Valuing Private Equity Investments Strip by Strip*

Arpit Gupta — NYU Stern

Stijn Van Nieuwerburgh — CBS, NBER, and CEPR

December 12, 2018

Abstract

We propose a new valuation method for private equity investments. First, we construct a cash-flow replicating portfolio for the private investment, using cash-flows on various listed equity and fixed income instruments. The second step values the replicating portfolio using a flexible asset pricing model that accurately prices the systematic risk in listed equity and fixed income instruments of different horizons. The method delivers a measure of the risk-adjusted profit earned on a PE investment as well as a time series for the expected return on PE funds. We apply the method to real estate, infrastructure, buyout, and venture capital funds, and find only modest average risk-adjusted profit but with substantial cross-sectional variation.

JEL codes: G24, G12

*Gupta: Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; agupta3@stern.nyu.edu; Tel: (212) 998-0305; 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 Ralph Koijen, Morten Sorensen, and Neng Wang for insightful discussions. This work is supported by the NYU Stern Infrastructure Initiative.

1 Introduction

Private equity investments have risen in importance over the past twenty-five years, relative to public equity. Private funds account for $4.7 trillion in assets under management, of which real estate funds comprise $800 billion (Preqin). Large institutional investors now allocate substantial fractions of their portfolios to such alternative investments. For example, the celebrated Yale University endowment has a portfolio weight of over 50% in alternative investments. Pension funds and sovereign wealth funds have also ramped up their allocations to alternatives. As the fraction of overall wealth that is held in the form of private investment grows, so does the importance of developing appropriate valuation methods. The non-traded nature of the assets and their infrequent cash-flows makes this a challenging problem.

As with any investment, the price of a private equity (PE) investment equals the present discounted value of its cash-flows. The general partner (GP, the fund manager) deploys the capital committed by the limited partners (LPs, investors). The risky projects may throw off some interim cash-flows that are distributed back to the LPs, but the bulk of the cash flows arise when the GP sells the projects, and distributes the proceeds net of fees (carry, promote) to the LPs. Our main question is how to adjust the profits the investor earns for the systematic risk inherent in the PE investment. Industry practice is to report the ratio of distributions to capital contributions and the internal rate of return. However, neither metric takes into account the riskiness of the cash-flows.

We propose a novel methodology that centers on the nature and the timing of cash-flow risk for PE investments. It proceeds in two steps. In a first step, we estimate the exposure of the PE fund's cash-flows to the cash-flows a set of liquid, publicly listed securities. In our application, we consider pay-offs to Treasury bonds, the stock market, and the publicly-traded real estate (REIT) and infrastructure markets; but the method easily accommodates additional publicly-traded factors. By "stripping" the sequence of PE cash-flows into individual cash-flows by horizon, and estimating the exposure to maturity-matched cash-flows on listed securities, our method decomposes the risk of a PE cash-flow into its different horizon components. We allow these horizon exposure profiles to depend on the vintage of the PE funds, and to differ by broad investment category.

In a second step, we use a flexible, no-arbitrage asset pricing model that prices the term structure of Treasury bonds, stocks, REITs, and infrastructure stocks. It postulates the main sources of systematic risk and estimates the prices of risk that the market assigns to these risk exposures. With those market price of risk estimates in hand, we can price strips, which are claims to a single risky cash-flow at each date and at each horizon in the bond, stock, REIT, and infrastructure market, inspired by Lettau and Wachter (2011) and van Binsbergen, Brandt, and Koijen (2012). We use the shock price elasticities of Borovička and Hansen (2014) to understand how risk prices change with horizon in the model. The model closely matches the time series of bond yields across maturities, stock price-dividend ratios, as well as stock and bond risk premia. It also matches the risk premium on short-maturity dividend futures from 2003-2014, calculated in the data by van Binsbergen, Hueskes, Koijen, and Vrugt (2013), and van Binsbergen and Koijen (2017), as well as the time series of the price-dividend ratio on 2-year cumulative dividend strips and the share of the overall stock market they represent from 1996-2009 as backed out from options data by van Binsbergen, Brandt, and Koijen (2012).

Combining the cash-flow replicating portfolio of strips obtained from the first step with the prices for these strips from the asset pricing model estimated in the second step, we obtain the fair price for the PE-replicating portfolio in each vintage and category. Each PE investment is scaled to represent \$1 of capital committed. Therefore, the replicating portfolio must also deploy \$1 of capital. Times in which the prices of stock and bond strips are high are times in which a \$1 replicating portfolio can buy a smaller quantity of stock and bond strips. All else equal, that will make it easier for the cash flows of PE funds started at that same time (i.e., of that vintage) to exceed those on the replicating portfolio. Of course, the assets the PE funds acquire may be more expensive in those times as well, so that out-performance is an empirical question. The risk-adjusted profit (RAP) of an individual PE fund is the net present value of the idiosyncratic PE cash flows, which are the difference between the realized cash flows on the PE fund and the realized cash flows on the replicating portfolio in that vintage-category. Under the null hypothesis that the asset pricing model is correct and that the average PE manager has no asset selection skill or no timing skill in when to deploy capital and harvest assets, the risk-adjusted profit is zero.

The model also delivers a time series for the expected return in each PE category. At each point in time , it reflects the systematic risk exposure of the PE funds in that category started at that point in time (vintage). That expected return can be broken down into its various horizon components, and, at each horizon, into its exposures to the various systematic risk factors. It helps us understand how the expected return on PE investments changes with the state of the economy. Our method can also be used to ask what the expected return is on all outstanding PE investments (vintages). It can be used to calculate the residual net asset value (NAV) of PE funds at each point during their life cycle. With expected returns, variance, and covariances with traded securities in hand, the approach opens the door to optimal portfolio analysis with alternative investments.

Now that trade in stock and bond strips has become feasible, our approach can be used by investors who have no access to PE funds (say, because they are too small) to mimic the return generating process of PE funds, possibly at lower cost. The argument is similar to that made for hedge-fund return replicating strategies. Conversely, in the absence of a full menu of dividend strips (say, REIT strips are not available), PE funds may be a trading strategy that provides an investor with access to exposure to the factor mimicking systematic cash flow risk (in say, real estate).

We use data from Preqin on all private equity funds with non-missing cash-flow information that were started between 1980 and 2017. Cash-flow data until March 2018 are used in the analysis. Our sample includes 4,221 funds in seven investment categories. The largest categories are Buyout, Real Estate, and Venture Capital. We are particularly interested in the categories Real Estate and Infrastructure, but report results for all categories. The main text reports results for these four categories and relegates the results for the other three to the appendix. Like in other papers in the literature, the PE data (here, taken from Preqin) are usually subject to some degree of selection bias.

One of our key findings is that the risk-adjusted profit to investors in PE funds is centered slightly above zero, about 3 cents per $1 invested, but with a large cross-sectional variation around the average. Using either a one or four-factor model, we find higher profitability for Real Estate and Buyout funds. Infrastructure, Debt Funds, and Restucturing funds also have positive (though small) risk-adjusted profits; we find little evidence for positive excess returns in Venture Capital. The intuition for this result is that PE cash-flows are well-spanned by public markets, and so their cash-flows are well-replicated ex-post through an estimated synthetic portfolio. However, we also find substantial cross-sectional variation in the profitability of funds. In the time-series, we generally find that more recent vintages perform worse than those in the 1990s. Despite the high apparent profitability of recent funds, our approach would suggest that these funds deliver little excess profit given their factor exposure and the performance of those publicly traded factors.

Related Literature This paper contributes to a large empirical literature on performance evaluation in private equity funds, such as Kaplan and Schoar (2005), Cochrane (2005), Korteweg and Sorensen (2017), Harris, Jenkinson, and Kaplan (2014), Phalippou and Gottschalg (2009), Robinson and Sensoy (2011), and many other papers cited above. Most of this literature focuses on Buyout and Venture Capital funds, though recent work in valuing privately-held real estate assets includes Peng (2016) and Sagi (2017). Ammar and Eling (2015) have studied infrastructure investments. This literature has found mixed results regarding PE fund outperformance and persistence of performance, depending on the data set and period in question. Our replicating portfolio approach results in a relatively low estimate of risk-adjusted profits for PE funds compared with the literature.

While performance evaluation in private equity is still often expressed as an internal rate of return or a ratio of distributions to capital committed, several important papers have incorporated risk into the analysis. The public market equivalent (PME) approach of Kaplan and Schoar (2005) compares the private equity investment to an appropriate public market benchmark (the aggregate stock market) with the same magnitude and timing of cash-flows. Sorensen and Jagannathan (2015) assess the PME approach from a SDF perspective. The closest antecedent to our paper is Korteweg and Nagel (2016), who propose a generalized PME approach that relaxes the assumption that the beta of PE funds to the stock market is one. This is particularly important in their application to venture capital funds. These approaches avoid making strong assumptions on the return-generating process of the PE fund, because they work directly with the cash-flows. See for example Cochrane (2005) and Korteweg and Sorensen (2010). Much of the literature assumes linear beta-pricing relationships, e.g., Ljungqvist and Richardson (2003), Driessen, Lin, and Phalippou (2012).

Several other papers have estimated beta exposures of PE funds with respect to the stock market, particularly for categories of buyout and venture capital. These include Gompers and Lerner (1997); Ewens, Jones, and Rhodes-Kropf (2013); Peng (2001); Woodward (2009). This literature has generally estimated stock market exposures of buyout funds above one, and even higher estimates for Venture Capital funds. Our work contributes to this literature estimating the risk exposure of PE funds by allowing for a flexible estimation approach across horizon and vintage; and in estimating fund exposures to a more expansive set of publicly listed securities. We also allow for our risk exposure estimates to differ by category, and examine a broader set of PE categories than typically examined in this literature. Our results for VC funds have implications for the returns on entrepreneurial activity (Moskowitz and Vissing-Jorgensen, 2002). Finally, we connect the systematic risk exposures of funds to a rich asset pricing model, which allows us to estimate risk-adjusted profits and time-varying expected returns.

Like Korteweg and Nagel (2016), we estimate a stochastic discount factor (SDF) from public securities. Our SDF contains additional risk factors, but more importantly, richer risk price dynamics. Those dynamics are central for generating realistic risk premia on bond and stock strips, which are the building blocks of our PE valuation method. The SDF model extends earlier work by Lustig, Van Nieuwerburgh, and Verdelhan (2013) who value a claim to aggregate consumption to help guide the construction of consumption-based asset pricing models. The asset pricing model combines a vector auto-regression model for the state variables as in Campbell (1991, 1993, 1996) with a no-arbitrage model for the (SDF) as in Duffie and Kan (1996); Dai and Singleton (2000); Ang and Piazzesi (2003). The SDF model needs to encompass the sources of aggregate risk that the investor has access to in public securities markets and that PE funds are exposed to. The question of performance evaluation then becomes whether, at the margin, PE funds add value to a portfolio that already contains these traded assets.

In complementary work, Ang, Chen, Goetzmann, and Phalippou (2017) 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 private equity returns, it does recover a time series of expected private equity returns. At each point in time, the asset pricing model can be used to revalue the replicating portfolio for the PE fund. Since it does not require a difficult Bayesian estimation step, our approach is 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 rest of the paper is organized as follows. Section 2 describes our methodology. Section 3 sets up and solves the asset pricing model. Section 4 presents the main results on the risk-adjusted profits and expected returns of PE funds. Section 5 concludes. The appendix provides derivations and additional results.

2 Methodology

Private equity investments are finite-horizon strategies, typically around ten to fifteen years in duration. Upon inception of the fund, the investor (LP) commits $1 to the fund manager (GP). The GP deploys that capital at his discretion, but typically within the first 2-4 years. Intermediate cash-flows may accrue from the operations of the assets, for example, net operating income from renting out an office building. Towards the end of the life of the fund (typically in years 5-10), the GP "harvests" the assets and distributes the proceeds to the investor after subtracting fees (including the carry or promote). These distribution cash-flows are risky, and understanding (and pricing) the nature of the risk in these cash-flows is the key question in this paper.

Denote this sequence of net-of-fees cash-flow distributions for fund i by $\{X^i_{t+h}\}^T_{h=0}$. Time t is the inception date of the fund, the vintage. The horizon h indicates the number of quarters since fund inception. We allocate all funds started in the same calendar year to the same vintage. The maximum horizon H is set to 60 quarters to allow for "zombie" funds that continue past their projected life span of 10 years. All cash flows observed after quarter H are allocated to quarter H. Monthly fund cash-flows are aggregated to the quarterly frequency. All PE cash-flows are reported for a \$1 investor commitment.

2.1 Two-Step Approach

In a first step, we find the cash-flow replicating portfolio of risk-free and risky securities for each of the PE cash-flow distributions. Let the $K \times 1$ vector $\boldsymbol{F}_{t+h}$ be the vector of cash flows on the securities in the replicating portfolio. The first element of $\boldsymbol{F}_{t+h}$ is a constant equal to 1. This is the cash-flow on a nominal zero-coupon U.S. Treasury bond that pays \$1 at time $t+h$. All other elements of $\boldsymbol{F}_{t+h}$ denote risky cash-flow realizations at time $t+h$. They are the payoffs on "zero coupon equity" or "dividend strips" (Lettau and Wachter, 2011; van Binsbergen, Brandt, and Koijen, 2012). They pay one (risky) cash-flow at time $t+h$ and nothing at any other date. We scale the risky dividend at $t+h$ by the cash flow at fund inception time t. For example, a risky cash-flow of $F_{t+h}(k) = \frac{D_{t+h}(k)}{D_t(k)} = 1.05$ implies that there was a 5% realized cash-flow growth rate between periods t and $t+h$ on the k^{th} asset in the replicating portfolio. This scaling gives the strips a "face value" around 1, comparable to that of the bond. It makes the exposures comparable in magnitude across assets.

Let the cash flow on the replicating portfolio be denoted by $\boldsymbol{\beta}^i_{t+h}\boldsymbol{F}_{t+h}$, where the $1 \times K$ vector $\boldsymbol{\beta}^i_{t+h}$ denotes the exposure of fund i to the assets in the replicating portfolio. We estimate the exposures from a projection of PE cash-flow at time $t+h$ on the cash-flows of the risk-free and risky strips:

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

where e denotes the idiosyncratic cash-flow component, orthogonal to F_{t+h}. The vector β^i_{t+h} describes how many units of each strip are in the replicating portfolio for the fund cash-flows. Equation (1) is estimated combining all funds in a given category, all vintages t, and all horizons h. We impose cross-equation restrictions, as explained below.

In the second step, we use our asset pricing model, spelled out in the next section, to price the zero coupon bond and equity strips. Denote the $K \times 1$ price vector for strips of horizon h by $\boldsymbol{P}_{t,h}$. The first element of this price vector is the price of the zero-coupon nominal bond which we denote by $P_{t,h}(1) = P^{\$}_{t,h}$. Let the one-period stochastic discount factor (SDF) be M_{t+1}, then the h-period SDF is:

$$M^h_{t+h} = \prod_{k=1}^{h} M_{t+k}.$$

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

$$\boldsymbol{P}_{t,h} = \mathbb{E}_t[M^h_{t+h}\boldsymbol{F}_{t+h}]$$

Budget Feasibility The first place where we use the asset pricing model is in making sure that the replicating portfolio for the PE fund is *budget feasible*. Since one dollar is available to buy a portfolio of bond and stock strips, no more than one dollar can be spent on the replicating portfolio. We impose that no less than one dollar should be spent. This is without loss of generality since the replicating portfolio can always invest in a one-period risk-free bond, which is equivalent to keeping cash. The replicating portfolio positions $\boldsymbol{\beta}^i_{t+h}$, estimated from equation (1), do not yet impose this budget feasibility requirement. Therefore, we define a vector of rescaled portfolio positions, $\boldsymbol{q}^i$, that costs exactly one dollar to buy

$$\boldsymbol{q}^i_{t,h} = \frac{\boldsymbol{\beta}^i_{t,h}}{\sum_{h=1}^{H} \boldsymbol{\beta}^i_{t,h}\boldsymbol{P}_{t,h}} \quad \Rightarrow \quad \sum_{h=1}^{H} \boldsymbol{q}^i_{t,h}\boldsymbol{P}_{t,h} = 1.$$

Since the strip prices change over time, each vintage has its own rescaling. With the budget feasible replicating portfolio in hand, we redefine the idiosyncratic component of fund cash-flows as v^i:

$$v^i_{t+h} = X^i_{t+h} - \boldsymbol{q}^i_{t,h}\boldsymbol{F}_{t+h}.$$

Under the joint null hypothesis of the asset pricing model and no fund outperformance, the expected present discounted value of fund cash-flow distributions must equal the $1 initially paid in by the investor:

$$\mathbb{E}_t\left[\sum_{h=1}^{H} M^h_{t+h} X^i_{t+h}\right] = \mathbb{E}_t\left[\sum_{h=1}^{H} M^h_{t+h}\boldsymbol{q}^i_{t,h}\boldsymbol{F}_{t+h}\right] = \sum_{h=1}^{H}\boldsymbol{q}^i_{t,h}\boldsymbol{P}_{t,h} = 1, \tag{2}$$

where the first equality follows from the fact that the idiosyncratic cash-flow component v^i is uncorrelated with the SDF since all priced cash-flow shocks are included in the vector F under the null hypothesis.

Expected Returns The second place where we use the asset pricing model is to calculate the expected return on the PE investment over the life of the investment. It equals the expected return on the replicating portfolio of strips:

$$\mathbb{E}_t\left[R^i\right] = \sum_{h=1}^{H}\sum_{k=1}^{K} \boldsymbol{w}^i_{t,h}(k)\mathbb{E}_t\left[\boldsymbol{R}_{t+h}(k)\right] \tag{3}$$

where $\boldsymbol{w}^i$ is a $1 \times KH$ vector of replicating portfolio weights with generic element $w^i_{t,h}(k) = q^i_{t,h}(k)P_{t,h}(k)$, in which $q^i_{t,h}(k)$ denotes the k^{th}element of the $1 \times K$ vector $q^i_{t,h}$ and $P_{t,h}(k)$ denotes the k^{th}element of the $K \times 1$ vector $P_{t,h}$. The $KH \times 1$ vector $\mathbb{E}_t[R]$ denotes the expected returns on the K traded asset strips at each horizon h. The asset pricing model provides the expected returns on these strips. Equation (3) decomposes the risk premium into compensation for exposure to the various risk factors, horizon by horizon.

Risk-Adjusted Profit Performance evaluation of PE funds requires quantifying the LP's profit on a particular PE investment, after taking into account its riskiness. 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 replicating portfolio, the PE cash-flows consist of one component that reflects compensation for risk and a risk-adjusted profit (RAP) equal to the discounted value of the idiosyncratic cash-flow component. The latter component for fund i in vintage t equals::

$$RAP_t^i = \sum_{h=1}^{H} P_{t,h}^{\$} v_{t+h}^i, \tag{4}$$

Since the idiosyncratic cash-flow components are orthogonal to the priced cash-flow shocks, they are to be discounted at the risk-free interest rate. The null hypothesis is that $\mathbb{E}[RAP_t^i] = 0$.

A fund with strong asset selection skills, which picks investment projects with payoffs superior to the cash flows on the traded assets, will have a positive risk-adjusted profit. Additionally, a fund with market timing skills, which invests at the right time (typically constrained to a fairly narrow investment period) and sells at the right time (within the harvesting period) will have positive risk-adjusted profit.[1] When calculating our profit measure, we exclude vintages after 2010, for which we are still missing a substantial fraction of the cash-flows.

We first discuss the relationship of our approach to other well-known approaches when valuing PE cash-flows. The rest of this section discusses implementation issues.

2.2 Connection to GPME and PME

Korteweg and Nagel (2016) define their *realized* GPME measure for fund i as:

$$\begin{aligned} GPME_t^i &= \sum_{h=0}^{H} M_{t+h}^h X_{t+h}^i \qquad (5) \\ &= \sum_{h=0}^{H} M_{t+h}^h \left\{ \boldsymbol{q}_{t,h}^i \boldsymbol{F}_{t+h} + v_{t+h}^i \right\} \\ &= RAP_t^i + \sum_{h=0}^{H} M_{t+h}^h \left\{ \boldsymbol{q}_{t,h}^i \left(\boldsymbol{F}_{t+h} - \mathbb{E}_t[\boldsymbol{F}_{t+h}] \right) \right\} \qquad (6) \end{aligned}$$

If the SDF model is correct, $\mathbb{E}_t[GPME^i] = 0$. The difficulty with computing (5) is that it contains the *realized* SDF which is highly volatile. In KN's implementation, the SDF is a function of only the market return: $\left(\Pi_{k=0}^h M_{t+k}\right) = \exp(0.088h - 2.65 \sum_{k=0}^h r_{t+k}^m)$. If the realized market return over a 10 year period is 100%, the realized SDF is 0.17. If the stock

[1] The fund's horizon is endogenous because it is correlated with the success of the fund. As noted by Korteweg and Nagel (2016), this endogeneity does not pose a problem as long as cash-flows are observed. "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."

return is 30%, the SDF is 1.08. Using our SDF, which is more volatile than one considered in KN, this approach leads to unrealistically low PE valuations, on average. Our risk-adjusted profit avoids using the realized SDF and instead relies on strip prices, which are *expectations* of SDFs multiplied by cash-flows.

A second difference between the two approaches is that the realized GPME can be high (low) because the factor payoffs F_{t+h} are unexpectedly high (low); the second term in equation (6). Our risk-adjusted profit measure does not credit the GP for this unexpected systematic cash-flow component. It removes a "factor timing" component of performance that is due to taking risk factor exposure. Like our approach, the simple PME does not credit the GP with factor timing.

Third, our approach credits the GP for "investment timing" skill while the GPME approach does not. Because it assumes that the replicating portfolio deploys the entire capital right away, a manager who successfully waits a few periods to invests will have a positive RAP. If the GP harvests at a more opportune time than the replicating portfolio, whose harvesting timing is determined by the average PE fund in that vintage and category, this also contributes to the RAP. The GPME as well as the simple PME approach do not credit the manager for investment timing because they assume that the replicating portfolio follows the observed sequence of PE capital calls and distributions.

Fourth, our approach accommodates heterogeneity in systematic risk exposure across PE funds that differ by vintage and category. In the standard PME approach, the market beta of each fund is trivially the same and equal to 1. In the GPME approach, PE funds are allowed to have a market beta that differs from 1, but the beta is the same for all funds. We allow for multiple risk factors, and the exposures differ for each vintage nd for each fund category. Appendix D provides more detail on the KN approach and more discussion on the points of differentiation.

2.3 Identifying and Estimating Cash-Flow Betas

The replicating portfolio must be rich enough that it spans all priced (aggregate) 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 beta profile leads to parameter proliferation and lack of identification. We impose cross-equation restrictions to aid identification.

One-factor Model To fix ideas, we start with a simple model in which all private equity cash-flows are assumed to only have interest rate risk. We refer to this as the one-factor model. The empirical model assumes that the cash-flows X of all funds i in the same category c (category superscripts are omitted for ease of notation) and vintage t have the same bond betas at each horizon h:

$$X_{t+h}^{i \in c} = \beta_{t,h}^{b} + e_{t+h}^{i} = q_{t,h}^{b} + v_{t+h}^{i} = a_t b_h + v_{t+h}^{i}. \tag{7}$$

We estimate (7) using a random effects model, category by category, with vintage times horizon effects as dependent variables. The estimation equally weights all funds. We scale the resulting bond exposures in each vintage so that they are budget-feasible. The resulting bond positions q_{t+h}^{b} are the product of a *vintage effect* a_t and a *horizon effect* b_h. With T vintages and H horizons, we estimate $T + H$ parameters using $N_f \times T \times H$ observations. Identification is achieved both from the cross-section (N_f funds in category c) and from the time series (T). Specifically, the vintage effects (the "a's") are only allowed to shift the horizon effects (the "b"-profiles) up and down in parallel fashion. The vintage effects are normalized to be 1 on average across vintages, and the horizon effects correspondingly rescaled. We include all available vintages that have at least eight quarters of cash flows because the extra information may be useful to better identify the first few elements of b_h.

One can test the null hypothesis that the vintage effects are the same: $H_0 : a_t = a_s, \ \forall t \neq s$. If the null cannot be rejected, imposing the equality of vintage effects may lead to efficiency gains in estimation.

With the portfolio weights $q_{t,h}^{b} P_{t,h}^{\$}$ in hand, we can calculate the expected return on the PE investment as a weighted average of the expected returns on the bond strips per equation (3). We can also calculate the remaining value of the PE fund at any point in its life by valuing the remaining replicating portfolio of zero-coupon bonds, using the bond prices prevailing at that point. Finally, we can calculate the risk-adjusted profit from equation (4), using the residuals from equation (7).

Four-factor Model Our main model is a four-factor model where we add to the bond factor a stock market factor, a traded real estate factor, and a traded infrastructure factor. We price these four assets in the model described in section 3. The model can easily be extended to an arbitrary number of $K > 4$ risk factors.

The empirical model assumes that the cash-flows X of all funds i in the same category

c and vintage t have the same vector of cash-flow betas at each horizon h. The betas on the bond, stock, real estate, and infrastructure factors are allowed to be different from one another, and to shift in different ways across vintages:

$$\begin{aligned} X_{t+h}^{i\in c} &= q_{t,h}^{b} + q_{t,h}^{mkt} F_{t+h}^{m} + q_{t,h}^{reit} F_{t+h}^{reit} + q_{t,h}^{infra} F_{t+h}^{infra} + v_{t+h}^{i} \\ &= a_t^1 b_h^1 + a_t^2 b_h^2 F_{t+h}^{m} + a_t^3 b_h^3 F_{t+h}^{reit} + a_t^4 b_h^4 F_{t+h}^{infra} + v_{t+h}^{i}. \end{aligned} \quad (8)$$

We estimate the $4T$ vintage effects $\left\{a_t^1, a_t^2, a_t^3, a_t^4\right\}_{t=1}^{T}$ and the $4H$ horizon profiles $\left\{b_h^1, b_h^2, b_h^3, b_h^4\right\}_{h=1}^{H}$ using a random effects model which contains vintage times horizon times cash-flows ($\boldsymbol{F}$) as dependent variables. In total, there are $4T + 4H$ coefficients to estimate using $N_f \times T \times H$ observations. The key identifying assumption is that the each vintage effect only shifts the corresponding horizon profile up and down in parallel fashion, rather than allowing for arbitrary shifting.

One can test the null hypothesis that the vintage effects are the same across risky assets: $H_0 : a_t^1 = a_t^2 = a_t^3 = a_t^4$. If the null cannot be rejected, imposing this condition may lead to efficiency gains in estimation.

3 Asset Pricing Model

The second main challenge is to price the replicating portfolio. If the only sources of risk were the risks inherent in the term structure of interest rates, this step would be straightforward. After all, on each date t, we can infer the prices of zero-coupon bonds of all maturities j from the observed yield curve. However, interest rate risk is not the only and (arguably) not even the main source of risk in the cash-flows of private equity funds. If stock market risk were the only other source of aggregate risk, then we could in principle use price information from dividend strips. Those prices can either be inferred from options and stock markets (van Binsbergen, Brandt, and Koijen, 2012) or observed directly from dividend strip futures markets (van Binsbergen, Hueskes, Koijen, and Vrugt, 2013). However, the available time series is too short for our purposes. Moreover, these strips are not available in one-quarter horizon increments. Third, there are no dividend strip data for publicly listed real estate or infrastructure assets, two additional traded factors we wish to include in our analysis given our special interest in real estate and infrastructure funds. Fourth, we do not observe expected returns on those strips, only realized excess returns over relatively short time series, and would need to obtain expected returns from an asset pricing model anyway. For all those reasons, we rely on a standard asset pricing model to obtain the time series of $\boldsymbol{P}_{t,h}$ and the corresponding expected returns. But we insist that this asset pricing model be consistent with the available dividend strip evidence.

The asset pricing model must correctly price the assets whose cash-flows $\boldsymbol{F}$ are used in the cash-flow replication. We propose a reduced-form asset pricing model rather than a structural model that starts from preferences, since it is more important for our purposes to price the replicating portfolio of publicly traded assets correctly than to understand the deeper sources of macro-economic risk that underly the prices of stocks and bonds. Our approach builds on Lustig et al. (2013), who price a claim to aggregate consumption and study the properties of the price-dividend ratio of this claim, the wealth-consumption ratio.

As emphasized by Korteweg and Nagel (2016), the objective is not to test the asset pricing model, but rather to investigate whether a potential PE investment adds value to an investor who already has access to securities whose sources of risk are captured by the SDF.

3.1 Setup

3.1.1 State Variable Dynamics

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

$$z_t = \boldsymbol{\Psi} z_{t-1} + \boldsymbol{\Sigma}^{\frac{1}{2}} \boldsymbol{\varepsilon}_t, \tag{9}$$

with shocks $\boldsymbol{\varepsilon}_t \sim i.i.d.\mathcal{N}(0, I)$ whose variance is the identity matrix. The companion matrix Ψ is a $N \times N$ matrix. The vector z is demeaned. The covariance matrix of the innovations to the state variables is Σ; the model is homoscedastic. We use a Cholesky decomposition of the covariance matrix, $\Sigma = \Sigma^{\frac{1}{2}} \Sigma^{\frac{1}{2}\prime}$, 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.

For now, we note that the (demeaned) one-month 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.

3.1.2 SDF

We specify an exponentially affine stochastic discount factor (SDF), similar in spirit to the no-arbitrage term structure literature (Ang and Piazzesi, 2003). The nominal SDF $M^\$_{t+1} = \exp(m^\$_{t+1})$ is conditionally log-normal:

$$m^\$_{t+1} = -y^\$_{t,1} - \frac{1}{2}\mathbf{\Lambda}'_t\mathbf{\Lambda}_t - \mathbf{\Lambda}'_t\varepsilon_{t+1}. \tag{10}$$

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

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

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

3.1.3 Bond Pricing

Proposition 1 in Appendix A shows that nominal bond yields of maturity τ are affine in the state variables:

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

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 will price a large cross-section of nominal bonds that differ by maturity, paying special attention to the one- and twenty-quarter bond yields since those are part of the state vector.

3.1.4 Equity Pricing

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

$$\frac{P_t^m}{D_t^m} = \sum_{\tau=1}^{\infty} P_{t,\tau}^d \tag{11}$$

$$\exp\left(\overline{pd} + e'_{pd^m} z_t\right) = \sum_{\tau=0}^{\infty} \exp\left(A_\tau^m + B_\tau^{m\prime} z_t\right), \tag{12}$$

where $P_{t,\tau}^d$ denotes the price of a τ-period dividend strip divided by the current dividend. Proposition 2 in Appendix A shows that the log price-dividend ratio on each dividend strip is affine in the state vector and provides recursions for the coefficients $\left(A_\tau^m, \boldsymbol{B}_\tau^m\right)$. Since the log price-dividend ratio on the stock market is an element of the state vector, it is affine in the state vector by assumption. Equation (12) restates the present-value relationship from equation (11). 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. Analogous present value restrictions holds for traded real estate and infrastructure strips, whose price-dividend ratios are also included in the state vector, and are also imposed on the estimation.

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 dividend strips' dividend-price ratios would equal yields on real bonds with the coupon adjusted for deterministic dividend growth. In this special case, 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.

3.1.5 Dividend Futures

The model readily implied the price of a futures contract that received 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:

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

The appendix 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. Appendix A provides the expressions and details.

3.2 Estimation

3.2.1 State Vector Elements

The state vector contains the following demeaned variables, in order of appearance: (1) GDP price inflation, (2) real per capita 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,[2] (5) the log price-dividend ratio on the CRSP stock market, (6) the log real dividend growth rate on the CRSP stock market, (7) the log price-dividend ratio on the NAREIT All Equity REIT index of publicly listed real estate companies, (8) the corresponding log real dividend growth rate on REITs, (9) the log price-dividend ratio on a listed infrastructure index, and (10) the corresponding log real dividend growth rate of infrastructure stocks:

$$z_t = [\pi_t, x_t, y_{t,1}^{\$}, y_{t,20}^{\$} - y_{t,1}^{\$}, pd_t^m, \Delta d_t^m, pd_t^{reit}, \Delta d_t^{reit}, pd_t^{infra}, \Delta d_t^{infra}]'.$$

This state vector is observed at quarterly frequency from 1974.Q1 until 2017.Q4 (176 observations). This is the longest available time series for which all variables are available.[3] Our PE cash flow data starts shortly thereafter in the early 1980s. While most of our PE fund data are 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.

[2] All yields we use are the average of daily Constant Maturity Treasury yields within the quarter.

[3] The first observation for REIT dividend growth is in 1974.Q1. We seasonally adjust dividends, which means we lose the first 8 quarters of data in 1972 and 1973. The seasonal adjustment is the same for the overall stock market and the infrastructure stock index.

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 B.1 contains the resulting point estimates for Ψ and $\Sigma^{\frac{1}{2}}$.

3.2.2 Market Prices of Risk

The state vector contains both priced sources of risk as well as predictors of bond and stock returns. We estimated 8 parameters in the constant market price of risk vector Λ_0 and 40 elements of the matrix Λ_1 which governs the dynamics of the risk prices. The point estimates are listed in Appendix B.2. We use the following target moments to estimate the market price of risk parameters.

First, 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 $T = 176$ quarters.[4]

Second, we impose restrictions that we exactly match the average five-year bond yield and its dynamics. This delivers 11 additional restrictions:

$$\begin{aligned} -A^{\$}_{20}/20 &= y^{\$}_{0,20} \\ -B^{\$}_{20}/20 &= [0, 0, 1, 1, 0, 0, 0, 0, 0, 0] \end{aligned}$$

Because the 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 and fourth place and zeroes everywhere else.

Third, we match the time-series of log price-dividend ratios on stocks, real estate, and infrastructure which are included in the state vector. The model-implied price-dividend ratios are built up from 3,500 quarterly dividend strips according to equation (11). Thus, we impose the present-value relationship for all three stock prices at each date. They constitute $3 \times T$ moments.

Fourth, we impose that the equity risk premia for the overall stock market, REITs, and infrastructure in the model match those given by the VAR, both in terms of the unconditional average and the dependence on the state variables. As usual, the expected excess return in logs (including a Jensen adjustment) must equal minus the conditional covariance between the log SDF and the log return. For example, for the overall stock

[4]The 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.

market:

$$E_t\left[r_{t+1}^{m,\$}\right]-y_{t,1}^{\$}+\frac{1}{2}V_t\left[r_{t+1}^{m,\$}\right] = -Cov_t\left[m_{t+1,}^{\$}r_{t+1}^{m,\$}\right]$$

$$r_0^m+\pi_0-y_0^{\$}(1)+\left[(e_{divm}+\kappa_1^m e_{pd}+e_\pi)'\Psi-e_{pd}'-e_{yn}'\right]z_t$$

$$+\frac{1}{2}\left(e_{divm}+\kappa_1^m e_{pd}+e_\pi\right)'\Sigma\left(e_{divm}+\kappa_1^m e_{pd}+e_\pi\right) = \left(e_{divm}+\kappa_1^m e_{pd}+e_\pi\right)'\Sigma^{\frac{1}{2}}\Lambda_t$$

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). Similar restrictions apply to the expected excess return for REITS and Infrastructure. This provides 29 additional restrictions. These moments dictate which elements of the 6th, 8th, and 10th rows of Λ_1 must be non-zero.

Fifth, we impose that the model match the realized price-dividend ratio time series on short-maturity strips on the stock market as well as the share of the overall market value that these short-maturity strips represent. Specifically, we consider a claim that pays the first eight quarters of realized nominal dividends. This claim can be priced in the model as the sum of the prices to the first eight dividend strips. Data for the price-dividend ratio and the share in the overall stock market (S&P500) on this claim are obtained from van Binsbergen, Brandt, and Koijen (2012) for the period 1996.Q1-2009.Q3 (55 quarters). This delivers 2×55 moments. We also want to make sure our model is consistent with the high average realized returns on short-horizon dividend futures, first documented by van Binsbergen, Hueskes, Koijen, and Vrugt (2013). Table 1 in van Binsbergen and Koijen (2017) reports the observed average monthly return on one- through seven-year U.S. SPX dividend futures over the period Nov 2002 - Jul 2014 . That average portfolio return is 0.726% per month or 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 average the realized return on this dividend futures portfolio between 2003.Q1 and 2014.Q2, and annualize it. We target 8.71% for this return. This results in one additional restriction. We free up the market price of risk associated with the market price-dividend ratio (fifth element of Λ_0 and first six elements of the fifth row of Λ_1) to help match the dividend strip evidence.

Sixth, 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).

Seventh, 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 growth, which is 1.65% per year in our sample. Long-run nominal yields must exceed long-run real yields by 2% inflation. These constraints are satisfied at the final solution.

Not counting the regularity conditions, we have 1,750 moments to estimate 48 parameters. Thus, the estimation is massively over-identified.

3.2.3 Model Fit

Figure 1 plots the bond yields on bonds of maturities 3 months, 1 year, 5 years, and 10 years. Those are the most relevant horizons for the private equity cash-flows. The model matches the time series of bond yields in the data closely. It matches nearly perfectly the one-quarter and 5-year bond yield which are part of the state space.

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.

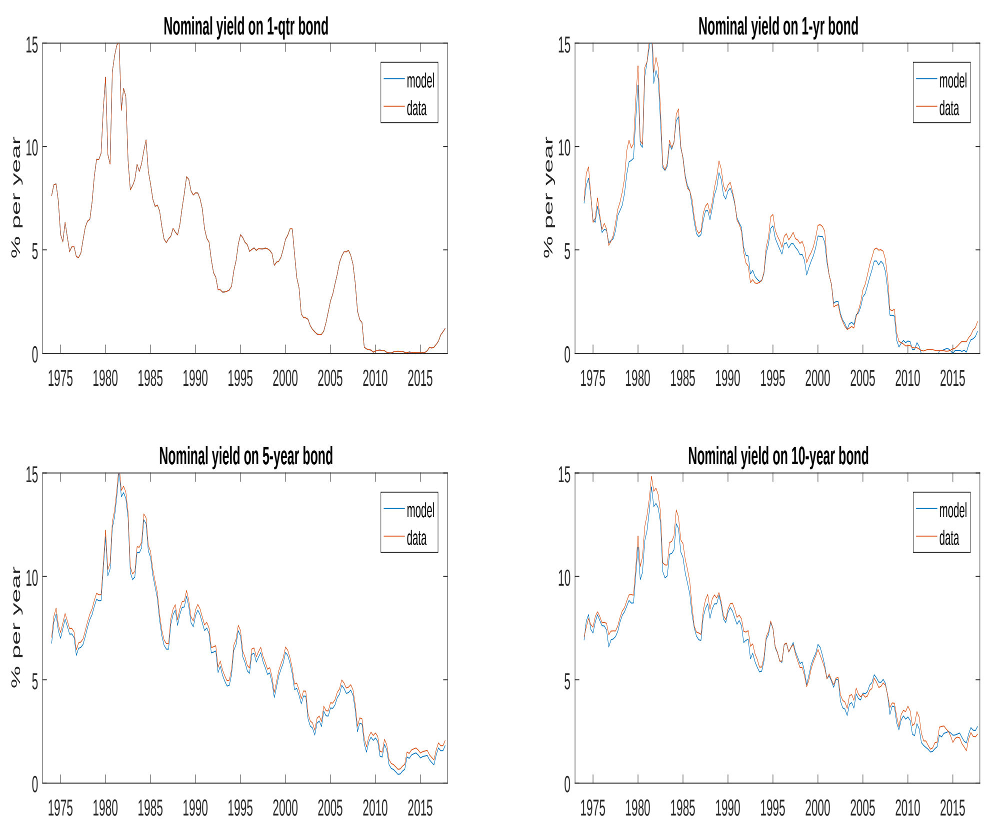

The top panels of Figure 2 show the model's implications for the average nominal (left panel) and real (right panel) yield curves at longer maturities. These long-term yields are well behaved. The bottom left panel shows that the model matches the dynamics of the nominal bond risk premium, defined as the expected excess return on five-year nominal bonds, nearly perfectly. The bottom right panel shows a decomposition of the nominal bond yield on a five-year bond into the five-year real bond yield, annual expected inflation inflation over the next five years, and the five-year inflation risk premium. On average, the 6.0% five-year nominal bond yield is comprised of a 1.7% real yield, a 3.3% expected inflation rate, and a 1.0% inflation risk premium. The graph shows that the importance of these components fluctuates over time. The inflation risk premium has been shrinking over time, consistent with the findings in the term structure literature.

Figure 2: 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.

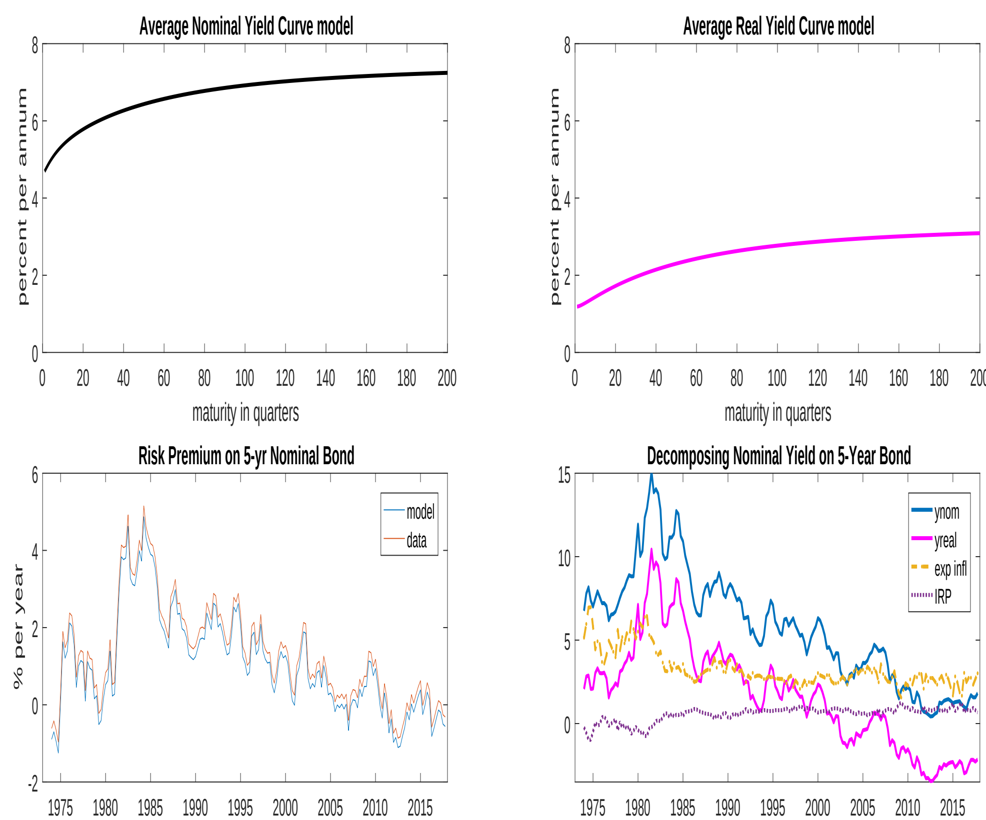

Figure 3 shows the equity risk premium, the expected excess return, in the left panels and the price-dividend ratio in the right panels. The top row is for the overall stock market, the middle row for REITs, and the bottom row for infrastructure stocks. The dynamics of the risk premia in the data are dictated by the VAR. The model chooses the market prices of risk to fit these risk premium dynamics as closely as possible.[5] The price-dividend ratios in the model are formed from the price-dividend ratios on the strips of maturities ranging from 1 to 3500 quarters, as explained above. The figure shows a good fit for price-dividend levels and 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. The model generates a 5.4% equity risk premium on the market with a standard deviation of the risk premium of 8.1%. The VAR implies an equity risk premium of 6.8% with a standard deviation of 11.2%. By its nature, the VAR may imply excessive movement in equity risk premia, especially on the stock market as a whole. A model that does not match every wiggle in the risk premium dynamics may in fact be the more realistic one.

3.3 Temporal Pricing of Risk

Zero-Coupon Bond and Zero-Coupon Equity Prices The first key output from the model, and input in the private equity valuation exercise, is a nominal bond price for zero-coupon bonds with maturities ranging from one to approximately 60 quarters. The second key output from the model, and input in the private equity valuation exercise, is the price for dividend strips with maturities ranging from one to 60 quarters. We scale this price by the current dividend. Figure 4 plots the time series for prices of nominal zero-coupon bonds and dividend strips on the overall stock market, REITs, and infrastructure stocks. We plot three maturities: one-month, five-years, and ten-years. We use these prices to value the replicating portfolio for private equity cash-flows in our main valuation equation (3). The model implies substantial time variation in dividend strip prices over time, as well as across risky assets.

As part of the estimation, the model fits the observed time series of the price-dividend ratio on a claim to the first 8 quarters of dividends, as well as the share of the total stock

[5] The quarterly risk premia are annualized by multiplying them by 4 for presentational purposes only. We note that the VAR does not restrict risk premia to remain positive. The VAR-implied equity risk premium is negative in 21% of the quarters. For REITS this is 10% and for infrastructure only 5% of quarters. The most negative value of the risk premium on the overall stock market is -8% quarterly. For REITS, the most negative value for the risk premium is -2.9% quarterly, while it is -1.2% for infrastructure.

Figure 3: Equity Risk Premia and Price-Dividend Ratios

The figure plots the observed and model-implied equity risk premium on the overall stock market, REIT market, and infrastructure sector, in the left panels, as well as the price-dividend ratio in the right panels.

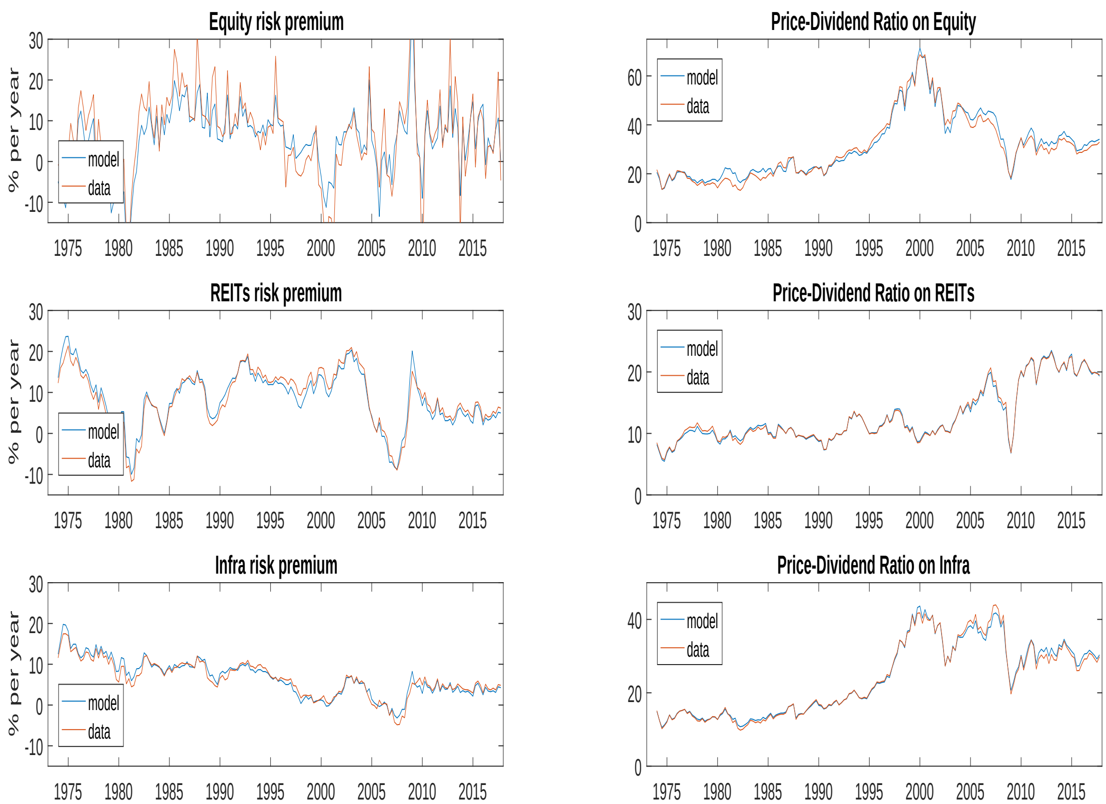

market value that these first eight quarters of dividends represent. Figure 5 shows the result for the price-dividend ratio in the left panel and the share in the right panel. The model generates the right level for the price-dividend ratio for the short-horizon claim. The average in the model, for the 55 quarters for which the data are available, is 7.74. The average in the data is 7.65. The model also generates the entire drop in the PD ratio in the financial crisis. However, it misses the equally large drop in 2000.Q4-2001.Q1. The model implies that the first 8 quarters of dividends represent 4.5% of the overall stock market value (the claim to all future quarters) over the period in which we have data. That same number in the data is 3.4%. The model mimics the observed dynamics of the share quite well, including the sharp decline in 2000.Q4-2001.Q1 when the short-term strips falls by more than the overall stock market. This reflects the market's perception that the

Figure 4: 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, REIT market, and infrastructure sector in the next three 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.

recession would be short-lived. In contrast, the share of short-term strips increases in the Great Recession, both in the data and in the model in recognition of the persistent nature of the crisis.

The model generates rich patterns in the temporal pricing of risk. Figure 6 plots the average risk premium on nominal zero coupon bond yields (top left panel) and on dividend strips (other three panels) of various maturities ranging from 1 to 60 quarters. Those are the relevant maturities for the private equity investments. Risk premia on nominal bonds are increasing with maturity from 0 to 3.5%. The second panel shows the risk premia on dividend strips on the overall stock market (solid blue line). It also plots the dividend futures risk premium. The difference between the dividend strip and the dividend futures risk premium is approximately equal to the nominal bond risk premium. The unconditional dividend futures risk premium (pink line with circles) is downward sloping in maturity at the short end, consistent with the empirical findings of van Binsbergen, Hueskes, Koijen, and Vrugt (2013), and van Binsbergen and Koijen (2017). The graph also

Figure 5: Short-run Cumulative Dividend Strips

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

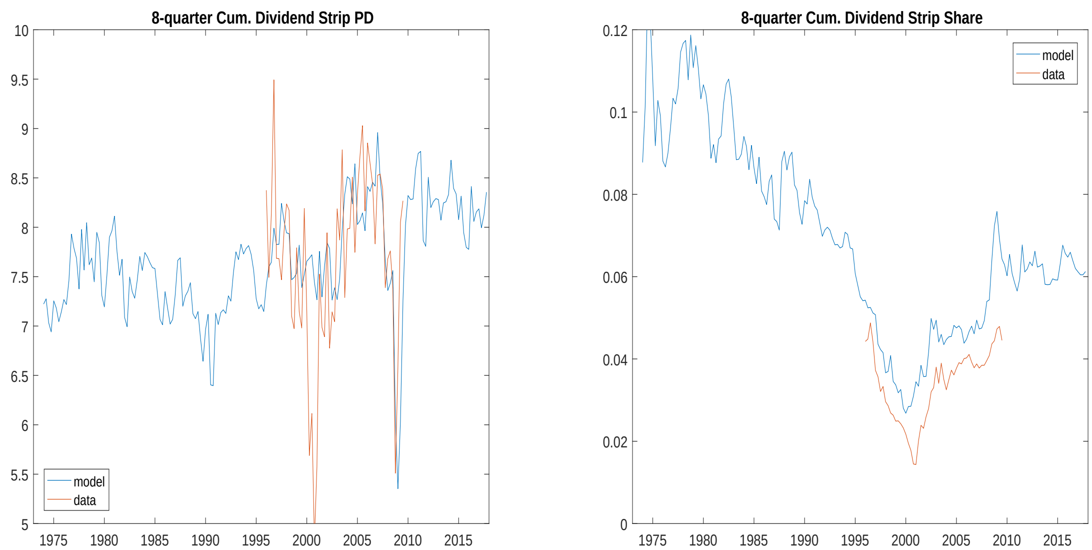

plots the dividend futures risk premium, averaged over the period 2003.Q1-2014.Q2 (yellow dash-dotted line). It is substantially more downward sloping at the short end than the risk premium averaged over the entire 1974-2017 sample. Indeed, the model does a good job matching the *realized* portfolio return on dividend futures of maturities 1-7 years over the period 2003.Q1-2014.Q2, which was 8.71% in the data and 8.70% in the model.[6] Risk premia on dividend futures peak at maturity of 70 quarters (not shown), and slowly decline thereafter stabilizing at 0.5% per year.

The average term structure of dividend strip risk premia for REITs (bottom left panel) starts at 6% at the short end, and is uniformly upward sloping , with risk premia growing to 10% per annum at the 60-quarter horizon. The futures risk premia are downward sloping at shorter horizons and essentially flat thereafter, suggesting that the upward slope in the spot risk premia is inherited from the nominal bond risk premia. The risk premium curve on public infrastructure assets (bottom right panel) shows a small hump at the short end and is much flatter than for the other two risky assets. It also is lower on average. The dividend futures risk premia on infrastructure are downward sloping,

[6] As an aside, the conditional risk premium, which is the expected (as opposed to realized) return on the dividend futures portfolio over the 2003.Q1-2014.Q2 period is 4.32% per year in the model. The unconditional risk premium on the dividend futures portfolio (over the full sample) is 3.81%.

similar to the overall stock market.

The difference in the level and the horizon-dependence of the three strip risk premia will generate differences in the risk premia on private equity investments if their cash-flows display differential exposure to the traded asset cash-flows.

Figure 6: Zero Coupon Bond Prices and Dividend Strip Prices

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, REIT market, and infrastructure sector in the next three panels, for maturities ranging from 1 to 180 months.

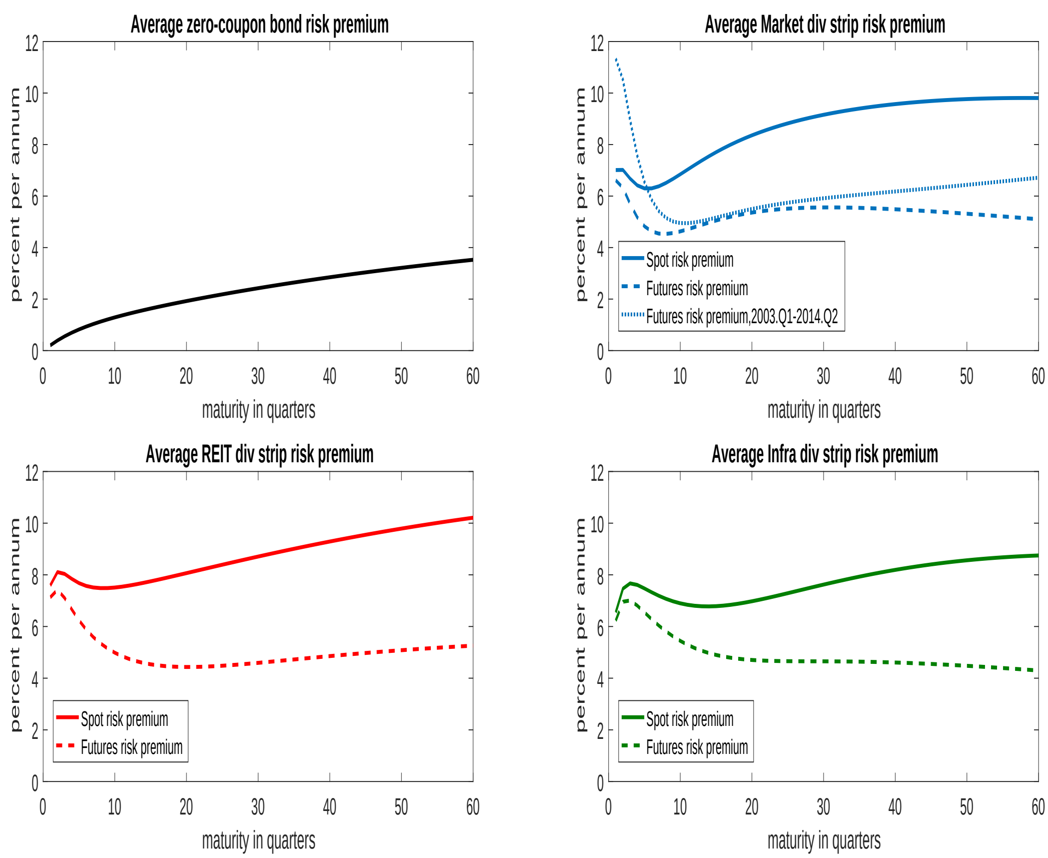

Shock Exposure and Shock-Price Elasticities Borovička and Hansen (2014), building on earlier work by Hansen and Scheinkman (2009), provide a *dynamic value decomposition*, the asset pricing counterpart to an impulse response function. It allows a researcher to decompose the risk premium an investor requires for exposure to a shock into the product of the exposure (shock exposure elasticity) and the risk price (shock price elasticity), horizon by horizon. Appendix C applies their analysis to our VAR setting.

Figure 7 plots the shock-exposure elasticities of dividend growth on the market (blue), dividend growth of REITs (red), and dividend growth of infrastructure stocks (green) to a one-standard deviation shock to inflation (top left), to real per capital GDP growth (top second), to the short rate (top third), to the slope factor (top right), to the dividend growth rate on the market (bottom left), to the dividend growth rate on REITs (bottom middle), and to the dividend growth rate on infrastructure (bottom right). The shock exposure elasticities are essentially impulse-responses to the original (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 the dividends on stocks, REITs, and infrastructure, the PE cash flow exposures to the VAR shocks will be linear combinations of the shock exposure elasticities of these three dividend growth rates. There is interesting heterogeneity in the three cash flow exposures to the various VAR shocks. For example, the top left panel shows that dividend growth on infrastructure responds positively to an inflation shock while the dividend growth responses for the aggregate stock market and especially for REITS are negative. This points to the inflation hedging potential of infrastructure assets and the inflation risk exposure of REITS and the market as a whole. The second panel shows that REIT dividend growth responds positively to GDP growth, while cash flow growth on the market responds negatively. All cash flows respond negatively to an increase in interest rates in the long-run, but the response of REIT cash flows is positive for the first five years. REIT cash flows are rents which can be adjusted upwards when rates increase, which typically occurs in a strong economy (see the GDP panel). The market dividend growth shows a substantial positive response to a steepening yield curve. The bottom three panels show that the dividend growth shock on the market is nearly permanent, while the other two cash flow shocks are mean reverting. A positive shock to infrastructure cash flows also has a positive effect on REIT cash flows, which grows over time (bottom right).

Figure 8 plots the shock-price elasticities to a one-standard deviation shock to each of the same (non-orthogonalized) VAR innovations. These shock price elasticities are a property of the (cumulative) SDF process, and therefore depend on the estimated market price of risk parameters. They quantify the compensation investors demand for horizon-dependent risk exposure. The price of inflation risk is negative, consistent with increases in inflation being bad states of the world. GDP growth risk is naturally priced positively, and more so at longer horizons. Level risk is negatively priced, consistent with standard results in the term structure literature that consider high interest rate periods bad states

Figure 7: Shock Exposure Elasticities

The figure plots the shock-exposure elasticities of dividend growth on the market, dividend growth of REITs, and dividend growth of infrastructure shocks to a one-standard deviation shock to the inflation factor (top left), real GDP growth (top second), the short rate (top third), the slope factor (top right), the price-dividend ratio on the market (bottom left), dividend growth rate on the market (bottom second), the dividend growth rate on REITs (bottom third), and the dividend growth rate on infrastructure (bottom right).

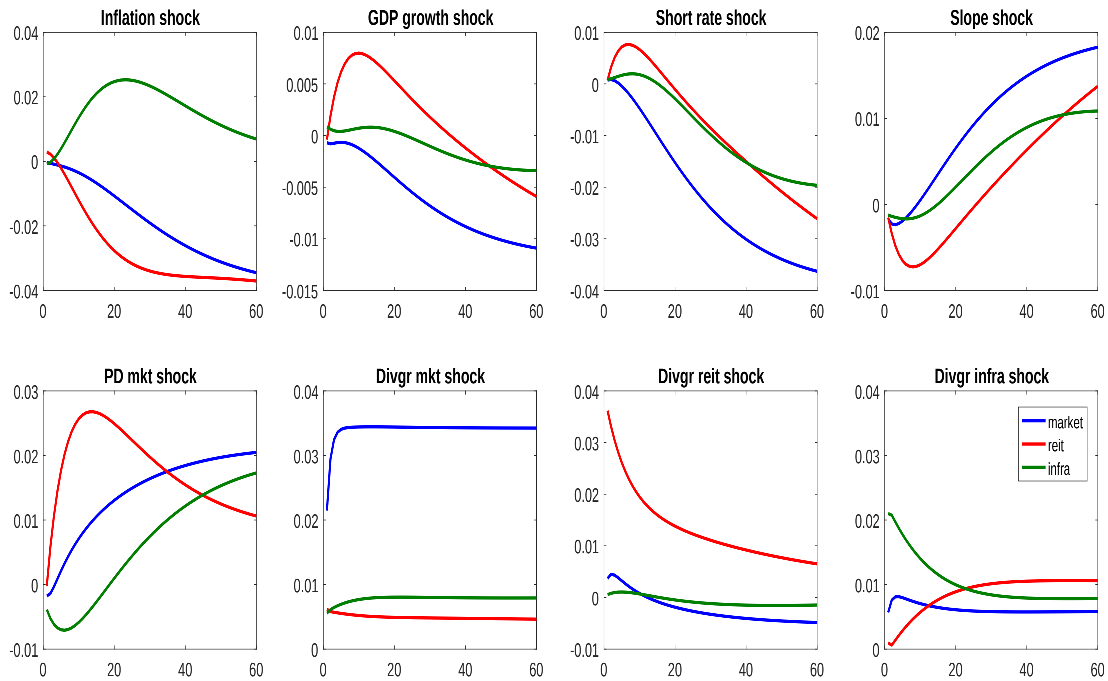

of the world. The price of level risk becomes less negative at longer horizons. The price of slope risk is positive. All cash-flow shocks in the bottom three panels naturally have positive risk prices since increases in cash-flow growth are good shocks to the representative investor. Those three dividend shock risk price elasticities are nearly flat across horizons.

Figure 8: 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 (top right), the price-dividend ratio on the market (bottom left), dividend growth rate on the market (bottom second), the dividend growth rate on REITs (bottom third), and the dividend growth rate on infrastructure (bottom right). The shocks whose risk prices are plotted are the (non-orthogonalized) VAR innovations $\Sigma^{\frac{1}{2}}\varepsilon$.

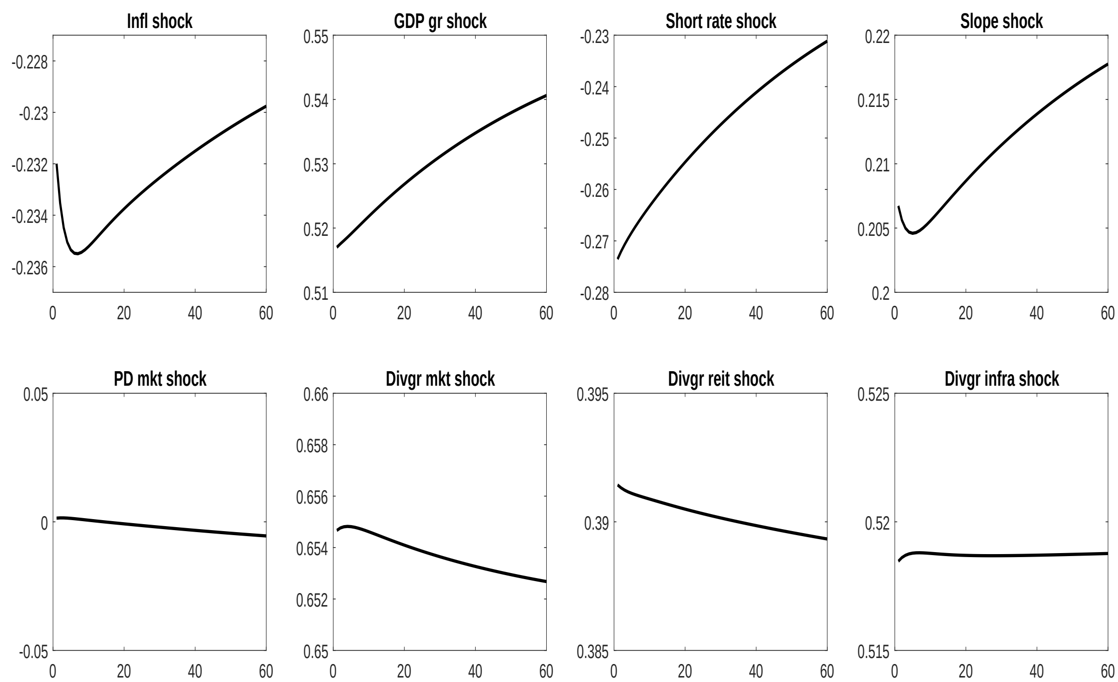

4 Expected Returns and Risk-adjusted Profits on Private Equity Funds

In this section, we combine the cash-flow exposures from section 2 with the asset prices from section 3 to obtain risk-adjusted profits on private equity funds.

4.1 Summary Statistics

Our fund data cover the period January 1981 until Dec 2017. The data source is Preqin. We group private equity funds into seven categories: Buyout (LBO), Venture Capital (VC), Real Estate (RE), Infrastructure (IN), 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. One may be able to further enrich the analysis by defining categories more granularly. For example, real estate strategies are often subdivided into opportunistic, value-add, core plus, and core funds. Infrastructure could be divided into greenfield and brownfield, etc. Additionally, one may be able to estimate a firm fixed effect as (larger) firms often have several funds in the data set.

We include all funds with non-missing cash-flow information. We group funds also by their vintage, the year in which they first appear in the data set. The last vintage we consider in the analysis is the 2017 vintage. Table 1 reports the number of funds and the aggregate AUM in each vintage-category pair. In total, we have 4,215 funds in our analysis and an aggregate of \$4.1 trillion in assets under management. There is clear business cycle variation in when funds funds get started as well as in their size (AUM). Buyouts are the largest category by AUM, followed by Real Estate, and then Venture Capital.

Figure 9 shows the average cash-flow profile in each category for distribution events, pooling all funds and vintages together and equally weighting them. We combine all monthly cash-flows into one quarterly cash-flow for each fund. Quarter zero is the quarter in which the first capital call takes place. The last bar is for the last quarter of year 15 (quarter 60). For the purposes of making this figure and in the cash-flow beta estimation, we include all observed cash-flows until 2017.Q4. Thus, the 2010 vintage funds have at most 32 quarters of cash-flows. Cash-flows arriving after quarter 60 are included in the last quarter under a separately highlighted color (green). The distribution cash-flows fol-

Table 1: Summary Statistics

Panel A: Fund Count

Vintage	Buyout	Debt Fund	Fund of Funds	Infrastructure	Real Estate	Restructuring	Venture Capital	Total
1981	0	0	0	0	0	0	1	1
1982	0	0	0	0	0	0	3	3
1983	0	0	1	0	0	0	1	2
1984	0	0	0	0	0	0	3	3
1985	4	0	0	0	0	0	6	10
1986	2	0	3	0	0	0	7	12
1987	5	0	0	0	0	0	5	10
1988	7	0	1	0	0	0	4	12
1989	3	0	0	1	0	0	5	9
1990	7	0	2	0	0	1	8	18
1991	2	0	1	0	0	2	4	9
1992	9	0	0	1	1	2	12	25
1993	9	0	2	1	0	0	11	23
1994	15	0	1	2	1	1	12	32
1995	14	0	5	1	2	0	17	39
1996	20	0	1	3	3	3	21	51
1997	23	0	4	2	6	2	26	63
1998	40	0	11	4	3	1	32	91
1999	31	1	9	1	2	3	47	94
2000	33	2	17	0	6	3	84	145
2001	21	0	19	1	2	5	55	103
2002	24	1	13	3	3	3	29	76
2003	18	1	12	2	7	4	20	64
2004	27	1	19	6	11	2	33	99
2005	55	2	33	5	19	6	49	169
2006	71	1	51	8	32	11	59	233
2007	72	1	48	13	35	13	72	254
2008	65	5	61	11	31	11	62	246
2009	27	2	30	9	12	8	26	114
2010	39	4	38	18	28	8	39	174
2011	53	2	64	22	46	12	49	248
2012	63	3	54	19	36	12	44	231
2013	63	15	66	19	59	20	53	295
2014	68	12	67	26	46	16	67	302
2015	72	16	72	21	73	18	76	348
2016	91	13	78	32	59	10	72	355
2017	47	20	38	22	60	7	58	252

Panel B: Fund AUM ($m)

Vintage	Buyout	Debt Fund	Fund of Funds	Infrastructure	Real Estate	Restructuring	Venture Capital	Total
1981	0	0	0	0	0	0	0	0
1982	0	0	0	0	0	0	55	55
1983	0	0	75	0	0	0	0	75
1984	0	0	0	0	0	0	189	189
1985	1,580	0	0	0	0	0	74	1,654
1986	59	0	1,510	0	0	0	293	1,862
1987	1,608	0	0	0	0	0	1,061	2,669
1988	2,789	0	0	0	0	0	463	3,252
1989	761	0	0	160	0	0	305	1,226
1990	2,457	0	1,906	0	0	153	1,125	5,641
1991	242	0	0	0	0	329	431	1,002
1992	1,150	0	0	184	0	59	1,320	2,713
1993	3,192	0	597	54	0	0	1,438	5,281
1994	6,882	0	140	1,519	488	93	1,413	10,535
1995	9,169	0	1,172	205	523	0	2,645	13,714
1996	7,435	0	242	1,114	1,851	1,600	3,820	16,062
1997	23,633	0	1,337	480	3,642	1,700	6,308	37,100
1998	38,956	0	10,879	3,933	3,461	52	8,441	65,722
1999	34,297	109	9,248	42	2,293	3,133	17,093	66,215
2000	56,299	730	13,570	0	6,757	3,320	37,802	118,478
2001	23,856	0	11,942	1,375	3,225	7,461	23,852	71,711
2002	23,705	100	10,043	1,795	5,507	1,844	7,391	50,385
2003	31,264	366	8,767	884	3,435	5,105	6,670	56,491
2004	32,855	215	5,808	5,446	6,269	2,580	8,419	61,592
2005	102,231	412	23,252	6,353	25,523	5,830	16,280	179,881
2006	215,875	778	41,728	9,726	42,442	22,928	37,285	370,762
2007	176,849	400	45,660	21,568	44,657	40,486	24,198	353,818
2008	165,982	4,697	39,493	27,228	37,789	26,158	31,839	333,186
2009	38,738	195	11,544	12,108	9,451	11,170	8,880	92,086
2010	30,612	1,684	20,662	19,891	20,125	11,855	21,457	126,286
2011	97,669	1,720	26,915	21,581	48,951	19,237	23,747	239,820
2012	94,732	733	40,384	34,467	23,417	23,943	29,566	247,242
2013	88,022	16,333	20,013	32,762	65,703	25,660	26,026	274,519
2014	126,182	5,638	37,191	50,504	32,206	21,792	35,863	309,376
2015	115,376	12,421	52,430	33,903	75,627	34,845	31,294	355,896
2016	165,779	10,721	47,705	53,388	45,416	11,368	36,433	370,810
2017	112,820	15,942	18,513	20,385	52,750	9,609	24,840	254,859

low a hum-shaped pattern. While a majority of distribution cash-flows occur between years 5 and 10, there meaningful cash-flow distributions before year 5 and after year 10. This is especially true for IN and VC funds.

Figure 9: Distribution Cash-flow Profiles

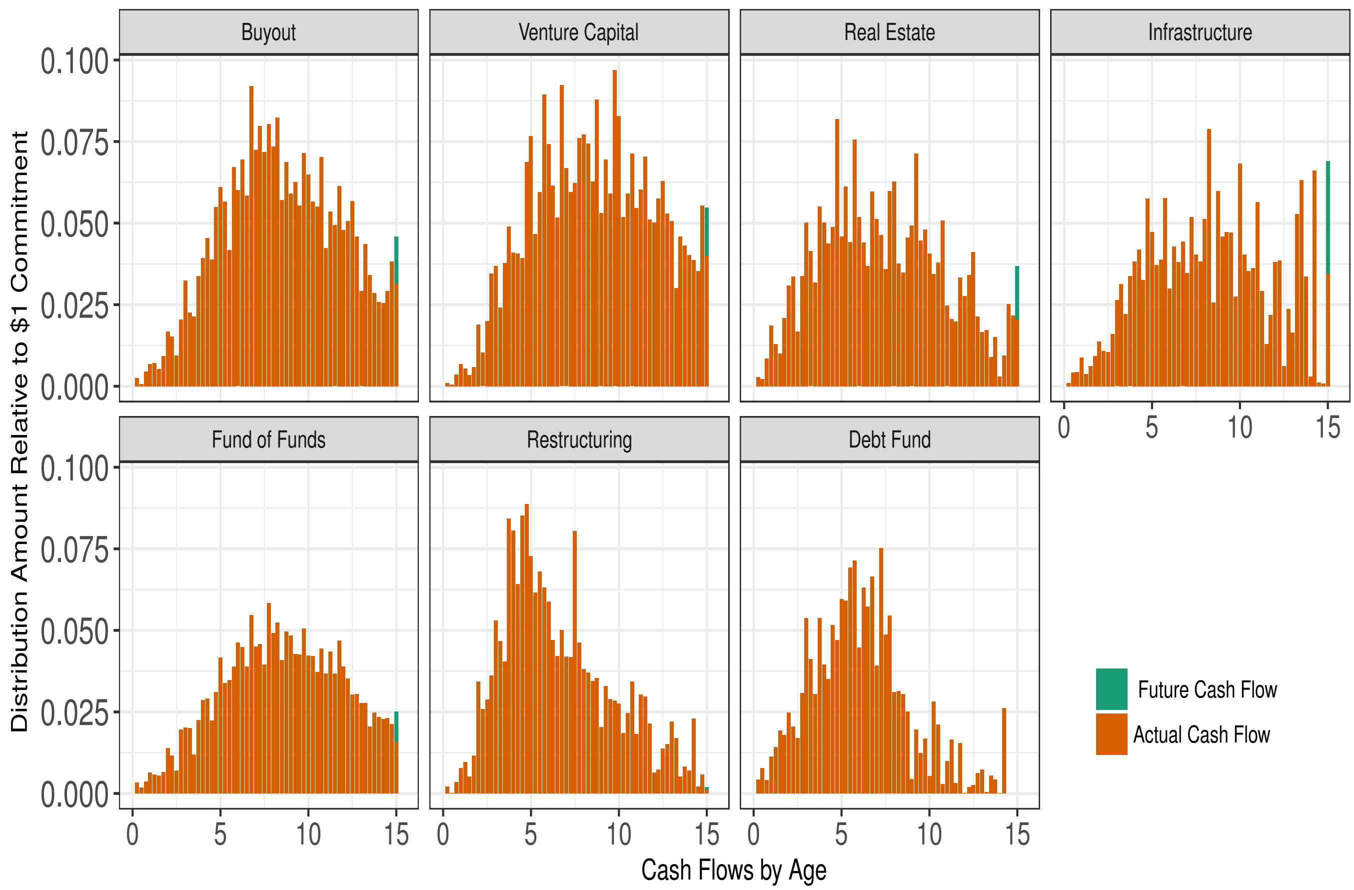

Figure 10 zooms in on the four investment categories of most interest to us: LBO, VC, RE, and IN. The figure shows the average cash-flow profile for each vintage. Since there are few LBO and VC funds prior to 1990 and few RE and IN funds prior to 2000, we start the former two panels with vintage year 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. Appendix Figures A.1 shows cash-flow profiles for the remaining categories.

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 the stock market also had very high cash-flow realizations in the year 2000 and a sharp drop thereafter, this type of variation will help the model identify a high stock market beta for VC funds. This is an important distinction with other methods, such as the PME, which assume constant risk exposure and so would attribute high cash flow distributions in this period to excess returns.

Figure 10: Cash-flows by Vintage

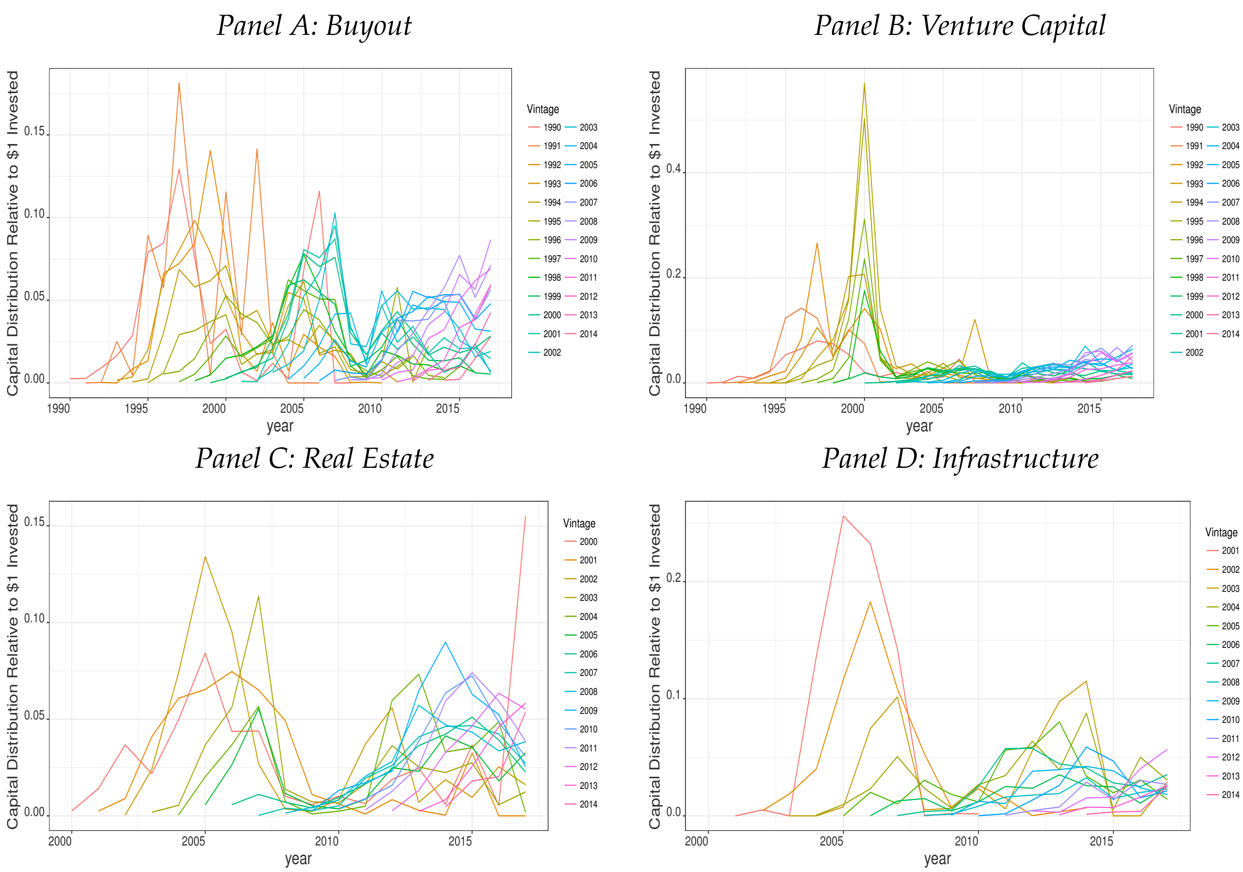

4.2 One-factor Model

We start with a discussion of the one-factor model, which assumes that the only risk that is priced is interest rate risk. It is straightforward to estimate and provides a useful point of comparison for our benchmark four-factor model. The estimation chooses parameters to match the average fund cash flows, for each category-vintage pair. Crucially, the resulting positions in bonds of various maturities are then scaled down (or up) to ensure that the replicating portfolio of bonds does not cost too much (too little). The high cash flows of a particular PE vintage-category may not be achievable/replicable with a budget-feasible bond portfolio, but only with a budget infeasible one. This will result in high "errors" v and high average risk-adjusted profits across the funds in that vintage-category.

Figure 11 show the estimated horizon effects $\widehat{b}_h$ in the left panels and the vintage effects $\widehat{a}_t$ in the right panels. Appendix Figure A.5 contains the same figure for other fund categories. Each row is for one of our four main investment categories. The plotted coefficients on the left are the positions that the replicating portfolio holds in bond strips (zero-coupon bonds) of the various horizons. The one-factor model explains 4.9% of the variation in cash-flows for LBO, 13.9% for VC, 5.2% for RE, and 6.2% for IN. There is substantial variation in the distribution of bond betas across categories of funds, both in terms of the horizon effects and vintage effects. The left panels show that IN funds have the highest average bond betas in years 10 and later, reflecting that they make distributions on average later in the life-cycle. RE funds have more short-horizon bond risk exposure, and also the highest exposure of all categories.

The right panels show how these exposure profiles from the left panel are shifted up or down for each vintage. On average, these vintage effects are 1 by construction. Values below one reflect lower bond risk. When bond risk is low, zero coupon bond prices are high and the $1 replicating portfolio can buy fewer bonds. The replicating portfolio therefore has lower average cash flows. Conversely, vintage effects above 1 reflect vintages with above-average risk and above-average cash flows on the replicating portfolio. For most categories the 2005-2007 vintages jump out as having the highest risk exposure. The mid-1990s also resulted in higher average systematic cash flows on RE and LBO.[7]

Expected Return With the replicating portfolio of zero-coupon bonds in hand, we can calculate the expected return on PE funds in each investment category as in equation (3). The expected returns on the traded strips are given by equation (A.12) in Appendix A. The left panels of Figure 12 plot the total expected return over the life of the investment, broken down into its horizon components, and averaged across all vintages. The expected return at each horizon is annualized in the left panels, and the fund's expected return is the weighted average of the horizon-specific returns, weighted by the portfolio weights over horizons. Since most of the cash-flows come later and later cash flows are riskier (higher bond beta), the risk premium is "backloaded." This backloading is more

[7]For these last three vintages, many funds have not yet reached their terminal date. Their vintage effect is thus estimated off the first 13-10 years of cash-flow data.

pronounced for IN and VC funds. Panels on the right plot the time-series of the expected return. Variation in the state variables of the VAR drive variation in the expected return of each zero-coupon bond, and therefore in the expected return of the replicating portfolio of the PE fund.

Performance Evaluation Next, we turn to performance evaluation in the one-factor model. The left panels of Figure 13 plots the histogram of risk-adjusted profits (RAP), pooling all vintages. The right panels plot the average risk-adjusted profit for each vintage. Appendix Figure A.8 reports on the remaining investment categories. If the only risk considered is interest rate risk, the average buyout fund has delivered a small negative RAP of -0.3%. The average VC fund has made a RAP of -0.9%. The average real estate and infrastructure fund have RAP of -0.3% each. More strikingly, there is large cross-sectional variation in profit across funds in the same category. Some funds gain 20 cents or more per dollar of committed capital while others lose 10 cents or more. Depending on the category, about 5% of PE funds have a RAP that exceeds 10%.

The right panels show interesting time series variation in average profits. Risk-adjusted profits are high for VC funds in the earlier part of the sample, confirming a key stylized fact in the VC literature. RE and IN funds have the highest profits for vintages in the early 2000s. Maybe surprisingly, the VC profits are not unusually high for the 1999-2000 vintages. The high factor exposures for those vintages are sufficient to explain the high average payouts. Also interesting is that the 2005-07 RE vintages performed poorly. The post-crisis vintages of PE funds tend to have weak average performance, with the exception of RE.[8]

[8] The profit for these vintages is calculated as the discounted value of the idiosyncratic cash-flow components v^i that are available through the end of the sample. Implicitly, the assumption is that the *non-systematic* cash-flow component of the average fund in those vintages will be zero in the remaining years for which no cash-flow data are available yet. In other words, the poor performance is not simply due to missing cash-flow information.

Figure 11: Replicating Portfolio Exposure by Feature

Panel A: Buyout

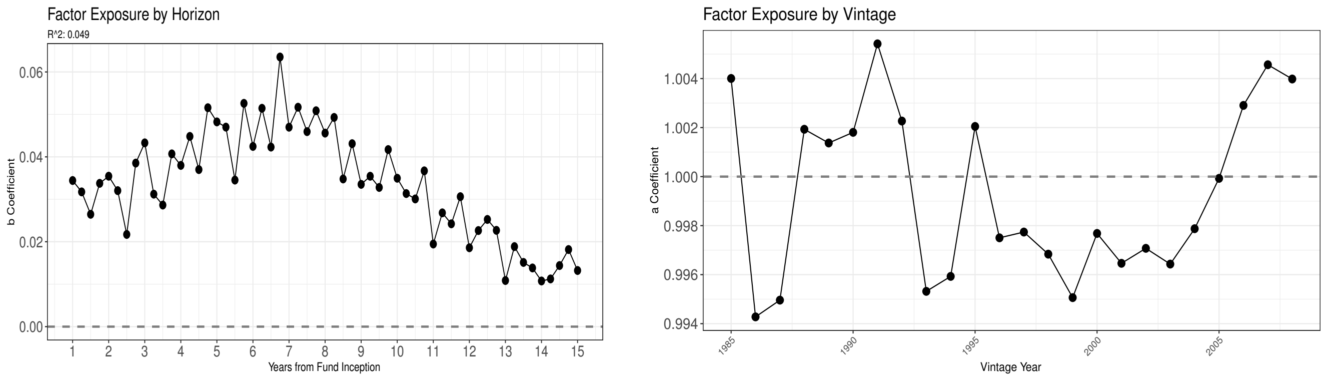

Panel B: Venture Capital

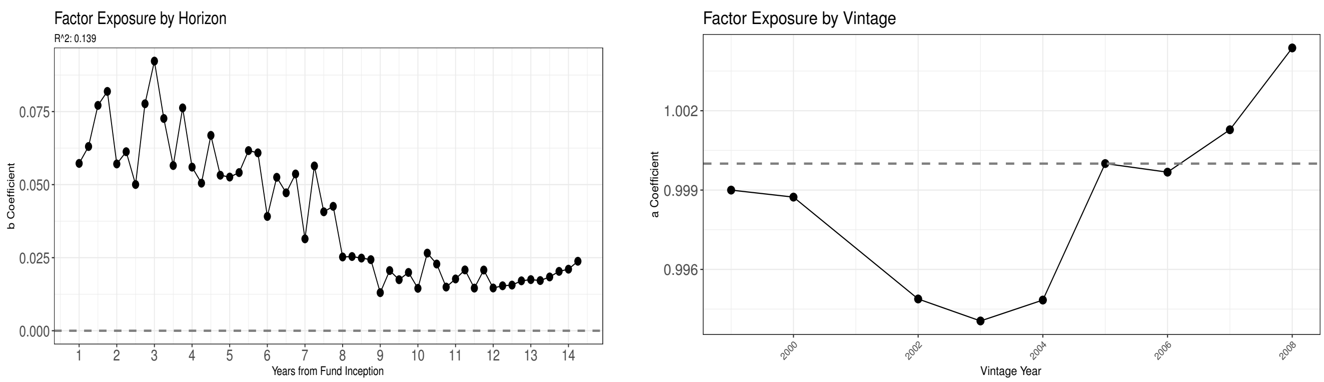

Panel C: Real Estate

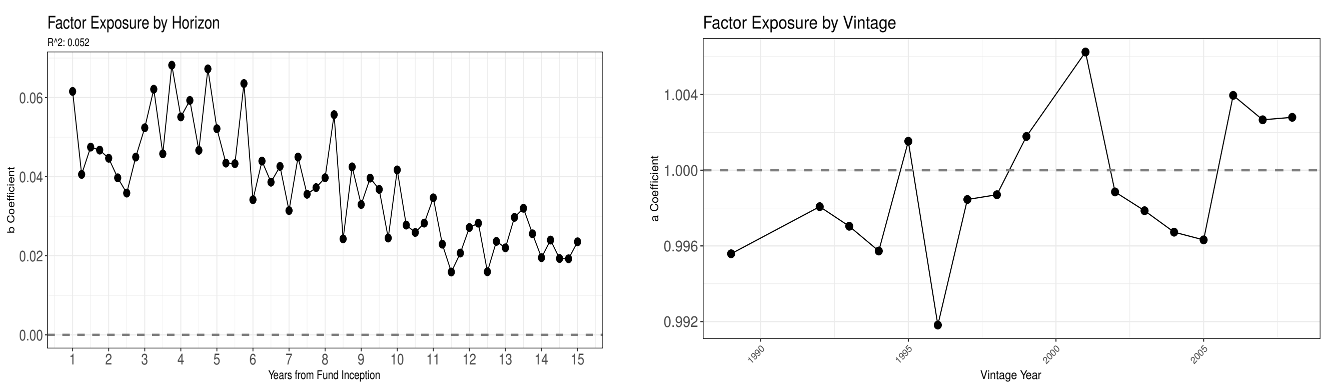

Panel D: Infrastructure

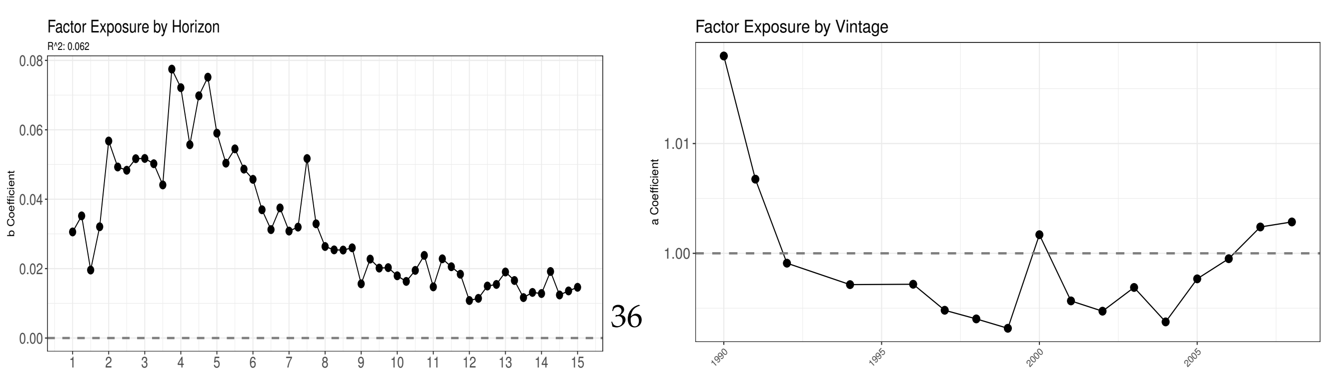

Figure 12: Expected Return

Panel A: Buyout

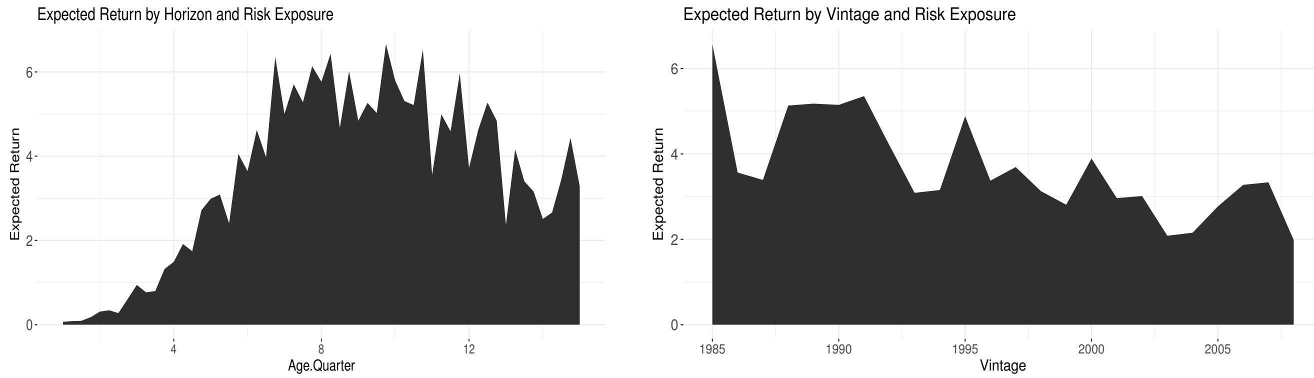

Panel B: Venture Capital

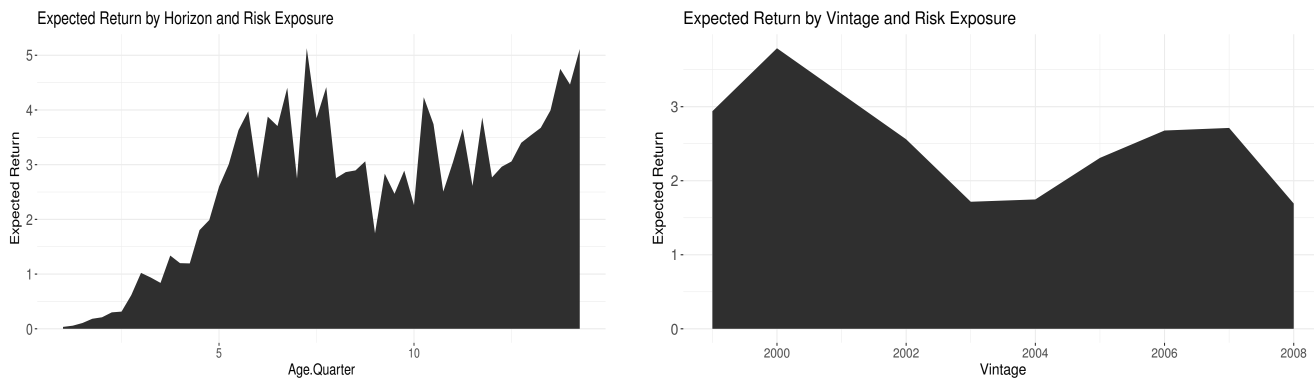

Panel C: Real Estate

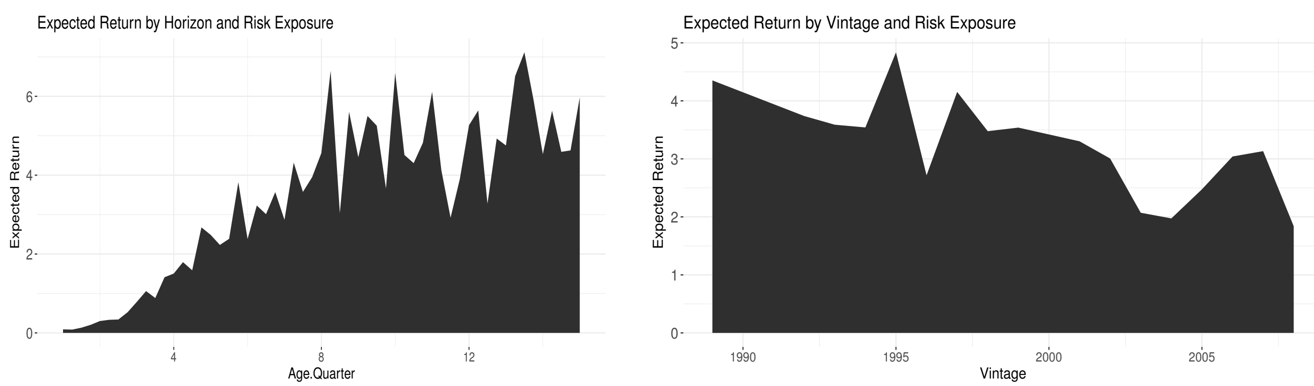

Panel D: Infrastructure

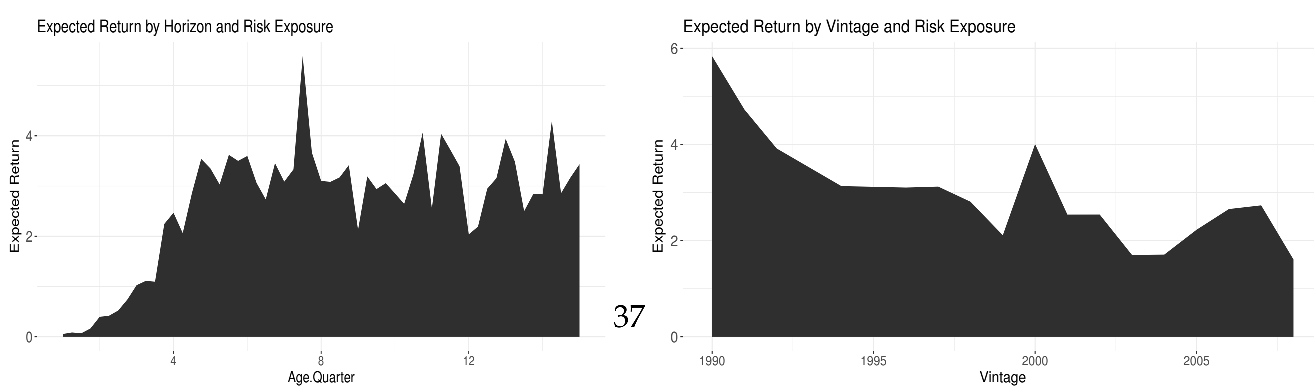

Figure 13: Risk-Adjusted Profits by Category

Panel A: Buyout

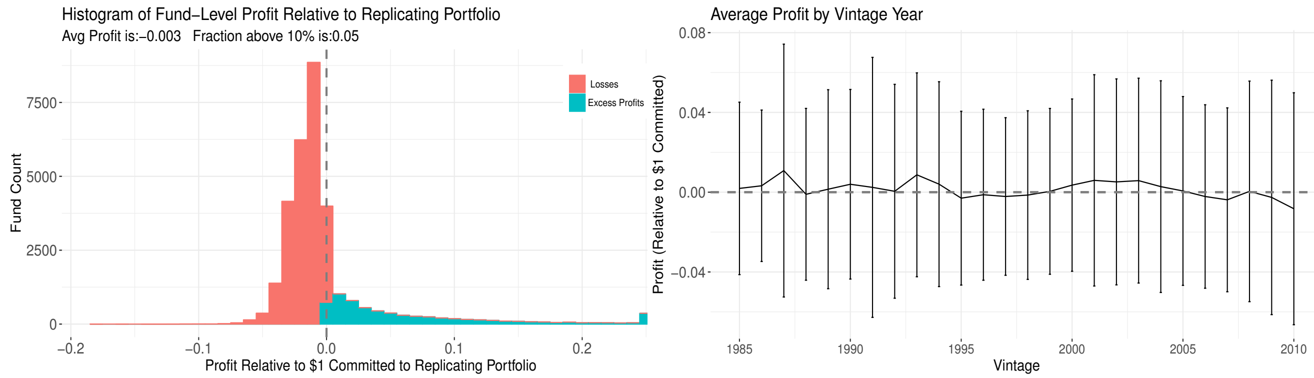

Panel B: Venture Capital

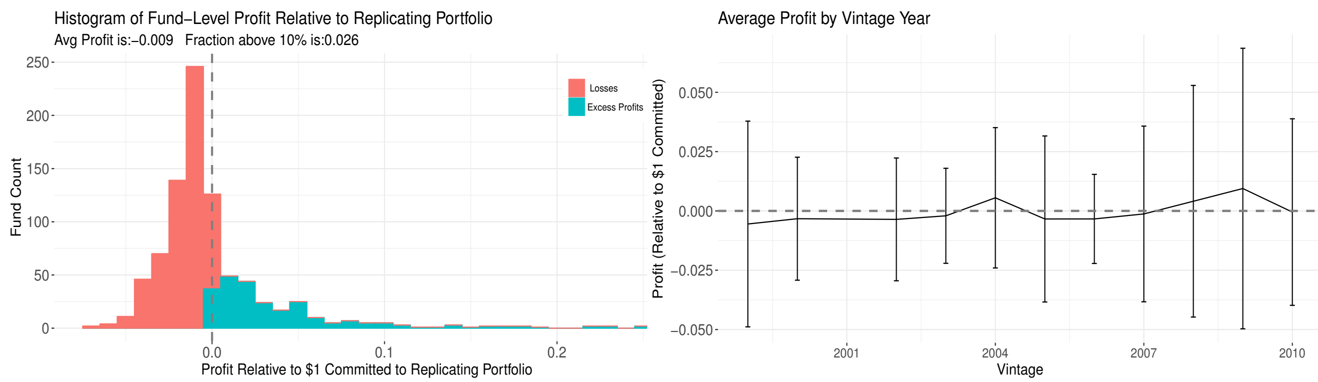

Panel C: Real Estate Funds

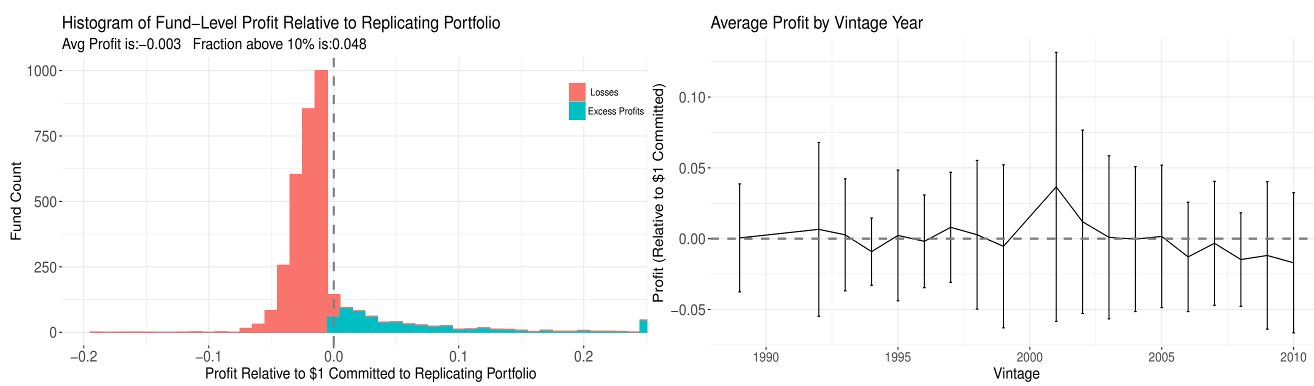

Panel D: Infrastructure Funds

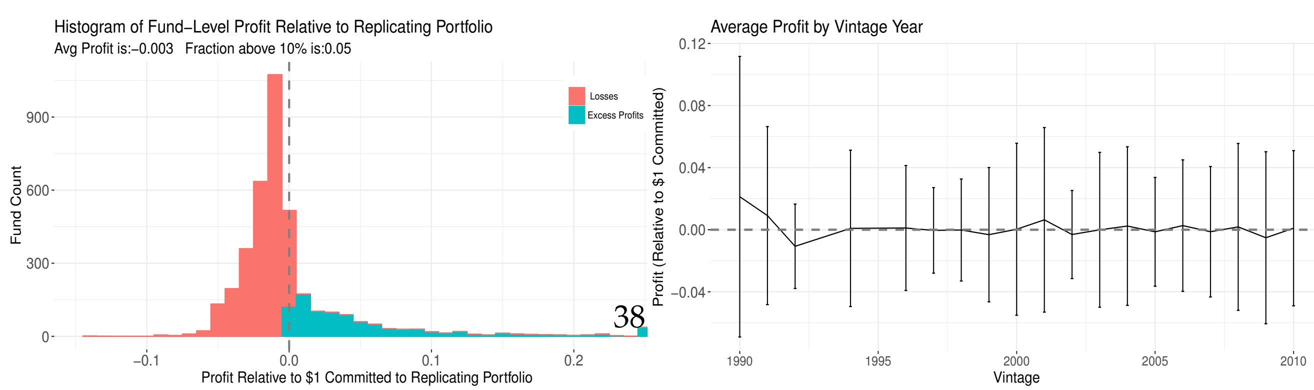

4.3 Four-factor Model

Next, we turn to our main results for the four factor model. Figure 14 show the estimated horizon effects $\widehat{b}_h$ in the left panels and the vintage effects $\widehat{a}_t$ in the right panels for the four main investment categories. The appendix contains the same plot for the three remaining categories. Our model illustrates a rich pattern of factor risk exposures. The overall R^2 model fit improves substantially from that of the one-factor model, with R^2 values between 9-19%. All PE fund categories tend to load strongly positively on REITs; all categories also have positive infrastructure risk exposure after year five. Offsetting the REIT and infrastructure exposure is negative stock market exposure. These patterns may be explained by REIT cash-flow distributions that have a strong pro-cyclical component, mirroring the cash flow pattern in PE distributions.

Expected Return The asset pricing model provides the expected return on each zero-coupon bond and each stock, REIT, and infra dividend strip. With the replicating portfolio in hand, we can calculate the expected return on PE funds in each investment category as in (3). Figure 15 breaks down the expected return by horizon, stacking the contribution from each security exposure. Each panel is for one investment category.

We find generally higher expected returns from the four-factor model than from the one-factor model, around 10%. Much of the expected return is earned for far out cash flows, despite those cash flows being relatively less important. The different patterns across horizons illustrates the value of breaking down PE returns strip-by-strip; while the variation across categories illustrates that private funds load very differently on risk exposures depending on industry focus. In the time series, we generally observe falling expected returns over vintages. The reason is low expected returns especially on REITs and Infrastructure stocks towards the end of the sample; see the bottom left panels of Figure 3. The differences in exposure across fund categories suggest limitations of existing measures in fund performance evaluation, such as the PME, which do not account for different risk exposures of different funds. Appendix Figure A.8 extends the analysis for the remaining three fund categories.

Performance Evaluation Figure 16 plots the histogram of risk-adjusted profits for the four-factor model, pooling all vintages, in the left panels. The right panels plot the average risk-adjusted profit for each vintage. The rows contain LBO, VC, RE, and IN funds. The appendix contains the same figures for the remaining investment categories. Like in the one-factor model, average profits are close to zero with considerable variation around them. Average profits in the later vintages are lower in the four- than in the one-factor model, underscoring the importance of incorporating the traded equity risk factors.The cross-sectional dispersion of profits is also more tightly concentrated around zero in these estimates, suggesting a tighter fit using the replicating portfolio.

All fund categories suggest negative but modest risk-adjusted profits under the four-factor model, averaging around -0.3 to -0.9%. Our estimates would suggest that apparent high returns in the VC sector primarily reflect high loadings on risk factors, and therefore high expected returns, as opposed to abnormal returns beyond the yield of a replicating portfolio.

Figure 14: Replicating Portfolio Exposure by Feature: Multi-Factor Model

Panel A: Buyout

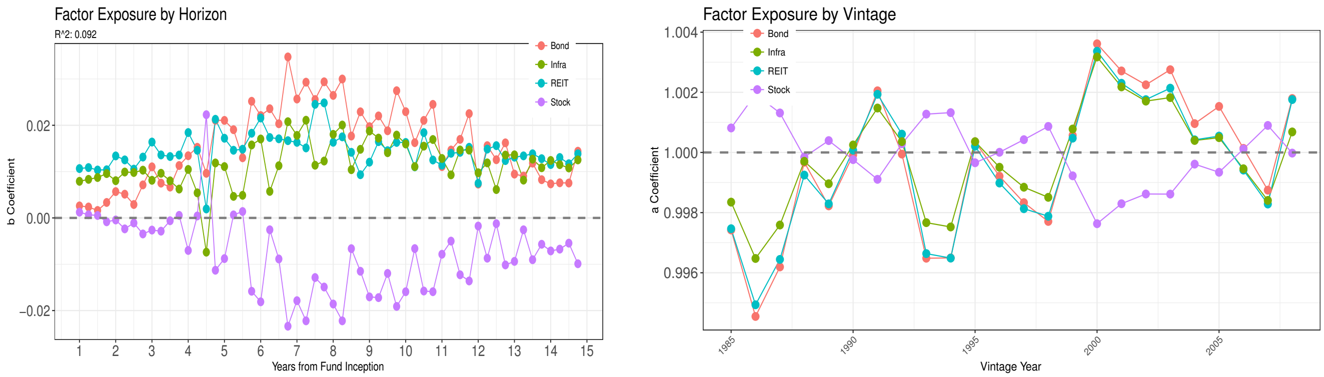

Panel B: Venture Capital

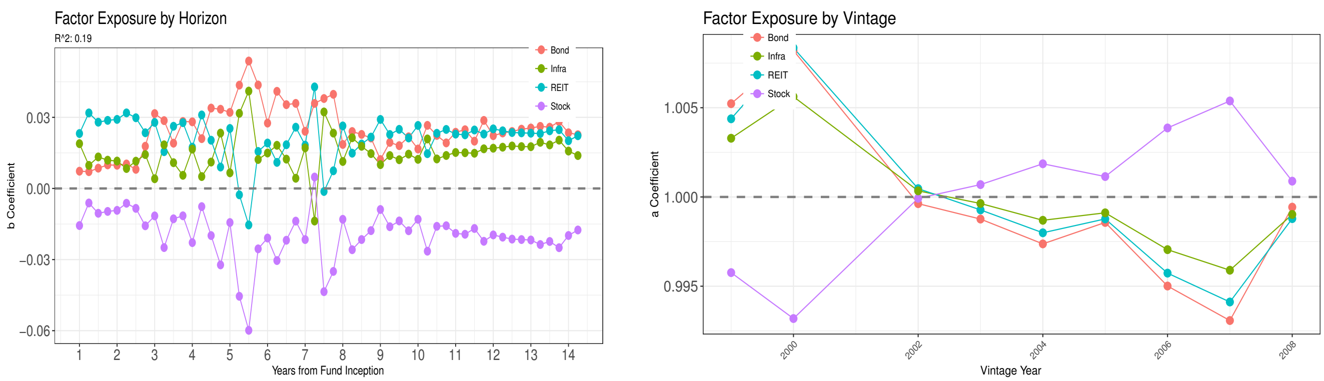

Panel C: Real Estate

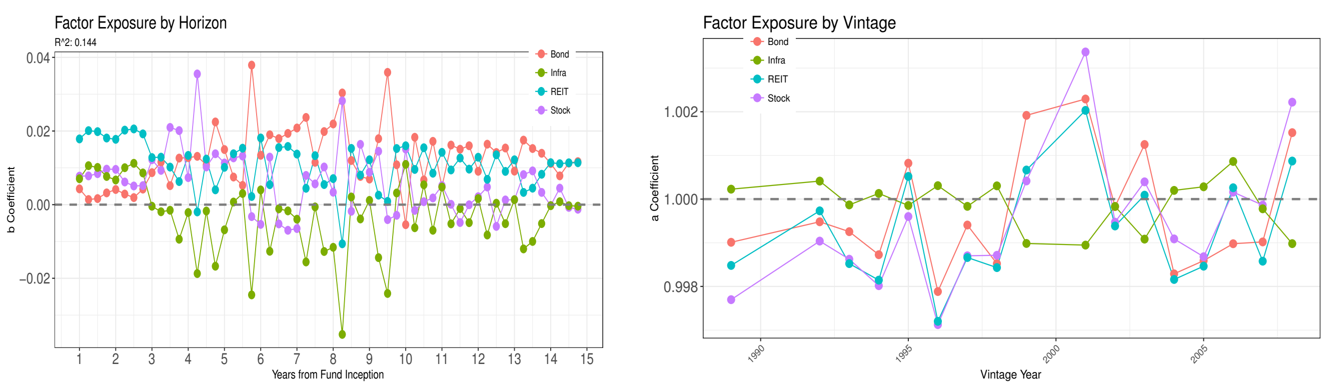

Panel D: Infrastructure

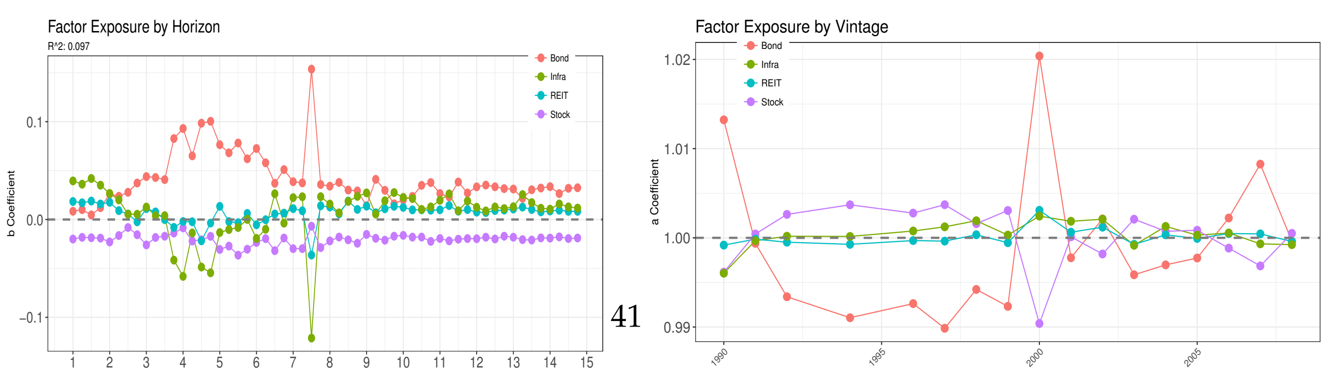

Figure 15: Expected Return

Panel A: Buyout

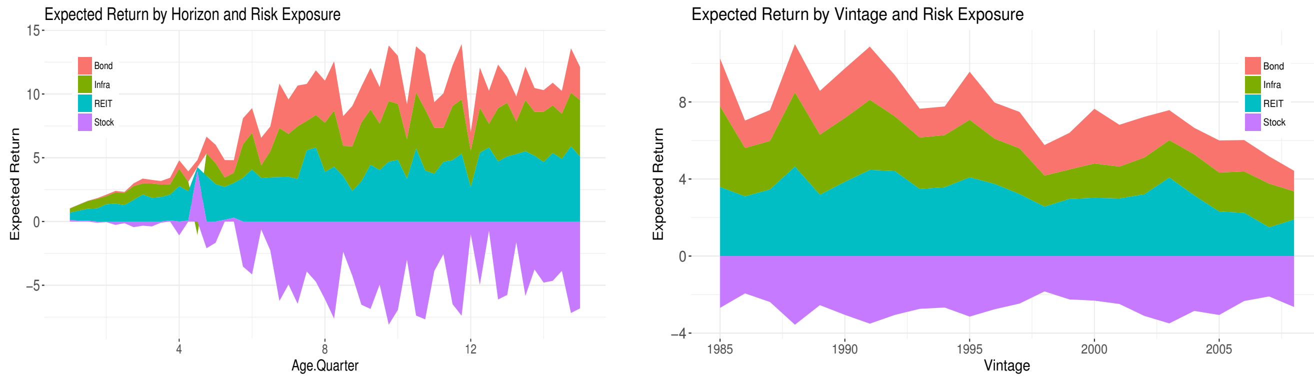

Panel B: Venture Capital

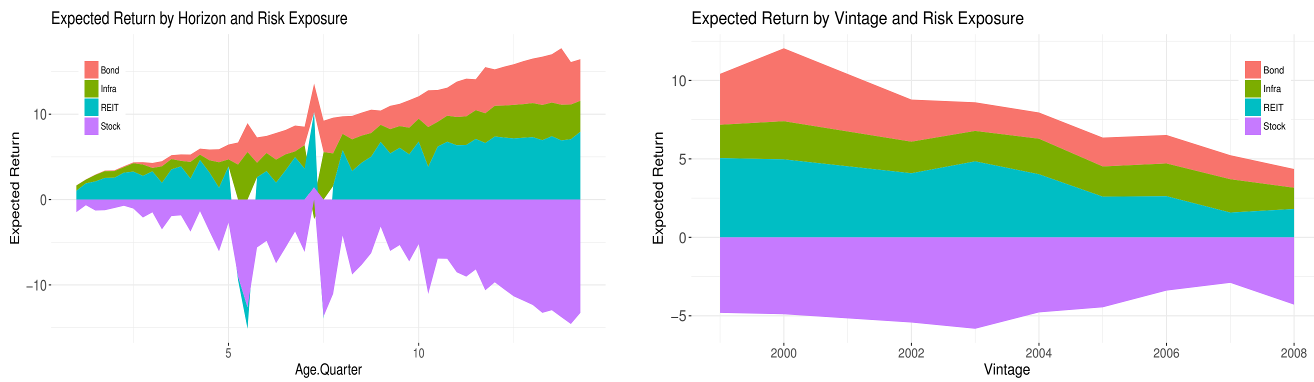

Panel C: Real Estate

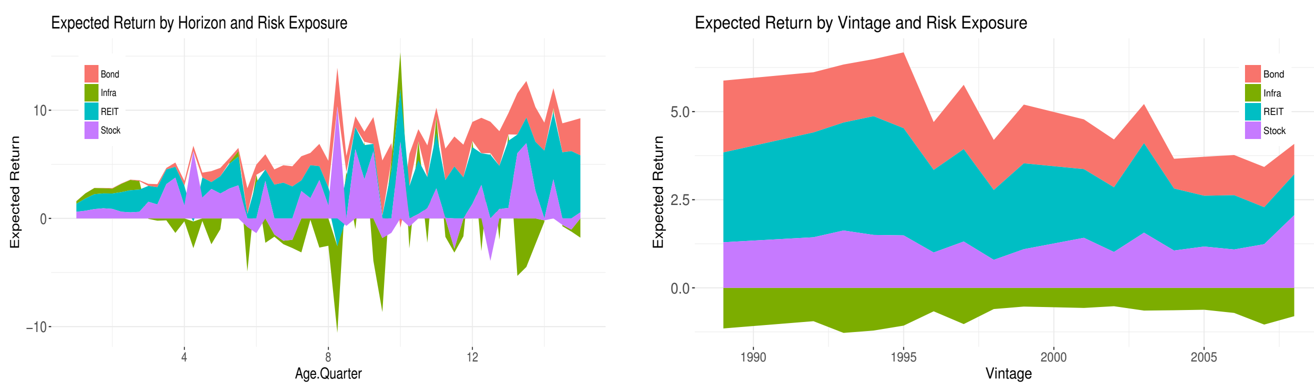

Panel D: Infrastructure

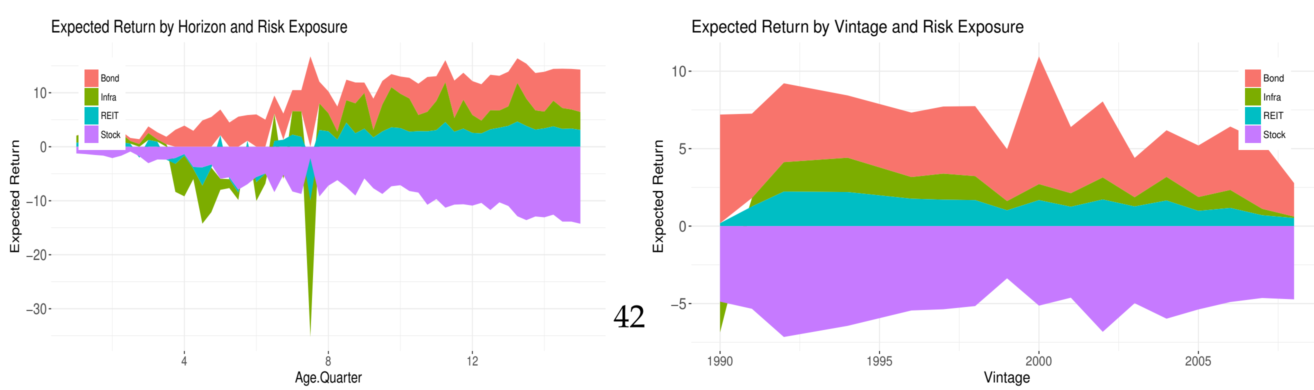

Figure 16: Risk-Adjusted Profits by Category: Multi-Factor Model

Panel A: Buyout

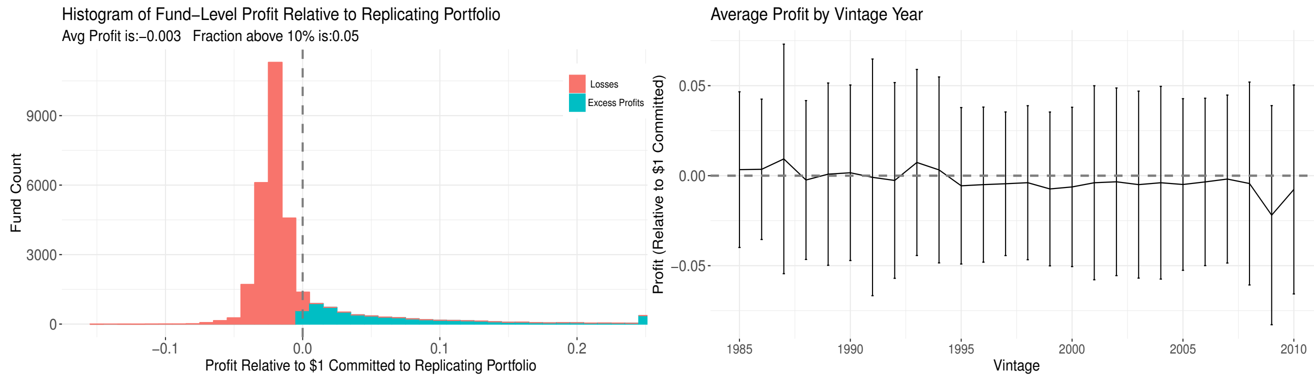

Panel B: Venture Capital

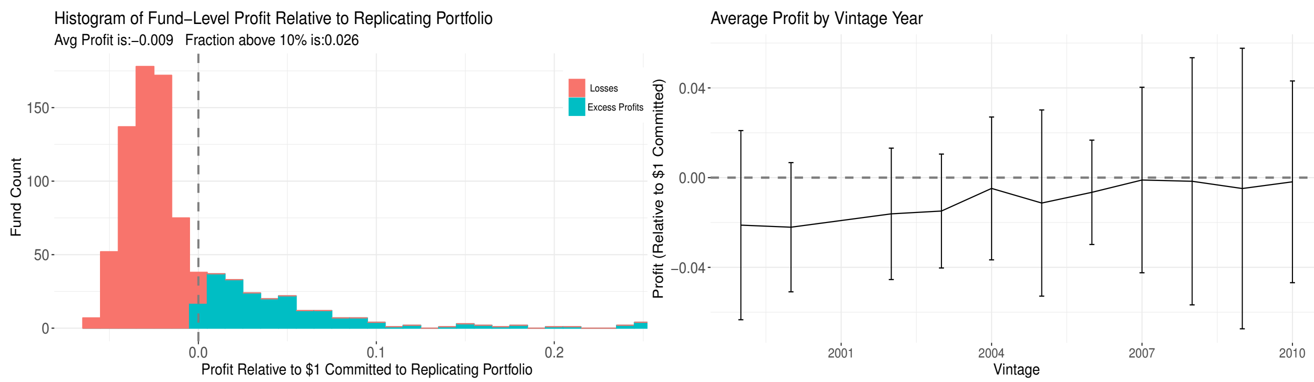

Panel C: Real Estate Funds

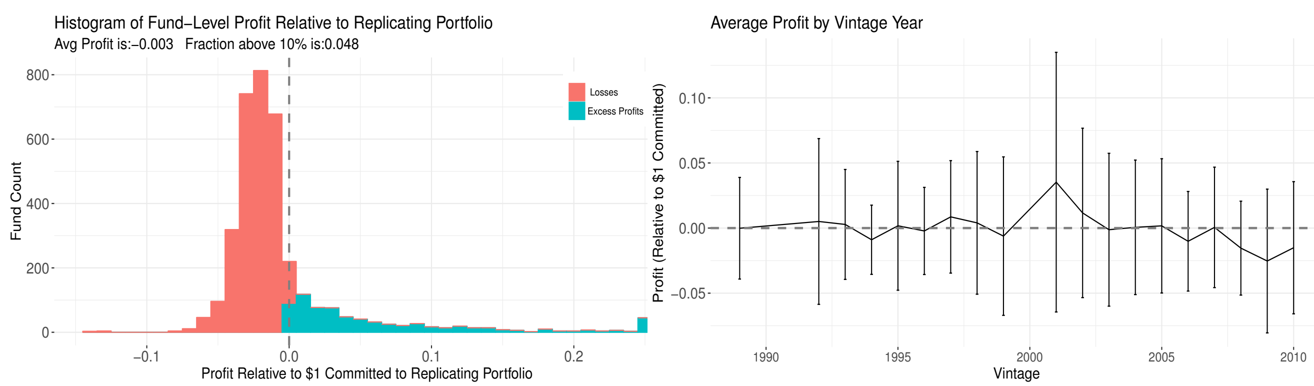

Panel D: Infrastructure Funds

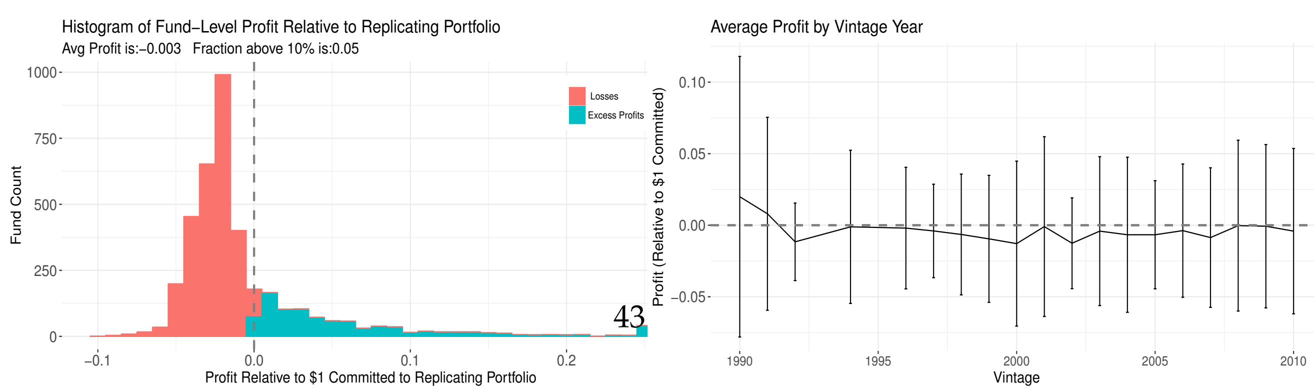

5 Conclusion

We provide a novel valuation method for private equity cash-flows that decomposes a private equity cash-flow at each horizon into a systematic component that reflects exposure to aggregate sources of risk, priced in listed securities markets, and an idiosyncratic component which is the risk-adjusted profit to the PE investor. The systematic component reflects the value of a portfolio of stock and bond strips, which pay safe or risky cash flows at horizons that match the horizon 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, fitting the time series bond yields and stock prices, including dividend strips, closely. The method improves on existing PE fund valuation techniques by considering exposure to multiple risk factors, decomposing risk into horizon-dependent components, and allowing for heterogeneity in systematic risk across PE vintages and categories.

We find that the average private equity fund generates little outperformance across most of the categories we consider. This is because private fund cash-flows can be replicated, at least on average, through a basket of publicly traded equivalents. However, we also document rich heterogeneity across horizons, in the cross-section, and in the time-series in terms of fund performance and expected returns. Average fund performance trends downward over time.

References

Ammar, S. B., and M. Eling, 2015, "Common risk factors of infrastructure investments," *Energy Economics*.

Ang, A., B. Chen, W. N. Goetzmann, and L. Phalippou, 2017, "Estimating Private Equity Returns from Limited Partner Cash Flows," *Netspar Discussion Paper No. 06/2014-021; Sad Business School WP 2014-8*.

Ang, A., and M. Piazzesi, 2003, "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables," *Journal of Monetary Economics*, 50, 745–787.

Borovička, J., and L. P. Hansen, 2014, "Examining Macroeconomic Models through the Lens of Asset Pricing," *Journal of Econometrics*, 183(1), 67–90.

Campbell, J. Y., 1991, "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.

——— , 1993, "Intertemporal Asset Pricing Without Consumption Data," *American Economic Review*, 83(3), 487–511.

——— , 1996, "Understanding Risk and Return," *The Journal of Political Economy*, 104(2), 298–345.

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

Cochrane, J. H., and J. Saa-Requejo, 2000, "Beyond Arbitrage: 'Good deal' asset price bounds in incomplete markets," *Journal of Political Economy*, 108, 79–119.

Dai, Q., and K. J. Singleton, 2000, "Specification Analysis of Affine Term Structure Models," *Journal of Finance*, 55, 1943–1978.

Driessen, J., T.-C. Lin, and L. 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(3), 511–535.

Duffie, D., and R. Kan, 1996, "A Yield Factor Model of Interest Rates," *Mathematical Finance*, 6, 379–406.

Ewens, M., C. M. Jones, and M. Rhodes-Kropf, 2013, "The Price of Diversifiable Risk in Venture Capital and Private Equity," *Review of Financial Studies*, 26, 1854–1889.

Gompers, P. A., and J. Lerner, 1997, "Risk and Reward in Private Equity Investments: The Challenge of Performance Assessment," *Journal of Private Equity*, 1, 5–12.

Hansen, L. P., and J. Scheinkman, 2009, "Long-Term Risk: An Operator Approach," *Econometrica*, 77 (1), 177–234.

Harris, R. S., T. Jenkinson, and S. N. Kaplan, 2014, "Private Equity Performance: What Do We Know?," *Journal of Finance*, 69(5).

Kaplan, S. N., and A. Schoar, 2005, "Private equity performance: Returns, persistence, and capital flows," *Journal of Finance*, 60(4), 1791–1823.

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

Korteweg, A., and M. Sorensen, 2010, "Risk and return characteristics of venture capital-backed entrepreneurial companies," *Review of Financial Studies*, 23(10), 3738–3772.

——— , 2017, "Skill and Luck in Private Equity Performance," *Journal of Financial Economics*, 0.

Lettau, M., and J. Wachter, 2011, "The Term Structures of Equity and Interest Rates," *Journal of Financial Economics*, 101 (1), 90–113.

Ljungqvist, A., and M. Richardson, 2003, "The cash flow, return and risk characteristics of private equity," *Working Paper, NYU Stern*.

Lustig, H., S. Van Nieuwerburgh, and A. Verdelhan, 2013, "The Wealth-Consumption Ratio," *Review of Asset Pricing Studies*, 3(1), 38–94., Review of Asset Pricing Studies.

Metrick, A., and A. Yasuda, 2010, "The economics of private equity funds," *Review of Financial Studies*, 23(6), 2303–2341.

Moskowitz, T. J., and A. Vissing-Jorgensen, 2002, "The returns to entrepreneurial investment: A private equity premium puzzle?," *American Economic Review*, 92, 745–778.

Peng, L., 2001, "Building a Venture Capital Index," *Working Paper*.

———, 2016, "The risk and return of commercial real estate: A property level analysis," *Real Estate Economics*, 44(3), 555–583.

Phalippou, L., and O. Gottschalg, 2009, "The performance of private equity funds," *Review of Financial Studies*, 22(4), 1747–1776.

Robinson, D. T., and B. A. Sensoy, 2011, "Cyclicality, Performance Measurement, and Cash Flow Liquidity in Private Equity," *Working Paper*.

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

Sorensen, M., and R. Jagannathan, 2015, "The Public Market Equivalent and Private Equity Performance," *Financial Analysts Journal*, 71(4), 43–50.

Sorensen, M., N. Wang, and J. Yang, 2014, "Valuing private equity," *Review of Financial Studies*, 27(7), 1977–2021.

van Binsbergen, J., M. Brandt, and R. Koijen, 2012, "On the Timing and Pricing of Dividends," *American Economic Review*, 102(4), 1596–1618.

van Binsbergen, J. H., W. H. Hueskes, R. S. J. Koijen, and E. B. Vrugt, 2013, "Equity Yields," *Journal of Financial Economics*, 110(3), 503–519.

van Binsbergen, J. H., and R. S. J. Koijen, 2017, "The Term Structure of Returns: Facts and Theory," *Journal of Financial Economics*, 124(1), 1.

Wachter, J., 2005, "Solving Models with External Habit," *Finance Research Letters*, 2, 210–226.

Woodward, S., 2009, "Measuring Risk for Venture Capital and Private Equity Portfolios," *Working Paper*.

A Appendix: Asset Pricing Model

A.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:

$$\begin{aligned}
y_{t,1} &= y^{\$}_{t,1} - E_t[\pi_{t+1}] - \frac{1}{2}e'_{\pi}\Sigma e_{\pi} + e'_{\pi}\Sigma^{\frac{1}{2}}\Lambda_t \\
&= y_0(1) + \left[e'_{yn} - e'_{\pi}\Psi + e'_{\pi}\Sigma^{\frac{1}{2}}\Lambda_1\right] z_t. &\text{(A.1)}\\
y_0(1) &\equiv y^{\$}_{0,1} - \pi_0 - \frac{1}{2}e'_{\pi}\Sigma e_{\pi} + e'_{\pi}\Sigma^{\frac{1}{2}}\Lambda_0. &\text{(A.2)}
\end{aligned}$$

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.

A.2 Nominal and real term structure

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

$$y^{\$}_t(\tau) = -\frac{A^{\$}_{\tau}}{\tau} - \frac{B^{\$\prime}_{\tau}}{\tau} z_t,$$

where the coefficients $A^{\$}_{\tau}$ and $B^{\$}_{\tau}$ satisfy the following recursions:

$$\begin{aligned}
A^{\$}_{\tau+1} &= -y^{\$}_{0,1} + A^{\$}_{\tau} + \frac{1}{2}\left(B^{\$}_{\tau}\right)'\Sigma\left(B^{\$}_{\tau}\right) - \left(B^{\$}_{\tau}\right)'\Sigma^{\frac{1}{2}}\Lambda_0, &\text{(A.3)}\\
\left(B^{\$}_{\tau+1}\right)' &= \left(B^{\$}_{\tau}\right)'\Psi - e'_{yn} - \left(B^{\$}_{\tau}\right)'\Sigma^{\frac{1}{2}}\Lambda_1, &\text{(A.4)}
\end{aligned}$$

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:

$$\log(P^{\$}_{t+1,\tau}) = A^{\$}_{\tau} + \left(B^{\$}_{\tau}\right)' z_{t+1}$$

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

$$\begin{aligned}
P^{\$}_{t,\tau+1} &= E_t[\exp\{m^{\$}_{t+1} + \log\left(P^{\$}_{t+1,\tau}\right)\}] \\
&= E_t[\exp\{-y^{\$}_{t,1} - \frac{1}{2}\Lambda_t'\Lambda_t - \Lambda_t'\varepsilon_{t+1} + A^{\$}_{\tau} + \left(B^{\$}_{\tau}\right)' z_{t+1}\}] \\
&= \exp\{-y^{\$}_{0,1} - e_{yn}' z_t - \frac{1}{2}\Lambda_t'\Lambda_t + A^{\$}_{\tau} + \left(B^{\$}_{\tau}\right)' \Psi z_t\} \times \\
&\quad E_t\left[\exp\{-\Lambda_t'\varepsilon_{t+1} + \left(B^{\$}_{\tau}\right)' \Sigma^{\frac{1}{2}}\varepsilon_{t+1}\}\right].
\end{aligned}$$

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

$$\begin{aligned}
P^{\$}_{t,\tau+1} &= \exp\left\{-y^{\$}_{0,1} - e_{yn}' z_t + A^{\$}_{\tau} + \left(B^{\$}_{\tau}\right)' \Psi z_t + \frac{1}{2}\left(B^{\$}_{\tau}\right)' \Sigma \left(B^{\$}_{\tau}\right)\right. \\
&\quad \left. - \left(B^{\$}_{\tau}\right)' \Sigma^{\frac{1}{2}}(\Lambda_0 + \Lambda_1 z_t)\right\}.
\end{aligned}$$

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

$$\log\left(P^{\$}_{t,\tau+1}\right) = A^{\$}_{\tau+1} + \left(B^{\$}_{\tau+1}\right)' z_t,$$

where $A^{\$}_{\tau+1}$ satisfies (A.3) and $B^{\$}_{\tau+1}$ satisfies (A.4). The relationship between log bond prices and bond yields is given by $-\log\left(P^{\$}_{t,\tau}\right)/\tau = y^{\$}_{t,\tau}$. □

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

$$r^{b,\$}_{t+1,\tau} = \log\left(P^{\$}_{t+1,\tau}\right) - \log\left(P^{\$}_{t,\tau+1}\right)$$

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:

$$\begin{aligned}
E_t\left[r^{b,\$}_{t+1,\tau}\right] - y^{\$}_{t,1} + \frac{1}{2}V_t\left[r^{b,\$}_{t+1,\tau}\right] &= -Cov_t\left[m^{\$}_{t+1}, r^{b,\$}_{t+1,\tau}\right] \\
&= \left(B^{\$}_{\tau}\right)' \Sigma^{\frac{1}{2}}\Lambda_t
\end{aligned}$$

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

$$
\begin{aligned}
A_{\tau+1} &= -y_{0,1} + A_\tau + \frac{1}{2}B'_\tau \Sigma B_\tau - B'_\tau \Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime} e_\pi\right), && \text{(A.5)} \\
B'_{\tau+1} &= -e'_{yn} + (e_\pi + B_\tau)' \left(\Psi - \Sigma^{\frac{1}{2}} \Lambda_1\right). && \text{(A.6)}
\end{aligned}
$$

For $\tau = 1$, we recover the expression for the risk-free rate in (A.1)-(A.2).

A.3 Stock Market

We define the real return on equity as $R^m_{t+1} = \frac{P^m_{t+1} + D^m_{t+1}}{P^m_t}$, where P^m_t is the end-of-period price on the equity market. A log-linearization delivers:

$$
r^m_{t+1} = \kappa^m_0 + \Delta d^m_{t+1} + \kappa^m_1 pd^m_{t+1} - pd^m_t. \tag{A.7}
$$

The unconditional mean log real stock return is $r^m_0 = E[r^m_t]$, the unconditional mean dividend growth rate is $\mu^m = E[\Delta d^m_{t+1}]$, and $\overline{pd^m} = E[pd^m_t]$ is the unconditional average log price-dividend ratio on equity. The linearization constants κ^m_0 and κ^m_1 are defined as:

$$
\kappa^m_1 = \frac{e^{\overline{pd^m}}}{e^{\overline{pd^m}} + 1} < 1 \ \text{ and } \ \kappa^m_0 = \log\left(e^{\overline{pd^m}} + 1\right) - \frac{e^{\overline{pd^m}}}{e^{\overline{pd^m}} + 1}\overline{pd^m}. \tag{A.8}
$$

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

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

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 $\{\Delta d^m_t, pd^m_t\}$. Rewriting (A.7):

$$
\begin{aligned}
r^m_{t+1} - r^m_0 &= \left[(e_{divm} + \kappa^m_1 e_{pd})'\Psi - e'_{pd}\right] z_t + \left(e_{divm} + \kappa^m_1 e_{pd}\right)' \Sigma^{\frac{1}{2}} \varepsilon_{t+1}. \\
r^m_0 &= \mu^m + \kappa^m_0 - \overline{pd^m}(1 - \kappa^m_1).
\end{aligned}
$$

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:

$$
\begin{aligned}
1 &= E_t\left[M_{t+1}\frac{P^m_{t+1}+D^m_{t+1}}{P^m_t}\right] = E_t\left[\exp\{m^\$_{t+1}+\pi_{t+1}+r^m_{t+1}\}\right] \\
&= E_t\left[\exp\left\{-y^\$_{t,1}-\frac{1}{2}\Lambda_t'\Lambda_t-\Lambda_t'\varepsilon_{t+1}+\pi_0+e_\pi' z_{t+1}+r^m_0+(e_{divm}+\kappa^m_1 e_{pd})'z_{t+1}-e_{pd}'z_t\right\}\right] \\
&= \exp\left\{-y^\$_0(1)-\frac{1}{2}\Lambda_t'\Lambda_t+\pi_0+r^m_0+\left[(e_{divm}+\kappa^m_1 e_{pd}+e_\pi)'\Psi-e_{pd}'-e_{yn}'\right]z_t\right\} \\
&\quad \times E_t\left[\exp\{-\Lambda_t'\varepsilon_{t+1}+\left(e_{divm}+\kappa^m_1 e_{pd}+e_\pi\right)'\Sigma^{\frac{1}{2}}\varepsilon_{t+1}\right] \\
&= \exp\left\{r^m_0+\pi_0-y^\$_0(1)+\left[(e_{divm}+\kappa^m_1 e_{pd}+e_\pi)'\Psi-e_{pd}'-e_{yn}'\right]z_t\right\} \\
&\quad \times \exp\left\{\frac{1}{2}\left(e_{divm}+\kappa^m_1 e_{pd}+e_\pi\right)'\Sigma\left(e_{divm}+\kappa^m_1 e_{pd}+e_\pi\right)-\left(e_{divm}+\kappa^m_1 e_{pd}+e_\pi\right)'\Sigma^{\frac{1}{2}}\Lambda_t\right\}
\end{aligned}
$$

Taking logs on both sides delivers:

$$
\begin{aligned}
r^m_0+\pi_0-y^\$_0(1)+\left[(e_{divm}+\kappa^m_1 e_{pd}+e_\pi)'\Psi-e_{pd}'-e_{yn}'\right]z_t & \\
+\frac{1}{2}\left(e_{divm}+\kappa^m_1 e_{pd}+e_\pi\right)'\Sigma\left(e_{divm}+\kappa^m_1 e_{pd}+e_\pi\right) &= \left(e_{divm}+\kappa^m_1 e_{pd}+e_\pi\right)'\Sigma^{\frac{1}{2}}\Lambda_t \\
E_t\left[r^{m,\$}_{t+1}\right]-y^\$_{t,1}+\frac{1}{2}V_t\left[r^{m,\$}_{t+1}\right] &= -Cov_t\left[m^\$_{t+1},r^{m,\$}_{t+1}\right]
\end{aligned}
\tag{A.9}
$$

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:

$$
\begin{aligned}
E_t\left[r^m_{t+1}\right]-y_{t,1}+\frac{1}{2}V_t\left[r^m_{t+1}\right] &= -Cov_t\left[m_{t+1},r^m_{t+1}\right] \\
r^m_0-y_0(1) &+ \left[(e_{divm}+\kappa^m_1 e_{pd}+e_\pi)'\Psi-e_{pd}'-e_{yn}'-e_\pi'\Sigma^{1/2}\Lambda_1\right]z_t \\
+\frac{1}{2}(e_{divm}+\kappa^m_1 e_{pd})'\Sigma(e_{divm}+\kappa^m_1 e_{pd}) &= \left(e_{divm}+\kappa^m_1 e_{pd}\right)'\Sigma^{1/2}(\Lambda_t-\left(\Sigma^{1/2}\right)'e_\pi)
\end{aligned}
$$

We combine the terms in Λ_0 and Λ_1 on the right-hand side and plug in for $y_0(1)$ from (A.2) to get:

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

This recovers equation (A.9).

A.4 Dividend Strips

A.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:

$$
p_{t,\tau}^d = \log\left(P_{t,\tau}^d\right) = A_\tau^m + B_\tau^{m\prime} z_t,
$$

where the coefficients A_τ^m and B_τ^m follow recursions:

$$
\begin{aligned}
A_{\tau+1}^m &= A_\tau^m + \mu_m - y_0(1) + \frac{1}{2}\left(e_{divm} + B_\tau^m\right)' \Sigma \left(e_{divm} + B_\tau^m\right) \\
&\quad - \left(e_{divm} + B_\tau^m\right)' \Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime} e_\pi\right), \qquad \text{(A.10)} \\
B_{\tau+1}^{m\prime} &= \left(e_{divm} + e_\pi + B_\tau^m\right)' \Psi - e_{yn}' - \left(e_{divm} + e_\pi + B_\tau^m\right)' \Sigma^{\frac{1}{2}} \Lambda_1, \qquad \text{(A.11)}
\end{aligned}
$$

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:

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

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

$$\begin{aligned} P^d_{t,\tau+1} &= \exp\{-y^{\$}_0(1) + \pi_0 + \mu_m + A^m_\tau + \left[(e_{divm} + e_\pi + B^m_\tau)' \Psi - e'_{yn}\right] z_t \\ &+ \frac{1}{2}\left(e_{divm} + e_\pi + B^m_\tau\right)' \Sigma \left(e_{divm} + e_\pi + B^m_\tau\right) \\ &- \left(e_{divm} + e_\pi + B^m_\tau\right)' \Sigma^{\frac{1}{2}} (\Lambda_0 + \Lambda_1 z_t)\} \end{aligned}$$

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

$$p^d_{t,\tau+1} = A^m_{\tau+1} + B^{m\prime}_{\tau+1} z_t,$$

where:

$$\begin{aligned} A^m_{\tau+1} &= A^m_\tau + \mu_m - y^{\$}_0(1) + \pi_0 + \frac{1}{2}\left(e_{divm} + e_\pi + B^m_\tau\right)' \Sigma \left(e_{divm} + e_\pi + B^m_\tau\right) \\ &- \left(e_{divm} + e_\pi + B^m_\tau\right)' \Sigma^{\frac{1}{2}} \Lambda_0, \\ B^{m\prime}_{\tau+1} &= \left(e_{divm} + e_\pi + B^m_\tau\right)' \Psi - e'_{yn} - \left(e_{divm} + e_\pi + B^m_\tau\right)' \Sigma^{\frac{1}{2}} \Lambda_1. \end{aligned}$$

We recover the recursions in (A.10) and (A.11) after using equation (A.2). □

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

$$\begin{aligned} \mathbb{E}_t\left[r^{d,\$}_{t+1,\tau}\right] - y^{\$}_{t,1} + \frac{1}{2} V_t\left[r^{d,\$}_{t+1,\tau}\right] &= -Cov_t\left[m^{\$}_{t+1}, r^{d,\$}_{t+1,\tau}\right] \\ &= \left(e_{divm} + e_\pi + B^m_\tau\right)' \Sigma^{\frac{1}{2}} \Lambda_t \end{aligned}$$

The risky strips for REITs and infrastructure are defined analogously.

A.4.2 Strip Expected Holding Period Return over k-horizons

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

$$\begin{aligned} \mathbb{E}_t[R_{t\to t+k}] &= \frac{\mathbb{E}_t\left[\frac{D^{\$}_{t+k}}{D^{\$}_t}\right]}{P^d_{t,k}} - 1 \\ &= \exp\left(-A^m_k - B^{m\prime}_k z_t + \mathbb{E}_t\left[\sum_{s=1}^{k} \Delta d_{t+s} + \pi_{t+s}\right] + \frac{1}{2} V_t\left[\sum_{s=1}^{k} \Delta d_{t+s} + \pi_{t+s}\right]\right) - 1 \end{aligned}$$

$$
\begin{aligned}
&= \exp\left(-A_k^m - B_k^{m\prime} z_t + k(\mu_m + \pi_0) + (e_{divm} + e_\pi)' \left[\sum_{s=1}^{k} \Psi^s\right] z_t \right. \\
&\quad \left. + \frac{k}{2}(e_{divm} + e_\pi)' \Sigma (e_{divm} + e_\pi)\right) - 1 \qquad \text{(A.12)}
\end{aligned}
$$

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

A.4.3 Strip Forward 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:

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

where $P^d_{t,\tau}$ 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:

$$
\frac{F^d_{t,\tau}}{D^\$_t} = \exp\left(A^m_\tau - A^\$_\tau + (B^m_\tau - B^\$_\tau)' z_t\right),
$$

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

$$
R^{fut,d}_{t+1,\tau} = \frac{F^d_{t+1,\tau-1}}{F^d_{t,\tau}} - 1 = \frac{F^d_{t+1,\tau-1}/D^\$_{t+1}}{F^d_{t,\tau}/D^\$_t} \frac{D^\$_{t+1}}{D^\$_t} - 1
$$

It can be written as:

$$
\begin{aligned}
\log\left(1 + R^{fut,d}_{t+1,\tau}\right) &= A^m_{\tau-1} - A^\$_{\tau-1} - A^m_\tau + A^\$_\tau + \mu_m + \pi_0 \\
&\quad + (B^m_{\tau-1} - B^\$_{\tau-1} + e_{divm} + e_\pi)' z_{t+1} - (B^m_\tau - B^\$_\tau)' z_t
\end{aligned}
$$

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:

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

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

A.4.4 Decomposing Strip Price into Deterministic and Risky Components

A claim to the dividend on a risky claim at time $t+k$, D_{t+k}, can be decomposed in the value of a claim to the deterministic component $D^{det}_{t+k} = D_t \exp(k\mu_m)$ and a risky component $D^{risky}_{t+k} = D_{t+k} - D^{det}_{t+k} = D_{t+k} - D_t \exp(k\mu_m)$. By no arbitrage, the price-dividend ratio of the risky claim must be the difference between the price-dividend ratio of the entire claim and the price-dividend ratio of the deterministic claim. We conjecture and verify that the price-dividend ratio of the deterministic claim is also affine in the state vector and solve for the coefficients from the Eurler equation.

Proof. We conjecture the affine structure for the price-dividend ratio and solve for the coefficients $A^{det}_{\tau+1}$ and $B^{det}_{\tau+1}$ in the process of verifying this conjecture using the Euler equation:

$$\begin{aligned}
P^{det}_{t,\tau+1} &= \mathbb{E}_t\left[M_{t+1}P^{det}_{t+1,\tau}\frac{D^{det,m}_{t+1}}{D^m_t}\right] = E_t\left[\exp\{m^{\$}_{t+1} + \pi_{t+1} + \mu^m + p^d_{t+1}(\tau)\}\right] \\
&= \mathbb{E}_t\left[\exp\{-y^{\$}_{t,1} - \frac{1}{2}\Lambda_t'\Lambda_t - \Lambda_t'\varepsilon_{t+1} + \pi_0 + e_\pi' z_{t+1} + \mu^m + A^{det}_\tau + B^{det\prime}_\tau z_{t+1}\}\right] \\
&= \exp\{-y^{\$}_0(1) - e_{yn}' z_t - \frac{1}{2}\Lambda_t'\Lambda_t + \pi_0 + e_\pi'\Psi z_t + \mu_m + A^{det}_\tau + B^{det\prime}_\tau \Psi z_t\} \\
&\quad \times \mathbb{E}_t\left[\exp\{-\Lambda_t'\varepsilon_{t+1} + (e_\pi + B^m_\tau)'\,\Sigma^{\frac{1}{2}}\varepsilon_{t+1}\right].
\end{aligned}$$

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

$$\begin{aligned}
P^{det}_{t,\tau+1} &= \exp\{-y^{\$}_0(1) + \pi_0 + \mu_m + A^{det}_\tau + \left[\left(e_\pi + B^{det}_\tau\right)'\Psi - e_{yn}'\right] z_t \\
&\quad + \frac{1}{2}\left(e_\pi + B^{det}_\tau\right)'\Sigma\left(e_\pi + B^{det}_\tau\right) - \left(e_\pi + B^{det}_\tau\right)'\Sigma^{\frac{1}{2}}(\Lambda_0 + \Lambda_1 z_t)\}
\end{aligned}$$

Taking logs and collecting terms, we obtain a log-linear expression for $p^{det}_t(\tau+1)$:

$$p^{det}_{t,\tau+1} = A^{det}_{\tau+1} + B^{det\prime}_{\tau+1} z_t,$$

where:

$$\begin{aligned} A^{det}_{\tau+1} &= A^{det}_{\tau} + \mu_m - y^{\$}_0(1) + \pi_0 + \frac{1}{2}\left(e_\pi + B^{det}_{\tau}\right)'\Sigma\left(e_\pi + B^{det}_{\tau}\right) \\ &\quad - \left(e_\pi + B^{det}_{\tau}\right)'\Sigma^{\frac{1}{2}}\Lambda_0, \\ B^{det\prime}_{\tau+1} &= \left(e_\pi + B^{det}_{\tau}\right)'\Psi - e'_{yn} - \left(e_\pi + B^{det}_{\tau}\right)'\Sigma^{\frac{1}{2}}\Lambda_1. \end{aligned}$$

After using equation (A.2), we get the final expressions:

$$A^{det}_{\tau+1} = A^{det}_{\tau} + \mu_m - y_{0,1} + \frac{1}{2}\left(B^{det}_{\tau}\right)'\Sigma\left(B^{det}_{\tau}\right) - \left(B^{det}_{\tau}\right)'\Sigma^{\frac{1}{2}}\left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime}e_\pi\right), \tag{A.13}$$

$$B^{det\prime}_{\tau+1} = \left(e_\pi + B^{det}_{\tau}\right)'\Psi - e'_{yn} - \left(e_\pi + B^{det}_{\tau}\right)'\Sigma^{\frac{1}{2}}\Lambda_1. \tag{A.14}$$

□

Expressions (A.13) and (A.14) are identical to the ones in (A.10) and (A.11), expect that they do not contain the terms e_{divm} which capture the cash flow risk. The price-dividend ratio on the risky component of the cash flow strip is the difference between the price-dividend ratio on the entire strip and the price-dividend ratio on the deterministic strip.

The expected return on the deterministic strip is given by:

$$\begin{aligned} \mathbb{E}_t[R_{t\to t+k}] &= \frac{\mathbb{E}_t\left[\frac{D^{\$,det}_{t+k}}{D^{\$}_t}\right]}{P^{det}_{t,k}} - 1 \qquad \text{(A.15)} \\ &= \exp\left(-A^{det}_k - B^{det\prime}_k z_t + k(\mu_m + \pi_0) + e'_\pi\left[\sum_{s=1}^{k}\Psi^s\right]z_t + \frac{k}{2}e'_\pi\Sigma e_\pi\right) - 1 \end{aligned}$$

The expected return on the risky part of the equity strips is the difference between the expected return on the entire strip given in (A.12) and the expected return on the deterministic strip given in (A.15).

B Point Estimates Baseline Model

B.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 is listed below.

$$\widehat{\Psi} = \begin{bmatrix} 0.9015 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.2116 & 0 & 0 & 0.0171 & 0 & 0.0078 & 0 & -0.0188 & 0 \\ 0.0694 & 0 & 0.9449 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.8309 & 0 & 0 & 0 & 0 & 0 & 0 \\ -4.7698 & 0 & 0 & 0 & 0.9573 & 0 & -0.0711 & 0 & 0 & 0.6467 \\ 0 & 0 & 0 & 0 & 0 & 0.3696 & 0.0179 & 0 & 0 & 0 \\ 0 & 0 & -3.5856 & 0 & 0 & 0 & 0.8938 & 0 & 0 & 0 \\ 0 & 0 & 1.9302 & 0 & 0.0694 & 0 & 0.0931 & 0 & -0.0720 & 0 \\ 0 & 0 & 0 & 0 & 0.0792 & 0 & 0 & 0 & 0.9156 & 0 \\ 0 & 0 & 0 & 0 & -0.0527 & 0 & 0 & 0 & 0.0466 & 0 \end{bmatrix}$$

The resulting variance-covariance matrix of the VAR residuals is Σ. It's Cholesky decomposition is $\Sigma = \Sigma^{\frac{1}{2}}\Sigma^{\frac{1}{2}\prime}$. After multiplying by 100, we obtain:

$$\widehat{\Sigma^{\frac{1}{2}}} = \begin{bmatrix} 0.2515 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.0406 & 0.6688 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.0350 & 0.0634 & 0.1725 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -0.0056 & -0.0119 & -0.0827 & 0.0968 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1.4238 & 1.1058 & -1.2234 & -0.6322 & 8.0735 & 0 & 0 & 0 & 0 & 0 \\ -0.0344 & -0.0679 & 0.1129 & -0.1378 & -0.1744 & 2.1303 & 0 & 0 & 0 & 0 \\ -0.8195 & 0.9114 & -1.1875 & -0.6028 & 5.4196 & 0.7971 & 7.4634 & 0 & 0 & 0 \\ 0.2899 & -0.0576 & 0.0378 & -0.1595 & 0.0315 & 0.6164 & -1.5111 & 3.2112 & 0 & 0 \\ -0.5775 & 0.7041 & -0.9354 & -0.7745 & 6.3894 & 0.9705 & 0.3429 & 0.5091 & 4.4511 & 0 \\ -0.0775 & 0.0886 & 0.0716 & -0.0964 & -0.4243 & 0.5170 & -0.0798 & -0.0679 & -0.5938 & 1.8896 \end{bmatrix}$$

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.

B.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 is estimated to be:

$$\widehat{\Lambda_0} = \begin{bmatrix} -0.2298 & 0.5320 & -0.4446 & -0.0542 & -0.1840 & 0.6771 & 0 & 0.3469 & 0 & 0.3413 \end{bmatrix}'$$

The matrix that governs the time variation in the market price of risk is estimated to be:

$$\widehat{\Lambda_1} = \begin{bmatrix} 60.52 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 10.99 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -54.39 & -248.61 & 0 & 0 & 0 & 0 & 0 & 0 \\ 54.97 & -6.94 & 0.01 & 0.33 & 0 & 0 & 0 & 0 & 0 & 0 \\ -5.91 & 3.95 & 5.91 & 9.10 & -0.22 & 4.60 & 0 & 0 & 0 & 0 \\ -104.33 & -17.40 & -89.14 & -161.67 & -0.95 & -0.00 & -1.88 & 0 & 0 & 7.46 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 101.04 & -3.48 & -70.54 & -31.87 & 2.33 & -7.63 & -0.46 & 0 & -1.78 & -3.15 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 162.10 & -5.76 & -5.98 & -7.05 & 1.59 & -10.47 & 1.20 & 0 & -1.50 & -5.10 \end{bmatrix}$$

The first four elements of Λ_0 and the first four rows of Λ_1 mostly govern the dynamics of bond yields and bond returns. The price of inflation risk is allowed to move with the inflation rate. The estimation shows that the price of inflation risk is negative on average ($\widehat{\Lambda_0}(1) = -0.23$), indicating that high inflation states are bad states of the world. The market price of inflation risk becomes larger (less negative) when inflation is higher than average ($\widehat{\Lambda_1}(1,1) = 60.52$). The price of real GDP growth risk is positive ($\widehat{\Lambda_0}(2) = 0.53$), indicating that high growth states are good states of the world. The price of growth risk increases when GDP growth is above average ($\widehat{\Lambda_1}(2,2) = 10.99$). The price of level risk (the shock to short rates that is orthogonal to inflation and real GDP growth) is estimated to be negative ($\widehat{\Lambda_0}(3) = -0.44$), consistent with e.g., the Cox-Ingersoll-Ross or Vasicek models. The price of level risk is allowed to change with both the level of interest rates, as in those simple term structure models, and also with the slope factor to capture the fact that bond excess returns are predictable by the slope of the yield curve (Campbell and Shiller). When interest rate levels are unusually high and the term structure steepens, the price of level risk becomes more negative, and expected future bond returns increase. The positive association between the slope and future bond returns is consistent with the bond return predictability evidence. The price of (orthogonal) slope risk is estimated to be slightly negative on average (-0.05). Since the spread between the five-year bond yield and the short rate is the fourth element of the state vector, and the short rate is the third element of the state vector, the five year bond yield can be written as:

$$y^{\$}_{t,20} = y^{\$}_{0,20} + (e_{yn} + e_{yspr})' z_t = -\frac{A^{\$}_{20}}{20} - \frac{B^{\$\prime}_{20}}{20} z_t$$

A necessary and sufficient condition to match the five-year bond yield dynamics is to allow for the first four elements of the fourth row of Λ_1 to be non-zero.

The last six elements of Λ_0 and last six rows of Λ_1 govern the stock pricing. We assume that the market prices of risk associated with the price-dividend ratios are zero, since those variables only play a role as predictors. The only exception is the price-dividend ratio on the stock market. The evidence from dividend strip spot and futures prices and the evidence on strip future returns helps us identify the market prices of risk associated with the pd ratio (fifth element of Λ_t).

The risk prices in the 6th, 8th, and 10th rows of Λ_t are chosen to match the observed mean and dynamics of the equity risk premium in model, as shown in appendix A, and data, as implied by the VAR. We only free up those elements of the 6th, 8th, and 10th rows of Λ_1 that are strictly necessary to allow the equity risk premium in the model to move with the same state variables as it does in the VAR. These rows of of Λ_t are also influenced heavily by our insistence on matching the entire time series of the price-dividend ratio on the stock market, REITS, and infrastructure.

C 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.

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:

$$\log(D^i_{t+1}) - \log(D^i_t) = \Delta d^i_{t+1} = A^i_0 + A^i_1 z_t + A^i_2 \varepsilon_{t+1},$$

where $A^i_0 = \mu_m$, $A_1 = e'_{divi}\Psi$, and $A^i_2 = e'_{divi}\Sigma^{\frac{1}{2}}$.

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

$$y_{t+1} - y_t = \Gamma_0 + \Gamma_1 z_t + z'_t \Gamma_3 z_t + \Psi_0 \varepsilon_{t+1} + z'_t \Psi_1 \varepsilon_{t+1} \tag{A.16}$$

with coefficients $\Gamma_0 = A^i_0$, $\Gamma_1 = A^i_1$, $\Gamma_3 = 0$, $\Psi_0 = A^i_2$, 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 (A.16), with coefficients $\Gamma_0 = A^i_0 + B_0$, $\Gamma_1 = A^i_1 + B_1$, $\Gamma_3 = B_3$, $\Psi_0 = A^i_2 + B_2$, 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:

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

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

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

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

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

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 (A.16) 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 Korteweg-Nagel Details

They propose:

$$m_{t+1} = a - br^m_{t+1},$$

whereby the coefficients a and b are chosen so that the Euler equation $1 = E[M_{t+1}R_{t+1}]$ holds for the public equity market portfolio and the risk-free asset return. More specifically, they estimate $a = 0.088$ and $b = 2.65$ using a GMM estimator:

$$\min_{a,b} \left(\frac{1}{N} \sum_i u_i(a,b) \right)' W \left(\frac{1}{N} \sum_i u_i(a,b) \right)$$

where

$$u_i(a,b) = \sum_{j=1}^{J} M_{t+h(j)}(a,b)[X_{if,t+h(j)}, X_{im,t+h(j)}],$$

N is the number of funds, and W is a 2×2 identity matrix. The T-bill benchmark fund cash-flow, X_{if}, and the market return benchmark cash-flow, X_{im}, are the cash-flows on a T-bill and stock market investment, respectively, that mimic the timing and magnitude of the private equity fund i's cash-flows. The $t + h(j)$ are the dates on which the private equity fund pays out cash-flow $j = 1, \cdots, J$. Date t is the date of the first cash-flow into the fund, so that $h(1) = 0$. For each of the two benchmark funds, the inflows are identical in size and magnitude as the inflows into the PE fund. If PE fund i makes a payout at $t + h(j)$, the benchmark funds also make a payout. That payout consists of two components. The first component is the return on the benchmark since the last cash-flow date. The second component is a return of principal, according to a preset formula which returns a fraction of the capital which is larger, the longer ago the previous cash-flow was.

A special case of this model is the public market equivalent of Kaplan and Schoar (2005), which sets $a = 0$ and $b = 1$. This is essentially the log utility model. The simple PME model is rejected by Korteweg and Nagel (2016), in favor of their generalized PME model.

There are several key differences between our method and that of Korteweg and Nagel (2016). First, we do not use SDF realizations to discount fund cash-flows. Rather, we use bond prices and dividend strip prices, which are conditional expectations. Realized SDFs are highly volatile. Second, the KN approach does not take into account heterogeneity in the amount of systematic risk of the funds. All private equity funds are assumed to have a 50-50 allocation to the stock and bond benchmark funds. Our model allows for different funds to have different stock and bond exposure. Third, the KN approach uses a preset capital return policy which is not tailored to the fund in question. For example, a fund may be making a modest distribution in year 5, say 10%, and a large distribution in year 10 (90%). Under the KN assumption, the public market equivalent fund would sell 50% in year 5 and the other 50% in year 10. There clearly is a mismatch between the risk exposure of the public market equivalent fund and that of the private equity fund. In other words, the KN approach does not take into the account the magnitude of the fund distributions, only their timing. Fourth, we use additional risk factors beyond those considered in KN.

To study just the importance of the last assumption, we can redo our calculations using a much simplified state vector that only contains the short rate, inflation, and the stock market return. This model has constant risk premia.

E Additional Results

Figure A.1: cash-flows by Vintage

(a) Fund of Funds

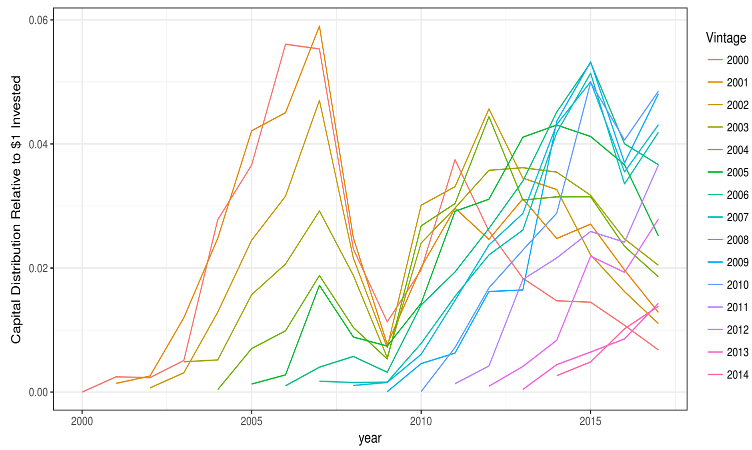

(b) Debt Funds

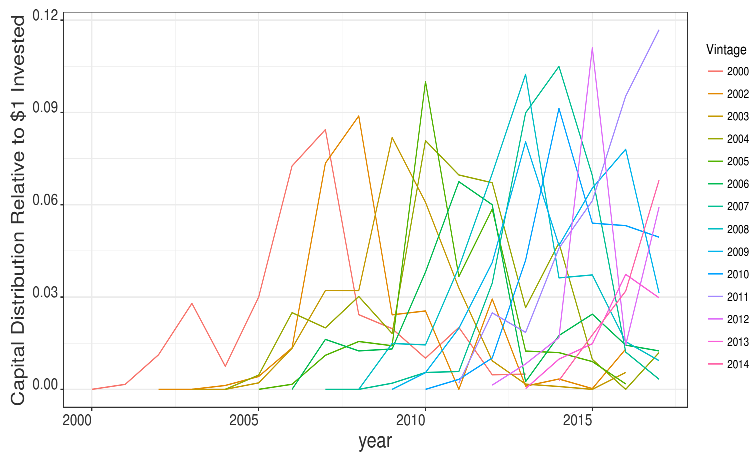

(c) Restructuring

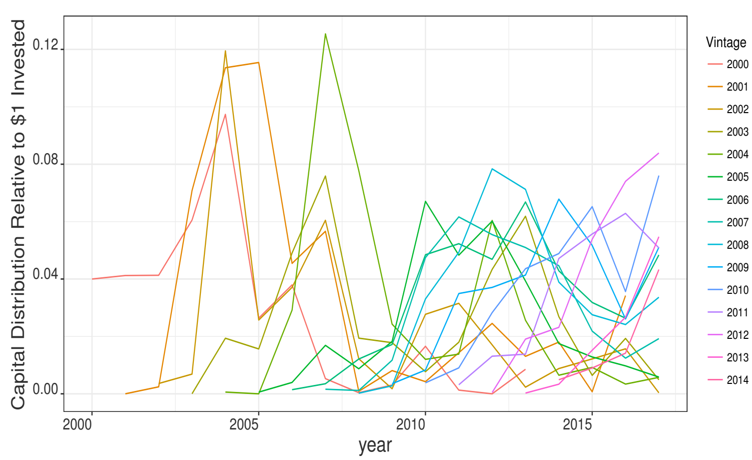

Figure A.2: One-Factor Fund Exposures for Other Categories

Panel A: Fund of Funds

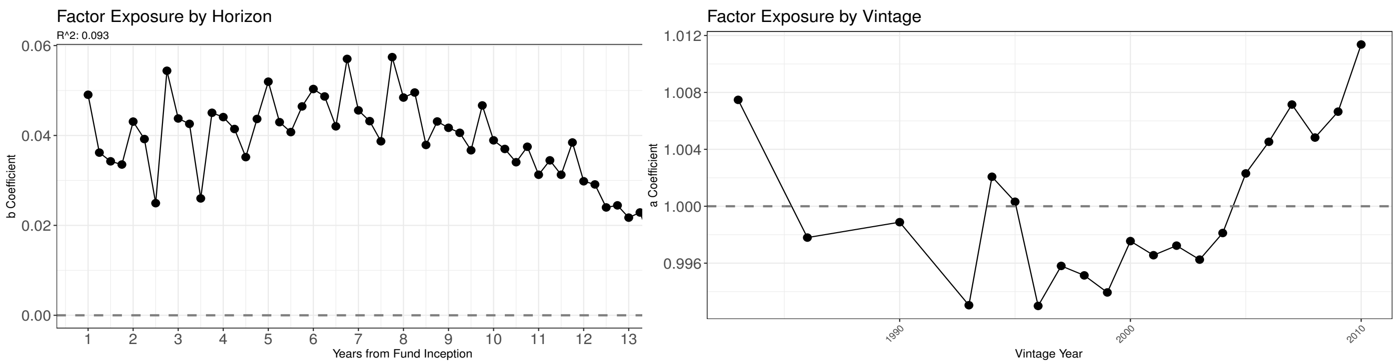

Panel B: Debt Funds

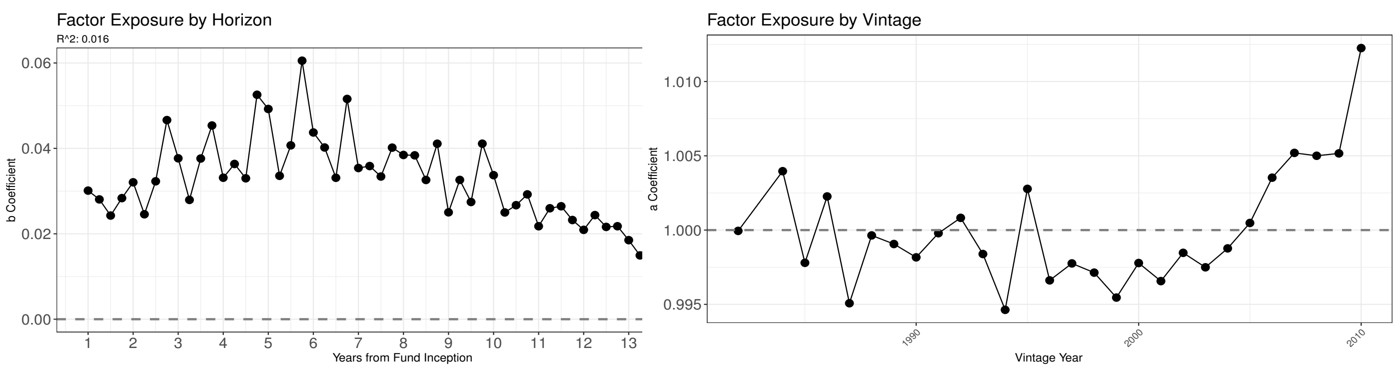

Panel C: Restructuring

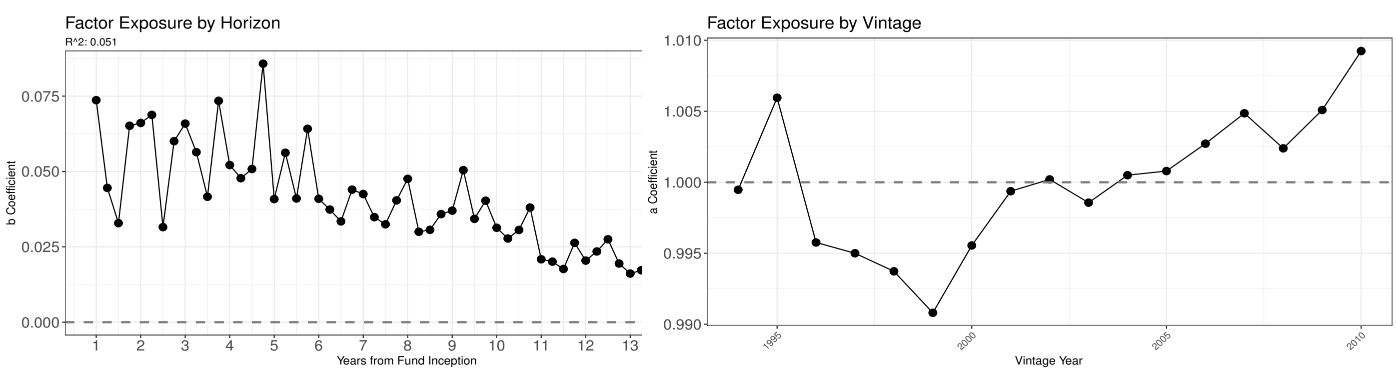

Figure A.3: One Factor Risk-Adjusted Profits by Category

Panel A: Fund of Funds

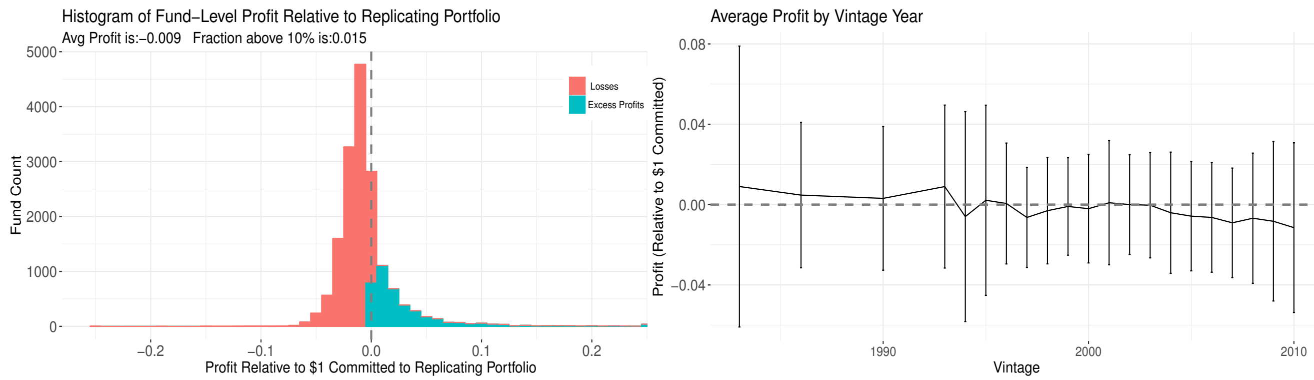

Panel B: Debt Funds

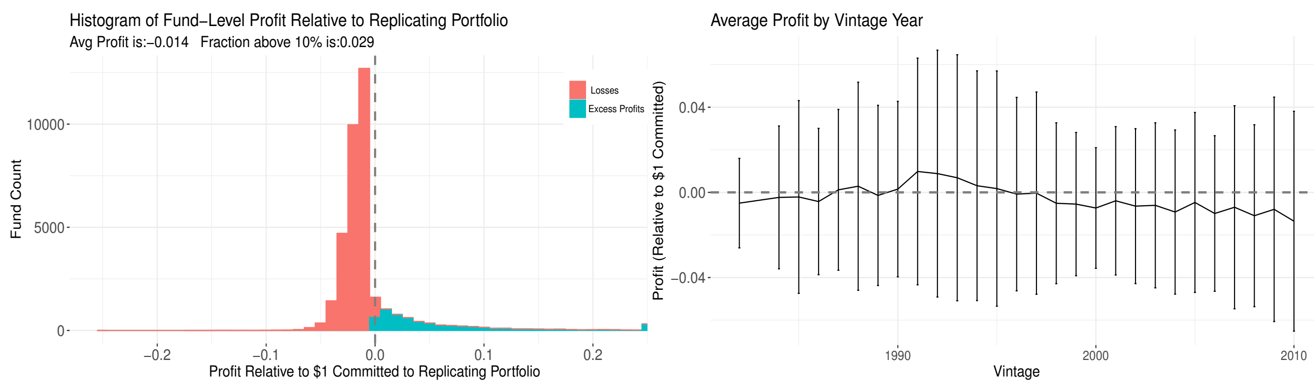

Panel C: Restructuring

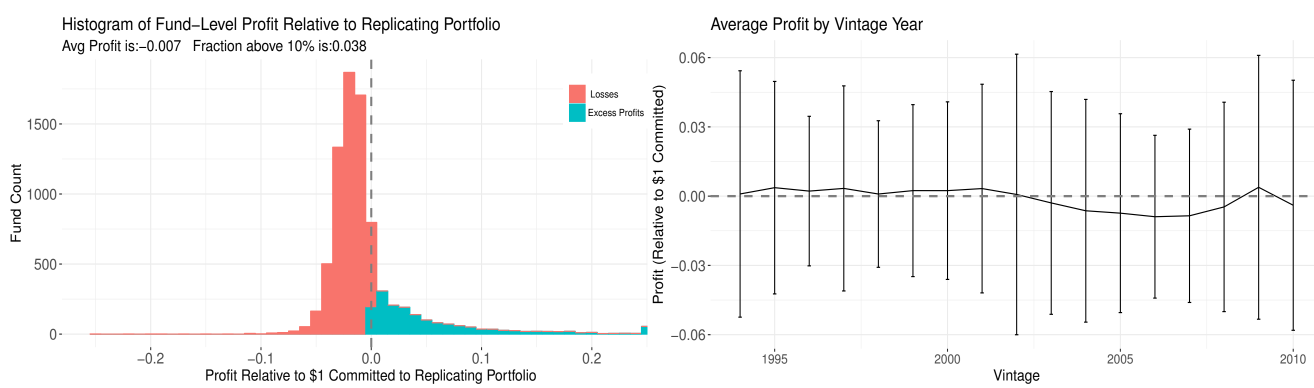

Figure A.4: One Factor Expected Return

Panel A: Fund of Funds

Panel B: Debt Funds

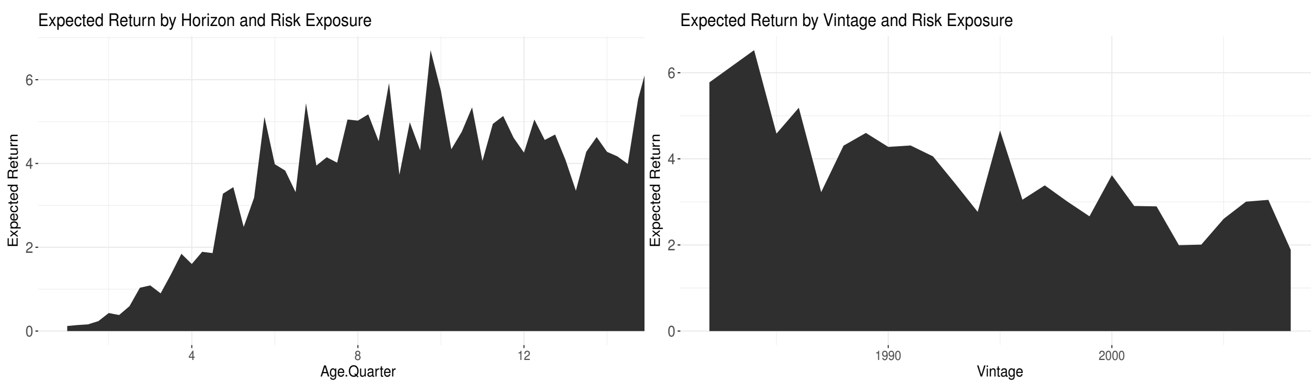

Panel C: Restructuring

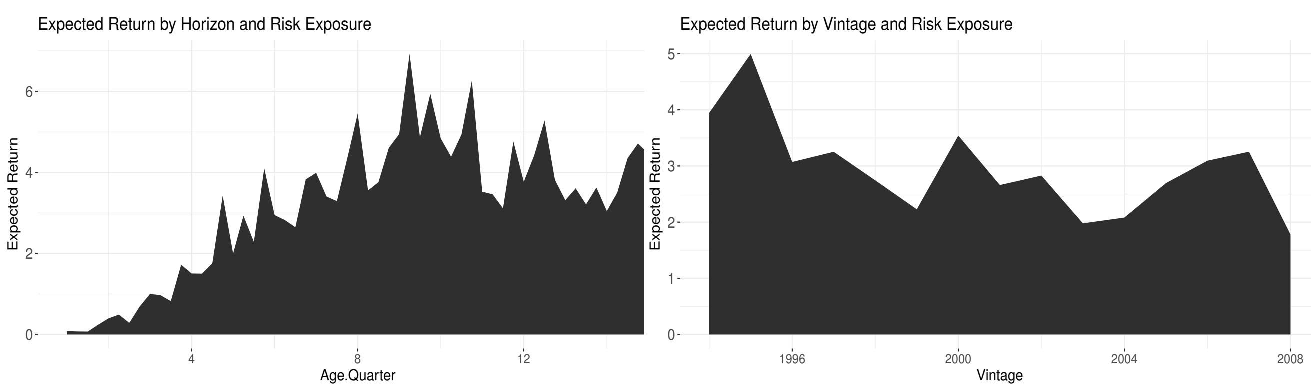

Figure A.5: Four-Factor Fund Exposures for Other Categories

Panel A: Fund of Funds

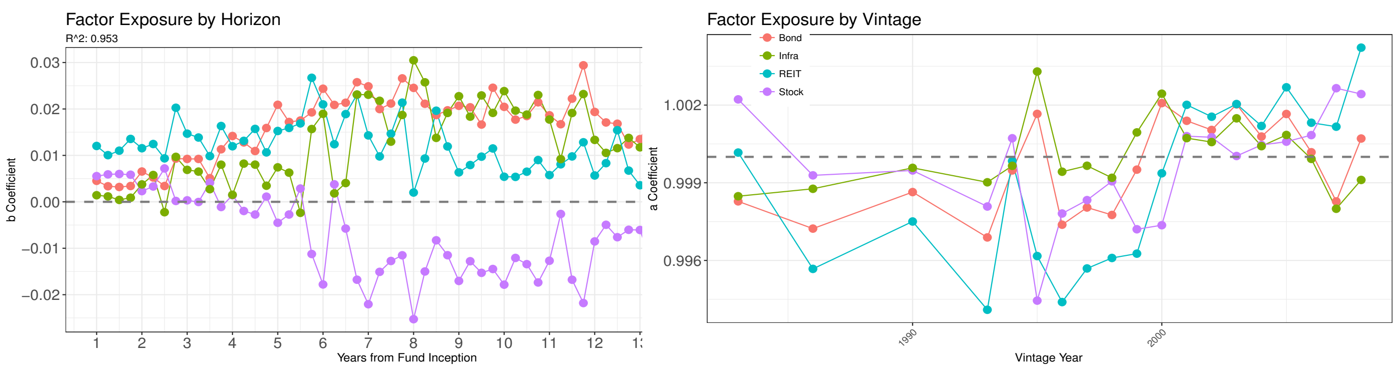

Panel B: Debt Funds

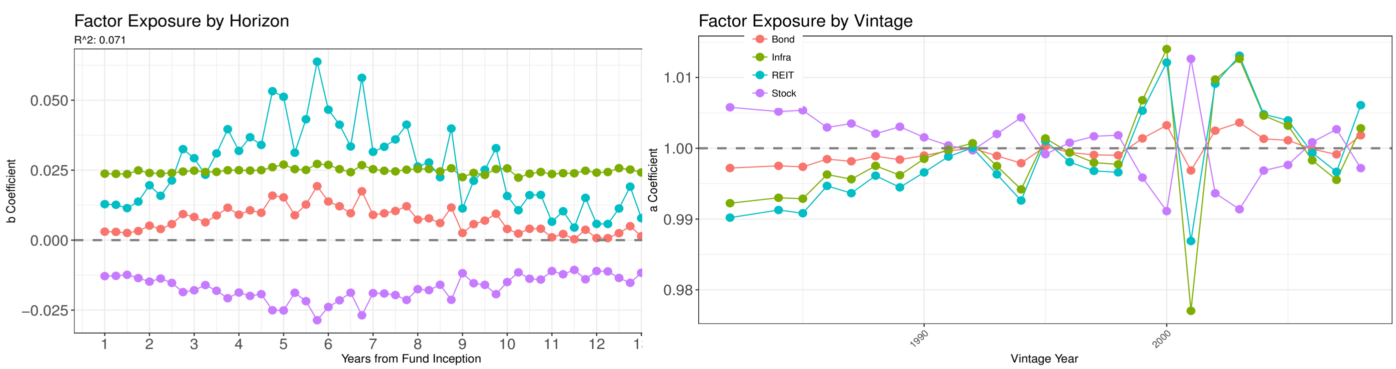

Panel C: Restructuring

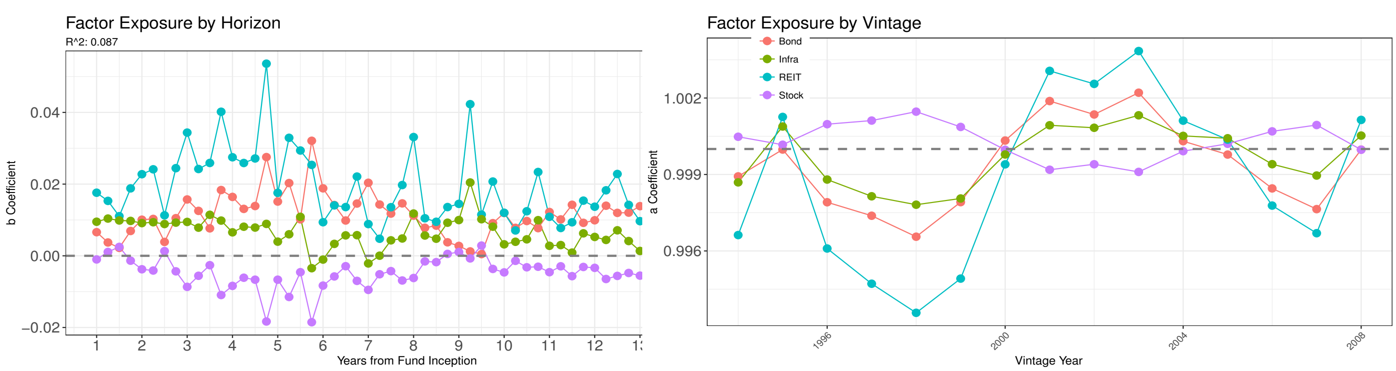

Figure A.6: Replicating Portfolio Exposure by Feature

Panel A: Buyout

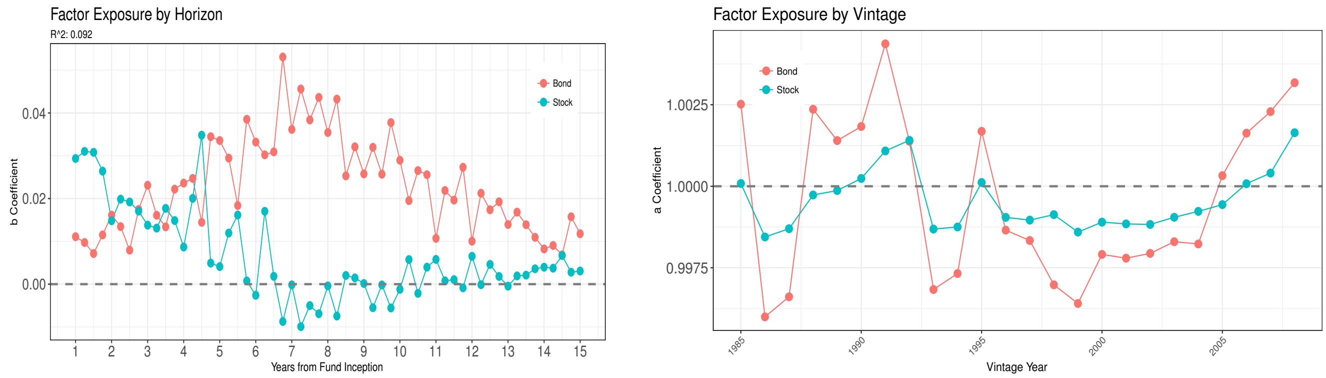

Panel B: Venture Capital

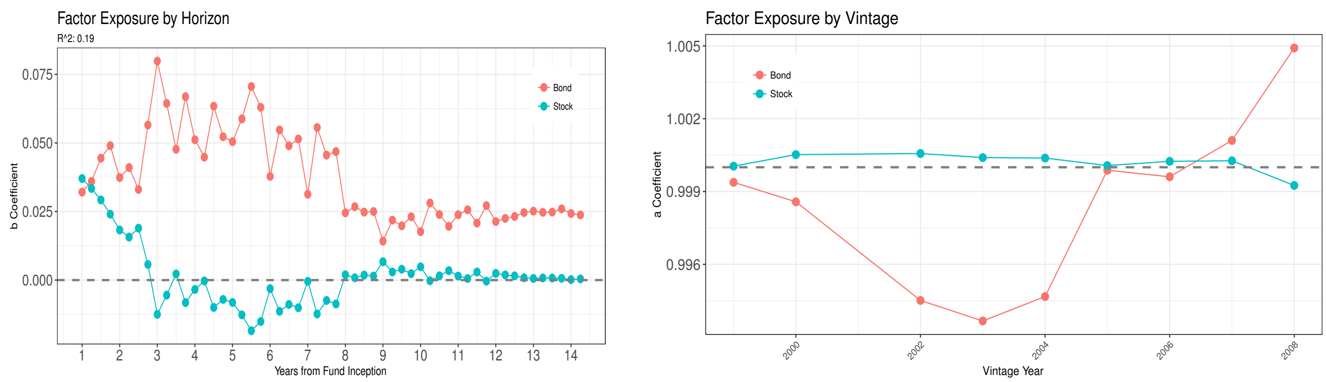

Panel C: Real Estate

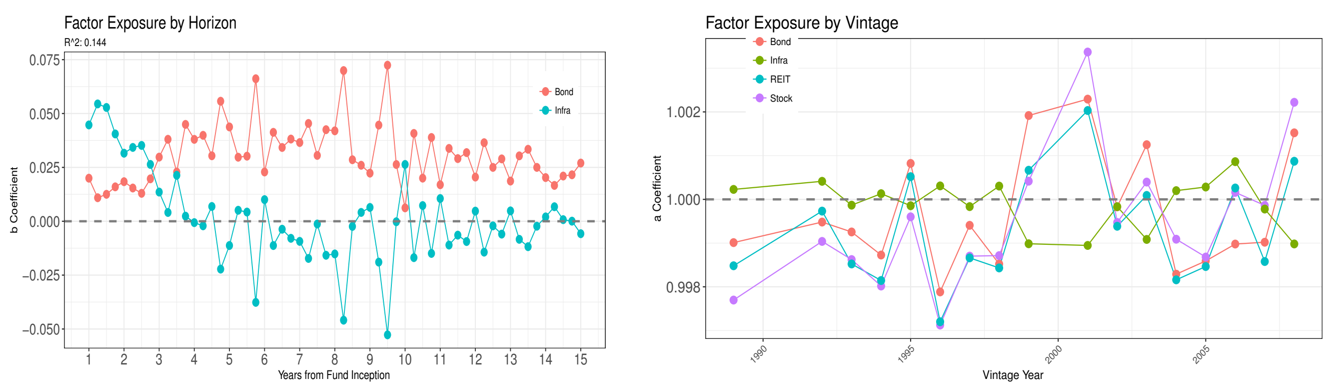

Panel D: Infrastructure

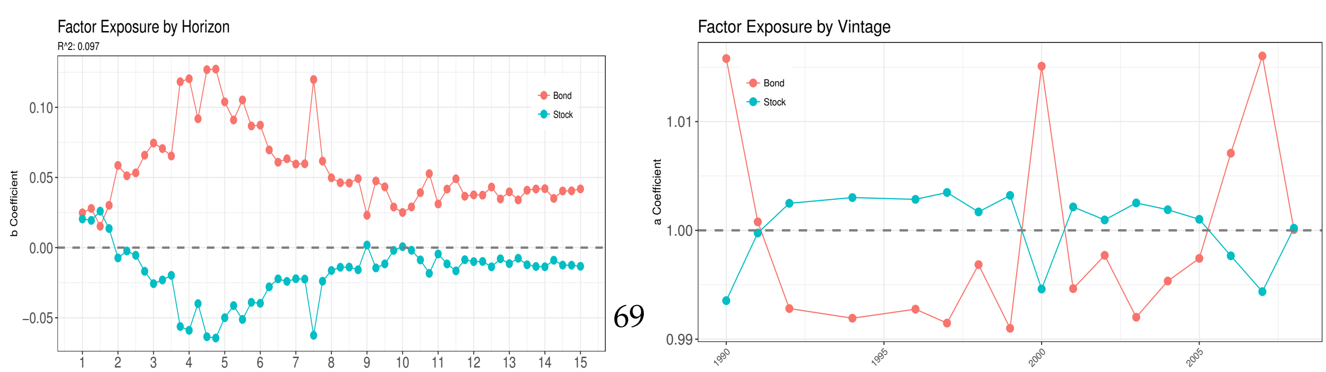

Figure A.7: Four Factor Risk-Adjusted Profits by Category

Panel A: Fund of Funds

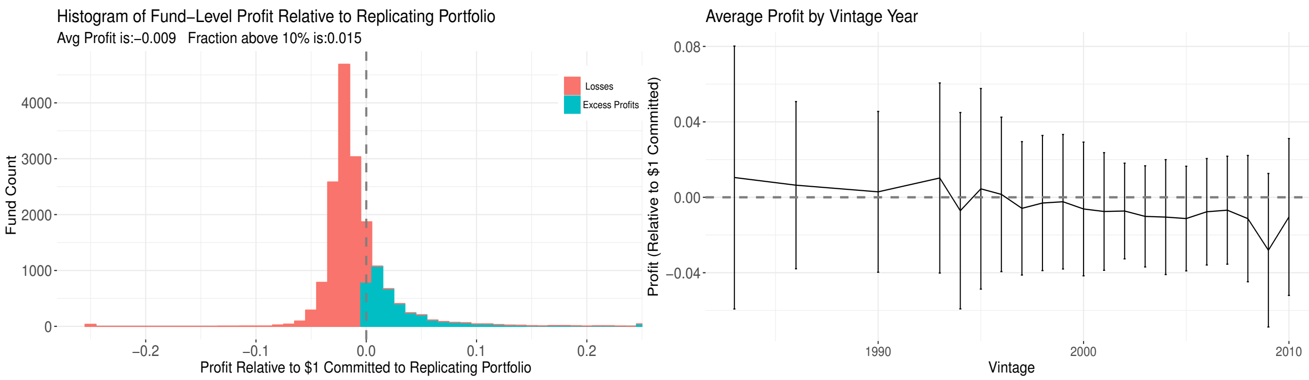

Panel B: Debt Funds

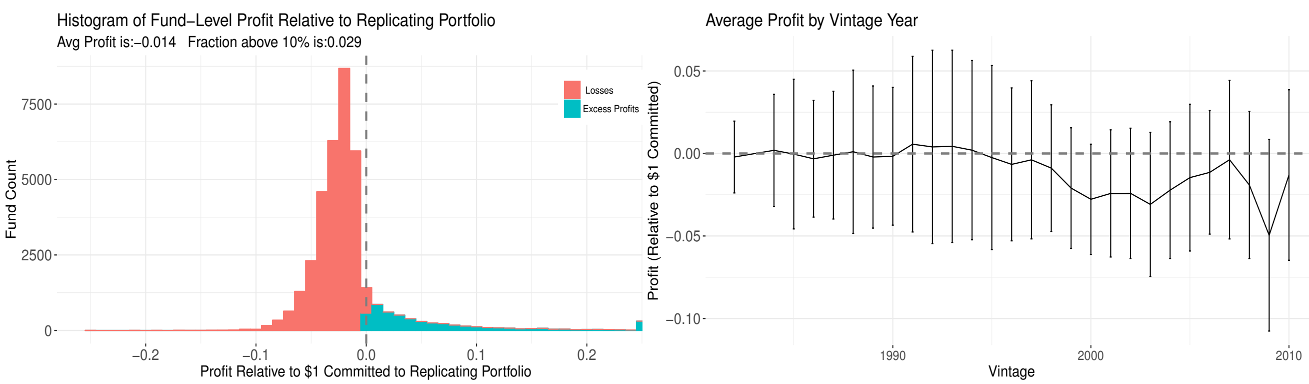

Panel C: Restructuring

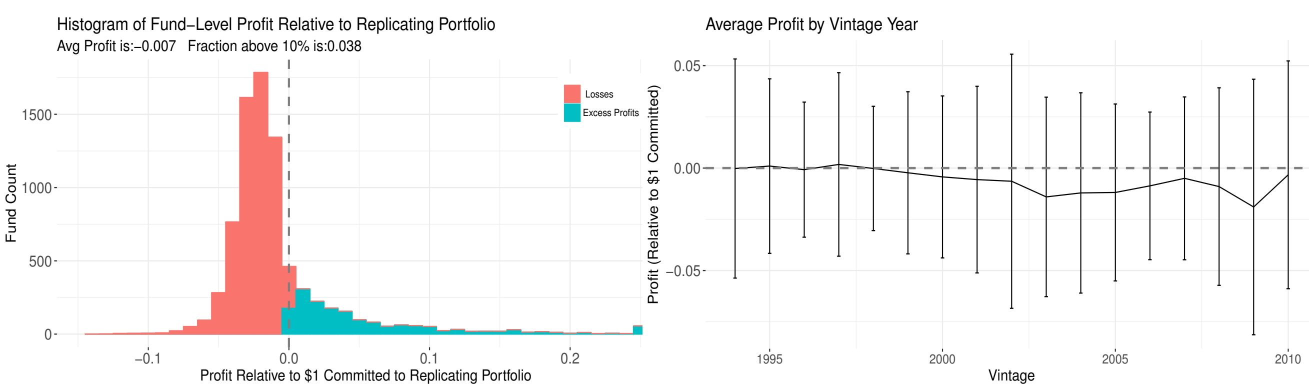

Figure A.8: Four Factor Expected Return

Panel A: Fund of Funds

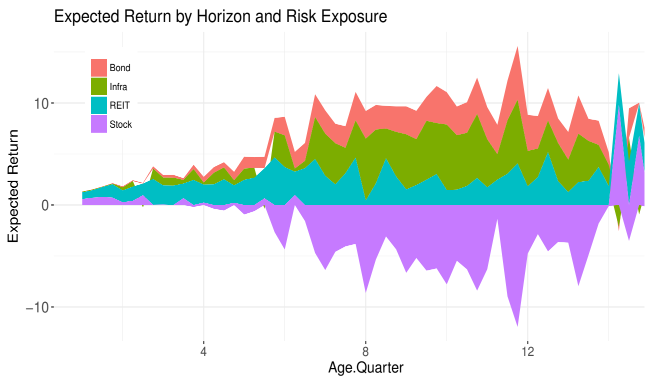

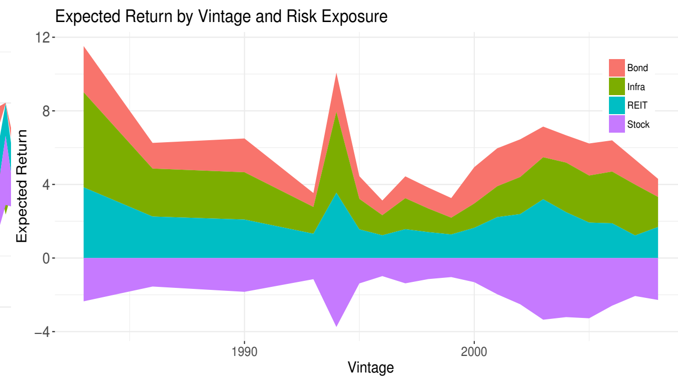

Panel B: Debt Funds

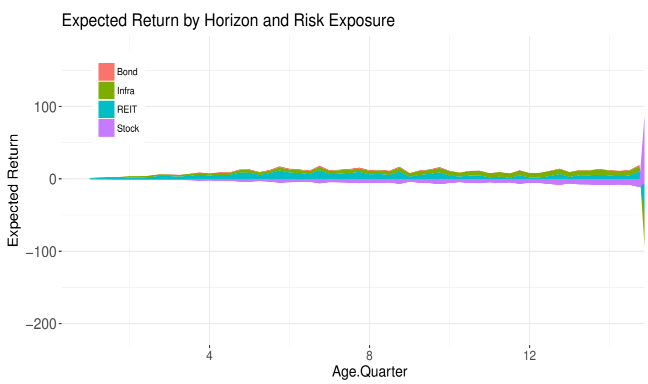

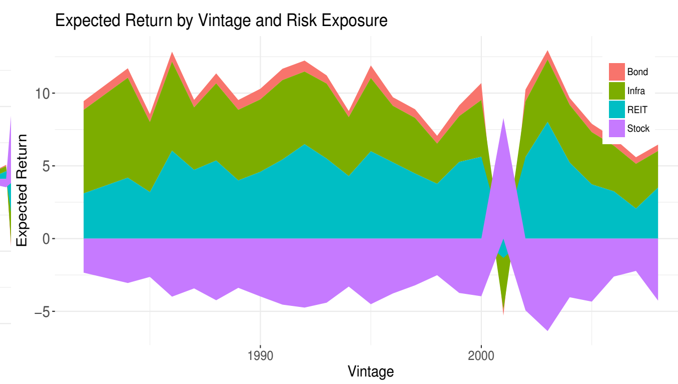

Panel C: Restructuring