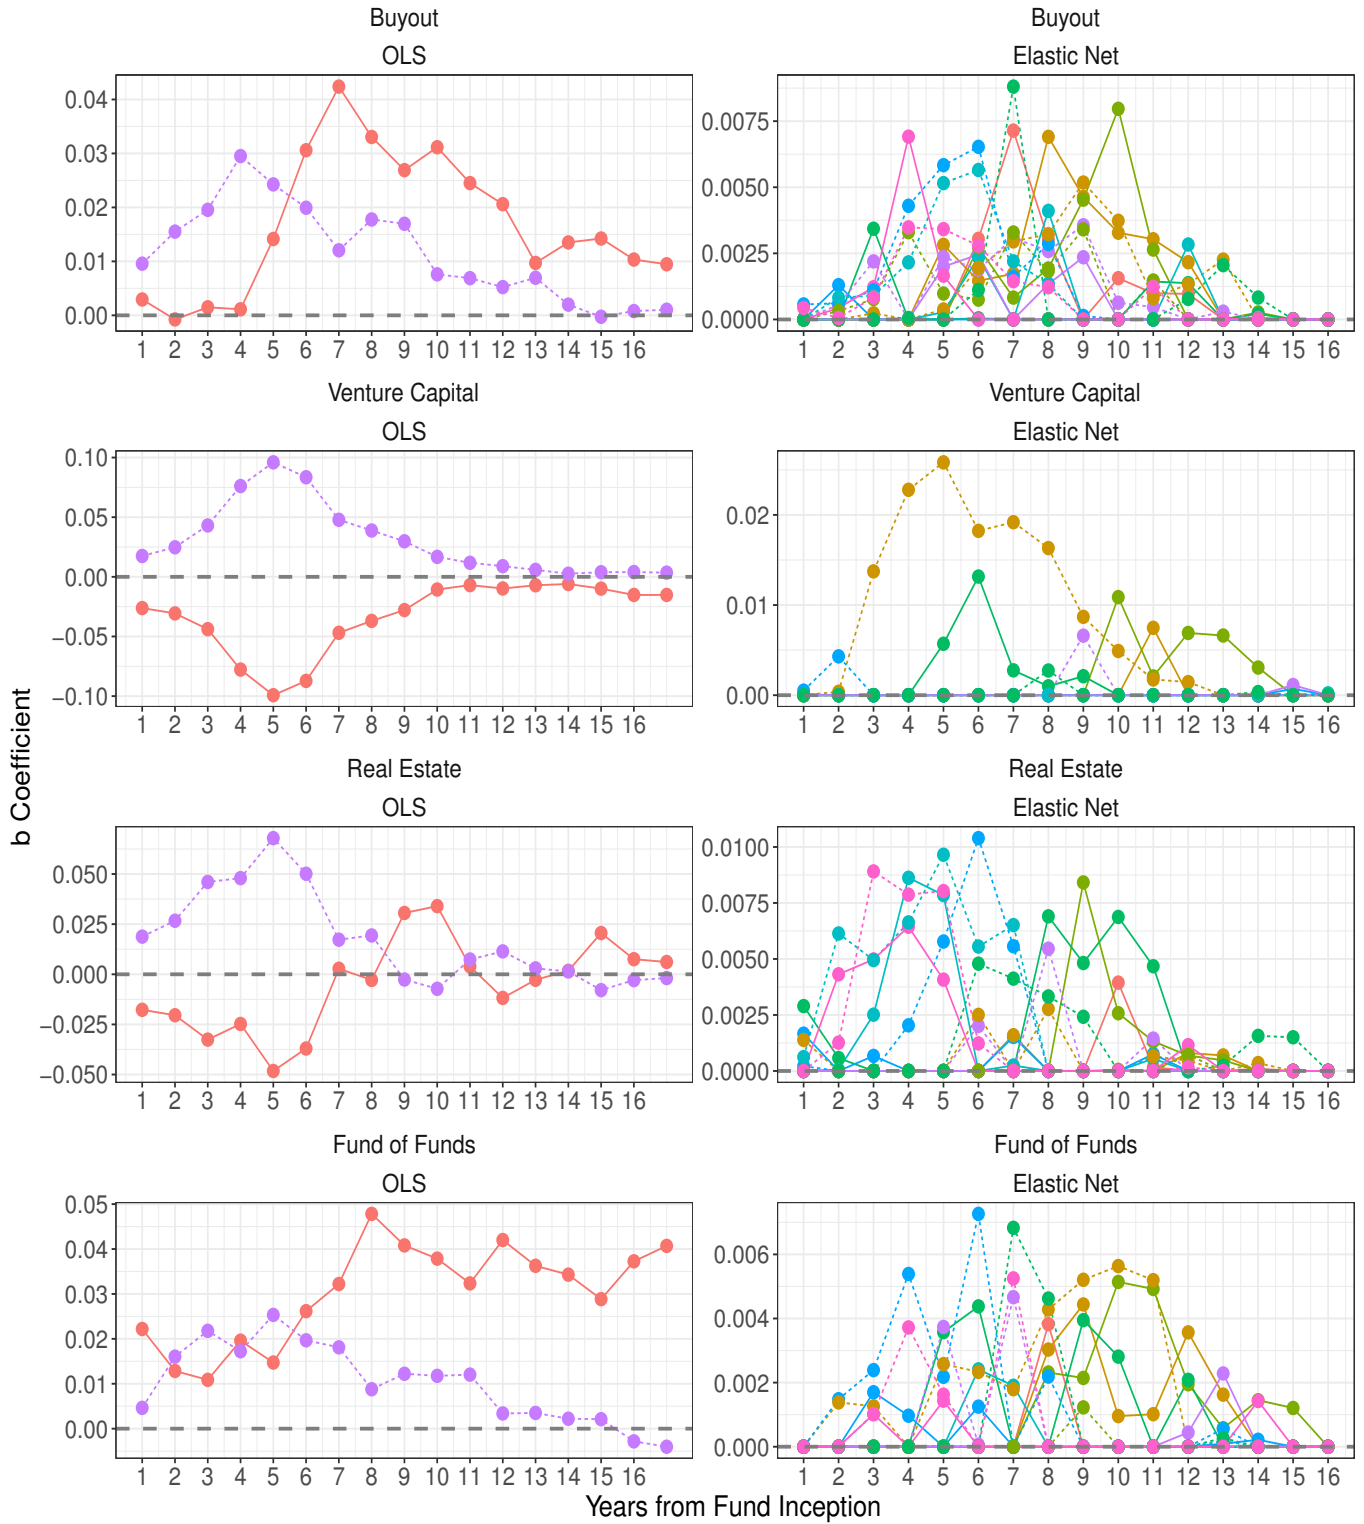




Figure 13. Expected Returns by Vintage



2-Factor

Elastic Net



Panel A: Buyout



Panel B: Venture Capital



Panel C: Real Estate



Panel D: Fund of Funds



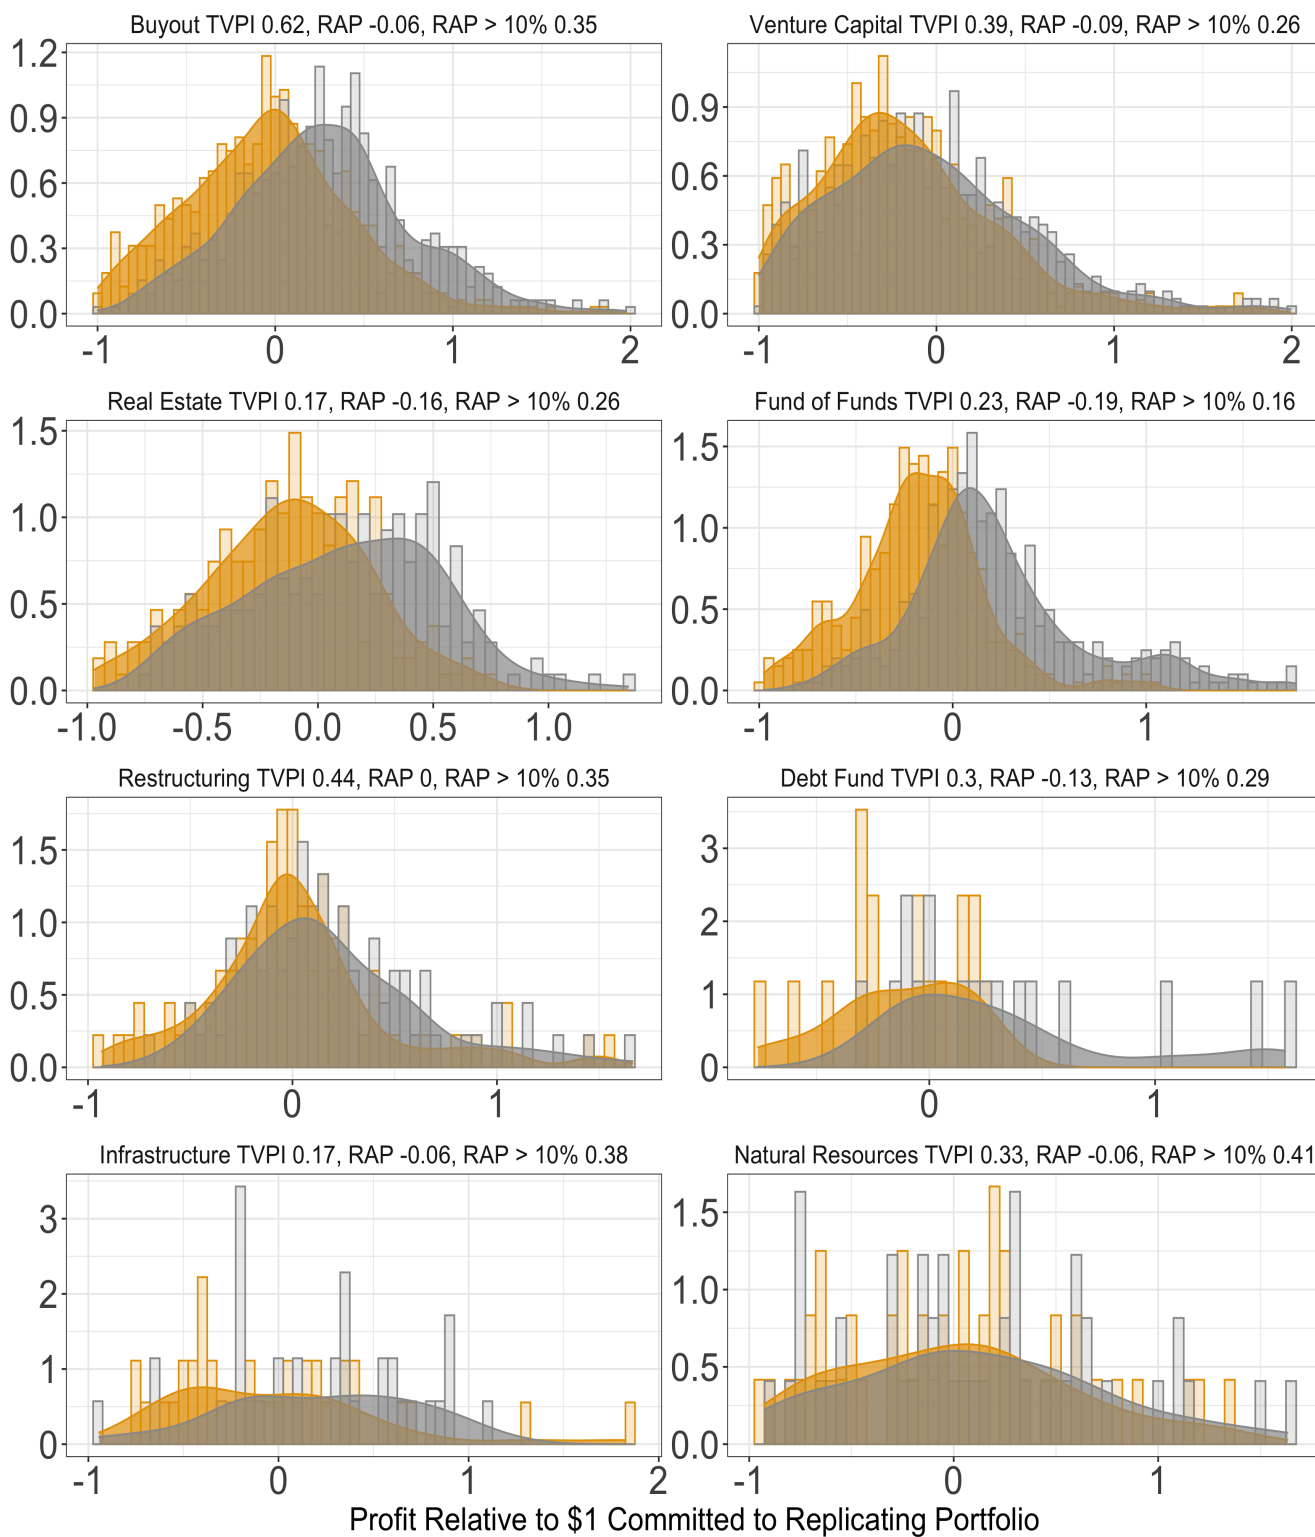


Figure 14. Fund Distribution of Risk-adjusted Profits

Histogram of Fund-Level Profit Relative to Replicating Portfolio



■ Elastic Net Full Factor ■ OLS Two Factors

Figure 15. Profits Over Time

Average Fund-level profit by Vintage

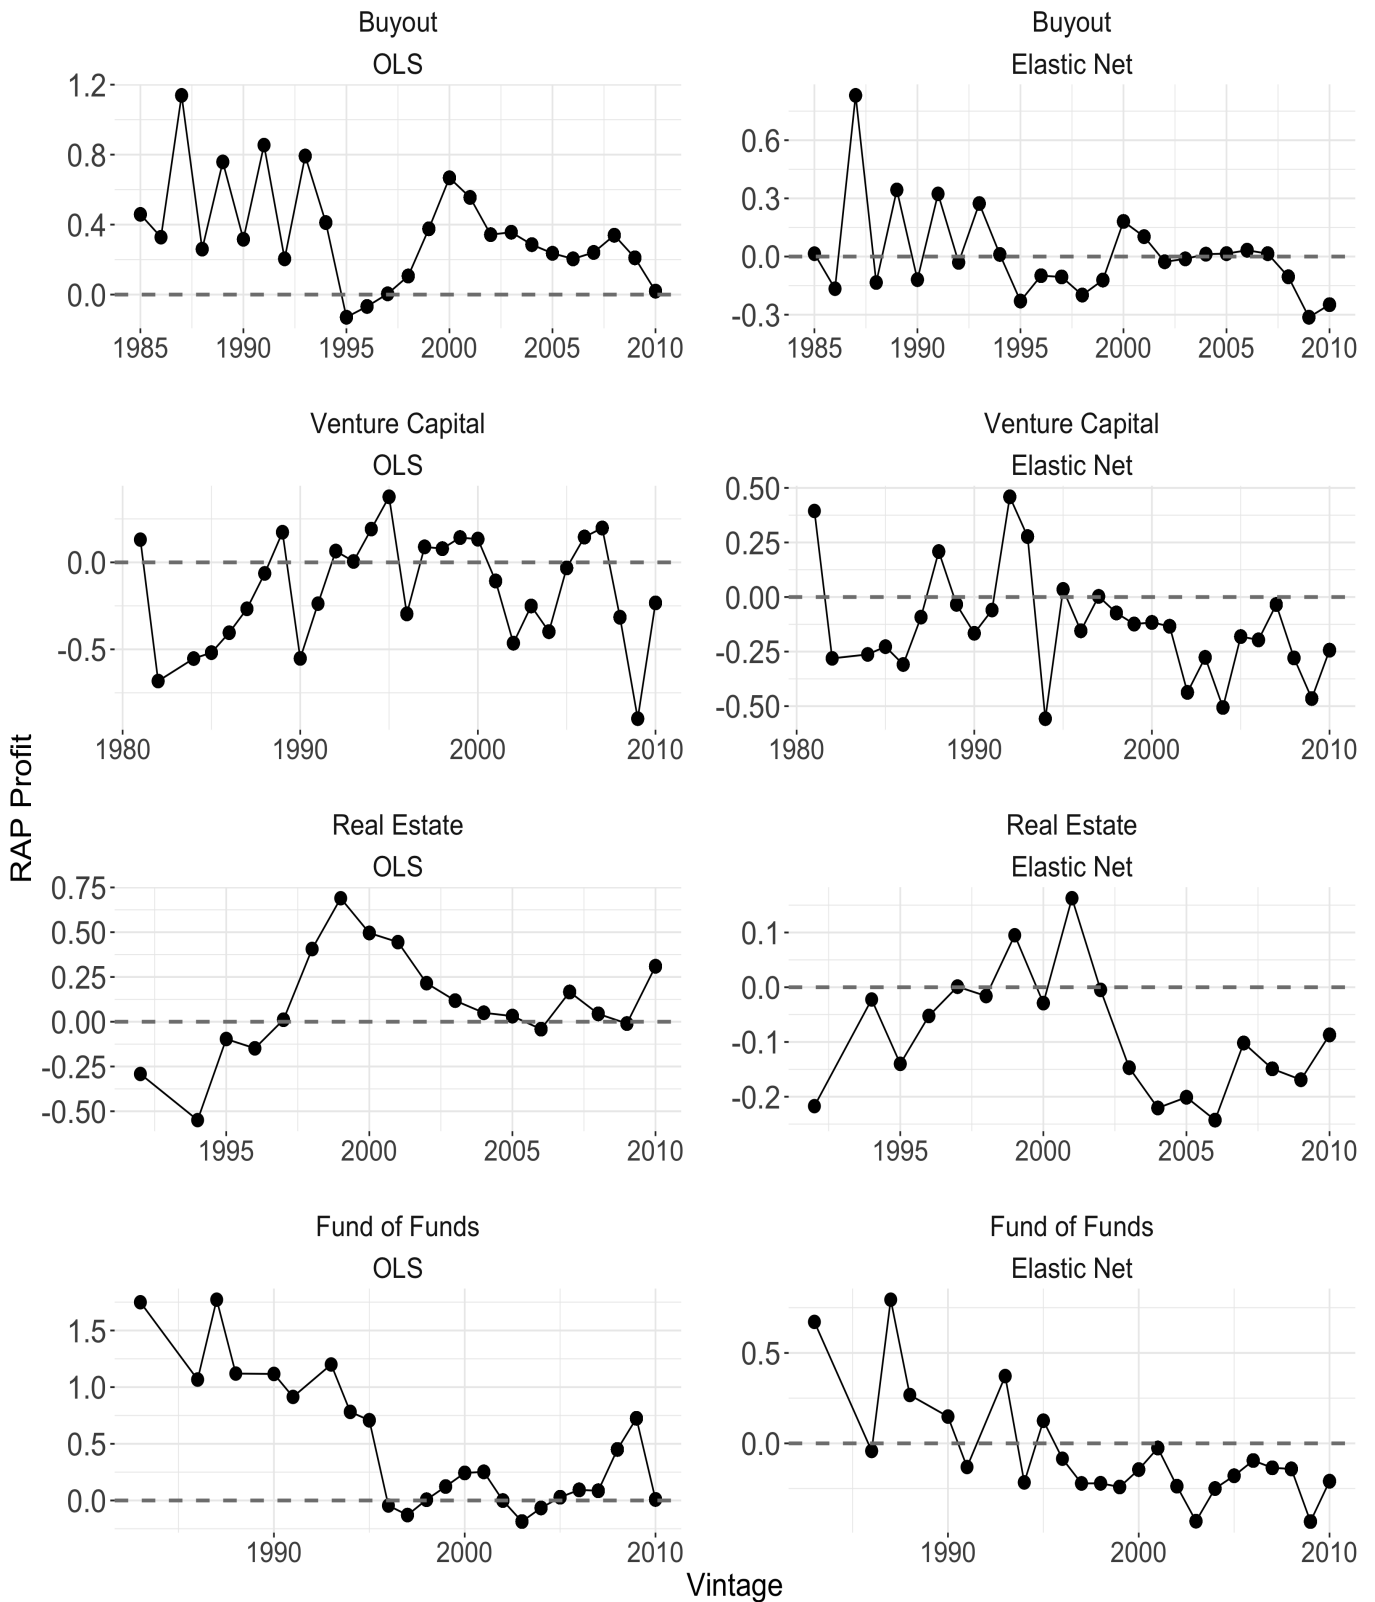




Figure 16. Comparing RAP to IRR and PME


Panel A: Buyout



Panel B: Venture Capital



Panel C: Real Estate



Panel D: Fund of Funds




Table I. Summary Statistics

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	0	1	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	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	0	1	25	2
1993	9	11	0	0	0	2	0	0	0	23	2
1994	16	11	1	1	1	2	0	0	1	33	2.09
1995	15	17	2	0	0	4	0	0	1	39	2.82
1996	21	24	4	1	3	1	0	0	0	56	3.32
1997	23	23	5	0	2	5	0	0	1	60	4
1998	40	33	3	1	1	12	0	0	3	93	4
1999	31	50	2	0	3	9	1	0	1	97	4
2000	35	87	7	0	3	17	1	0	0	150	4
2001	20	52	2	0	5	19	0	0	1	99	4
2002	23	30	2	1	4	13	1	0	2	76	4
2003	18	20	7	1	4	13	1	0	1	65	3.99
2004	28	34	11	4	2	24	1	0	2	106	3.99
2005	56	48	20	0	6	35	2	0	5	172	3.74
2006	76	60	35	5	10	54	0	0	4	244	3.96
2007	74	69	37	6	14	49	1	0	7	257	3.55
2008	65	63	39	4	11	72	5	0	8	268	2.55
2009	31	31	14	6	8	36	2	0	4	132	2.49
2010	43	41	34	9	9	41	3	0	8	189	3.5
2011	55	52	50	9	10	70	2	0	9	260	2.94
2012	64	46	41	8	13	55	4	0	11	245	3
2013	69	50	60	11	18	70	14	0	8	302	3
2014	67	65	55	15	16	77	12	0	15	322	2.99
2015	77	79	83	14	20	75	18	0	8	375	2.78
2016	96	79	59	16	11	92	13	0	16	383	3
2017	53	76	65	14	14	49	29	0	11	312	3
Total:	1145	1210	639	126	193	904	110	0	130	4474	-

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

The universe is all funds in Prequin that have cash-flow data. Not all funds have AUM data, particularly in the early years of the sample.

Table II. Model Comparison

	<i>Panel A: Main Categories</i>							
	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)
	<i>Panel B: Additional Categories</i>							
	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)

The 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.

Table III. Persistence in Performance

	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

The table reports various fund persistence measures for TVPI, IRR, PME-1, the 2-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 that 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 that the second fund in the pair is also in the bottom quartile.

Internet Appendix

A Monte Carlo Validation Study

This appendix conducts a Monte Carlo simulation analysis that studies the internal validity of two estimation approaches to determine the added-value of PE funds: the strip-based approach at the core of this paper and an alternative approach that uses the realized SDF.

A.1 Strip-based Approach

We address three main questions. First, can the strip-based estimation method reliably recover the true underlying risk exposures? Second, does the risk-adjusted profit (RAP) measure properly reflect the absence of out-performance when in truth there is no outperformance? Third, does the RAP measure recover skill (heterogeneity) when there is skill (heterogeneity), and what else does the cross-sectional dispersion in RAP reflect?

To answer these questions, we build a panel of fictitious “private equity” cash flow data of the form:

$$X_{t,t+h}^i = c^i + (a_t^1 + b_h^1) + (a_t^2 + b_h^2) \frac{D_{t+h}^m}{D_t^m} + (a_t^3 + b_h^3) \frac{D_{t+h}^{reit}}{D_t^{reit}} + v_{t+h}^i, \quad (15)$$

where $X_{t,t+h}^i$ is the distribution cash flow of a fund i that was started in period t (vintage) in quarter $t+h$, (a_t^1, a_t^2, a_t^3) are vintage effects associated with each of the risk factor exposures, (b_h^1, b_h^2, b_h^3) are the corresponding horizon effects, c^i is fund-specific abnormal cash flow, and v_{t+h}^i is an idiosyncratic cash flow realization.

For ease of discussion, we refer to c^i as a measure of “skill,” with the understanding that fund managers (GPs) may have skill but capture all the rents in the form of fees rather than passing them through to the investors (Berk and Green, 2004). Since our entire analysis takes the perspective of the investor (LP), c^i captures “net alpha” (after fees) rather than “gross alpha.” c^i is measured in terms of quarterly abnormal cash flow per \$1 initial investment. The assumption in (15) is that the systematic component of fund cash flow depends on exposure to three listed factors: zero-coupon bonds (who have a pay-off of \$1 per unit at maturity), the aggregate stock market dividend normalized by the market dividend at the time of fund formation, and the normalized REIT dividend.

Like in the actual estimation, we assume that funds started at different times may have differential riskiness, and that the vintage effects a_t depend on what quartile of the stock market price-dividend dis-

tribution we are in at the time of fund formation. We make up the shape of the vintage effects a and the horizon effects b . We choose these profiles to be different from one another—to make sure the estimation can handle different shapes—and to be at least somewhat realistic. These profiles are shown in Figure A1 with the label “truth.”

Like in the real data, we simulate one time-series for the $N \times 1$ aggregate state vector z by drawing the structural VAR innovations $\varepsilon_t \sim \mathcal{N}(0, I)$. We recursively form the VAR z_t using the companion matrix Ψ and innovation covariance matrix Σ estimated from the data. This preserves the correlation structure of the aggregate shocks to the one observed in the data, and generates the same persistence for the state variables as in the real data. Since dividend growth on the market and on REITS are both elements of the state vector, the simulation delivers the time series for the dividends $\{D_t^m, D_t^{reit}\}$ that enter in (15). The simulation also provides the price-dividend ratio on the stock market, another element of the VAR. We create pd-quartile indicator variables. Given our estimates for the market prices of risk from the real-data estimation (Λ_0, Λ_1) , taken here as truth, we also form a time-series for the SDF, zero-coupon bond prices $P_{t,h}^\$$, and dividend strip prices, $P_{t,h}$. The cash flow exposures are scaled such that the replicating portfolio costs \$1 on average across periods (vintages). When calculating RAP, we compute the actual cost C_t in each period (for each vintage).

We assume that the idiosyncratic PE cash flow shocks v_{t+h}^i are i.i.d. across funds and time, and normally distributed with mean zero and standard deviation σ_v . We create a balanced panels of N_f funds in each of N_v vintages, for a total of $S = N_f \times N_v$ funds. Each of these funds is observed for H quarters. We set $N_f = 60$ and $N_v = 100$ to match the panel dimension of the Buyout category, where we observe a median of 60 funds and have complete cash flow data for 100 quarterly vintages. Thus the artificial fund panel contains $S = 6,000$ funds. Each of these funds has $H = 64$ quarterly cash flows. We set the idiosyncratic cash flow volatility to $\sigma_v = 0.02$. This level of idiosyncratic risk generates an R-squared for the cash-flow replication regression of 18%, roughly consistent with what we find in the data.

Estimating Strip Exposures We start by assuming that funds do not generate any abnormal cash flows: $c^i = 0, \forall i$.

Having generated this realistic panel of fictitious PE cash flow data, we first attempt to recover the exposure coefficients (a_t^1, a_t^2, a_t^3) and (b_h^1, b_h^2, b_h^3) by estimating equation (15) by OLS. All $S \times H = 384,000$ cash flows are stacked and regressed on pd-quartile indicator variables, age indicator variables, and both pd-quartile and horizon indicators interacted with both market dividends and REIT dividends. Only three of the four vintage effects are identified, so we normalize the one associated with the second pd-quartile to zero. With $K = 3$ factors and $H = 64$ horizons, there are $K \times H + K \times 3 = 201$ coefficients to be estimated.

Figure A1 shows that this estimation “works.” It overlays the estimated horizon effects and vintage effects for each of the three factors on the true (hypothesized) exposures. The model retrieves these cash-flow exposures successfully, for each of the different “shapes” of the hypothesized exposures. This is despite the large amount of idiosyncratic cash flow noise (the v_t^i) and for a realistic assumption on the amount of PE cash-flow data that are available to the researcher.

The one place where the estimation struggles a bit in the early part of the life-cycle. This is because (i) we assume that the true exposures on the market and REIT dividend factors are small at that stage of the life cycle, and (ii) there is less variation in the scaled dividend growth rates $\left(\frac{D_{t+h}^m}{D_t^m}, \frac{D_{t+h}^{reit}}{D_t^{reit}}\right)$ at short horizons $h < 10$. A lack of variation in dividend strip payoffs at short horizons across vintages make it difficult to distinguish a small exposure to the dividend strips from a constant exposure (to bonds).

This answers affirmatively the first question of whether the strip-based method can reliably recover the true underlying risk exposures.

Estimating RAPs in the Absence of Skill Second, we ask what the estimated RAP distribution looks like when in truth PE funds have no skill. By assumption, none of the funds generate any abnormal cash flows for their investors ($c^i = 0$), and so all RAPs should be zero. But because of idiosyncratic cash flows and some estimation error in the estimated exposures, we expect a non-degenerate estimated RAP distribution.

The histogram of the estimated RAP is plotted in Figure A2. There are two important findings. First, the estimated mean of the RAP distribution is essentially unbiased. The mean of the distribution is -0.002, close to the truth of 0.

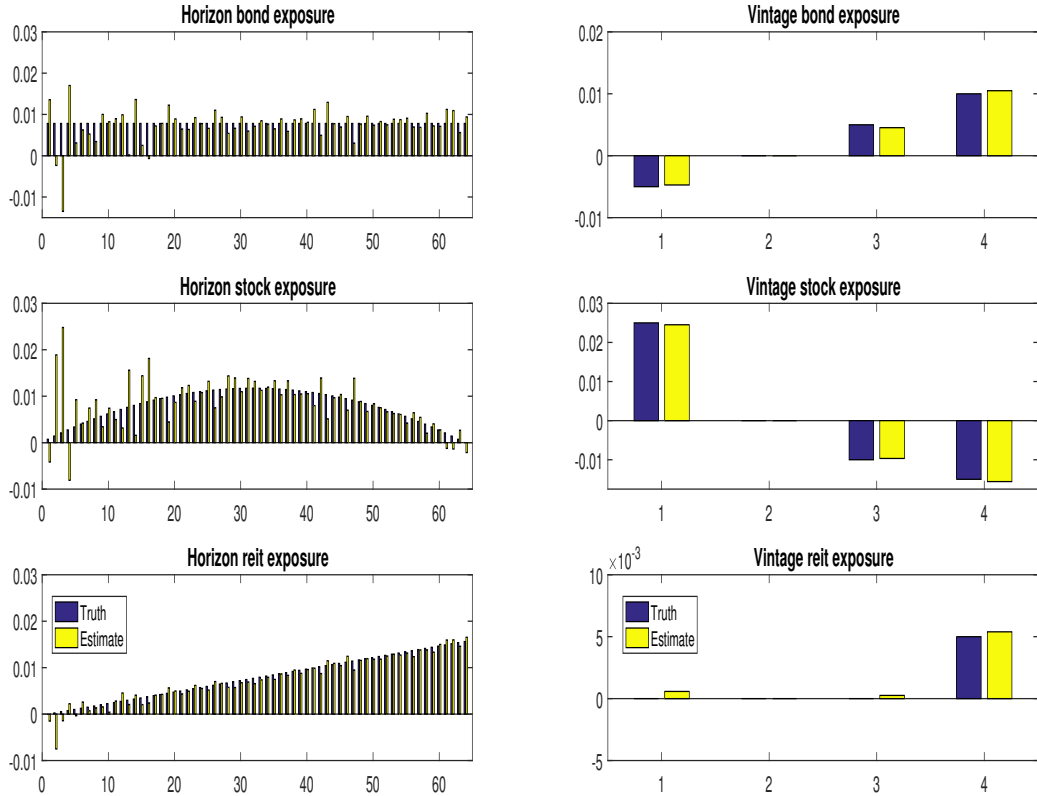
Second, the cross-sectional dispersion of RAP is small: 0.032. Of this dispersion, 95% comes from the residuals, the first term in the RAP equation (5), while 5% comes from estimation error in the exposures, the second term in the RAP equation. The [5th,95th] percentiles of the estimated RAP distribution are [-0.053,0.050]. To summarize, it would be rare to see a fund with an estimated RAP larger than 5% in absolute value (5 cents per \$1 invested), if PE funds had no skill (zero net alpha).

This answers affirmatively the second question of whether the estimated RAP accurately recovers a true absence of out-performance.

Estimating RAP in the Presence of Skill Third, we ask the question: If (some) funds had non-zero net alpha, would the estimation be able to recover it? To do so, we set $c^i \neq 0$. We consider three cases: homogenous positive-mean skill across managers, heterogeneous mean-zero skill, and heterogeneous positive-mean skill.

Figure A1. Monte Carlo Exercise: True and Estimated Exposures

The figure plots the hypothesized true exposures of a panel of simulated “PE funds” as well as the estimated exposures. The true model is a 3-factor model with zero-coupon bonds, stock market dividend strips, and REIT dividend strips. There are 60 funds in each of 100 quarterly vintages, each with 64 quarters of cash flow distributions.



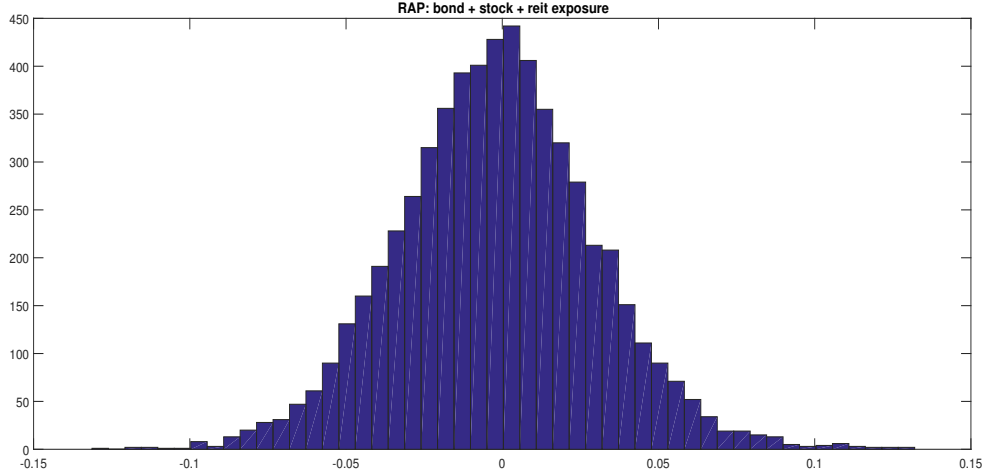
We find that even a small amount of homogenous out-performance of $c^i = c = 0.001, \forall i$ can be recovered precisely. The true mean RAP for this value of c^i is 0.0522 with a standard deviation of 0.008. These numbers are calculated by zeroing out the idiosyncratic risk $\sigma_v = 0$. There is still some dispersion in RAP in the absence of idiosyncratic risk because there is aggregate risk; funds in different vintages span different 64-quarter stretches of the time series.

At the baseline value of $\sigma_v = 0.02$, the estimated mean RAP is 0.0503, which is close to the true 0.0522. In other words, the estimated mean RAP is nearly unbiased. The dispersion in RAP is 0.032, the same value as in the case of $c = 0$. Again, this dispersion mostly reflects idiosyncratic cash-flow risk.

The way our method recovers the skill is in the form of estimated bond exposures \hat{b}_h^1 that exceed the true exposure b_h^1 by the value of $c = 0.001$. The intercept in equation (15) picks up both bond exposure and skill

Figure A2. Monte Carlo Exercise: RAP Distribution

The figure plots the estimated RAP distribution in a panel of simulated “PE funds.” The true RAP is zero for each fund. The true model is a 3-factor model with zero-coupon bonds, stock market dividend strips, and REIT dividend strips. There are 60 funds in each of 100 quarterly vintages, each with 64 quarters of cash flow distributions.



level: $\hat{b}_h^1 = b_h^1 + c$. This is intuitive. A skilled PE fund is equivalent to a passive replicating portfolio with a larger position in risk-free bonds. Since bond prices are high, the replicating portfolio is extra valuable. The second term of the RAP expression (5) picks up the value of that extra bond position.

The second case features a non-trivial distribution of skill (net alpha). We assume that $c^i \sim \mathcal{N}(0, 0.002)$, so that there is a lot of cross-fund dispersion in out-performance around a zero mean (no average skill). In the absence of idiosyncratic risk, we obtain a nearly unbiased mean RAP estimate of -0.0069 and a RAP dispersion of 0.101. With idiosyncratic risk, we obtain a similar mean RAP estimate of -0.0087 and a RAP dispersion of 0.106. Hence, we continue to estimate the mean skill accurately, and find that the bulk of the dispersion in estimated RAP reflects the true skill heterogeneity.

Finally, we consider a case where $c^i \sim \mathcal{N}(0.001, 0.002)$. We find that the estimated RAP distribution has a mean of 0.0434 and standard deviation of 0.106. Turning off idiosyncratic risk, we get a true mean RAP of 0.0453 with a standard deviation of 0.101. Hence, estimated mean RAP is again nearly unbiased, and true skill dispersion contributes nearly all of the estimated dispersion in RAP.

In conclusion, our method reliably detects skill (heterogeneity) when there is skill (heterogeneity) and detects no skill when there is none. This affirmatively answers the third question.

Robustness and Extensions We have performed several additional exercises, omitted here for brevity. First, the same results obtain if we replace the market and REIT exposures by other factors. Second, results

are robust to adding more factors to the stock market and REIT exposures. Third, we find the same results when dividend strips and capital gain strips are used in conjunction.

Fourth, ignoring the vintage effects in estimation when they are truly present results in severely distorted estimates of the age effects. This suggests that vintage effects are identified, and that ignoring vintage effects could be problematic.

Fifth, we have verified that if the true exposure to an additional risk factor is zero, but that risk factor is (erroneously) included in the estimation, the estimated exposures to that factor are close to zero. This is the case even when that spurious factor is correlated with the true factors that drive fund cash flows, using similar correlations to the ones we observe among factors in the data.

Sixth, we have used this framework to understand how RAP is biased when a true risk factor is erroneously omitted. The true model is the three factor model discussed above, but the researcher erroneously estimates a two-factor model that includes only bond strips and aggregate stock market strips. This will lead the researcher to over-estimate the number of bonds and the number of market dividend strips in the replicating portfolio, and to underestimate the number of reit dividend strips in the replicating portfolio since the REIT strips have mistakenly been left out. The over-estimation of bond and market dividend strip positions makes the replicating portfolio too expensive. However, the under-estimation of the reit dividend strip exposures makes the portfolio too cheap. While bond prices are higher than REIT dividend strip stripes, it is not clear ex-ante which way RAP will be biased. It depends on the magnitude of the loadings of the PE funds on the omitted risk factor relative to the loading on the included factors, which in turn depends on the covariance structure of the included and omitted risk factors, as well as on the relative strip prices. In this case, we find that the replicating portfolio is too cheap and the bias in the RAP is large and negative. The mean RAP estimate is -0.24, the median is -0.16, and 86% of simulated funds have a negative RAP estimate, when in truth all funds have a RAP of 0.

A.2 Using Realized SDF

An alternative method to computing the compute the risk-adjusted value-added of a PE fund is to discount the PE cash flows by the realized SDF. We refer to this metric as $VAsdf$ in equation (8).

Benchmark Asset Pricing Model In a first Monte Carlo exercise we form the realized SDF $M_{t+1} = \exp \{m_{t+1}\}$, where m_{t+1} is given by (10), repeated below:

$$M_{t+1} = \exp \left\{ -r_t^f - \frac{1}{2} \Lambda_t' \Lambda_t - \Lambda_t \varepsilon_{t+1} \right\}.$$

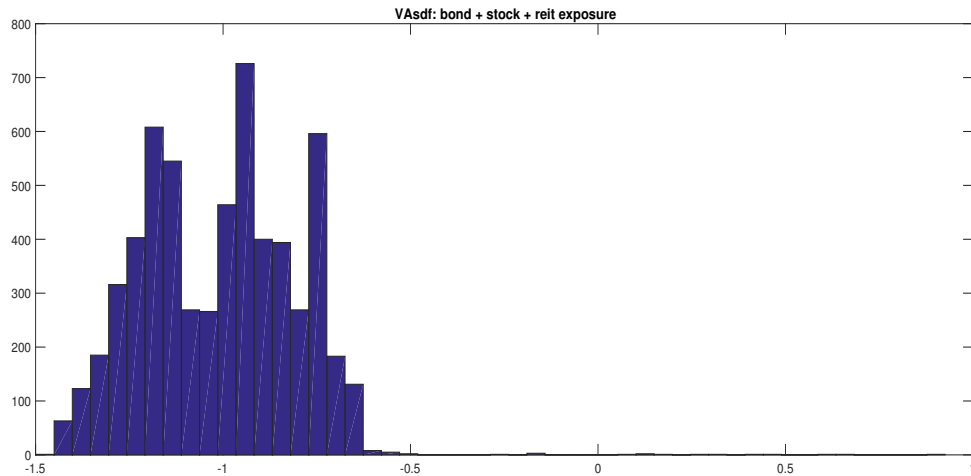
We use the same exact draws for the structural VAR innovations $\varepsilon_t \sim \mathcal{N}(0, I)$ and hence the same state vector sample path $\{z_t\}$ from the previous exercise. We form the market prices of risk Λ_t as the estimated market prices of risk coefficients (Λ_0, Λ_1) from the benchmark asset pricing model estimation and the state vector z_t : $\Lambda_t = \Lambda_0 + \Lambda_1' z_t$.

We using the exact same $S \times H$ panel of artificial PE fund cash flow realizations, assumed to be driven by the same 3-factor model of listed cash flows under $c^i = 0, \forall i$. We form each fund's value-added using the realized SDF $VAsdf_t^i = \sum_h M_{t,t+h} X_{t+h}^i - C_t$. We can compute the true expected risk-adjusted discounted value of the cash-flow stream C_t given that we know the true (assumed) exposures of each fund.

Figure A3 plots the histogram of $VAsdf_t^i$ across all funds and vintages. It is directly comparable to the RAP distribution plotted in Figure A2. While the true mean of $VAsdf$ is zero, the estimated average across funds is -1.01, and the estimated median is -1.00. This estimate implies that the median fund loses the entire initial investment of \$1 after accounting for risk, a -100% return on investment. Second, the $VAsdf$ has a large standard deviation of 0.25, eight times larger than that of RAP. The [5th,95th] percentiles of its distribution are [-1.33,-0.71]. Nearly the entire distribution of $VAsdf$ lies to the left of \$0. This illustrates the problem with using the realized SDF to evaluate fund performance.

Figure A3. Monte Carlo Exercise: VAsdf Distribution

The figure plots the distribution of the VAsdf metric across funds, where $VAsdf_t^i = \sum_h M_{t,t+h} X_{t+h}^i - C_t$, where C_t is the true expected risk-adjusted discounted value of the cash-flow stream. The true model is a 3-factor model with zero-coupon bonds, stock market dividend strips, and REIT dividend strips. There are 60 funds in each of 100 quarterly vintages, each with 64 quarters of cash flow distributions.



What is causing this problem? [Martin \(2012\)](#) shows that this is a generic issue in the valuation of long-lived assets. Define a risk-adjusted cumulative payoff $Y_t \equiv M_1 X_1 - M_2 X_2 + \dots - M_t X_t$ then $\mathbb{E}Y_t = 1$ for all t and Y_t is a non-negative martingale. This risk-adjusted cumulative payoff Y_t tends to zero almost

surely. But then, why is its expected value equal to 1? Because this asset value is driven by a small number of extreme sample paths along which Y_t explodes. $\mathbb{E} \max Y_t = \infty$ and for any $\nu > 0$, $\mathbb{E}[Y_t^{1+\nu}] \rightarrow \infty$ as $t \rightarrow \infty$. Martin calls these exploding risk-adjusted returns aggregate disasters. He shows that they are rare, but not very rare in the sense that for any $\nu > 0$, we can find an arbitrarily large N such that $\text{Prob}(\max Y_t \geq Y_t) > 1/N^{1+\nu}$.

The *VAsdf* contains the cumulative SDF over h periods $M_{t,t+h} = M_t M_{t+1} \cdots M_{t+h}$, where each term $M_{t+j} = e^{-r_{t+j-1}^f Z_{t+j}}$ and $Z_{t+j} = \exp \left\{ -\frac{1}{2} \Lambda'_{t+j-1} \Lambda_{t+j-1} - \Lambda_{t+j-1} \varepsilon_{t+j} \right\}$. This Z_{t+j} is a random variable with expectation $\mathbb{E} Z_{t+j} = 1$. The same is true for $Y_{t+h} = Z_{t+1} Z_{t+2} \cdots Z_{t+h}$: $\mathbb{E} Y_{t+h} = 1$. As Martin (2012) argued, $h \rightarrow \infty$, Y_{t+h} tends to zero almost surely. We have a finite h , but when the market prices of risk $|\Lambda_t|$ are high enough, the problem kicks in for the cash flow horizons that are relevant for PE analysis ($h = 20 - 60$ quarters). Hence, in the *VAsdf* calculations, we are multiplying positive PE cash flows X_{t+h} by SDF realizations $M_{t,t+h}$ that are usually close to zero. The *VAsdf* ends up close to $-C_t = -1$.

One could conjecture that there may be an issue with the specific sequence of SDF shocks ε_{t+j} that happened to materialize in the historical sample or that happened to be drawn in this Monte Carlo simulation. Extensive simulation work that draws many different paths for these shocks delivers similar results. In other words, the specific aggregate shocks that underly the results above are typical.

The main idea behind our RAP approach is to not use the SDF realizations, but use strip prices instead, which are *conditional expectations* of the SDF multiplied by a risky cash flow. These strip prices are well behaved, even in rich models of risk and return.


Simpler Asset Pricing Model Using the realized SDF generates similar issues even in much simpler SDF models. To illustrate this, we consider a special case of our model, namely the conditional CAPM. First, the state space only contains the first six elements of z_t . The companion matrix Ψ recursively zeros out all insignificant elements, as we do for the full model. Second, the SDF only features one priced source of risk, shocks to aggregate dividend growth, which is the sixth element of the VAR. That is, $L_0 = 0$ except for $L_0(6)$ and $L_1 = 0$ except for $L_1(6, 1)$, $L_1(6, 3)$, $L_1(6, 5)$, and $L_1(6, 6)$. The five risk price parameters are chosen to exactly match the unconditional equity risk premium and its dynamics as implied by the VAR to the corresponding objects in the SDF model. This equalization boils down to a linear system of five equations in five unknowns. With the VAR and the market prices of risk in hand, we draw a sequence of aggregate shocks ε_{t+j} of the same length T as the data.

Next, we generate an artificial panel data set of PE fund cash flows, exactly as described above ($\sigma_v = 0.02$, $c^i = 0$). The true model is a 3-factor model where funds are exposed to bonds, market dividend strips, and market gain strips. As before, our approach recovers the true exposures reliably. The mean RAP

estimate is -0.001 , which is close to the true mean of zero, with a small cross-sectional dispersion of 0.023 . In sharp contrast, the estimated average $VAsdf$ across funds is -0.83 , substantially the true mean of 0 , with a large dispersion of 0.83 . Figure A4 plots the histograms of $VAsdf$ and RAP side by side; they are obtained for the exact same fund cash flow realizations. Fully 98% of $VAsdf$ realizations are below 0 .

Figure A4. Monte Carlo Exercise: $VAsdf$ and RAP Distribution Under Conditional CAPM

The figure plots the distribution of the $VAsdf$ metric (left panel) and RAP (right panel) across funds, where $VAsdf_t^i = \sum_h M_{t,t+h} X_{t+h}^i - C_t$, where C_t is the true expected risk-adjusted discounted value of the cash-flow stream. The true model is a 3-factor model with zero-coupon bonds, stock market dividend strips, and stock market gain strips. There are 60 funds in each of 100 quarterly vintages, each with 64 quarters of cash flow distributions.



We have explored several intermediate models between the simple conditional CAPM and the full SDF model, all of which share the same issue with the $VAsdf$ statistic. This result should not come as a surprise. Result 4 in [Martin \(2012\)](#) shows that the issue $Y_\infty = 0$ arises even in the static CAPM as long as the maximum Sharpe ratio exceeds the volatility of the market return, a condition that is satisfied in the data. Even though we do not have $h = \infty$, the problem is already severe when we consider only values of h up to 64. The larger the market prices of risk, the sooner (the lower the first h for which) the problem kicks in. Even when simulating the static CAPM with (return innovations that have a) quarterly Sharpe ratio of $\Lambda_0 = 0.17$ (0.34 annually), the distribution of cumulative SDF realizations $M_{t,t+h}$ is left-skewed, dispersed, and with a large fraction of realizations close to zero. The resulting distribution of $VAsdf$ for the typical fund will have substantial mass close to $-C_t$.

The dispersed and left-skewed distribution of $VAsdf$ implies that it is highly likely that the actual sample an econometrician gets to work with could have $VAsdf$, or averages of $VAsdf$ across funds and vintages, that are very different from $E[VAsdf]$.

B Asset Pricing Model Results

B.1 Risk-free rate

The real short yield $y_{t,1}$, or risk-free rate, satisfies $E_t[\exp\{m_{t+1} + y_{t,1}\}] = 1$. Solving out this Euler equation, we get:

$$\begin{aligned} y_{t,1} &= y_{t,1}^{\$} - \mathbb{E}_t[\pi_{t+1}] - \frac{1}{2}e'_{\pi}\Sigma e_{\pi} + e'_{\pi}\Sigma^{\frac{1}{2}}\Lambda_t \\ &= y_0(1) + \left(e'_{y_n} - e'_{\pi}\Psi + e'_{\pi}\Sigma^{\frac{1}{2}}\Lambda_1\right)z_t. \end{aligned} \quad (16)$$

$$y_0(1) \equiv y_{0,1}^{\$} - \pi_0 - \frac{1}{2}e'_{\pi}\Sigma e_{\pi} + e'_{\pi}\Sigma^{\frac{1}{2}}\Lambda_0. \quad (17)$$

where we used the expression for the real SDF

$$\begin{aligned} m_{t+1} &= m_{t+1}^{\$} + \pi_{t+1} \\ &= -y_{t,1}^{\$} - \frac{1}{2}\Lambda'_t\Lambda_t - \Lambda'_t\varepsilon_{t+1} + \pi_0 + e'_{\pi}\Psi z_t + e'_{\pi}\Sigma^{\frac{1}{2}}\varepsilon_{t+1} \\ &= -y_{t,1} - \frac{1}{2}e'_{\pi}\Sigma e_{\pi} + e'_{\pi}\Sigma^{\frac{1}{2}}\Lambda_t - \frac{1}{2}\Lambda'_t\Lambda_t - \left(\Lambda'_t - e'_{\pi}\Sigma^{\frac{1}{2}}\right)\varepsilon_{t+1} \end{aligned}$$

The real short yield is the nominal short yield minus expected inflation minus a Jensen adjustment minus the inflation risk premium.

B.2 Nominal and Real Term Structure

Proposition 1. Nominal bond yields are affine in the state vector:

$$y_t^{\$}(\tau) = -\frac{A_{\tau}^{\$}}{\tau} - \frac{B_{\tau}^{\$'}}{\tau}z_t,$$

where the coefficients $A_{\tau}^{\$}$ and $B_{\tau}^{\$}$ satisfy the following recursions:

$$A_{\tau+1}^{\$} = -y_{0,1}^{\$} + A_{\tau}^{\$} + \frac{1}{2}\left(B_{\tau}^{\$}\right)'\Sigma\left(B_{\tau}^{\$}\right) - \left(B_{\tau}^{\$}\right)'\Sigma^{\frac{1}{2}}\Lambda_0, \quad (18)$$

$$\left(B_{\tau+1}^{\$}\right)' = \left(B_{\tau}^{\$}\right)'\Psi - e'_{y_n} - \left(B_{\tau}^{\$}\right)'\Sigma^{\frac{1}{2}}\Lambda_1, \quad (19)$$

initialized at $A_0^{\$} = 0$ and $B_0^{\$} = 0$.

Proof. We conjecture that the $t + 1$ -price of a τ -period bond is exponentially affine in the state:

$$\log(P_{t+1,\tau}^{\$}) = A_{\tau}^{\$} + (B_{\tau}^{\$})' z_{t+1}$$

and solve for the coefficients $A_{\tau+1}^{\$}$ and $B_{\tau+1}^{\$}$ in the process of verifying this conjecture using the Euler equation:

$$\begin{aligned} P_{t,\tau+1}^{\$} &= \mathbb{E}_t[\exp\{m_{t+1}^{\$} + \log(P_{t+1,\tau}^{\$})\}] \\ &= \mathbb{E}_t[\exp\{-y_{t,1}^{\$} - \frac{1}{2}\Lambda_t'\Lambda_t - \Lambda_t'\varepsilon_{t+1} + A_{\tau}^{\$} + (B_{\tau}^{\$})' z_{t+1}\}] \\ &= \exp\{-y_{0,1}^{\$} - e'_{yn}z_t - \frac{1}{2}\Lambda_t'\Lambda_t + A_{\tau}^{\$} + (B_{\tau}^{\$})' \Psi z_t\} \times \\ &\quad \mathbb{E}_t \left[\exp\{-\Lambda_t'\varepsilon_{t+1} + (B_{\tau}^{\$})' \Sigma^{\frac{1}{2}}\varepsilon_{t+1}\} \right]. \end{aligned}$$

We use the log-normality of ε_{t+1} and substitute for the affine expression for Λ_t to get:

$$P_{t,\tau+1}^{\$} = \exp \left\{ -y_{0,1}^{\$} - e'_{yn}z_t + A_{\tau}^{\$} + (B_{\tau}^{\$})' \Psi z_t + \frac{1}{2} (B_{\tau}^{\$})' \Sigma (B_{\tau}^{\$}) - (B_{\tau}^{\$})' \Sigma^{\frac{1}{2}} (\Lambda_0 + \Lambda_1 z_t) \right\}.$$

Taking logs and collecting terms, we obtain a linear equation for $\log(p_t(\tau + 1))$:

$$\log(P_{t,\tau+1}^{\$}) = A_{\tau+1}^{\$} + (B_{\tau+1}^{\$})' z_t,$$

where $A_{\tau+1}^{\$}$ satisfies (18) and $B_{\tau+1}^{\$}$ satisfies (19). The relationship between log bond prices and bond yields is given by $-\log(P_{t,\tau}^{\$})/\tau = y_{t,\tau}^{\$}$. \square

Define the one-period return on a nominal zero-coupon bond as:

$$r_{t+1,\tau}^{b,\$} = \log(P_{t+1,\tau}^{\$}) - \log(P_{t,\tau+1}^{\$})$$

The nominal bond risk premium on a bond of maturity τ is the expected excess return corrected for a Jensen term, and equals negative the conditional covariance between that bond return and the nominal SDF:

$$\begin{aligned} \mathbb{E}_t \left[r_{t+1,\tau}^{b,\$} \right] - y_{t,1}^{\$} + \frac{1}{2} \mathbf{V}_t \left[r_{t+1,\tau}^{b,\$} \right] &= -\mathbf{C}_t \left[m_{t+1}^{\$}, r_{t+1,\tau}^{b,\$} \right] \\ &= (B_{\tau}^{\$})' \Sigma^{\frac{1}{2}} \Lambda_t \end{aligned}$$

Real bond yields, $y_{t,\tau}$, denoted without the \$ superscript, are affine as well with coefficients that follow similar recursions:

$$A_{\tau+1} = -y_{0,1} + A_\tau + \frac{1}{2}B'_\tau \Sigma B_\tau - B'_\tau \Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}'} e_\pi \right), \quad (20)$$

$$B'_{\tau+1} = -e'_{y\pi} + (e_\pi + B_\tau)' \left(\Psi - \Sigma^{\frac{1}{2}} \Lambda_1 \right). \quad (21)$$

For $\tau = 1$, we recover the expression for the risk-free rate in (16)-(17).

B.3 Stock Market

We define the real return on equity as $R_{t+1}^m = \frac{P_{t+1}^m + D_{t+1}^m}{P_t^m}$, where P_t^m is the end-of-period price on the equity market. A log-linearization delivers:

$$r_{t+1}^m = \kappa_0^m + \Delta d_{t+1}^m + \kappa_1^m pd_{t+1}^m - pd_t^m. \quad (22)$$

The unconditional mean log real stock return is $r_0^m = \mathbb{E}[r_t^m]$, the unconditional mean dividend growth rate is $\mu^m = \mathbb{E}[\Delta d_{t+1}^m]$, and $\overline{pd^m} = \mathbb{E}[pd_t^m]$ is the unconditional average log price-dividend ratio on equity. The linearization constants κ_0^m and κ_1^m are defined as:

$$\kappa_1^m = \frac{e^{\overline{pd^m}}}{e^{\overline{pd^m}} + 1} < 1 \quad \text{and} \quad \kappa_0^m = \log \left(e^{\overline{pd^m}} + 1 \right) - \frac{e^{\overline{pd^m}}}{e^{\overline{pd^m}} + 1} \overline{pd^m}. \quad (23)$$

Our state vector z contains the (demeaned) log real dividend growth rate on the stock market, $\Delta d_{t+1}^m - \mu^m$, and the (demeaned) log price-dividend ratio $pd_t^m - \overline{pd^m}$.

$$\begin{aligned} pd_t^m(\tau) &= \overline{pd^m} + e'_{pd} z_t, \\ \Delta d_t^m &= \mu^m + e'_{divm} z_t, \end{aligned}$$

where e'_{pd} (e_{divm}) is a selector vector that has a one in the fifth (sixth) entry, since the log pd ratio (log dividend growth rate) is the fifth (sixth) element of the VAR.

We define the log return on the stock market so that the return equation holds exactly, given the time series for $\{\Delta d_t^m, pd_t^m\}$. Rewriting (22):

$$\begin{aligned} r_{t+1}^m - r_0^m &= \left[(e_{divm} + \kappa_1^m e_{pd})' \Psi - e'_{pd} \right] z_t + \left(e_{divm} + \kappa_1^m e_{pd} \right)' \Sigma^{\frac{1}{2}} \varepsilon_{t+1}. \\ r_0^m &= \mu^m + \kappa_0^m - \overline{pd^m} (1 - \kappa_1^m). \end{aligned}$$

The equity risk premium is the expected excess return on the stock market corrected for a Jensen term. By the Euler equation, it equals minus the conditional covariance between the log SDF and the log return:

$$\begin{aligned}
1 &= \mathbb{E}_t \left[M_{t+1} \frac{P_{t+1}^m + D_{t+1}^m}{P_t^m} \right] = E_t \left[\exp\{m_{t+1}^{\$} + \pi_{t+1} + r_{t+1}^m\} \right] \\
&= \mathbb{E}_t \left[\exp \left\{ -y_{t,1}^{\$} - \frac{1}{2} \Lambda_t' \Lambda_t - \Lambda_t' \varepsilon_{t+1} + \pi_0 + e'_{\pi} z_{t+1} + r_0^m + (e_{divm} + \kappa_1^m e_{pd})' z_{t+1} - e'_{pd} z_t \right\} \right] \\
&= \exp \left\{ -y_0^{\$}(1) - \frac{1}{2} \Lambda_t' \Lambda_t + \pi_0 + r_0^m + [(e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Psi - e'_{pd} - e'_{yn}] z_t \right\} \\
&\quad \times \mathbb{E}_t \left[\exp \left\{ -\Lambda_t' \varepsilon_{t+1} + (e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Sigma^{\frac{1}{2}} \varepsilon_{t+1} \right\} \right] \\
&= \exp \left\{ r_0^m + \pi_0 - y_0^{\$}(1) + [(e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Psi - e'_{pd} - e'_{yn}] z_t \right\} \\
&\quad \times \exp \left\{ \frac{1}{2} (e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Sigma (e_{divm} + \kappa_1^m e_{pd} + e_{\pi}) - (e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Sigma^{\frac{1}{2}} \Lambda_t \right\}
\end{aligned}$$

Taking logs on both sides delivers:

$$\begin{aligned}
r_0^m + \pi_0 - y_0^{\$}(1) + [(e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Psi - e'_{pd} - e'_{yn}] z_t & \tag{24} \\
+ \frac{1}{2} (e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Sigma (e_{divm} + \kappa_1^m e_{pd} + e_{\pi}) &= (e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Sigma^{\frac{1}{2}} \Lambda_t \\
\mathbb{E}_t [r_{t+1}^{m,\$}] - y_{t,1}^{\$} + \frac{1}{2} \mathbb{V}_t [r_{t+1}^{m,\$}] &= -\mathbf{C}_t [m_{t+1}^{\$}, r_{t+1}^{m,\$}]
\end{aligned}$$

The left-hand side is the expected excess return on the stock market, corrected for a Jensen term, while the right-hand side is the negative of the conditional covariance between the (nominal) log stock return and the nominal log SDF. We refer to the left-hand side as the equity risk premium in the data, since it is implied directly by the VAR. We refer to the right-hand side as the equity risk premium in the model, since it requires knowledge of the market prices of risk.

Note that we can obtain the same expression using the log real SDF and log real stock return:

$$\begin{aligned}
\mathbb{E}_t [r_{t+1}^m] - y_{t,1} + \frac{1}{2} \mathbb{V}_t [r_{t+1}^m] &= -\mathbf{C}_t [m_{t+1}, r_{t+1}^m] \\
r_0^m - y_0(1) + [(e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Psi - e'_{pd} - e'_{yn} - e'_{\pi} \Sigma^{1/2} \Lambda_1] z_t & \\
+ \frac{1}{2} (e_{divm} + \kappa_1^m e_{pd})' \Sigma (e_{divm} + \kappa_1^m e_{pd}) &= (e_{divm} + \kappa_1^m e_{pd})' \Sigma^{1/2} (\Lambda_t - (\Sigma^{1/2})' e_{\pi})
\end{aligned}$$

We combine the terms in Λ_0 and Λ_1 on the right-hand side and plug in for $y_0(1)$ from (17) to get:

$$\begin{aligned}
r_0^m + \pi_0 - y_{0,1}^{\$} + \frac{1}{2} e'_{\pi} \Sigma e_{\pi} &+ \frac{1}{2} (e_{divm} + \kappa_1^m e_{pd})' \Sigma (e_{divm} + \kappa_1^m e_{pd}) + e'_{\pi} \Sigma (e_{divm} + \kappa_1^m e_{pd}) \\
+ [(e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Psi - e'_{pd} - e'_{yn}] z_t &= (e_{divm} + \kappa_1^m e_{pd})' \Sigma^{1/2} \Lambda_t + e'_{\pi} \Sigma^{\frac{1}{2}} \Lambda_0 + e'_{\pi} \Sigma^{1/2} \Lambda_1 z_t
\end{aligned}$$

This recovers equation (24).

B.4 Dividend Strips

B.4.1 Affine Structure for Price-Dividend Ratio on Equity Strip

Proposition 2. Log price-dividend ratios on dividend strips are affine in the state vector:

$$p_{t,\tau}^d = \log \left(P_{t,\tau}^d \right) = A_\tau^m + B_\tau^{m'} z_t,$$

where the coefficients A_τ^m and $B_\tau^{m'}$ follow recursions:

$$\begin{aligned} A_{\tau+1}^m &= A_\tau^m + \mu_m - y_0(1) + \frac{1}{2} (e_{divm} + B_\tau^m)' \Sigma (e_{divm} + B_\tau^m) \\ &\quad - (e_{divm} + B_\tau^m)' \Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}'} e_\pi \right), \end{aligned} \quad (25)$$

$$B_{\tau+1}^{m'} = (e_{divm} + e_\pi + B_\tau^m)' \Psi - e'_{yn} - (e_{divm} + e_\pi + B_\tau^m)' \Sigma^{\frac{1}{2}} \Lambda_1, \quad (26)$$

initialized at $A_0^m = 0$ and $B_0^m = 0$.

Proof. We conjecture the affine structure and solve for the coefficients $A_{\tau+1}^m$ and $B_{\tau+1}^{m'}$ in the process of verifying this conjecture using the Euler equation:

$$\begin{aligned} P_{t,\tau+1}^d &= \mathbb{E}_t \left[M_{t+1} P_{t+1,\tau}^d \frac{D_{t+1}^m}{D_t^m} \right] = E_t \left[\exp \{ m_{t+1}^\$ + \pi_{t+1} + \Delta d_{t+1}^m + p_{t+1}^d(\tau) \} \right] \\ &= \mathbb{E}_t \left[\exp \left\{ -y_{t,1}^\$ - \frac{1}{2} \Lambda_t' \Lambda_t - \Lambda_t' \varepsilon_{t+1} + \pi_0 + e'_{\pi} z_{t+1} + \mu^m + e'_{divm} z_{t+1} + A_\tau^m + B_\tau^{m'} z_{t+1} \right\} \right] \\ &= \exp \left\{ -y_0^\$(1) - e'_{yn} z_t - \frac{1}{2} \Lambda_t' \Lambda_t + \pi_0 + e'_{\pi} \Psi z_t + \mu_m + e'_{divm} \Psi z_t + A_\tau^m + B_\tau^{m'} \Psi z_t \right\} \\ &\quad \times \mathbb{E}_t \left[\exp \left\{ -\Lambda_t' \varepsilon_{t+1} + (e_{divm} + e_\pi + B_\tau^m)' \Sigma^{\frac{1}{2}} \varepsilon_{t+1} \right\} \right]. \end{aligned}$$

We use the log-normality of ε_{t+1} and substitute for the affine expression for Λ_t to get:

$$\begin{aligned} P_{t,\tau+1}^d &= \exp \left\{ -y_0^\$(1) + \pi_0 + \mu_m + A_\tau^m + \left[(e_{divm} + e_\pi + B_\tau^m)' \Psi - e'_{yn} \right] z_t \right. \\ &\quad \left. + \frac{1}{2} (e_{divm} + e_\pi + B_\tau^m)' \Sigma (e_{divm} + e_\pi + B_\tau^m) \right. \\ &\quad \left. - (e_{divm} + e_\pi + B_\tau^m)' \Sigma^{\frac{1}{2}} (\Lambda_0 + \Lambda_1 z_t) \right\} \end{aligned}$$

Taking logs and collecting terms, we obtain a log-linear expression for $p_t^d(\tau + 1)$:

$$p_{t,\tau+1}^d = A_{\tau+1}^m + B_{\tau+1}^{m'} z_t,$$

where:

$$\begin{aligned} A_{\tau+1}^m &= A_{\tau}^m + \mu_m - y_0^{\$}(1) + \pi_0 + \frac{1}{2} (e_{divm} + e_{\pi} + B_{\tau}^m)' \Sigma (e_{divm} + e_{\pi} + B_{\tau}^m) \\ &\quad - (e_{divm} + e_{\pi} + B_{\tau}^m)' \Sigma^{\frac{1}{2}} \Lambda_0, \\ B_{\tau+1}^{m'} &= (e_{divm} + e_{\pi} + B_{\tau}^m)' \Psi - e'_{yn} - (e_{divm} + e_{\pi} + B_{\tau}^m)' \Sigma^{\frac{1}{2}} \Lambda_1. \end{aligned}$$

We recover the recursions in (25) and (26) after using equation (17). □

Like we did for the stock market as a whole, we define the strip risk premium as:

$$\begin{aligned} \mathbb{E}_t \left[r_{t+1,\tau}^{d,\$} \right] - y_{t,1}^{\$} + \frac{1}{2} \mathbb{V}_t \left[r_{t+1,\tau}^{d,\$} \right] &= -\mathbf{C}_t \left[m_{t+1}^{\$}, r_{t+1,\tau}^{d,\$} \right] \\ &= (e_{divm} + e_{\pi} + B_{\tau}^m)' \Sigma^{\frac{1}{2}} \Lambda_t \end{aligned}$$

The risky strips for REITs, infrastructure, small stocks, growth stocks, and natural resources are defined analogously.

B.4.2 Expected Holding Period Return over k-horizons on Dividend Strips

The expected return on a dividend strip that pays the realized nominal dividend k quarters hence and that is held to maturity is:

$$\begin{aligned} \mathbb{E}_t [R_{t \rightarrow t+k}] &= \frac{\mathbb{E}_t \left[\frac{D_{t+k}^{\$}}{D_t^{\$}} \right]}{P_{t,k}^d} - 1 \\ &= \exp \left(-A_k^m - B_k^{m'} z_t + \mathbb{E}_t \left[\sum_{s=1}^k \Delta d_{t+s} + \pi_{t+s} \right] + \frac{1}{2} \mathbb{V}_t \left[\sum_{s=1}^k \Delta d_{t+s} + \pi_{t+s} \right] \right) - 1 \\ &= \exp \left(-A_k^m - B_k^{m'} z_t + k(\mu_m + \pi_0) + (e_{divm} + e_{\pi})' \left[\sum_{s=1}^k \Psi^s \right] z_t + \frac{k}{2} (e_{divm} + e_{\pi})' \Sigma (e_{divm} + e_{\pi}) \right) - 1 \end{aligned} \quad (27)$$

These are the building blocks for computing the expected return on a PE investment.

B.4.3 Dividend Futures: Price and Return

The price of a dividend futures contract which delivers one quarter worth of nominal dividends at quarter $t + \tau$, divided by the current dividend, is equal to:

$$\frac{F_{t,\tau}^d}{D_t^\$} = P_{t,\tau}^d \exp\left(\tau y_{t,\tau}^\$\right),$$

where $P_{t,\tau}^d$ is the spot price-dividend ratio. Using the affine expressions for the strip price-dividend ratio and the nominal bond price, it can be written as:

$$\frac{F_{t,\tau}^d}{D_t^\$} = \exp\left(A_\tau^m - A_\tau^\$ + (B_\tau^m - B_\tau^\$)'z_t\right),$$

The one-period holding period return on the dividend future of maturity τ is:

$$R_{t+1,\tau}^{fut,d} = \frac{F_{t+1,\tau-1}^d}{F_{t,\tau}^d} - 1 = \frac{F_{t+1,\tau-1}^d/D_{t+1}^\$}{F_{t,\tau}^d/D_t^\$} \frac{D_{t+1}^\$}{D_t^\$} - 1$$

It can be written as:

$$\begin{aligned} \log\left(1 + R_{t+1,\tau}^{fut,d}\right) &= A_{\tau-1}^m - A_{\tau-1}^\$ - A_\tau^m + A_\tau^\$ + \mu_m + \pi_0 \\ &\quad + (B_{\tau-1}^m - B_{\tau-1}^\$ + e_{divm} + e_\pi)'z_{t+1} - (B_\tau^m - B_\tau^\$)'z_t \end{aligned}$$

The expected log return, which is already a risk premium on account of the fact that the dividend future already takes out the return on an equal-maturity nominal Treasury bond, equals:

$$\begin{aligned} \mathbb{E}_t \left[\log\left(1 + R_{t+1,\tau}^{fut,d}\right) \right] &= A_{\tau-1}^m - A_{\tau-1}^\$ - A_\tau^m + A_\tau^\$ + \mu_m + \pi_0 \\ &\quad + \left[(B_{\tau-1}^m - B_{\tau-1}^\$ + e_{divm} + e_\pi)' \Psi - (B_\tau^m - B_\tau^\$)' \right] z_t \end{aligned}$$

Given that the state variable z_t is mean-zero, the first row denotes the unconditional dividend futures risk premium.

B.5 Capital Gain Strips

A capital gain strip is a strip that pays the realized ex-dividend stock price P_{t+k}^m at time $t + k$. For convenience, we scale this payout by the current stock price P_t^m . In other words, the claim pays off the realized cumulative capital gain between periods t and $t + k$, $\frac{P_{t+k}^m}{P_t^m}$.

By value additivity of the dividend strips, the time- t price of this claim is today's stock price minus the prices of the dividend strips of horizons $1, \dots, k$:

$$\frac{P_t^m - (P_{t,1}^d + \dots + P_{t,k}^d) D_t^m}{P_t^m} = 1 - \frac{P_{t,1}^d + \dots + P_{t,k}^d}{P_t^m / D_t^m} = 1 - \frac{\sum_{\tau=1}^k \exp\{A_\tau^m + (B_\tau^m)' z_t\}}{\exp\{pd^m + e'_{pdm} z_t\}}$$

The expected return on the capital gains strip is given by

$$\begin{aligned} \frac{\mathbb{E}_t \left[\frac{P_{t+k}^m}{P_t^m} \right]}{P_t^m - (P_{t,1}^d + \dots + P_{t,k}^d) D_t^m} &= \frac{\mathbb{E}_t \left[\frac{P_{t+k}^m}{P_t^m} \right]}{1 - \frac{\sum_{\tau=1}^k \exp\{A_\tau^m + (B_\tau^m)' z_t\}}{\exp\{pd^m + e'_{pdm} z_t\}}} = \frac{\mathbb{E}_t \left[\frac{P_{t+k}^m / D_{t+k}^m}{P_t^m / D_t^m} \frac{D_{t+k}^m}{D_t^m} \right]}{1 - \frac{\sum_{\tau=1}^k \exp\{A_\tau^m + (B_\tau^m)' z_t\}}{\exp\{pd^m + e'_{pdm} z_t\}}} \\ &= \frac{\mathbb{E}_t \left[\exp\{e'_{pdm} (z_{t+k} - z_t) + \Delta d_{t+k}^{\$} + \dots + \Delta d_{t+1}^{\$}\} \right]}{1 - \frac{\sum_{\tau=1}^k \exp\{A_\tau^m + (B_\tau^m)' z_t\}}{\exp\{pd^m + e'_{pdm} z_t\}}} \\ &= \frac{\mathbb{E}_t \left[\exp\{e'_{pdm} (z_{t+k} - z_t) + k(\mu_m + \pi_0) + \sum_{\tau=1}^k (e_{divm} + e_\pi)' z_{t+\tau}\} \right]}{1 - \frac{\sum_{\tau=1}^k \exp\{A_\tau^m + (B_\tau^m)' z_t\}}{\exp\{pd^m + e'_{pdm} z_t\}}} \\ &= \frac{\exp \left\{ k(\mu_m + \pi_0) + \left[e'_{pdm} (\Psi^k - I) + (e_{divm} + e_\pi)' \sum_{\tau=1}^k (\Psi)^\tau \right] z_t + \frac{1}{2} \mathcal{V} \right\}}{1 - \frac{\sum_{\tau=1}^k \exp\{A_\tau^m + (B_\tau^m)' z_t\}}{\exp\{pd^m + e'_{pdm} z_t\}}} \end{aligned}$$

where

$$\mathcal{V} = e'_{pdm} \left[\sum_{\tau=1}^k \Psi^{\tau-1} \Sigma (\Psi^{\tau-1})' \right] e_{pdm} + (e_{divm} + e_\pi)' \left[\sum_{\tau=1}^k \sum_{n=1}^{\tau} \Psi^{n-1} \Sigma \Psi^{n-1} \right] (e_{divm} + e_\pi)$$

B.6 Covariance PE with stock and bond returns

The covariance of PE fund returns with stock returns and with bond returns is given by the covariance of the PE fund's replication portfolio return with stock and bond returns. The return on the replicating portfolio is the weighted average of the return on the strips that make up the portfolio. The weights are described in equation (2), where w^i is a $1 \times HK$ vector with generic element $w_{t,h,k}^i = q_{t,h}^i(k) P_{t,h}(k)$. The weights sum to one. We focus on the one-period, conditional stock beta $\beta_{t,m}^i$ and bond beta $\beta_{t,b}^i$ of PE fund i . For the bond beta, we focus on the beta with the nominal five-year bond ($\tau = 20$).

$$\beta_{t,m}^i = \frac{\mathbf{C}_t(r_{t+1}^i, r_{t+1}^m)}{\mathbf{V}_t(r_{t+1}^m)} = \frac{\mathbf{C}_t(\sum_{h,k}^{HK} w_{t,h,k}^i r_{t+1,h,k}^i, r_{t+1}^m)}{\mathbf{V}_t(r_{t+1}^m)} = \frac{\sum_{h,k}^{HK} w_{t,h,k}^i \mathbf{C}_t(r_{t+1,h}(k), r_{t+1}^m)}{\mathbf{V}_t(r_{t+1}^m)}$$

$$\beta_{t,b}^i = \frac{\sum_{h,k}^{HK} w_{t,h,k}^i (e_{strip,k,h})' \Sigma (e_{divm} + e_{\pi} + \kappa_1^m e_{pd})}{\left(e_{divm} + e_{\pi} + \kappa_1^m e_{pd} \right)' \Sigma \left(e_{divm} + e_{\pi} + \kappa_1^m e_{pd} \right)}$$

$$\beta_{t,b}^i = \frac{\sum_{h,k}^{HK} w_{t,h,k}^i(k) (e_{strip,k,h})' \Sigma (B_{20}^{\$})}{\left(B_{20}^{\$} \right)' \Sigma \left(B_{20}^{\$} \right)}$$

For example, the three factor model for Real Estate PE funds implies that:

$$e_{strip,1,h} = B_h^{\$}$$

$$e_{strip,2,h} = e_{divm} + e_{\pi} + B_h^m$$

$$e_{strip,3,h} = e_{divreit} + e_{\pi} + B_h^{reit}$$

As another example, the three-factor model for VC funds implies:

$$e_{strip,1,h} = B_h^{\$}$$

$$e_{strip,2,h} = e_{divsmall} + e_{\pi} + B_h^{small}$$

$$e_{strip,3,h} = e_{divgrowth} + e_{\pi} + B_h^{growth}$$

Similar expressions are obtained for the other fund categories and for models with fewer or more risk factors.

C Point Estimates Baseline Model

C.1 VAR Estimation

In the first stage we estimate the VAR companion matrix by OLS, equation by equation. We start from an initial VAR where all elements of Ψ are non-zero. We zero out the elements whose t-statistic is less than 1.96. We then re-estimate Ψ and zero out the elements whose t-statistic is less than 1.96. We continue this procedure until the Ψ matrix no longer changes and all remaining elements have t-statistic greater than 1.96. The resulting VAR companion matrix estimate, $\hat{\Psi}$, is listed below.

0.81	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.28	0.00	0.00	0.01	0.00	0.00	0.00	-0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.11	0.00	0.91	0.19	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.71	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.98	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.52	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.96	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	1.78	0.00	0.07	0.48	0.06	0.10	-0.07	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.98	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.35	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	-5.70	0.00	0.00	0.00	-0.12	0.00	0.00	0.00	0.94	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	-3.92	0.00	0.60	0.00	0.00	0.00	0.00	0.00	0.29	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.97	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	-0.17	0.00	0.00	0.00	0.00	0.00	0.42	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.93	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.04	0.49	0.00	0.00
0.00	0.00	-3.76	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.93	-0.59
0.00	0.00	2.84	0.00	-0.23	0.00	0.00	0.00	0.00	0.51	0.00	0.00	0.12	0.00	0.00	0.19	0.14	0.54

The Cholesky decomposition of the residual variance-covariance matrix, $\Sigma^{\frac{1}{2}}$, multiplied by 100 for readability is given by:

0.25	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.04	0.67	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.03	0.09	0.23	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-0.01	-0.01	-0.12	-0.13	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-1.52	1.05	-1.00	0.37	8.38	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-0.06	0.05	0.07	0.10	-0.09	1.60	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-1.21	0.90	-1.34	-0.25	5.76	0.27	7.47	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.08	0.02	0.02	-0.13	0.08	0.50	-1.30	3.12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-0.75	0.83	-0.42	-0.07	6.56	1.57	0.01	0.29	4.10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-0.06	0.12	-0.04	-0.19	-0.35	0.22	0.09	-0.13	-0.72	1.73	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-1.98	1.47	-1.14	1.38	9.71	1.42	0.99	0.91	-0.90	0.76	6.90	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.37	0.04	0.27	0.46	-0.19	0.94	0.26	0.38	-0.11	-0.90	-3.13	4.29	0.00	0.00	0.00	0.00	0.00	0.00
-1.51	1.40	-0.67	-0.22	8.50	0.04	-1.34	-0.76	-0.83	-0.00	0.00	0.33	3.75	0.00	0.00	0.00	0.00	0.00
-0.16	-0.08	-0.03	0.11	-0.02	1.52	-0.11	0.17	0.09	-0.10	0.05	-0.41	-1.96	2.10	0.00	0.00	0.00	0.00
-0.07	0.70	0.12	-0.29	4.92	0.76	0.33	1.66	4.38	0.09	0.30	-0.64	-0.65	-1.44	6.92	0.00	0.00	0.00
0.05	0.04	-0.09	-0.30	-0.06	-0.32	0.33	-0.74	-0.56	0.85	-0.41	0.50	0.29	0.08	-1.76	3.47	0.00	0.00
-1.24	0.85	-1.70	1.13	8.02	0.25	1.82	0.79	0.50	-0.60	1.61	1.32	-1.66	-0.90	-0.69	0.62	6.05	0.00
0.09	0.15	0.76	0.37	0.16	1.46	-0.23	-0.24	-0.20	0.65	-0.54	0.17	0.86	-0.16	0.19	-0.55	-2.72	3.13

The diagonal elements report the standard deviation of the VAR innovations, each one orthogonalized to the shocks that precede it in the VAR, expressed in percent per quarter.

C.2 Market Price of Risk Estimates

The market prices of risk are pinned down by the moments discussed in the main text. Here we report and discuss the point estimates. Note that the prices of risk are associated with the orthogonal VAR innovations $\varepsilon \sim \mathcal{N}(0, I)$. Therefore, their magnitudes can be interpreted as (quarterly) Sharpe ratios. The constant in the market price of risk estimate $\widehat{\Lambda}_0$ is:

0.32	-0.19	-0.43	0.23	-0.07	1.18	0.00	0.37	0.00	0.01	0.00	-0.14	0.00	0.51	0.00	0.52	0.00	0.33
------	-------	-------	------	-------	------	------	------	------	------	------	-------	------	------	------	------	------	------

The matrix that governs the time variation in the market price of risk is estimated to be $\widehat{\Lambda}_1 =$:

71.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	-49.8	-148.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
82.0	-0.3	-19.3	7.4	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
25.0	5.3	-6.7	-16.3	0.1	0.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-40.1	-25.7	-51.8	-3.4	-2.1	21.8	0.7	0.0	0.0	0.0	-0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
19.5	-3.7	26.9	-29.9	2.5	8.4	0.4	0.4	-2.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
32.4	1.3	-0.7	26.2	1.5	-20.0	-0.7	-0.2	-1.4	16.2	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-36.0	3.6	-116.8	-76.2	0.1	-4.8	-3.4	-0.1	0.6	0.2	-1.4	6.7	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
24.1	-3.3	3.6	9.1	1.8	-15.4	-0.6	-3.2	-0.6	-3.2	0.0	0.2	-1.8	17.4	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-2.8	8.8	-15.0	26.5	-0.4	-0.3	0.1	-1.3	0.8	-12.5	-0.1	0.3	-0.3	4.2	-0.4	4.6	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-14.9	-5.2	47.8	39.9	-6.7	-17.2	1.1	-0.6	-0.0	6.2	0.7	-2.1	3.2	5.8	0.1	4.3	2.3	-0.8

C.3 Identification

The estimation proceeds in nine steps; those steps are also useful to explain the identification of the market price of risk parameters.

Step 1: Only Bonds The first four elements of Λ_0 and the first four rows of Λ_1 govern the dynamics of bond yields and bond returns. The state vector contains the one-quarter bond yield. The model matches the time series of the one-quarter bond yield automatically by virtue of its inclusion in the state space. This moment does not identify any market price of risk parameters since $A^\$(1) = 1$ and $B^\$(1) = [1, 1, \dots, 1]$ and do not depend on any market price of risk parameters.

Since the spread between the 20-quarter bond yield and the one-quarter bond yield is the fourth element of the state vector, and the short rate is the third element of the state vector, the 20-quarter bond yield can be written as:

$$y_{t,20}^\$ = y_{0,20}^\$ + (e_{yn} + e_{yspr})' z_t = -\frac{A_{20}^\$}{20} - \frac{B_{20}^{\$'}}{20} z_t$$

This restriction identifies one element in the constant Λ_0 , specifically

$$y_{0,1}^\$ + yspr_0 = -\frac{1}{20} A_{20}^\$$$

and N elements in the time-varying market price of risk matrix Λ_1 :

$$e'_{y1} + e'_{yspr} = -\frac{1}{20} (B_{20}^\$)'$$

The recursions for the coefficients in the affine term structure model are repeated here for convenience:

$$A_{\tau+1}^\$ = -y_{0,1}^\$ + A_\tau^\$ + \frac{1}{2} (B_\tau^\$)' \Sigma (B_\tau^\$) - (B_\tau^\$)' \Sigma^{\frac{1}{2}} \Lambda_0, \quad (28)$$

$$(B_{\tau+1}^\$)' = (B_\tau^\$)' \Psi - e'_{yn} - (B_\tau^\$)' \Sigma^{\frac{1}{2}} \Lambda_1, \quad (29)$$

initialized at $A_0^\$ = 0$ and $B_0^\$ = 0$. Define $\tilde{\Psi} = \Psi - \Sigma^{\frac{1}{2}} \Lambda_1$ to be the risk-neutral companion matrix. Then (19) can be written as:

$$-\frac{(B_{\tau+1}^\$)'}{20} = \frac{1}{20} e'_{yn} (I - \tilde{\Psi})^{-1} (I - \tilde{\Psi}^{\tau+1})$$

The restriction on Λ_1 can be written as:

$$e'_{y1} + e'_{yspr} = \frac{1}{20} e'_{yn} (I - \tilde{\Psi}^{\tau+1}) (I - \tilde{\Psi})^{-1}$$

The left-hand side is a $N \times 1$ vector with a 1 in elements 3 and 4 and a 0 in the other four positions. Hence, the same must be true of the right-hand side. This imposes $N = 18$ restrictions on Λ_1 which affects $\tilde{\Psi}$, given Ψ . There is one restriction on each of the columns of Λ_1 . It is a restriction on a linear combination of the elements in the first four rows of Λ_1 of that column. For example, if the elements in the 10th column and first three rows of Λ_1 are all zero, it is a simple restriction on the element in the fourth row and 10th column.

We allow for four non-zero elements in the first three rows of Λ_1 : $\Lambda_1(1, 1)$, $\Lambda_1(2, 2)$, $\Lambda_1(3, 3)$, $\Lambda_1(3, 4)$. The price of inflation risk is allowed to depend on the level of inflation, the price of GDP growth risk is allowed to depend on the level of GDP growth, and the price of interest rate risk is allowed to depend on the level of the interest rate (as in [Cox, Ingersoll, and Ross, 1985](#)) and on the slope (as in [Campbell and Shiller, 1991](#)). The positive association between the slope and future bond returns is also consistent with the bond return predictability evidence ([Cochrane and Piazzesi, 2006](#)). The constant market prices of risk associated with inflation, GDP growth, and the level factor are also allowed to differ from zero. These seven parameters are identified off the cross-section of nominal and real bond yields; we include five additional nominal yield maturities beyond the 1- and 20-quarter yields, and five real bond maturities. Given the structure imposed on the first three rows of Λ_1 , matching the 20-quarter bond yield requires non-zero elements in the first four columns of the fourth row.

To assure that the model ends up with reasonable implication for yields at long maturities not observed in the data, but nevertheless important for valuation, we impose the following regularity conditions on nominal and real bonds. First, we impose that the unconditional mean nominal bond yield be at least as high as the sample mean nominal GDP growth rate at maturities of 200, 400, 800, 1200, 1600, and 2000 quarters. Second, we impose that the unconditional mean real yield be at least as high as the sample mean real GDP growth rate at the same maturities. Third, we impose that the unconditional annual nominal yield be at least 2% higher than the unconditional real bond yield of the same maturity, at maturities of 200, 400, 800, 1200, and 2000 quarters. Fourth, we assume that the nominal and real yield curves flatten out at long maturities. Specifically, the difference between the annual yield at 400 and 200 quarters maturity must not exceed 1%; the difference between the annual yield at 800 and 600 quarters maturity must not exceed 0.5%; and the difference between the annual yield at 2000 and 1000 quarters maturity must not exceed 0.25%. These bond regularity conditions have strong bite in early iterations of the estimation algorithm. However, none of the restrictions bind at the final parameter estimates.

The estimation also imposes good deal bounds of the kind suggested by [Cochrane and Saa-Requejo \(2000\)](#). Specifically, we impose a severe penalty on large values for the unconditional standard deviation of

the log SDF (the maximum Sharpe ratio):

$$\mathcal{P}_1 = 100 \left(\exp [\max \{Std(m_{t+1}) - 1.5, 0\}]^2 - 1 \right) \quad (30)$$

We also include a penalty on large values for the market prices of risk Λ_t at each date t :

$$\mathcal{P}_2 = 10 \sum_{t=1}^T \left(\exp \left[\max \left\{ \sqrt{\Lambda_t' \Lambda_t} - 3, 0 \right\} \right]^2 - 1 \right) \quad (31)$$

Step 2: Only Aggregate Stock Market In the second step, we hold fixed the market price of risk parameters in the first four elements (rows) of Λ_0 (Λ_1), and estimate the market prices of risk in the fifth and sixth rows. Recall that the VAR models the dynamics of the price-dividend ratio (row 5) and the dividend growth rate (row 6). The VAR implies an expected excess log stock return including a Jensen adjustment given by the left-hand side of the following equation:

$$\begin{aligned} r_0^m + \pi_0 - y_{0,1}^s + \frac{1}{2} e_{\pi}' \Sigma e_{\pi} + \frac{1}{2} (e_{divm} + \kappa_1^m e_{pd})' \Sigma (e_{divm} + \kappa_1^m e_{pd}) + e_{\pi}' \Sigma (e_{divm} + \kappa_1^m e_{pd}) \\ + \left[(e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Psi - e_{pd}' - e_{yn}' \right] z_t = (e_{divm} + \kappa_1^m e_{pd})' \Sigma^{1/2} (\Lambda_0 + \Lambda_1' z_t) + e_{\pi}' \Sigma^{1/2} \Lambda_0 + e_{\pi}' \Sigma^{1/2} \Lambda_1 z_t \end{aligned} \quad (32)$$

The first term on the second line is the time-varying component of the expected excess stock return. This is just data, i.e. the companion matrix of the VAR Ψ , not asset pricing. The asset pricing is on the right-hand side of the above equation. It reports the equity risk premium, which is negative the conditional covariance of the the log stock return and the log SDF. It depends on the market prices of risk.

To replicate this time-variation, the sixth row of Λ_1 must be such that the time-varying components of the left- and right-hand sides of equation (32) are equalized:

$$(e_{divm} + \kappa_1^m e_{pd} + e_{\pi})' \Psi - e_{pd}' - e_{yn}' = (e_{divm} + \kappa_1^m e_{pd})' \Sigma^{1/2} \Lambda_1' + e_{\pi}' \Sigma^{1/2} \Lambda_1$$

This is a linear system of N equations in N unknowns, which uniquely pins down the N elements in the sixth row of Λ_1 . Given the structure of Ψ and the structure of the market prices of risk in the first four rows of Λ_1 , the sixth row of Λ_1 must have non-zero elements in the first through sixth columns, as well as in the seventh and eleventh columns. It must have zeros in all other columns.

We free up $\Lambda_0(5)$ and $\Lambda_1(5, 1 : 6)$ in order to match the dividend strip evidence described in the main text. We match the time series of the price-dividend ratio on the first 8 quarters of dividends, the time series of the share of the overall stock market value these first 8 strips represent, and the unconditional mean return on a portfolio of 1- through 7-year dividend futures.

The identification of the fifth and sixth elements of the market prices of risk is also aided by the moment that equates the price-dividend ratio in the data to the price-dividend ratio in the model, where the latter is calculated at the sum of the price-dividend ratios of the first 3600 dividend strips.

This estimation also imposes the bond yield regularity conditions and good deal bounds (30) and (31).

Step 3-8 In step 3 of the estimation, we hold fixed the market prices of risk in the first six elements (rows) of Λ_0 (Λ_1), and estimate the market prices of risk in the eight rows. Recall that the seventh and eighth row contain the log price-dividend ratio and log dividend growth on REITs. We only free up the market prices of risk in the eighth row, associated with REIT dividend growth. the constant $\Lambda_0(8)$ is chosen to match the equity risk premium on REITS. The eighth row to match the VAR-implied dynamics of the expected excess return on REITS, a condition analogous to equation (32). This implies N conditions on N unknowns. We also impose the moment that equates the price-dividend ratio on REITs in the data to the price-dividend ratio on REITs in the model, where the latter is calculated at the sum of the price-dividend ratios of the first 3600 REIT dividend strips. This estimation also imposes the bond yield regularity conditions and good deal bounds (30) and (31).

Steps 4-8 are analogous to step 3. Given the estimated market prices of risk before the step, we free up one additional element (row) of Λ_0 (Λ_1) to match the expected excess return and the price-dividend ratio on an additional equity risk factor. The expected excess return moment tells us exactly which elements must and must not equal zero.

Step 9 The last step of the estimation takes as starting values all market prices of risk obtained thus far. It re-estimates those 116 parameters by matching jointly all moments. This step is non-trivial because bond moments are predicted by and correlated with equity variables, and vice versa. Since inflation, GDP growth, level, and slope risk also affect all equity prices and expected equity returns, those equity moments help identify the market prices of risk in the first four rows (bond block). And vice versa, bond moments can help identify prices of risk in the equity block. Finally, regularity conditions on yields and good deal bounds can be satisfied in the stock and bond estimations, but violated in the joint estimation at the starting parameters. The regularity conditions are satisfied at the final parameter estimates and the maximum Sharpe ratio falls substantially.

D Shock-exposure and Shock-price Elasticities

Borovička and Hansen (2014) provide a *dynamic value decomposition*, the asset pricing counterparts to impulse response functions, which let a researcher study how a shock to an asset's cash-flow today affects future cash-flow dynamics as well as the prices of risk that pertain to these future cash-flows. What results is a set of *shock-exposure elasticities* that measure the quantities of risk resulting from an initial impulse at various investment horizons, and a set of *shock-price elasticities* that measure how much the investor needs to be compensated currently for each unit of future risk exposure at those various investment horizons. We now apply their analysis to our VAR setting.

D.1 Derivation

Recall that the underlying state vector dynamics are described by:

$$z_{t+1} = \Psi z_t + \Sigma^{\frac{1}{2}} \varepsilon_{t+1}$$

The log cash-flow growth rates on stocks, REITs, and infrastructure stocks are described implicitly by the VAR since it contains both log returns and log price-dividend ratios for each of these assets. The log real dividend growth rate on an asset $i \in \{m, reit, infra\}$ is given by:

$$\log(D_{t+1}^i) - \log(D_t^i) = \Delta d_{t+1}^i = A_0^i + A_1^i z_t + A_2^i \varepsilon_{t+1},$$

where $A_0^i = \mu_m$, $A_1 = e'_{divi} \Psi$, and $A_2^i = e'_{divi} \Sigma^{\frac{1}{2}}$.

Denote the cash-flow process $Y_t = D_t$. Its increments in logs can be written as:

$$y_{t+1} - y_t = \Gamma_0 + \Gamma_1 z_t + z_t' \Gamma_3 z_t + \Psi_0 \varepsilon_{t+1} + z_t' \Psi_1 \varepsilon_{t+1} \quad (33)$$

with coefficients $\Gamma_0 = A_0^i$, $\Gamma_1 = A_1^i$, $\Gamma_3 = 0$, $\Psi_0 = A_2^i$, and $\Psi_1 = 0$.

The one-period log real SDF, which is the log change in the real pricing kernel S_t , is a quadratic function of the state:

$$\log(S_{t+1}) - \log(S_t) = m_{t+1} = B_0 + B_1 z_t + B_2 \varepsilon_{t+1} + z_t' B_3 z_t + z_t' B_4 \varepsilon_{t+1}$$

where $B_0 = -y_0^{\$}(1) + \pi_0 - \frac{1}{2} \Lambda_0' \Lambda_0$, $B_1 = -e'_{yn} + e'_{\pi} \Psi - \Lambda_0' \Lambda_1$, $B_2 = -\Lambda_0' + e'_{\pi} \Sigma^{\frac{1}{2}}$, $B_3 = -\frac{1}{2} \Lambda_1' \Lambda_1$, and $B_4 = -\Lambda_1'$.

We are interested in the product $Y_t = S_t D_t$. Its increments in logs can be written as in equation (33),

with coefficients $\Gamma_0 = A_0^i + B_0$, $\Gamma_1 = A_1^i + B_1$, $\Gamma_3 = B_3$, $\Psi_0 = A_2^i + B_2$, and $\Psi_1 = B_4$.

Starting from a state $z_0 = z$ at time 0, consider a shock at time 1 to a linear combination of state variables, $\alpha'_h \varepsilon_1$. The shock elasticity $\epsilon(z, t)$ quantifies the date- t impact:

$$\epsilon(z, t) = \alpha'_h (I - 2\tilde{\Psi}_{2,t})^{-1} (\tilde{\Psi}'_{0,t} + \tilde{\Psi}'_{1,t}z)$$

where the $\tilde{\Psi}$ matrices solve the recursions

$$\begin{aligned}\tilde{\Psi}_{0,j+1} &= \hat{\Gamma}_{1,j}\Sigma^{1/2} + \Psi_0 \\ \tilde{\Psi}_{1,j+1} &= 2\Psi'\hat{\Gamma}_{3,j}\Sigma^{1/2} + \Psi_1 \\ \tilde{\Psi}_{2,j+1} &= (\Sigma^{1/2})'\hat{\Gamma}_{3,j}\Sigma^{1/2}\end{aligned}$$

The $\hat{\Gamma}$ and $\tilde{\Gamma}$ coefficients follow the recursions:

$$\begin{aligned}\tilde{\Gamma}_{0,j+1} &= \hat{\Gamma}_{0,j} + \Gamma_0 \\ \tilde{\Gamma}_{1,j+1} &= \hat{\Gamma}_{1,j}\Psi + \Gamma_1 \\ \tilde{\Gamma}_{3,j+1} &= \Psi'\hat{\Gamma}_{3,j}\Psi + \Gamma_3 \\ \hat{\Gamma}_{0,j+1} &= \tilde{\Gamma}_{0,j+1} - \frac{1}{2}\log(|I - 2\tilde{\Psi}_{2,j+1}|) + \frac{1}{2}\tilde{\Psi}_{0,j+1}(I - 2\tilde{\Psi}_{2,j+1})^{-1}\tilde{\Psi}'_{0,j+1} \\ \hat{\Gamma}_{1,j+1} &= \tilde{\Gamma}_{1,j+1} + \tilde{\Psi}_{0,j+1}(I - 2\tilde{\Psi}_{2,j+1})^{-1}\tilde{\Psi}'_{1,j+1} \\ \hat{\Gamma}_{3,j+1} &= \tilde{\Gamma}_{3,j+1} + \frac{1}{2}\tilde{\Psi}_{1,j+1}(I - 2\tilde{\Psi}_{2,j+1})^{-1}\tilde{\Psi}'_{1,j+1}\end{aligned}$$

starting from $\hat{\Gamma}_{0,0} = 0$, $\hat{\Gamma}_{1,0} = 0_{1 \times N}$, $\hat{\Gamma}_{2,0} = 0_{N \times N}$, and where I is the $N \times N$ identity matrix.

Let $\epsilon_g(z, t)$ be the shock-exposure elasticity (cash-flows $Y = D$) and $\epsilon_{sg}(z, t)$ the shock-value elasticity, then the shock-price elasticity $\epsilon_p(z, t)$ is given by

$$\epsilon_p(z, t) = \epsilon_g(z, t) - \epsilon_{sg}(z, t).$$

In an exponentially affine framework like ours, the shock price elasticity can also directly be derived by setting $Y_t = S_t^{-1}$ or $y_{t+1} - y_t = -m_{t+1}$, with coefficients in equation (33) equal to $\Gamma_0 = -B_0$, $\Gamma_1 = -B_1$, $\Gamma_3 = -B_3$, $\Psi_0 = -B_2$, and $\Psi_1 = -B_4$.

The shock-price elasticity quantifies implied market compensation for horizon-specific risk exposures. In our case, these risk compensations are extracted from a rich menu of observed asset prices matched by a

reduced form model, rather than by constructing a structural asset pricing model. The horizon-dependent risk prices are the multi-period impulse responses for the cumulative stochastic discount factor process.


D.2 Results

Figure D1 plots the shock-exposure elasticities of the seven dividend growth processes -on the market (blue), small stocks (cyan), growth (magenta), value (yellow), REITs (red), infrastructure (green), and natural resource stocks (black)- to a one-standard deviation shock to inflation (top left), real GDP growth (top second), the short rate (top third), the slope factor (top fourth), the price-dividend ratio on the market (top fifth), the dividend growth rate on the market (top right), the dividend growth rate on small stocks (bottom left), the dividend growth rate on growth stocks, on value stocks, on REITs, on infrastructure, and the dividend growth rate on natural resource stocks (bottom right). The shock exposure elasticities are essentially impulse-responses to the original (i.e., non-orthogonalized) VAR innovations. They describe properties of the VAR, not of the asset pricing model. Since our private equity cash-flows are linear combinations of these dividends, the PE cash flow exposures to the VAR shocks will be linear combinations of the plotted shock exposure elasticities of these dividends.

There is interesting heterogeneity in the cash flow exposures of the seven risky assets to the VAR shocks. For example, the top left panel shows that REIT and to a lesser extent small stock cash flows increase in response to a positive inflation shock, while the dividend growth responses for the aggregate stock market and especially for growth stocks and value stocks (at longer horizons) are negative. This points to the inflation hedging potential of REITS. The second panel shows that REIT and small stock dividend growth responds positively to a GDP growth shock, while cash flow growth on value stocks respond negatively. The response of REIT and small stock cash flows to higher interest rates is positive in the short run. REIT cash flows are rents which can be adjusted upwards when rates increase, which typically occurs in a strong economy (see the GDP panel). Value stocks have strong negative interest rate exposure at intermediate horizons. The market dividend growth shows a substantial positive response to a steepening yield curve. Slope exposure is even larger for value stocks and substantial for growth stocks. When the slope steepens, which tends to happen during recessions, NR cash flows fall. The top right panel shows that the dividend growth shock to the market is nearly permanent. Many other stock indices' cash flows also go up, with small and growth stocks showing the greatest exposure to the market cash flow shock, and NR a consistently negative exposure. A positive innovation in REIT cash flows is associated with a positive innovation in small-stock dividends. This makes sense given the small stock-like behavior of REITs (Van Nieuwerburgh, 2019). A positive shock to infrastructure cash flows has long-lasting positive effects on REIT cash flows. the other

Figure D1. Shock Exposure Elasticities

The figure plots the shock-exposure elasticities of dividend growth on the market, dividend growth of REITs, dividend growth of infrastructure stocks, dividend growth of small stocks, dividend growth of growth stocks, and dividend growth of natural resource stocks to a one-standard deviation shock to inflation (top left), real GDP growth (top second), the short rate (top third), the slope factor, the price-dividend ratio on the market, dividend growth rate on the market (top right), dividend growth rate on small stocks (bottom left), dividend growth on growth stocks (bottom second), dividend growth on value stocks, on REITs, on infrastructure stocks, and dividend growth of natural resource stocks (bottom right). The shock exposures are to the (non-orthogonalized) VAR innovations.

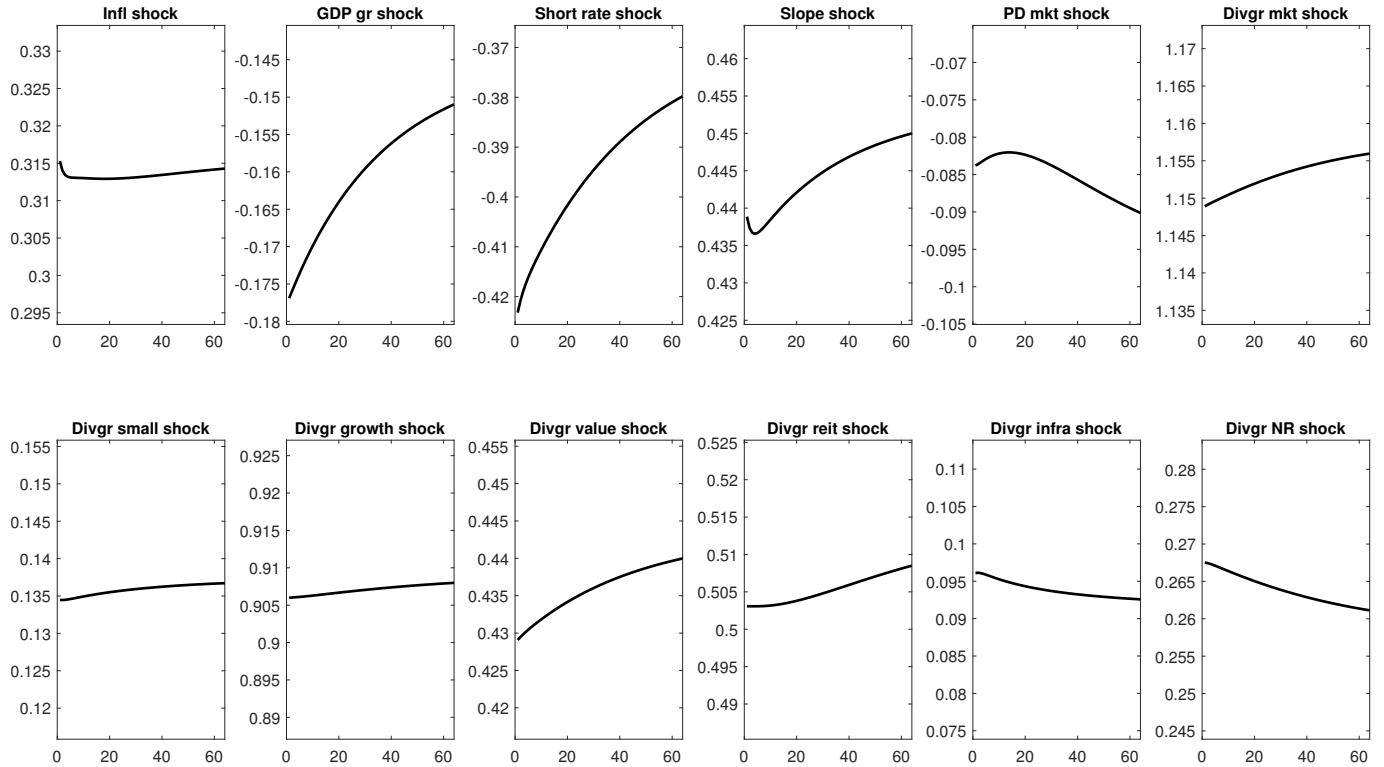


five cash flow series (bottom middle panel). Positive growth stock cash flow innovations have strongly persistent positive effects for market and small stock cash flows. NR cash flows have longer-run spillovers on REITs.

Figure D2 plots the shock-price elasticities to a one-standard deviation shock to each of the same (non-orthogonalized) VAR innovations. Shock price elasticities are properties of the SDF process, and therefore depend on the estimated market price of risk parameters. They quantify the compensation investors demand for horizon-dependent risk exposures. The price of inflation risk is positive, indicating that increases in inflation are good states of the world. This is a common finding in papers that include the recent sample when bonds became deflation hedges rather than inflation bets (Campbell, Sunderam, and Viceira, 2017).

Figure D2. Shock Price Elasticities

The figure plots the shock-price elasticities to a one-standard deviation shock to the inflation factor (top left), real GDP growth (top second), short rate (top third), the slope factor, the price-dividend ratio on the market, the dividend growth rate on the market (top right), the dividend growth rate on small stocks (bottom left), the dividend growth rate on growth stocks, on value stocks on REITs, on infrastructure stocks, and on natural resources stocks (bottom right). The shocks whose risk prices are plotted are the (non-orthogonalized) VAR innovations $\Sigma^{\frac{1}{2}}\epsilon$.



GDP growth risk is priced negatively, but less so at longer horizons. Rate level risk is negatively priced, consistent with standard results in the term structure literature that consider high interest rate periods bad states of the world. The price of level risk becomes less negative at longer horizons. The price of slope risk is positive, consistent with the findings in [Kojien, Lustig, and Van Nieuwerburgh \(2017\)](#). All cash-flow shocks in the bottom six panels naturally have positive risk prices since increases in cash-flow growth are good shocks to the representative investor. The highest risk price is associated with shocks to the aggregate stock market, followed by infrastructure shocks, then growth shocks. Compensation for exposure to the various dividend shocks varies little with the horizon.

E Additional Results

Figure E1. Cash-Flows by Vintage, Alternate Categories




Figure E2. Factor Exposure for other Categories by Fund Horizon

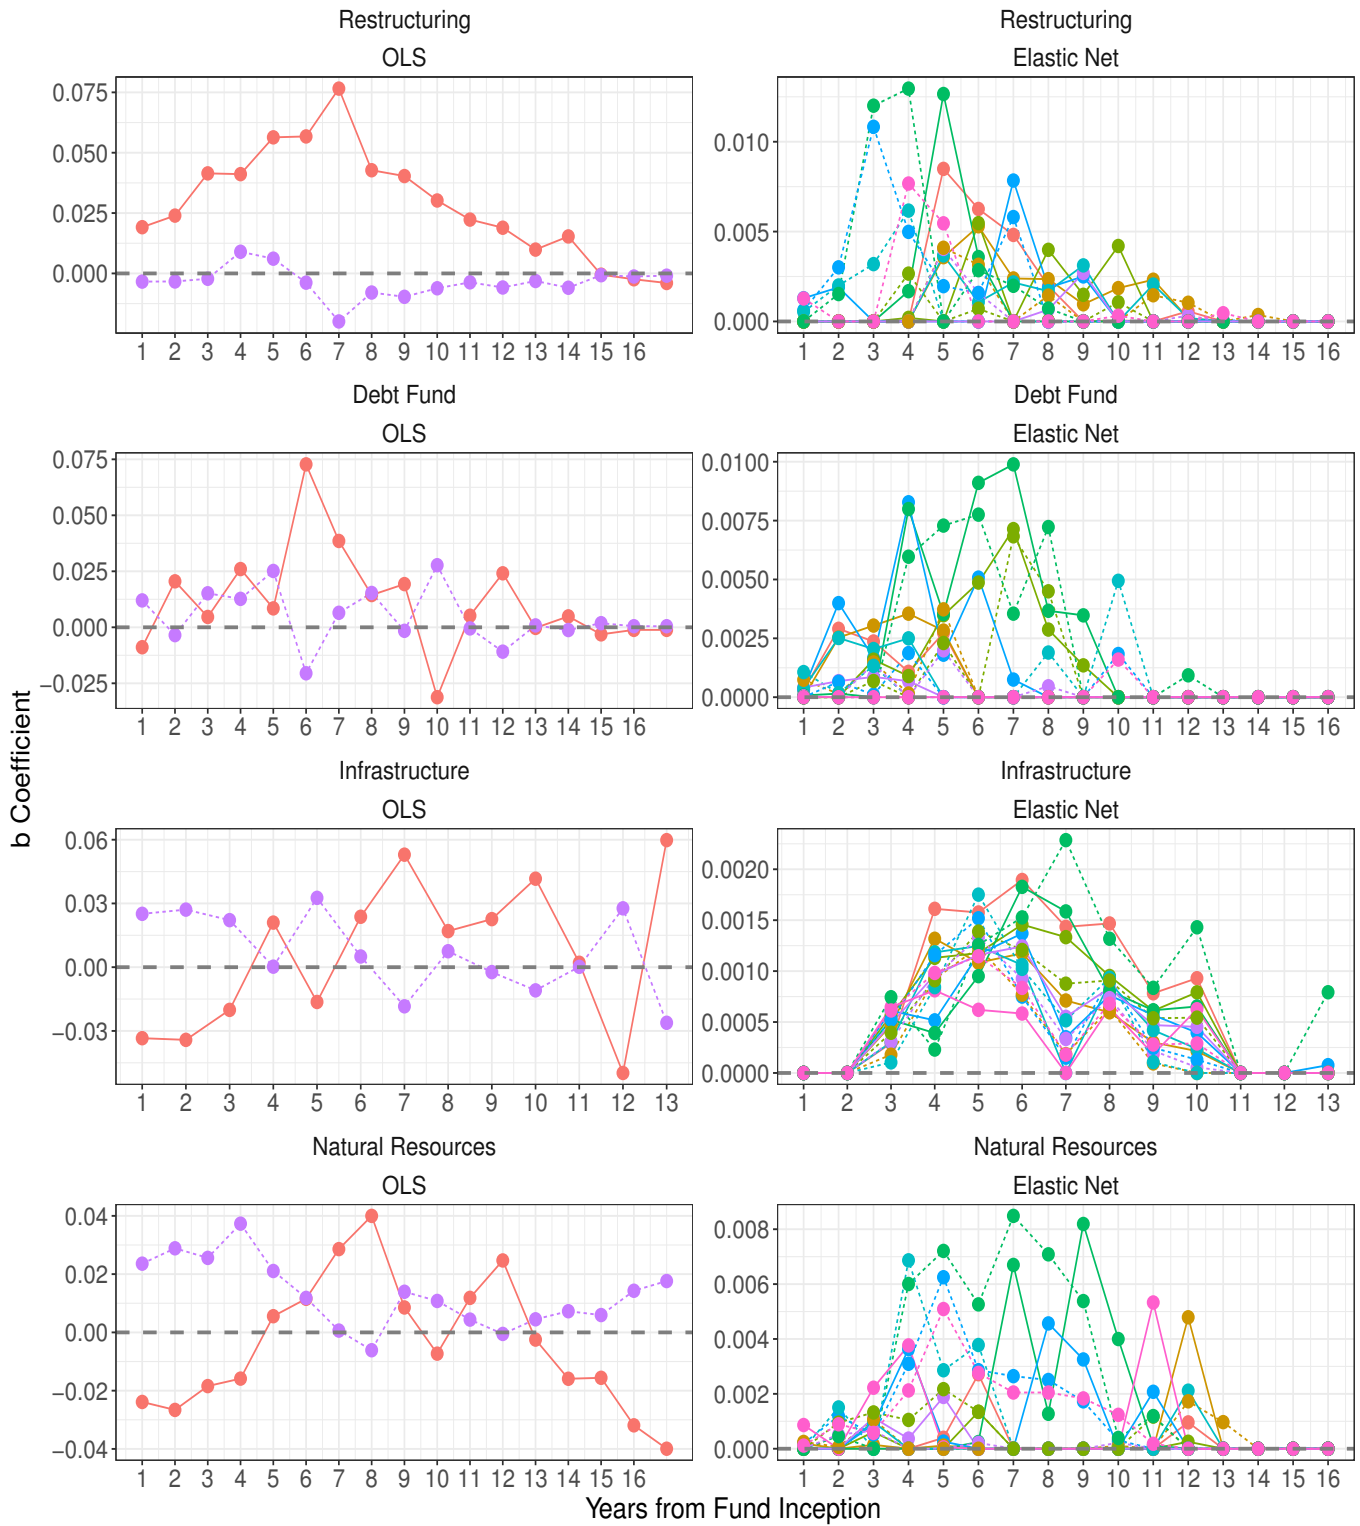
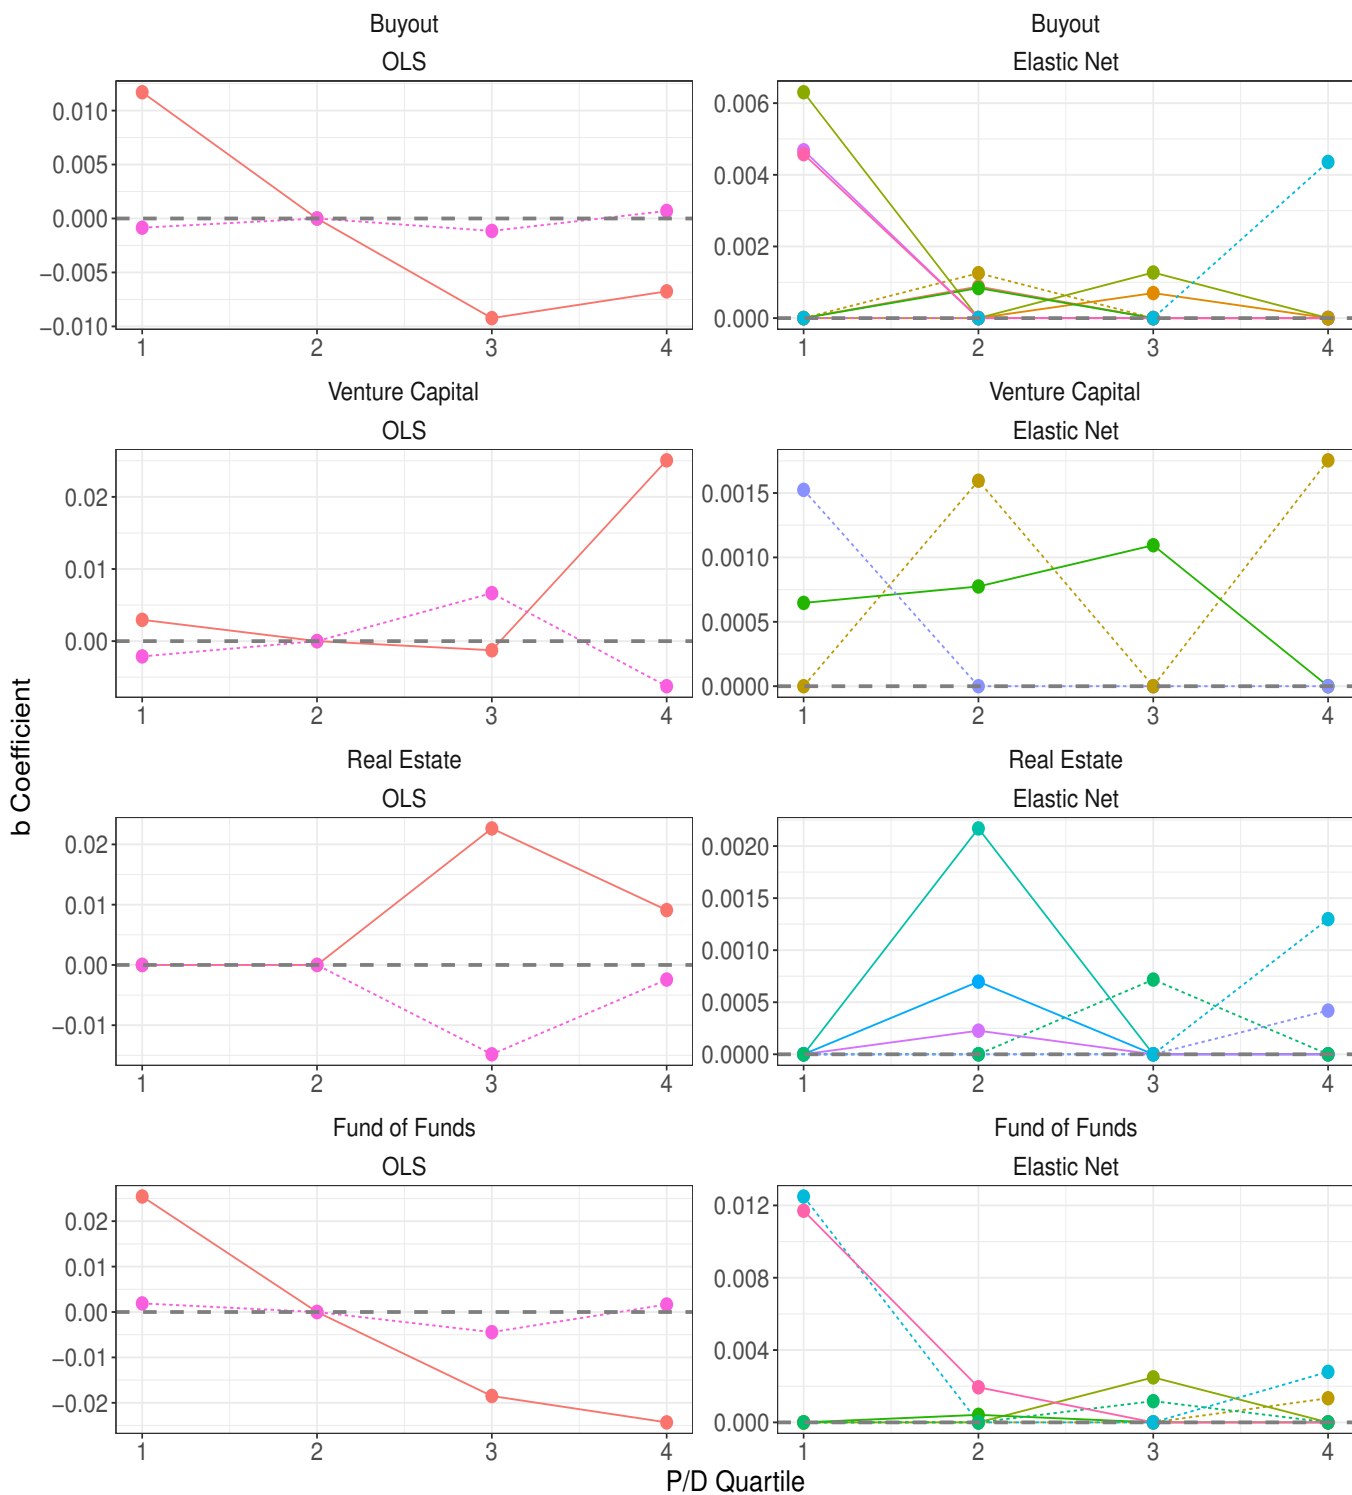


Figure E3. Factor Exposure by P/D Quartile



- Dividend ··· Gain
- Bond ● Infra
- Growth ● Natural Resources (NR)
- Growth Gain ● NR Gain
- REIT ● Small Gain
- REIT Gain ● Stock
- Small ● Stock Gain
- Value

Figure E4. Factor Exposure by P/D Quartile for Additional categories

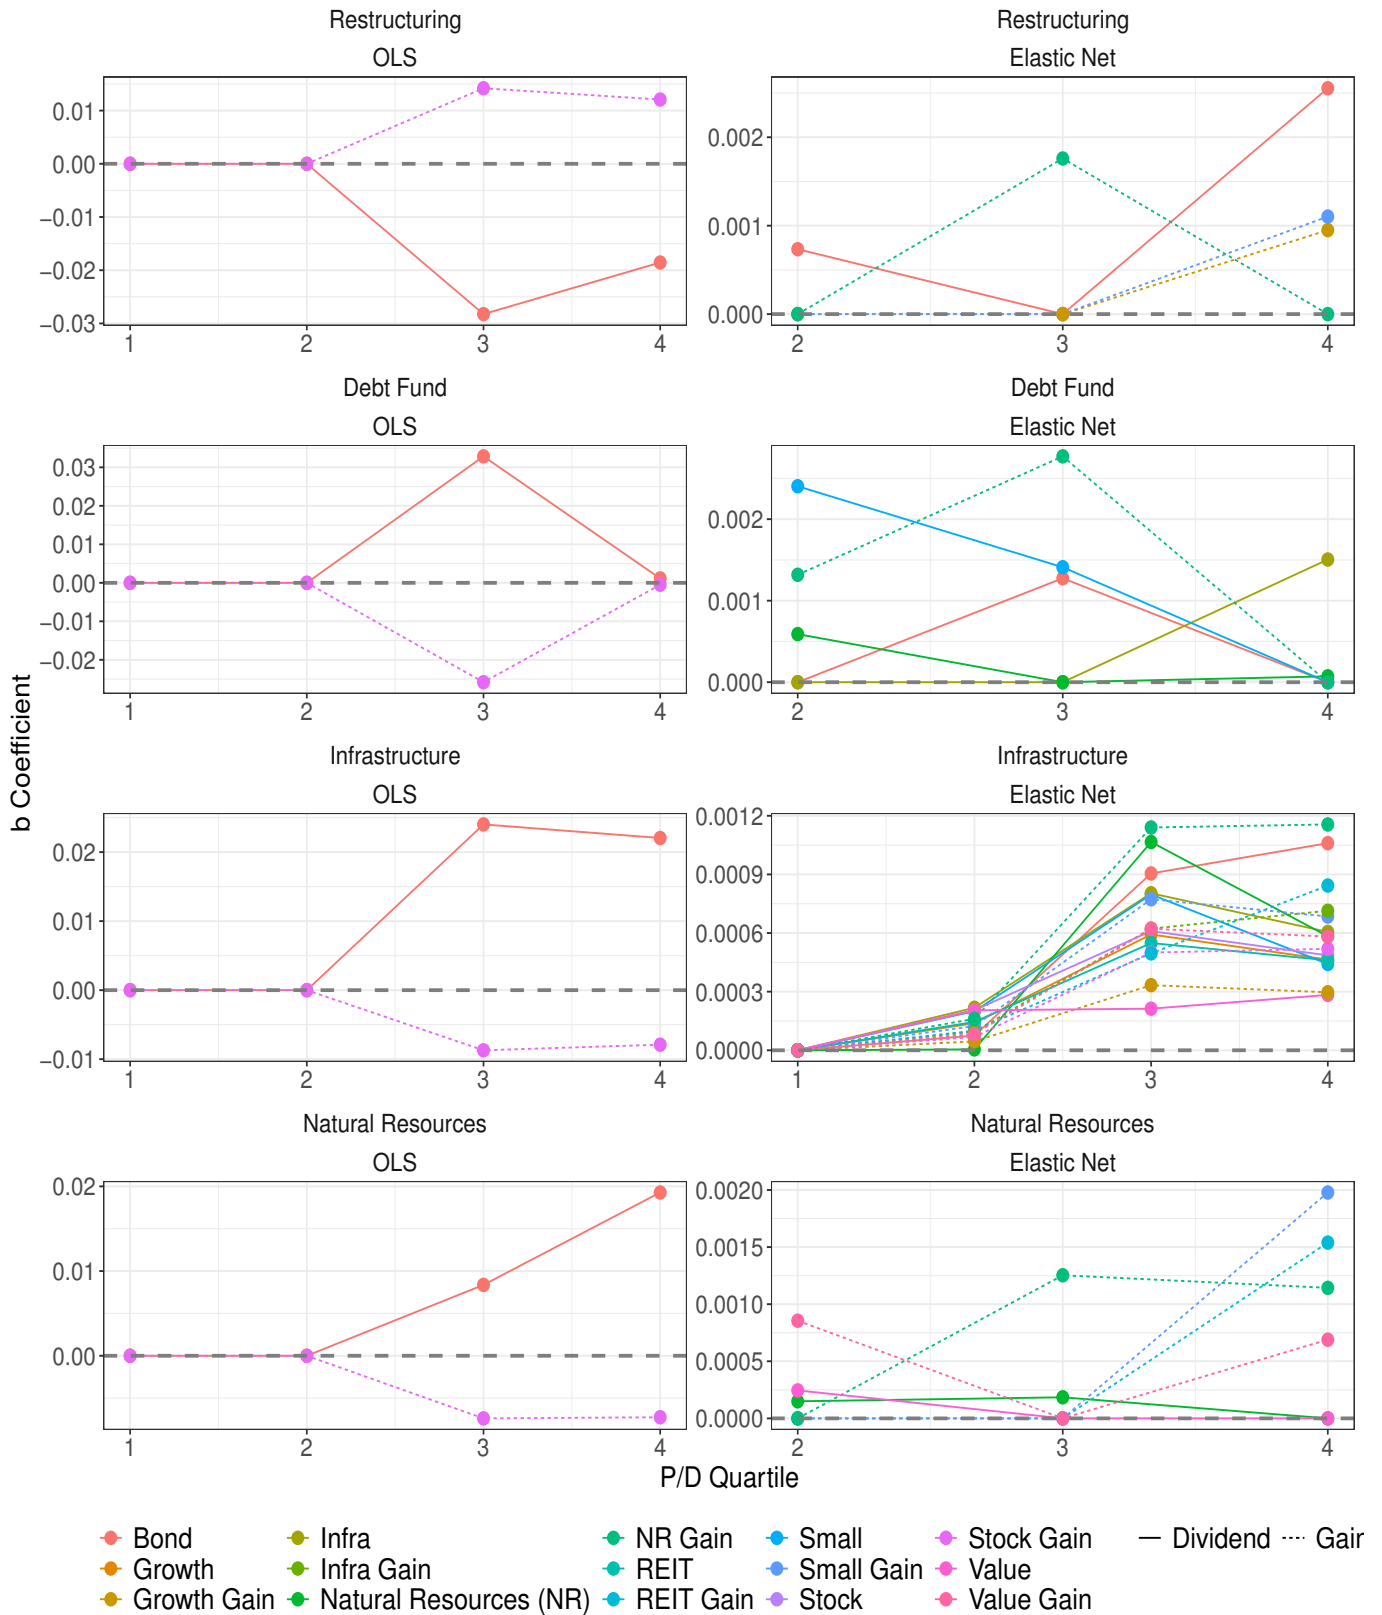




Figure E5. Expected Returns by Vintage for Additional Categories



2-Factor

Elastic Net



Panel A: Restructuring



Panel B: Debt Fund



Panel D: Infrastructure



Panel D: Natural Resources





Figure E6. Profits Over Time for Additional Categories

Average Fund-level profit by Vintage

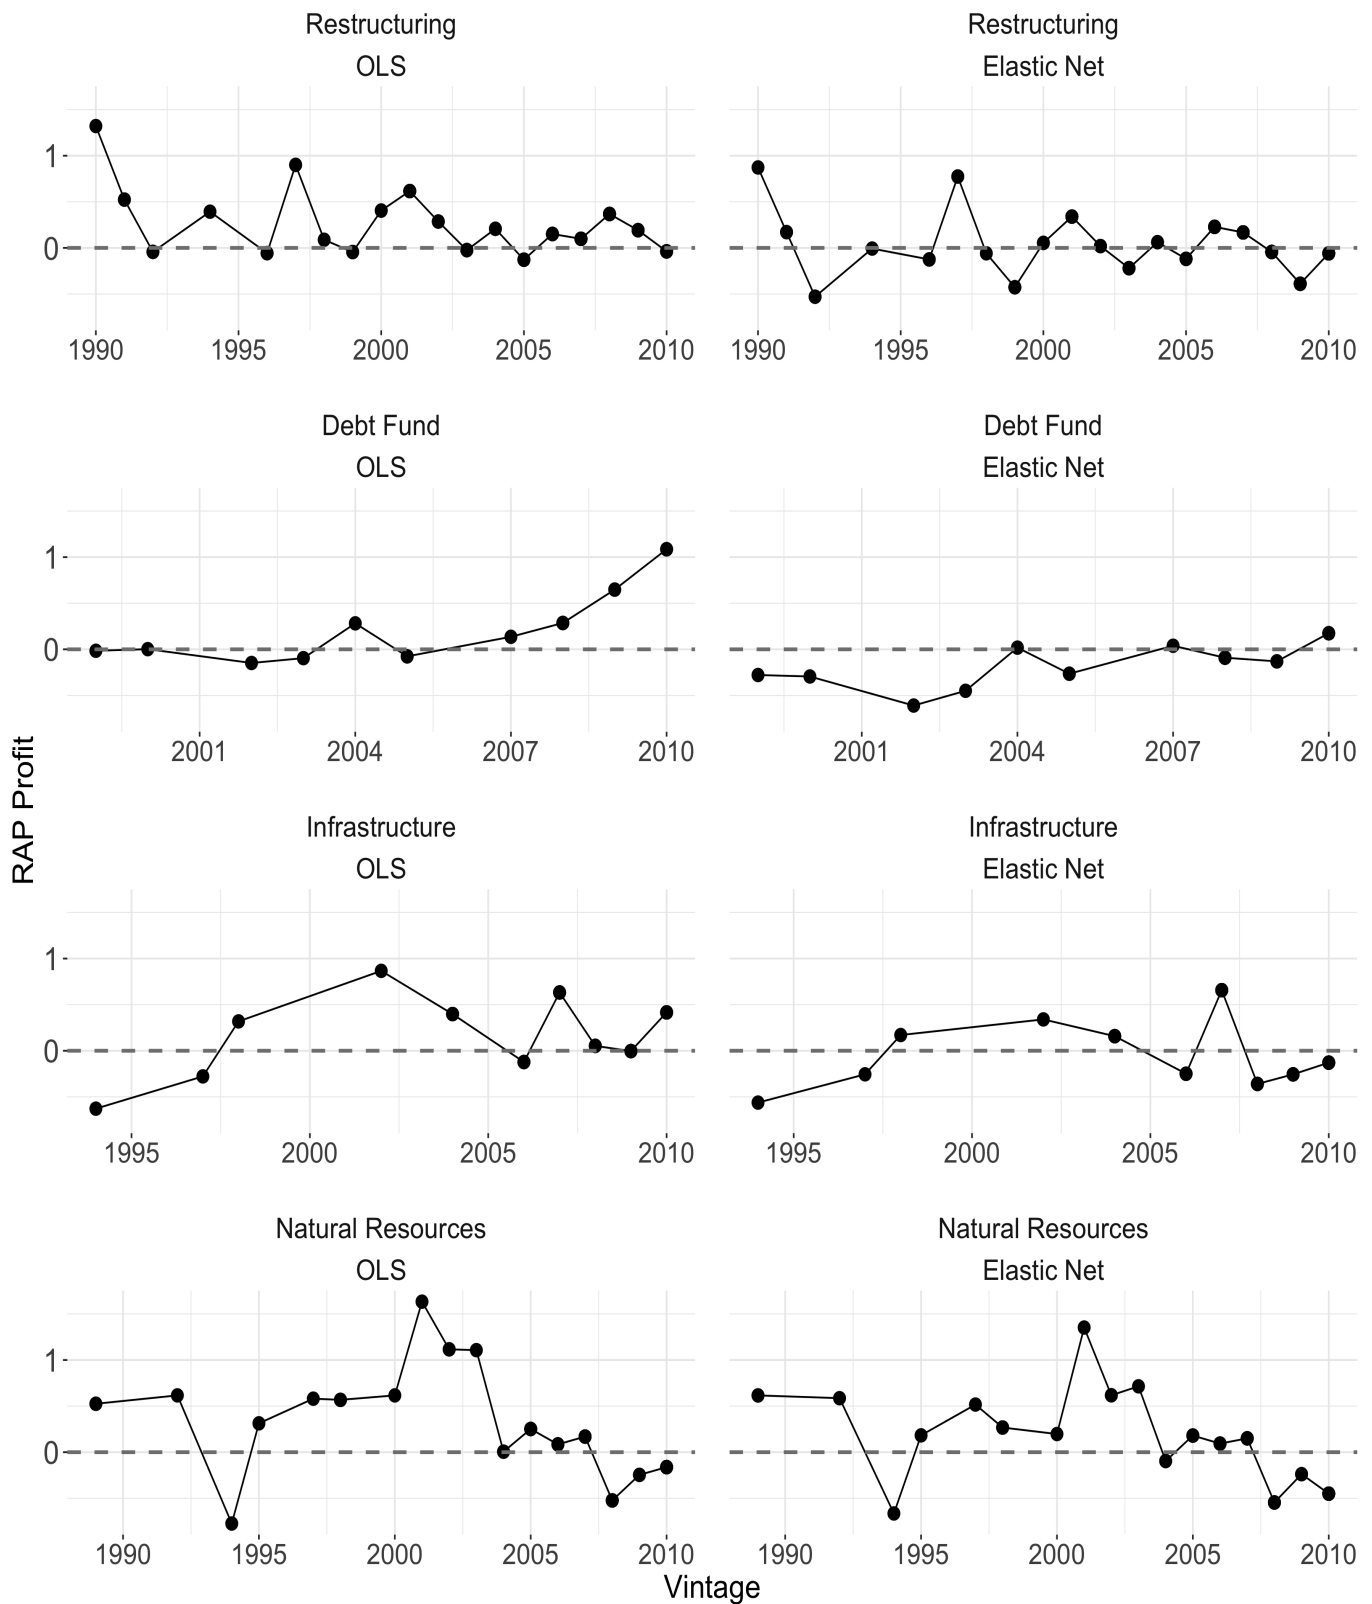




Figure E7. Elastic Net Model Comparison for Additional Categories


Panel A: Restructuring




Panel B: Debt Fund



Panel C: Infrastructure



Panel D: Natural Resources



F Validation on Public Equities

We conduct two exercises on public equities that serve as further model validation. The first exercise is an internal validity exercise, which also serves to illustrate how the typical cash flow life-cycle pattern of PE funds affects our inference. The second exercise is an external validity exercise that prices a cross-section of stocks not included in the asset pricing.

F.A REIT Exercise


Each quarter from 1974.Q1 until 2003.Q4, we form two new “REIT hedge funds.” Each hedge fund operates for 64 quarters. The first hedge fund’s strategy is to buy \$1 of the REIT index in the vintage-origination quarter, collect the dividends on the index, and sell 25% of the portfolio in each of quarters 61 through 64. Dividends in those last four quarters are proportional to the remaining position in the index. The second hedge fund’s strategy is to liquidate a share of the portfolio in each quarter, where the share follows a parabolic (inverted-U) shape that begins and ends at zero. Peak liquidation occurs in quarter 32. Dividends are received on whatever portion of the portfolio is not yet liquidated. In both strategies, the hedge fund is fully liquidated at the end of quarter 64.

The left panel of Figure F1 plots the cash flows on the first strategy for one particular vintage (1989.Q1); the right panel plots the second strategy for the same vintage. The total cash flows (“total CF”) in each case are the sum of dividends on the index (“cash div”) and liquidation proceeds (“liq div”). Since stocks and stock indices pay low dividend yields—even REITS which pay a lot of dividends for tax reasons,—most of the cash flows from the first strategy occur upon liquidation at the end of the life of the hedge fund. But the dividend cash flows are also substantial and growing over the life of the fund as the companies in the index grow their dividends. The second REIT hedge fund, in contrast, has declining dividends as the fund gradually liquidates, and substantial liq-

liquidation cash-flows in the middle of the life cycle when liquidations peak. The share of cash flows that comes from dividends falls over the life of the fund, and is lower on average as in the first strategy. In general, the first strategy offers much higher total cash flows than the second strategy since prices and dividends tend to rise over time in our sample period.

Figure F1. REIT Hedge Fund Cash Flows

The left panel plots the cash flows on a REIT hedge fund strategy that buys \$1 of the REIT index, collects dividends for 60 quarters, and liquidates the fund in the last four quarters. The right panel plots the cash flows on a second REIT hedge fund strategy that buys \$1 of the REIT index, gradually liquidates the shares in parabolic fashion over the next 64 quarters, and receives the dividends on the remaining portion of the portfolio in each quarter. The vintage plotted is 1989.Q1 in both panels.



We then estimate a three-factor model ($K = 3$), which consists of zero-coupon bonds, REIT dividend strips, and REIT capital gain strips. As we do with the actual PE data, we pool all hedge fund observations. In this case there are 120 funds, one for each vintage. We estimate $H = 64$ age effects for each factor, b_h^k , and we allow each of the three sets of age effects to be shifted by their own vintage effect, a_t^k . The vintage shifters depend on the quartile of the price-dividend ratio of the stock market at fund formation, the same assumption as we make in the PE fund estimation. There are a total of $K \times (H + 3) = 201$ exposures to be estimated.

Figure F2 shows the estimated exposures for the second hedge fund strategy which the parabolic liquidation pattern. This is the case that most closely mimics the pattern in the

PE data. The middle-left panel shows that the exposures to the REIT index dividend strips pick up the dividend cash flows of the hedge fund strategy. The declining shape in the exposures mirrors that in the hedge fund dividend cash flows shown in the right panel of Figure F1. The bottom-left panel of Figure F2 shows the REIT hedge fund exposure to REIT index gain strips. These gain strip exposures closely track the hump-shaped liquidation cash flows in the right panel of Figure F1. This illustrates the separate roles of dividend strip and gain strip factors. The former picks up dividend cash flows earned by the hedge fund, while the latter picks up asset dispositions. The top-left panel shows the bond exposures. These are generally small after 10 quarters. The vintage effects in the right column are generally estimated to be very small, which is correct given that the investment strategy does not depend on the pd ratio at fund origination. The 3-factor model explains 99.4% of the variation in the REIT hedge fund cash flows. In short, the replication works well, and the factors we would expect to pick up the dividend and the liquidation components of the REIT hedge fund cash flows do so appropriately and fully.

The only inference issue is the estimated non-zero bond exposures in the first 10 quarters. Those should be zero and reflected instead in higher REIT dividend strip exposures at ages 1-10 quarters. The reason the estimation confuses these two is because there is insufficient variation in short-horizon REIT dividend growth across vintages. Since short-horizon dividend growth is hard to distinguish from a constant, the estimation assigns those exposures to the bond rather than to the REIT dividend strip. This problem diminishes quickly with the horizon. This does not constitute a big problem for our PE analysis since PE funds in the data distribute very little cash flows in the first 10 quarters after the first capital call.

The average RAP for the REIT funds is 0.019 with a standard deviation (across vintages) of 0.032. In sum, RAP is close to zero for all vintages. Again, this confirms that the approach works. The small positive RAP comes from the inference issue with the bond exposures. The replicating portfolio interprets the early hedge fund cash flows as risk-

free rather than exposed to dividend strips. Since zero coupon bond prices are higher than REIT dividend strip prices, this assignment error increases the value of the replicating portfolio slightly above \$1. Again, we do not expect this to be much of an issue for PE funds since there is not much cash going out early in the life cycle.

Figure F2. REIT Hedge Fund 2 Exposures

The figure plots the estimated exposures to the bond factor (top row), the REIT dividend strip factor (middle row) and the REIT gain factor (bottom row). The left panels plot the estimated age effects b_h for cash flow horizons $h = 1, \dots, H$. The right panels plot the estimated vintage effects, associated with quartiles of the price-dividend ratio distribution on the stock market. The second-quartile effect is normalized to zero. The figure is for the second REIT hedge fund strategy that liquidates shares in parabolic fashion over the life of the fund.

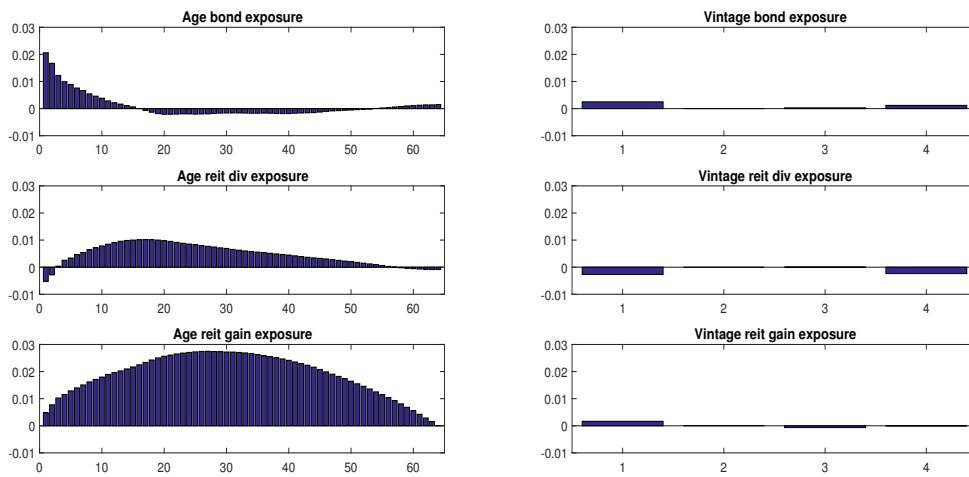
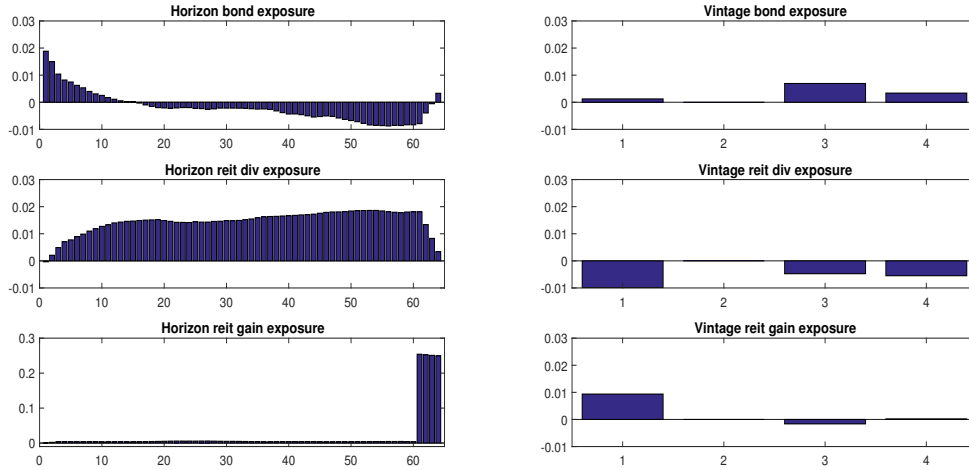


Figure F3 shows the estimated exposures for the first hedge fund strategy which does not liquidate until the end. Again, the approach “works.” The gain strip exposure picks up the large liquidations in the last four quarters of the life of the hedge fund. The dividend strip exposures pick up the slightly rising dividends of the hedge fund. The bond effects are small except for the first 10 quarters. The vintage effects are small. The R^2 is 100%. The average RAP of 0.027 is close to zero with small standard deviation of 0.053. This exercise shows that our approach is able to recover starkly different patterns in cash flows exemplified by the two hedge fund strategies.

Figure F3. REIT Hedge Fund 1 Exposures

The figure plots the estimated exposures to the bond factor (top row), the REIT dividend strip factor (middle row) and the REIT gain factor (bottom row). The left panels plot the estimated age effects b_h for cash flow horizons $h = 1, \dots, H$. The right panels plot the estimated vintage effects, associated with quartiles of the price-dividend ratio distribution on the stock market. The second-quartile effect is normalized to zero. The figure is for the first REIT hedge fund strategy that does not liquidate any shares until after quarter 60.



F.B Operating Profitability and Investment Portfolios

next, we use our method to price a set of portfolios whose risk factors are not included in the state space. We choose the 25 portfolios sorted by operating profitability (OP) and investment (INV). Fama and French (2015) show that these portfolios are not spanned by the market, size, and value portfolios. This lead them to propose a new 5-factor model with OP and INV factors added to the traditional 3-factor model. Having formed time series for the dividends and prices for each of the 25 portfolios and applied the same special dividend smoothing as for the other price and dividend series, we form a “hedge fund” that buys one of the portfolios, earns its cash dividends, and gradually liquidates it at the prevailing market prices in parabolic fashion, just like in the REIT validation exercise. After 64 quarters, the fund is fully liquidated. Total fund cash flows are the sum of cash dividends and liquidation proceeds. We start one such fund in each of the 120 quarters from 1974.Q1 until 2003.Q4.

We then estimate a five-factor model that contains the following factors: bonds, the

Table FI. RAP for 25 Operating Profitability and Investment Portfolios

	Panel A: Risk-Adjusted Profit					
	OP1	OP2	OP3	OP4	OP5	OP5-OP1
INV1	-0.15	0.16	-0.12	-0.23	0.02	0.17
INV2	-0.25	0.01	0.38	0.34	0.25	0.5
INV3	0.14	-0.2	0.08	0.08	0.15	0.01
INV4	-0.12	0.09	-0.05	0.36	0.55	0.67
INV5	-0.56	-0.29	-0.2	-0.04	0.44	1.00
INV5-INV1	-0.42	-0.45	-0.07	0.19	0.41	0.83
	Panel B: FF Three-Factor Alpha					
	OP1	OP2	OP3	OP4	OP5	OP5-OP1
INV1	3.37	3.56	5.61	8.30	8.86	5.49
INV2	1.54	4.41	6.25	6.99	6.96	5.42
INV3	3.56	2.24	4.54	4.46	6.47	2.91
INV4	2.52	4.45	4.44	5.02	6.92	4.41
INV5	-1.52	2.36	2.84	3.94	7.50	9.02
INV5-INV1	-4.88	-1.20	-2.77	-4.36	-1.36	3.52

aggregate stock market dividend and capital gain strips, and value stock dividend and gain strips. We find similar results for two seven-factor models that add either growth-stock strips or small-stock strips. These results are not reported for brevity. First, we find high R^2 values for all 25 portfolios, ranging from 82.1% to 97.5%. Second, we find that the 5-factor model generates a near-zero RAP of 0.034 on average across portfolios. Third, Panel A of Table FI shows that the model leaves a large positive RAP for an investment that goes long in high-OP stocks with high investment-minus-low investment and that goes short in low-OP stocks with high-INV-minus-low-INV. This is consistent with the 3-factor FF model leaving an alpha in the INV-OP-sorted portfolios. Highly profitable firms that invest a lot earn a high RAP, but low-profitability firms that invest a lot destroy value.

Reassuringly, there are clear patterns in the exposures of these 25 portfolios to the factors. For example, Panel A of Table FII shows the exposures to the value gain strip factor. Since we estimate an exposure at each age, we average across age for presenta-

Table FII. Exposure to Value Factor for 25 Operating Profitability and Investment Portfolios

	Panel A: Exposure to Value Gain Strips					
	OP1	OP2	OP3	OP4	OP5	OP5-OP1
INV1	0.45	0.79	0.83	1.31	0.41	-0.03
INV2	0.11	0.42	0.59	0.51	-0.01	-0.12
INV3	0.84	0.05	0.22	0.26	0.1	-0.74
INV4	0.49	0.87	0.18	0.48	-0.01	-0.5
INV5	-0.06	0.34	-0.06	-0.27	-0.57	-0.51
INV5-INV1	-0.51	-0.45	-0.88	-1.58	-0.98	-0.47
	Panel B: FF Three-Factor HML-Beta					
	OP1	OP2	OP3	OP4	OP5	OP5-OP1
INV1	0.19	0.44	0.28	0.41	0.23	0.04
INV2	0.23	0.36	0.33	0.18	0.08	-0.15
INV3	0.26	0.28	0.27	0.11	0.08	-0.17
INV4	0.21	0.21	0.13	0.05	-0.18	-0.39
INV5	-0.25	-0.13	-0.19	-0.24	-0.4	-0.14
INV5-INV1	-0.45	-0.56	-0.47	-0.65	-0.63	-0.18

tion purposes. As the last two indicates, high-investment firms tend to have negative exposures to the value gain factor while low-investment firms tend to have high exposure. Indeed, low-investment firms have low asset growth and are more like value firms while high-investment firms are more like growth firms. Hence, the value factor exposures are consistent with intuition. High OP-stocks tend to have lower exposure to value gain strips than low-OP shocks, especially among high-INV stocks. In sum, the value factor exposure clearly captures some of the cross-sectional variation in the 25 OP-INV portfolios.

These patterns in RAP and value exposures, obtained using our estimation methodology, match the patterns obtained from estimating a standard 3-factor model on monthly returns. Panel B of Table FI reports three-factor alphas for the 25 OP-INV portfolios; monthly alphas have been multiplied by 1200. Just like the RAPs, FF-alphas tend to be positive along the OP dimension and negative along the INV dimension. The corner portfolio has a large positive alpha, indicating that a high-investment strategy among high-OP

firms outperforms.

Panel B of Table [FII](#) shows the HML-beta from the three-factor model for the 25 OP-INV portfolios. Just like in our method, the standard 3-factor model finds consistently negative value betas for the high-INV minus low-INV portfolio strategy. It also finds (weaker) negative value exposure of a high-OP minus low-OP strategy, especially among higher-INV firms. Again, our method arrives at the same finding.

This exercise illustrates that our approach results in sensible result for listed securities that are not included in the state space.

G Cross-Validation

In this appendix we provide details on our cross-validation procedure and show robustness on the hyper-parameters governing the Elastic Net. First, recall that our Elastic Net estimation equation is given by:

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

We set $\lambda_0 = \infty$, to restrict our attention to positive exposure coefficients, corresponding to a long-only portfolio. We have two other hyper-parameters: α and λ . The parameter α is an elastic-net penalty. It is bounded by $\alpha = 1$, which corresponds to the special case of a lasso specification, and by $\alpha = 0$, which corresponds to the special case of a ridge regression penalty. The parameter λ determines the total penalty amount.

We tune these two parameters through a cross-validation procedure in which we divide our sample into 10 non-overlapping folds of similar size. We do so for each fund category separately. We leave out one fold and fit the model by choosing α and λ to minimize the mean-squared error (MSE) on the remaining 9 folds.

Table [G1](#) shows the MSE and fitted hyper-parameters for each of our fund categories, in comparison with the MSE for the 2-factor OLS model. Generally, we estimate α parameters close to 0—suggesting that the model prefers a ridge regression specification. Shrinkage is instead achieved through the λ parameter, the fitted value of which varies across categories.

In [Figure G1](#) we show the robustness of our key coefficients to different values of $\log(\lambda)$. We add up the cumulative fitted coefficients, or portfolio positions, across each of the factors across all horizons ($\sum_h \beta_{t,h}$). Each vertical slice represents a single model for a particular choice of λ , estimated at the optimal level of α . The vertical dotted line represents the optimal λ . Moving further to the right corresponds to a higher λ , implying a more constrained model. Across several categories, we find that the dominant exposure

across parameter values suggests a factor exposure consistent with the asset selection of the category. For instance, the dominant factor selection in Venture Capital across a broad range of factor choices is the growth gain strip exposure. It is a REIT gains strip in the Real Estate category, while natural resources (and natural resources gains) are dominant in the Natural Resources category. Natural resources and infrastructure tend to be the two largest components in the Infrastructure category. The pattern is more complex in some of the other categories, such as Buyout, Fund of Funds, and Restructuring. The asset selection in these categories tends to be broader than in the other categories we consider. Figure G2 performs a similar exercise varying our α coefficient (holding λ fixed at the optimal value). The message of these graphs is that our results result in similar exposure patterns are not very sensitive to hyper-parameter choices in the neighborhood of the one that minimizes the MSE.

Figure G3 shows heat maps illustrating how the MSE varies across combinations of α and λ for the four main PE categories. The optimal combination of α and λ is highlighted with a black box. This figure suggests that the models with lowest MSE inhabit a lower-center region of our state space—entailing low α and moderate λ . The graph also confirms that MSEs are fairly flat in the neighborhood of the chosen hyper-parameter, indicating robustness of the results.

Our model combines exposure estimates with asset strip prices to deliver risk-adjusted profits (RAPs). We investigate how different hyper-parameter choices translate into average RAP estimates in Figure G4. Again, we find that RAP estimates display little variation in a neighborhood around the optimal hyper-parameters. Far away from the optimal choices, on the far right of the graph, we find more negative RAPs. These correspond to highly sparse portfolios that turn out to be cheaper to replicate, and result in lower RAPs. Importantly, we observe that the entire region of the parameter space corresponds to negative profits in all four of the main PE categories. This suggests that our key finding—the negative risk-adjusted return of private equity funds—is highly robust and not sensitive

to the Elastic Net model specification.

Table GI. Model Fit and Baseline Hyper-Parameter Choices

Model	Buyout			VC			Real Estate			Fund of Funds		
	MSE	α	$\log(\lambda)$	MSE	α	$\log(\lambda)$	MSE	α	$\log(\lambda)$	MSE	α	$\log(\lambda)$
2-factor OLS	0.0067883	-	-	0.0208785	-	-	0.0053303	-	-	0.0019508	-	-
15 factor Elastic Net	0.0067486	0.064	-6.11	0.0209316	0.008	-6.21	0.0052946	0.064	-6.29	0.00194	0.008	-6.53
Model	Restructuring			Debt Fund			Infrastructure			Natural Resources		
	MSE	α	$\log(\lambda)$	MSE	α	$\log(\lambda)$	MSE	α	$\log(\lambda)$	MSE	α	$\log(\lambda)$
2-factor OLS	0.005513	-	-	0.0024353	-	-	0.0132752	-	-	0.0052202	-	-
15 factor Elastic Net	0.0054679	0	-4.85	0.0024268	0	-4.52	0.0133175	0	-1.12	0.0051197	0	-2.65

Figure G1. Variation in Coefficients Across Lambda

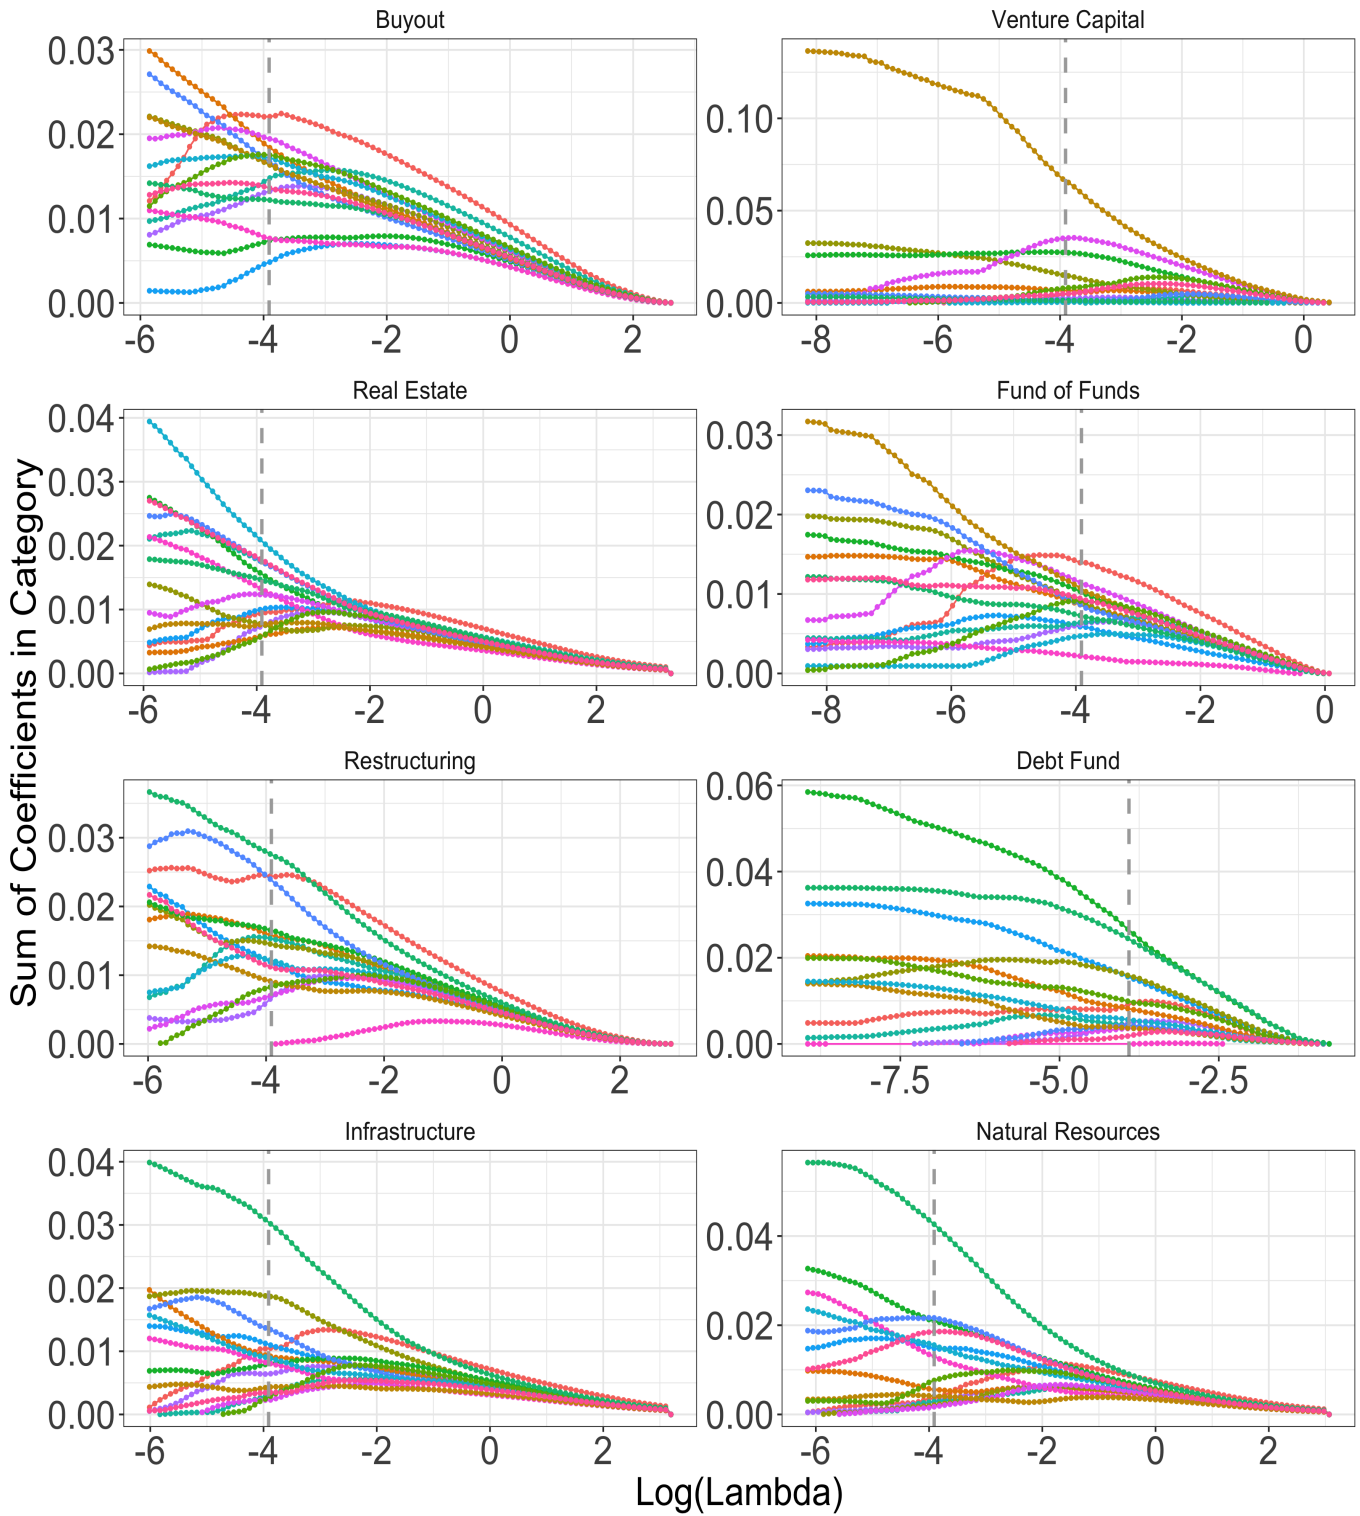


Figure G2. Variation in Coefficients Across Alpha

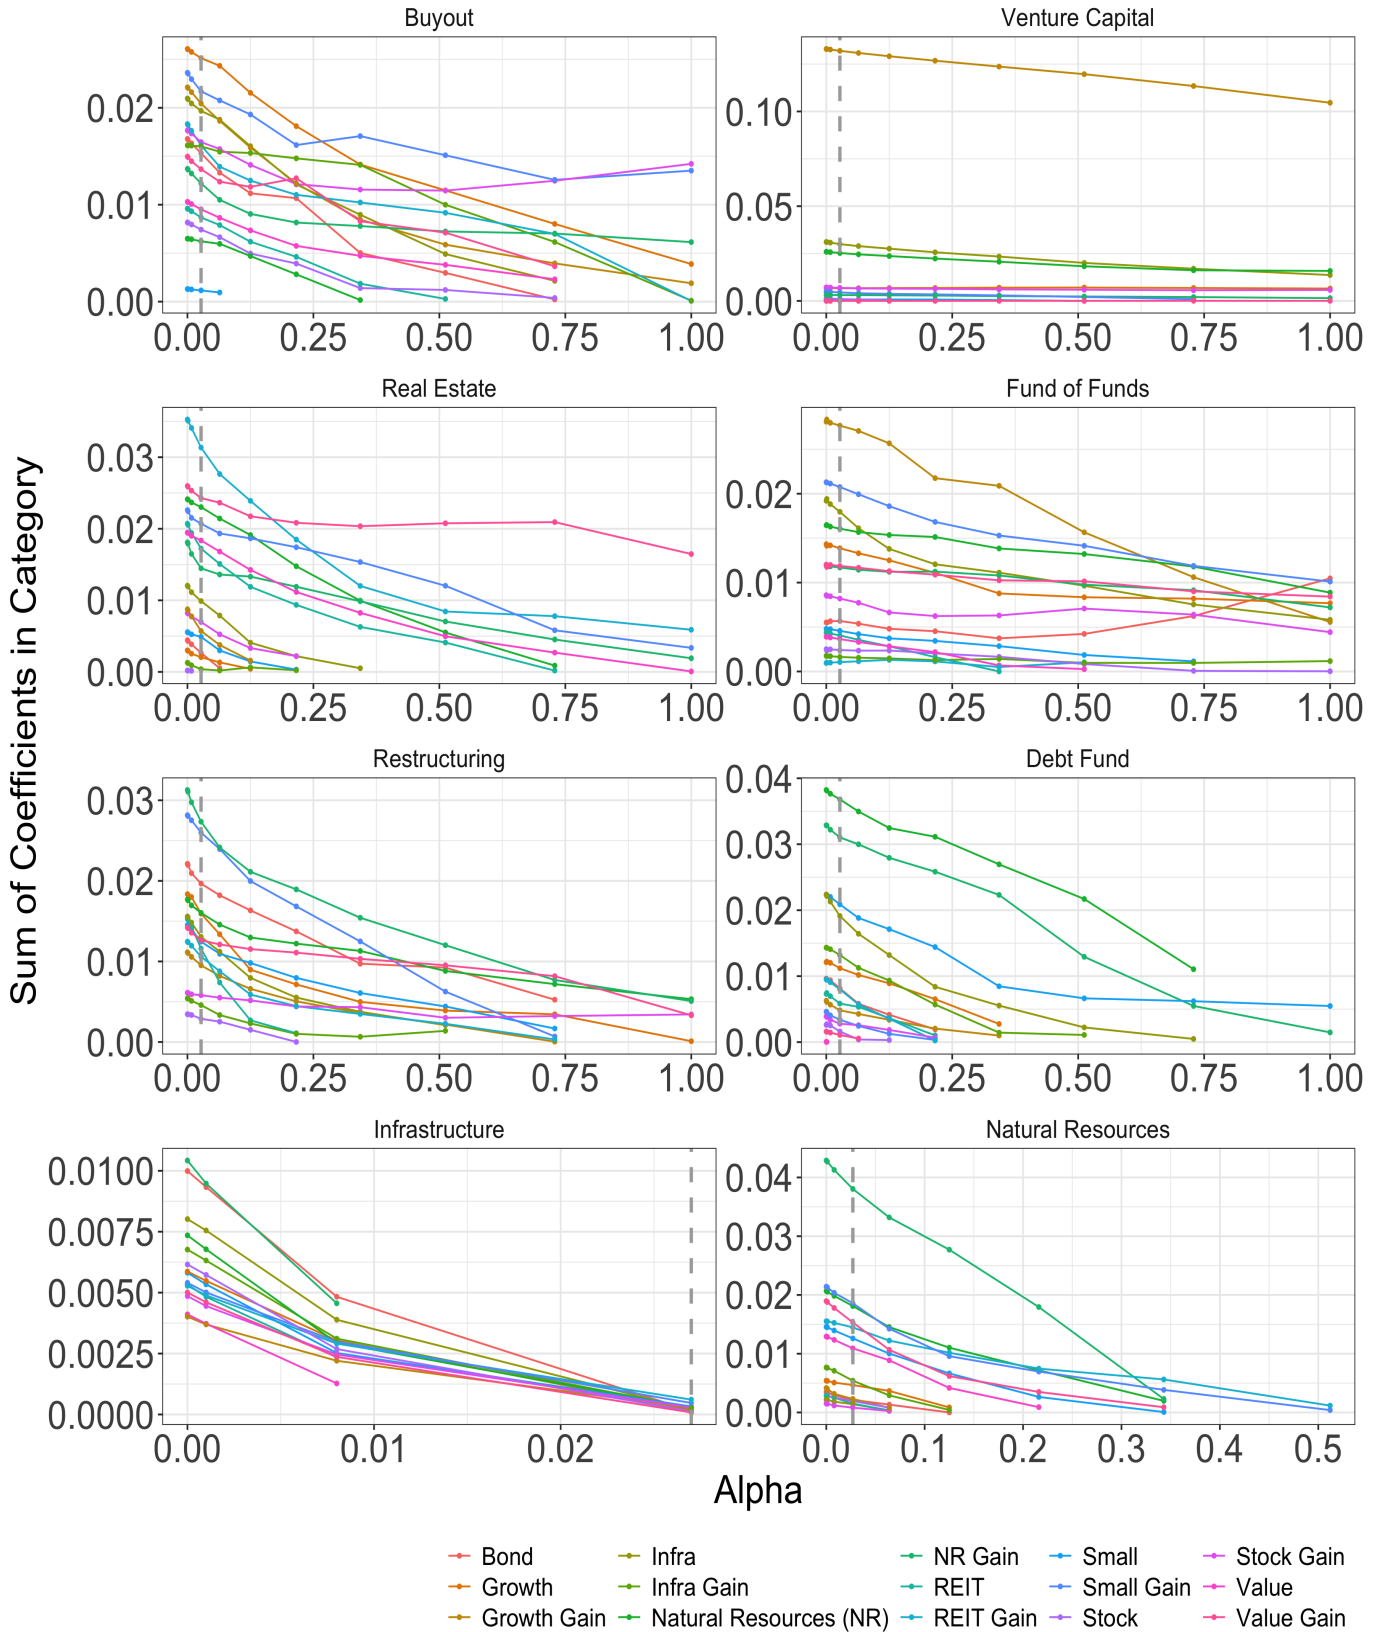


Figure G3. MSE Heatmaps

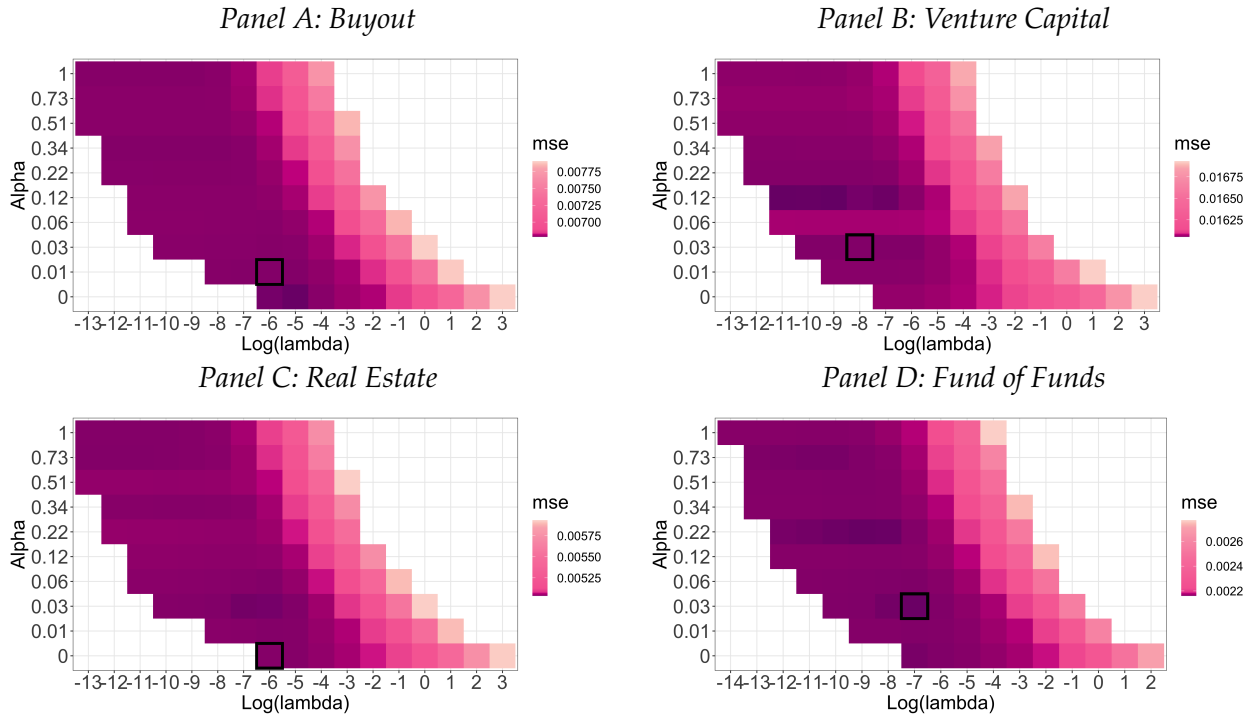
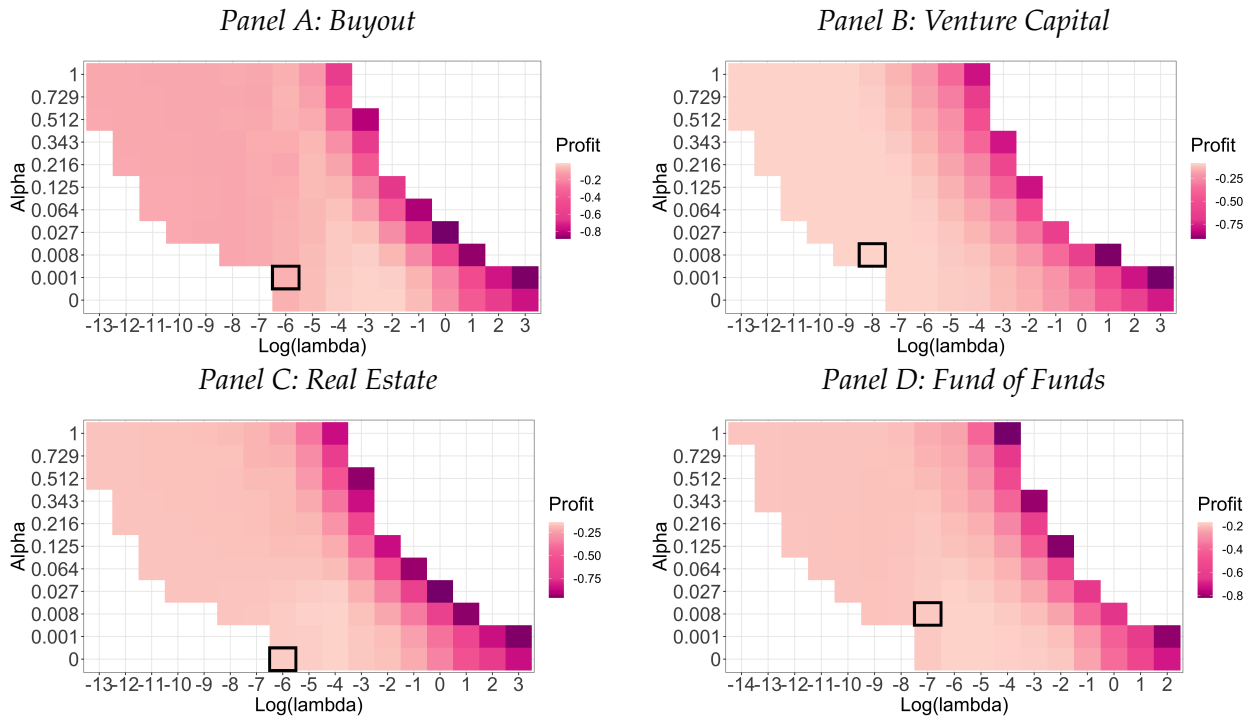


Figure G4. RAP Heatmaps



H Estimates on Burgiss Dataset

We repeat our analysis on the Burgiss private equity fund data set. As Table GII shows, this data set contains more funds than the Preqin data set. Figure G5 shows estimated age effects for Buyout, VC, RE, and IN are similar to our benchmark estimates. Table GII contains the average RAP estimates per fund category. The treatment of calls is the baseline one (NPV calls). The message is the same as in the main text. Adjusting for a large set of risk exposures substantially reduces risk-adjusted performance. Figure G6 plots the average RAP by vintage. It shows large positive RAP for VC for the mid-1990s vintages, and declining RAP in all categories on the later part of the sample.

Figure G5. Burgiss Factor Exposures

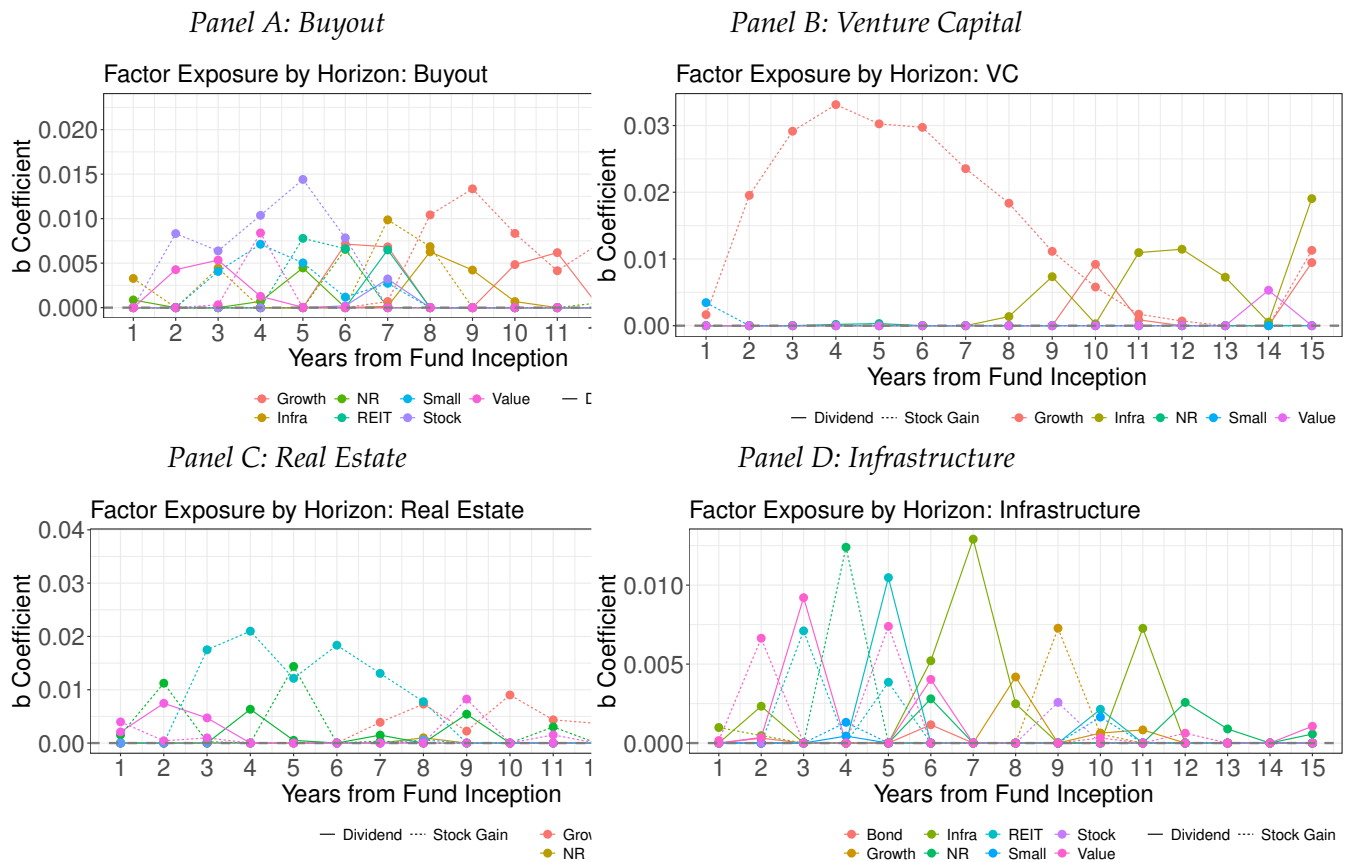



Table GII. Burgiss Model Comparison


<i>Panel A: Main Categories</i>			
	Buyout	VC	Real Estate
	Mean	Mean	Mean
TVPI	0.58	0.66	0.24
RAP 2-factor (NPV Calls)	0.34	-0.21	-0.32
RAP 15-factor (NPV Calls)	0.01	0.13	-0.06
Preqin Fund Count:	1,145	1,210	639
Burgiss Fund Count:	3,022	2,306	1,018
<i>Panel B: Additional Categories</i>			
	Restructuring	Debt Fund	Infrastructure
	Mean	Mean	Mean
TVPI	0.45	0.36	0.07
RAP 2-factor (NPV Calls)	0.43	0.21	-0.49
RAP 15-factor (NPV Calls)	0.08	-0.01	-0.16
Preqin Fund Count:	193	110	126
Burgiss Fund Count:	191	569	154

Figure G6. Burgiss Profit Over Time


Panel A: Buyout



Panel B: Venture Capital



Panel C: Real Estate



Panel D: Infrastructure

