NBER WORKING PAPER SERIES

EQUITY TERM STRUCTURES WITHOUT DIVIDEND STRIPS DATA

Stefano Giglio
Bryan T. Kelly
Serhiy Kozak

Working Paper 31119
http://www.nber.org/papers/w31119

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2023

We appreciate helpful comments from Jules Van Binsbergen, Niels Gormsen, Ralph Koijen, Jessica Wachter, participants at the Chicago Booth Asset Pricing Conference, SFS Finance Cavalcade, Virtual Derivatives Workshop, WFA, and seminar participants at Stockholm School of Business, Temple University, Tilburg University, Tsinghua University, University of Gothenburg, University of Hong Kong, University of Massachusetts at Amherst, and the Five Star conference. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Equity Term Structures without Dividend Strips Data
Stefano Giglio, Bryan T. Kelly, and Serhiy Kozak
NBER Working Paper No. 31119
April 2023
JEL No. G11,G12,G13

ABSTRACT

We use a large cross-section of equity returns to estimate a rich affine model of equity prices, dividends, returns and their dynamics. Using the model, we price dividend strips of the aggregate market index, as well as any other well-diversified equity portfolio. We do not use any dividend strips data in the estimation of the model; however, model-implied equity yields generated by the model match closely the equity yields from the traded dividend forwards reported in the literature. Our model can be used to extend the data on the term structure of aggregate (market) discount rates over time (back to the 1970s) and across maturities, since we are not limited by the maturities of actually traded dividend claims. Most importantly, the model generates term structures for any portfolio of stocks (e.g., small and value portfolios, high and low investment portfolios, etc). The novel cross-section of term structure data estimated by our model, covering a span of 45 years that includes several recessions, represents a rich set of new empirical moments that can be used to guide and evaluate asset pricing models, beyond the aggregate term structure of dividend strips that has been studied in the literature.

Stefano Giglio
Yale School of Management
165 Whitney Avenue
New Haven, CT 06520
and NBER
stefano.giglio@yale.edu

Bryan T. Kelly
Yale School of Management
165 Whitney Ave.
New Haven, CT 06511
and NBER
bryan.kelly@yale.edu

Serhiy Kozak
University of Maryland
7621 Mowatt Ln
4453
VMH 4453
College Park, MD 20742
https://www.serhiykozak.com
sekozak@umd.edu

1 Introduction

The term structure of discount rates for risky assets plays an important role in many fundamental economic contexts. For example, pricing an asset with a specific horizon of cash flows and evaluating an investment opportunity with a specific maturity requires knowing the maturity-specific discount rate. Investment in climate-change mitigation, where the maturity of the project is especially long and therefore the long end of the term structure is especially important, is one well-known case.

In this paper we specify and estimate a rich affine model of equity portfolios, that describes the prices, dividends, and excess returns of a large cross-section of portfolios, as well as their dynamics. While the model is driven by many parameters, we impose discipline on the model in several ways: by imposing pricing restrictions, by choosing appropriately the state vector that drives the dynamics of the economy, and by imposing parameter restrictions that reflect recent findings in the literature on returns predictability with large cross-sections. We then use our model to generate term structure of discount rates not only for aggregate cash flows (the S&P 500), but also for many other equity portfolios, obtaining a large panel of 102 term structures of discount rates, with arbitrary maturity going up to infinity, and with a time series going back to the 1970s. We validate the predictions of our model for the discount rates of risky cash flows by comparing the implied dividend strips from our model to the prices of actually traded dividend strips (from Bansal et al. (2017) and van Binsbergen and Koijen (2015)). We show that our model – estimated using no dividend strip data at all – manages to match well the prices of traded dividend strips on the S&P 500 of maturities 1, 2, 5 and 7 years, observed since 2004, along a variety of dimensions (average slope, time-series, and so on). After validating the model using observed strip data, we then use it to explore the properties of implied dividend strips extending the time series back to the 1970s, and studying the cross-section of term structures of different portfolios.

The asset pricing literature has tackled the task of estimating the term structure of discount rates for risky assets in a variety of ways. An earlier literature has proposed extracting information about the term structure from the cross-section of equity portfolios. The broad idea behind this method is that if some stocks are mostly exposed to long-term cash-flow shocks, and others are mostly exposed to short-term shocks, the difference in risk premia between the two types of stocks can be reconducted to a difference in how investors price shocks to cash flows of different maturities – that is, to the term structure of discount rates applied by investors (see Bansal et al. (2005), Lettau and Wachter (2007), Hansen et al. (2008), Da (2009)).

While this literature made substantial progress in understanding what the cross-section of equity portfolios implies for the term structure of discount rates, it has faced an important hurdle: the term structure of discount rates depends on the entire dynamics of cash flows as well as the risk preferences of investors (and their variation over time). That is, the term structure of discount rates is an equilibrium result that depends on the interaction of a large number of forces. To identify and estimate the model, the papers in this literature have imposed strong assumptions on the preferences (e.g., Epstein–Zin preferences, constant risk aversion, restrictions on which shocks are priced by investors, etc.), or on the dynamics of the economy and cash flows, or both.

New impetus in the study of term structures has come from the introduction of data on traded dividend claims (van Binsbergen et al. (2012b), Van Binsbergen et al. (2013) and van Binsbergen and Koijen (2015)). The ability to directly observe the returns of finite-maturity dividend claims gives a direct window into the risk premia investors require to hold risks of different maturities, and obviates the need to estimate the dynamics of the economy and the preferences of investors. Studying term structures with traded dividend strip data, however, has a few shortcomings of its own. First, the time series is quite limited, since data is available only starting around 2004, and it includes only one full recession with associated recovery, the Great Recession. Second, there is little cross-sectional data, since only the aggregate market dividend strips – for the US and other countries – start around 2004; much more limited data is instead available on individual firms. Third, only a portion of the term structure is typically observed (a few maturities up to 7 years). Fourth, there are concerns about liquidity of these contracts, which could potentially lead to measurement error.

In this paper, we return to the first approach to estimate term structures, based on equity portfolios alone; we use the model to effectively produce new (implied) term structure data that expands the existing (observed) data along each of those dimension. The term structures we generate cover a large number of cross-sectional portfolios, in addition to the S&P 500: value, size, profitability, momentum, etc., for the total of 102 portfolios. They have a long time series, starting in 1975 and therefore covering several recessions and booms. They have all possible maturities, including the very short and the very long ends of the term structure.

While closely related to the models that had been used in this literature to study term structures using equity portfolios, our model has a few distinct features that are crucial to generate realistic implied term structures that match the ones we observe from traded dividend claims. The model features rich dynamics which are motivated by recent empirical findings in the literature, and we believe are both economically reasonable and statistically parsimonious. In particular, consistent with Kozak et al. (2020) and Giglio and Xiu (2021), who show that a few dominant principal components (PCs) of a large cross-section of anomaly portfolio returns explain the cross-section of expected returns well, our state vector includes four factor returns (PCs), estimated from a large cross-section of 51 anomalies.

Our further motivation comes from Haddad et al. (2020), who demonstrate that valuation ratios strongly and robustly predict expected returns on these PCs—and their risk prices in the stochastic discount factor (SDF)—in the time-series. They argue that the resulting time-variation in risk prices is critical for adequately capturing dynamic properties of the pricing kernel. Chernov et al. (2018) echoes the importance of time-variation in prices of risk by proposing to test asset-pricing models using multi-horizon returns. Motivated by these findings, our specification also includes four factor yields (D/P ratios) associated with the factor returns—for a total of eight variables in the state vector—which allows us to capture the dynamics of conditional means and SDF risk prices.

Overall, our specification allows both dividend growth and risk premia to vary over time in minimally restricted ways; and has a general, affine specification for the stochastic discount factor in which shocks to factor returns are priced and their risk prices are captured by valuation ratios. The fact that our state variables are factor yields and factor returns means that, on the one hand, our state

variables are forward-looking and can be expected to contain information about the evolution of the economy; and on the other hand, it implies that the factors need to satisfy certain pricing restrictions, which help better pin down their dynamics. It is this balance of a rich model with appropriately-chosen restrictions that represents the core of our paper: it allows us to produce term structures of discount rates that match well the observed ones – and gives us confidence in extending them over time, maturities, and portfolios.

More specifically, we start with a large cross-section of test asset returns r_t on 102 portfolios. We proceed with specifying and estimating a homoskedastic affine model, in which the factors F_t that drive the dynamics of the test asset returns are chosen to be the linear combinations of assets' excess returns and their dividend yields, that is, $F_t = [f_{r,t}\ f_{y,t}]'$, where factor returns, $f_{r,t} = Q'(r_t - r_{f,t})$, and factor yields, $f_{y,t} = Q'y_t$, are the same linear combination of log excess returns and yields on the test assets, given by some matrix Q. We restrict the first factor to coincide with the excess log market return, and the rest to be based on the three largest principal components of 51 long-short anomaly portfolios constructed from the 102 long-only portfolios, as in Haddad et al. (2020); Kozak (2020). This choice of factors for the model has two main advantages. First, it captures very well the covariation of returns of the anomaly portfolios (together, the four factor returns explain 93.3% of the total variation in excess returns in the panel of 102 long portfolios). Second, the factor returns are well predicted by their own dividend yields, which, conveniently, are also part of our state vector F_t. So our state vector F_t contains variables that are useful to predict excess returns (and dividends, which are related by an identity to prices and returns).

We restrict this affine model in only two ways. First, we impose that the innovations in the stochastic discount factor in the economy depend only on the innovations in factor returns (and not on the innovations in factor yields): that is, we are assuming that the SDF innovations are fully spanned by the factor returns. Second, we impose that the four factor yields $f_{y,t}$ contain all available information about the future (so that lagged returns do not help predict future yields and returns after controlling for lagged yields). In practice, it is well known that dividend yields have much stronger predictive power for future dividends and returns than lagged returns, and we simply impose this by assumption in our statistical model.

We do not impose any other restrictions to the model, except of course the pricing restrictions that link the pricing kernel to prices, dividends and returns. These restrictions imply, as usual in these cases, that the process of dividend growth for each portfolio is fully pinned down by the (nonlinear) identity that links prices, returns and dividends.

We end up with an estimated model that, with its eight state variables, is rich enough to capture a variety of possible dynamics for prices, excess returns and dividends. The model immediately produces implied term structures of dividend strips and forwards, that is, spot or forward claims to a specific dividend at some point in the future; it produces a different term structure for each of the four portfolios that are part of the state vector F_t. But the model does more: because the four factors span extremely well the cross-section of returns and dividend yields of all the original 102 characteristic-sorted portfolios, it can easily build implied term structures of any of these portfolios.

Our estimated model delivers a variety of novel empirical results. First, we extend the study of the

term structure of aggregate dividend claims (on the S&P 500, as in van Binsbergen and Koijen (2015)) over time (back to the 70s) and across maturities. In the sample starting in 2004 that was used in van Binsbergen and Koijen (2015), we match the time series of forward dividend yields very closely, and therefore, mechanically, we also match the term structure of discount rates.[1] The term structure of forward risk premia appears in this sample mildly upward sloping (not significantly so). We also confirm that our implied term structure inverts during the financial crisis, just like the observed one does.[2]

Extending the sample to the 1970s allows us to include several additional recessions to our sample; at the same time, the Great Recession carries less overall weight in the sample. It is interesting to see that all the results of the post-2004 sample carry over to the longer sample. The term structure inverts in almost all of the additional recessions (for example, in the early 80s and 90s). And the term structure of forward discount rates is still close to flat; it is mildly upward sloping on average, but not significantly so.

In addition, a decomposition of the movements of prices of short-term dividend claims into expected dividend growth and expected risk premia shows that the former varies substantially over time: investors expected low dividend growth during the 1980 and 1990 recessions, as well as during the Great Recession – and this moved the prices of the short-term dividend claims substantially.

The most important and novel results of our paper are the estimated term structures of discount rates for different portfolios, like value and growth firms, and small and large firms. Our model generates interesting differences both in the average term structure across portfolios, and in the time series. For example, term structures of long-short portfolios such as SMB and HML are fundamentally different: the SMB portfolio appears to exhibit a downward-sloping term structure, while the HML portfolio has an upward-sloping term structure, despite the fact that both portfolios have positive unconditional risk premium in our sample. We show that other long-short portfolios, such as profitability, growth, momentum, idiosyncratic volatility, exhibit distinct term structures of expected returns. Theoretical models aiming to explain these different spreads (like the many that have been proposed to explain the size and value premia) can make use of these estimates to help refine and calibrate the economic mechanisms: our estimates of the term structure of different portfolios provide us with additional empirical moments, that can be compared to the corresponding ones in the models.

Finally, there are interesting patterns in the *time series* of slopes of the yield term structure of different portfolios. For example, the slopes of small and large stocks tend to move together; both term structures were upward sloping during the 1990s, and both were downward sloping during the Great Recession. Yet, only the term structure for small stocks inverted during the late 90s stock cycle, marking an important divergence between the two portfolios that lasted several years. On the contrary, no such divergence in the shape of the term structure can be seen for value and growth stocks

[1] We also verify that the ability of the model to match traded strip prices is robust to a variety of changes in our model specification.

[2] Incidentally, we note that because our implied dividend strip prices are obtained from equity portfolios, they are less susceptible to the potential criticism raised by Bansal et al. (2017) that traded dividend strips might be illiquid and that bid-ask spreads might affect the conclusions about the slope of the term structure. The fact that we actually match those prices very closely using only (very liquid) equity portfolios suggests that liquidity is not a driver of the findings of van Binsbergen and Koijen (2015) in the first place.

in that period – instead, the largest difference in that case occurred in the recovery from the financial crisis: after 2008, the term structure of value stock expected returns increased significantly, whereas this didn't happen for growth stocks.

To sum up: our model effectively processes a rich information set (the time-series and cross-sectional behavior of 102 portfolios spanning a wide range of equity risks) to produce "stylized facts" – the time series and cross-sectional behavior of implied dividend term structures – that summarize a dimension of the data that is particularly informative about our economic models. Similarly in spirit to the way in which the introduction of vector autoregressions (VAR) by Sims (1980) provided new moments against which to evaluate structural macro models (the impulse-response functions that were generated by the VARs), the objective of this paper is to produce realistic term structure of discount rates for different portfolios that closely resemble the actual dividend claims we observe in the data, and that can be used by asset pricing models as a moment for evaluation and guidance.

It is important to note that our estimates of equity yields, discount rates, and returns on specific dividend claims are subject to estimation uncertainty. In fact, there are two types of uncertainty that affect the final empirical estimates. First, there is uncertainty coming from the fact that the prices of the dividend strips that our model produces are not actually observed, but they are obtained from the model parameters, which are subject to estimation error. We refer to this source of uncertainty as "parameter uncertainty." Then, there is uncertainty coming from the fact that some moments of interest (for example the average slope of the dividend term structure) need to be estimated from averages taken over finite samples. This last type of uncertainty, that we call "sampling uncertainty," is present even when prices are observable. What our model brings to the table is, in a sense, a tradeoff between these two sources of uncertainty. Compared to using the actual observed prices of dividend strips, working with estimated dividend strips introduces parameter uncertainty. At the same time, our method allows to dramatically expand the time series available, reducing sampling uncertainty. Which of the two forces dominates total uncertainty depends on the specific context. In our empirical estimation, we find that the sampling uncertainty dominates, so that our method allows us to significantly reduce standard errors on many moments of interest. To help future researcher take explicitly into account this estimation uncertainty, we provide data on standard errors of our estimated equity yields along with their point estimates on our website.[3]

Our methodology lends itself to a variety of applications. As we mentioned above, one direct application is to test and calibrate asset pricing models, by using our estimated model to produce additional empirical moments. In the paper, we provide an illustration of this idea, using our estimated term structures to test workhorse asset pricing models (like Bansal and Yaron (2004), Campbell and Cochrane (1999), etc.). More interestingly, however, our rich set of term structures of different portfolios can be used to test further predictions of the models (about the cross-sectional heterogeneity in the shapes of the term structures across portfolios). In addition, our term structures can be used for the valuation of project with specific horizons. For example, Gupta and Van Nieuwerburgh (2019) use them to evaluate private equity investments. An alternative application is to climate change mitigation investments, where the long end of the term structure is especially important (see, for example, Giglio

[3] `https://www.serhiykozak.com/data`.

et al. (2015)).

Beyond the seminal literature using equity portfolios or dividend strips to calibrate and estimate empirical term structures, our paper also relates to a more recent literature that has also built on those approaches to improve our understanding of term structures. This literature for the most part focuses on the term structure of aggregate dividend claims. Some papers explore the joint behavior of the aggregate stock market and treasury bonds (Lettau and Wachter (2011), Ang and Ulrich (2012), Koijen et al. (2017)), whereas others use the traded dividend strip data in the estimation (Kragt et al. (2014), Gomes and Ribeiro (2019), Yan (2015)). Recently, Gupta and Van Nieuwerburgh (2019) use term structures of discount rates from a similar affine model to value private equity investment; given their different objective, they use specific portfolios in their model (small and large firms, REITs, and infrastructure firms). Our objective is instead to select the state variables that best describe the whole dynamics of the economy; so our choice of factors is determined by the ability to best represent a vast cross-section of portfolios, and our main objective is to produce realistic dividend strips, that best match the traded ones.

The paper also relates to a large number of studies that have explored the term structure of risky assets, in addition to that of equity market dividend claims. Among these, different studies have focused on the term structure of currency risk (Backus et al. (2018)), variance risk (Dew-Becker et al. (2017)), housing risk (Giglio et al. (2014)). Chernov et al. (2018) have stressed the usefulness of multi-period returns to test asset pricing models. Several papers have proposed models that aim to explain observed patterns in the term structure of discount rates, among which Croce et al. (2014), Gormsen (2018). Methodologically, the paper is also similar to Adrian et al. (2015) who propose a similar affine structure of the SDF but do not use it to explore the term structure of risky assets.

Finally, our paper relates to a third approach used in the literature to explore the term structure of discount rates using only equity portfolios. This approach, followed by Weber (2018), Gormsen and Lazarus (2019), and Gonçalves (2019), is based on estimating the duration of portfolios directly (instead of by modeling the dynamics of dividends and effectively estimating duration as exposures to dividend shocks of different horizons), and using it to back out implied discount rates at different horizons.

2 The Model

2.1 Motivation

In this section we introduce our reduced-form asset pricing model. The model specifies: 1) a stochastic discount factor (SDF) that depends on factors and their shocks; 2) a full specification of the factor dynamics; and 3) equations linking the prices and returns of any portfolios to those factors. This model is quite general, and to make estimation and identification feasible, we impose different restrictions when we bring it to the data.

The specification for the SDF and the choice of factors builds on the empirical evidence in Kozak et al. (2018) and Haddad et al. (2020), who construct asset pricing factors for the cross-section using principal components (PCs) of a large cross-section of test portfolios, and also discuss empirically

successful restrictions to the drivers of risk prices (specifically, the valuation ratios of each PC). This specification strikes a good balance between fit and parsimony, that was explored in the papers mentioned above. To this specification of the SDF, our model adds the dynamics for the factors, which then allows us to solve for the equilibrium prices of securities with arbitrary maturity.

Our model, therefore, integrates an existing, empirically successful, specification for the SDF with a vector autoregression (VAR) structure to capture predictability across multiple horizons and therefore obtain term-structure implications. When applied to pricing one-period returns, the model directly maps into the standard cross-sectional pricing literature, e.g., Kozak et al. (2018, 2020). Prices of risk and expected returns on factor portfolios, however, are time-varying, which maps directly into the framework of Haddad et al. (2020). Our paper thus generalizes and combines existing approaches aiming at explaining cross-sectional and time-series patterns, with the ultimate goal of extracting a cross-section of term structures for different portfolios. Specifically, the structure of the affine model allows us to make progress by decomposing dividend yields into discount rates and expected dividend growth at each horizon, and use the entire estimated model (both the dynamics and the SDF, which are estimated simultaneously) to construct the various term structures.

2.2 The setup

State space dynamics. We begin by specifying a general factor model with k factors F_t, whose identity we will discuss later, that follow linear dynamics:

$$\underbrace{F_{t+1}}_{k\times 1} = \underbrace{c}_{k\times 1} + \underbrace{\rho}_{k\times k} F_t + u_{t+1}, \tag{1}$$

and with $\text{var}_t\left(u_{t+1}\right) = \Sigma$ constant (i.e., we assume homoskedasticity). Further, we assume that shocks u_{t+1} are normally distributed.

SDF. Denote the risk-free rate as $r_{f,t}$. We assume a log-linear SDF, where the priced shocks are u_{t+1}, with time-varying risk prices λ_t,

$$m_{t+1} = -r_{f,t} - \frac{1}{2}\lambda_t'\Sigma\lambda_t - \lambda_t' u_{t+1}. \tag{2}$$

Risk prices λ_t are assumed to be affine in F_t,

$$\underbrace{\lambda_t}_{k\times 1} = \underbrace{\lambda}_{k\times 1} + \underbrace{\Lambda}_{k\times k}\underbrace{F_t}_{k\times 1}. \tag{3}$$

The price P_t, dividend D_t, and gross return $R_{t+1} = \frac{P_{t+1}+D_{t+1}}{P_t}$ of any test asset satisfy the Euler equation:

$$1 = \mathbb{E}_t\left[e^{m_{t+1}}\frac{P_{t+1}}{P_t}\left(1+\frac{D_{t+1}}{P_{t+1}}\right)\right] = \mathbb{E}_t\left[e^{m_{t+1}+r_{t+1}}\right]. \tag{4}$$

Here,

$$r_{t+1} = \log\left[R_{t+1}\right] = \log\left[\frac{P_{t+1}}{P_t}\left(1 + \frac{D_{t+1}}{P_{t+1}}\right)\right] = \Delta p_{t+1} + y_{t+1}, \tag{5}$$

where we denote as r_{t+1} the log returns on the asset, which we decompose into the sum of log price changes, $\Delta p_{t+1} = \log\left(\frac{P_{t+1}}{P_t}\right)$, and the dividend yield, which we define as

$$y_t \equiv \log\left(1 + \frac{D_t}{P_t}\right). \tag{6}$$

Note that, given this definition of the dividend yield, we do not use any approximations in the decomposition of log returns in Equation (5). In other words, the identity linking price changes, dividend yields, and returns is always satisfied exactly in our paper given our definition of the dividend yield.

Under log-normality, by taking the log of both sides of Equation (4), we obtain:

$$0 = \mathbb{E}_t\left[m_{t+1}\right] + \mathbb{E}_t\left[\Delta p_{t+1} + y_{t+1}\right] + \frac{1}{2}\operatorname{var}_t\left[m_{t+1} + \Delta p_{t+1} + y_{t+1}\right]. \tag{7}$$

Next, we consider the cross-section of n financial assets which can be priced using Equation (7).

Price dynamics. We directly specify the dynamics of log price changes, in excess of the risk-free rate on financial assets:

$$\Delta p_{t+1} - r_{f,t} = \gamma_0 + \gamma_1 F_t + \gamma_2 u_{t+1} + \epsilon_{p,t+1}, \tag{8}$$

where we assume that expected price changes are driven by state variables in F_t, and shocks include both shocks to the state vector, u_{t+1}, and asset-specific shocks $\epsilon_{p,t+1}$. Note that in macro-finance models, often the dynamics of *dividends* are specified first, together with the SDF, and prices are then obtained by combining the two through the Euler equation. Here we follow the alternative approach of specifying the dynamics of *prices* first, together with the SDF, and the dynamics of dividends can be backed out of the other equations (the Euler equation and the identity linking prices, returns, and dividends). This is similar to the setup of Campbell (1991), who eliminates consumption from the various equations and expresses the entire model in terms of the dynamics of wealth and returns. The relationship between of the two approaches is discussed in detail later in Section 2.3.

Dividend yields. The Euler equation (4) implies that the price-dividend ratio of any asset can be expressed as an affine function of the state vector and stock-specific residuals $\epsilon_{y,t}$:

$$y_t = b_0 + b_1 F_t + \epsilon_{y,t}, \tag{9}$$

where parameters b_0 and b_1 are of sizes $n \times 1$, and $n \times k$, respectively as we show in Appendix A.1.[4]

[4]In principle, y_t could also depend on u_t, in addition to F_t. In such a case one could express u_t in terms of F_t and F_{t-1} using Equation (1), substitute it in, and satisfy the assumption by expanding the state vector to include F_{t-1}. To reduce estimation errors associated with large state spaces, we instead rely on a specific empirical choice of the state vector, which we discuss below, to guarantee that this assumption holds up well in the data (our state vector spans the

As discussed below, the dividend yields and returns of some linear combination of the original assets will be part of the state vector F_t itself. For those portfolios, Equation (9) is automatically satisfied without residual, with an appropriate choice of coefficients b_0 (zeroes) and b_1 (linking each portfolio to its position in the state vector). For other assets, parameters b_0 and b_1 can be solved for using Equation (7), given the estimates of the risk price parameters.

Excess returns. Equations (5), (8), and (9) imply that excess returns are also affine in factors and shocks:

$$r_{t+1} - r_{f,t} = \beta_0 + \beta_1 F_t + \beta_2 u_{t+1} + \epsilon_{r,t+1}. \tag{10}$$

In this specification, $\mathbb{E}_t\left[r_{t+1} - r_{f,t}\right] = \beta_0 + \beta_1 F_t$ is the risk premium of the n assets, which satisfies the no-arbitrage condition in Equation (7); parameters β_0 and β_1 are of sizes $n \times 1$, and $n \times k$, respectively. u_{t+1} and $\epsilon_{r,t+1}$ are the systematic and idiosyncratic shocks, respectively, and β_2 is the $n \times k$ matrix of exposures of n assets to k systematic risk such that

$$\beta_0 = \gamma_0 + b_0 + b_1 c \tag{11}$$

$$\beta_1 = \gamma_1 + b_1 \rho \tag{12}$$

$$\beta_2 = \gamma_2 + b_1, \tag{13}$$

$$\epsilon_{r,t+1} = \epsilon_{p,t+1} + \epsilon_{y,t+1}. \tag{14}$$

Note that the factor dynamics, the SDF specification, and the Euler equation impose additional restrictions among the coefficients of these equations. For example, the Euler equation links risk premia β_0 and β_1 to risk exposures β_2. We specify the additional restrictions that lead us to an identified model below.

Deflated dividends. We infer log *deflated* dividend growth (defined as: log dividend growth net of the log risk-free rate, $\Delta d_{t+1} - r_{f,t}$) from the returns identity:

$$r_{t+1} - r_{f,t} = y_{t+1} + pd_{t+1} - pd_t + (\Delta d_{t+1} - r_{f,t}), \tag{15}$$

where the log price-dividend ratio, pd_t, is a non-linear function of y_t given by $pd_t = -\log\left(\exp\left(y_t\right) - 1\right)$.

Dividend dynamics are, in general, not linear, because of the nonlinearity of the relation between returns, dividends, and prices. Rather than imposing approximate log-linearity via a Campbell-Shiller loglinearization, we work directly with the deflated dividends and their nonlinear dynamics, expressing them in closed-form after the model has been fully solved using observable returns r and dividend yields y.

Specializing the state vector. We now take a specific stand on the state vector F_t. Our choice is motivated by the empirical findings of Kozak et al. (2020) and Haddad et al. (2020). In particular, Kozak et al. (2020) show that a few dominant principal components (PCs) of a large cross-section

yields on included financial assets with an average R-squared of 99%).

of anomaly portfolio returns explain the cross-section of expected returns well. Haddad et al. (2020) further demonstrate that valuation ratios strongly and robustly predict expected returns on these PCs—and their risk prices in the SDF—in the time-series. They argue that the resulting time variation in risk prices is critical for adequately capturing the dynamic properties of the pricing kernel.

We, therefore, choose a specification which is motivated by both results. We assume that the dynamics of the economy are fully captured by $p = k/2$ linear combinations of *excess log returns*, $f_{r,t} = Q'(r_t - r_{f,t})$ and p linear combinations of *yields*, $f_{y,t} = Q'y_t$ of the n assets, for some $n \times p$ matrix Q which constructs the same p linear combinations of these variables based on the assets:

$$F_t \equiv \begin{bmatrix} f_{r,t} \\ f_{y,t} \end{bmatrix} = \begin{bmatrix} Q'(r_t - r_{f,t}) \\ Q'y_t \end{bmatrix}. \tag{16}$$

We refer to the p linear combinations of excess log returns, $f_{r,t}$, as *factor (excess) returns* and to the p linear combinations of yields, $f_{y,t} = Q'y_t$, as *factor yields* (or factor dividend-price ratios).

Being a linear combination of log returns, the "factor returns" $f_{r,t}$ are *not* themselves log returns on a portfolio of assets with weights given by Q – that is, $f_{r,t}$ is not itself the log excess return of a tradable portfolio (except for the market which we treat separately and include as the first factor). However, because each of the n assets is tradable, the Euler equation has to hold for each of them, so that the linear combination (Q) of those Euler equations has to hold as well – a restriction we impose in solving the model.

Our parsimonious specification allows us to be consistent with both Kozak et al. (2020) and Haddad et al. (2020): the p factor returns explain the cross-section of expected returns well, while the p factor yields allow us to explain the cross-section of dividend yields to satisfy Equation (9), as well as capture the dynamics of conditional means and SDF risk prices.

This specification can be seen as an extension of the setup of Campbell (1991), in which the dynamics of the economy (also represented by a VAR) include returns and the dividend-price ratio of *one* portfolio (the market), plus additional predictors. Here, the state vector includes a pair $(f_{r,t}, f_{y,t})$ for p linear combinations of test assets, and no additional predictors.

Given that our SDF is based on that in Kozak et al. (2018) and Haddad et al. (2020), our model naturally performs equally well in pricing the cross-section of characteristic-sorted portfolios that we use as test assets. It is important to note, however, that this does not imply that our specification will mechanically be able to price *other* assets, including dividend strips on different underlying portfolios. The ability to price these other assets depends both on the SDF and on the structure of the dynamics in the model. This is why, after estimating our model using available stock portfolios, we evaluate the ability of the model to generate realistic term structures, by comparing the implied dividend term structures with the ones observed in the data. As we discuss below, we find that the model does a good job in matching traded S&P 500 dividend strip data.

Restrictions. We now introduce three restrictions on the model. First, we assume that there are only p priced risks, and they are fully spanned by our p return factors: so m_t loads only on the p innovations in return factors (the first p elements of u_{t+1}). This means that only the first p elements

of λ_t are non-zero. In turn, this implies that only the first p rows of λ and Λ can be nonzero. The remaining shocks in u_t drive the dynamics of the economy but are not priced by investors. This assumption is motivated by Kozak et al. (2018, 2020) who suggests that an SDF constructed from a small number of (p) diversified portfolio returns prices well the cross-section of returns. We then simply allow the dynamics to also include additional, non-priced shocks.

Second, we impose that only dividend yields (and not lagged returns) drive time variation in risk premia. This means that the matrix Λ will have the following structure:

$$\Lambda = \begin{bmatrix} 0_{p\times p} & \tilde{\Lambda} \\ 0_{p\times p} & 0_{p\times p} \end{bmatrix}$$

where the zeros in the second row are due to the fact that only returns shocks are priced, and the zeros in the top-left corner imply that risk premia variation is entirely driven by factor yields $f_{y,t}$; $\tilde{\Lambda}$ is a $p \times p$ matrix of risk price loadings on dividend yields.

As we show in greater detail below, these assumptions imply restrictions on the transition matrix ρ of the factors: they imply that expectations of factor returns are a function of lagged yields but not lagged realized returns. To this restriction, we add a third restriction on the conditional mean of the yields: we impose that it is also only a function of the lagged yields but not of the lagged returns. We therefore assume that

$$\rho = \begin{bmatrix} 0_{p\times p} & \rho_{r,y} \\ 0_{p\times p} & \rho_{y,y} \end{bmatrix},$$

where $\rho_{r,y}$ could potentially be further restricted to be a $p\times p$ diagonal matrix, based on the evidence in Haddad et al. (2020) that own valuation ratios are the strongest predictors of factor returns. We do not currently impose the latter restriction to remain as flexible as possible. The former restriction—that lagged returns forecast neither returns nor yields—is relatively mild, in our opinion, and consistent with voluminous literature documenting low autocorrelation in equity returns.

Note also that similar restrictions have been imposed in the term-structure literature, for instance in Cochrane and Piazzesi (2008). They specify an SDF in which only shocks to bond yields are priced, and their SDF risk prices are fully driven only by the Cochrane and Piazzesi (2005) factor.

Prices of risk. Using the definition of factor returns in (10), we use (16) to express them as:

$$f_{r,t+1} = Q'\left(r_{t+1} - r_{f,t}\right) = \beta_{f,0} + \beta_{f,1}F_t + \beta_{f,2}u_{t+1}, \tag{17}$$

where

$$\beta_{f,0} = Q'\beta_0 = c_r, \tag{18}$$

$$\beta_{f,1} = Q'\beta_1 = \left[\mathbf{0}_{p\times p}, \rho_{r,y}\right], \tag{19}$$

$$\beta_{f,2} = Q'\beta_2 = \left[I_{p\times p}, \mathbf{0}_{p\times p}\right], \tag{20}$$

where the first equality in each of the equations above is reflecting the fact that "factor returns" $f_{r,t+1}$ are linear combinations of test assets' log excess returns, and the second equality holds true because each factor is part of the state space vector F_t, which imposes specific restrictions on all of the parameters (which also leads to the absence of idiosyncratic shocks in (17)).

We now use these restrictions to show how the factors and the Euler equation can be used to link the parameters β to the prices of risk in the SDF.

We plug the expressions in (2), (8), and (9) into (7):

$$\mathbf{0}_{n\times 1} = \beta_0 + \beta_1 F_t - \beta_2 \Sigma (\lambda + \Lambda F_t) + \frac{1}{2} \operatorname{diag}\left[\beta_2 \Sigma \beta_2' + \Sigma_\epsilon\right]. \tag{21}$$

Pre-multiplying by Q', expressing in terms of $\beta_{f,\cdot}$, and matching coefficients on F_t, we get two equations which can be used to solve for λ and Λ, given the parameters $c_r, \rho_{r,y}$ and estimates of the test assets' variance terms:

$$0 = \beta_{f,1} - \beta_{f,2}\Sigma\Lambda, \tag{22}$$

$$0 = \beta_{f,0} - \beta_{f,2}\Sigma\lambda + \frac{1}{2} Q' \operatorname{diag}\left[\beta_2 \Sigma \beta_2' + \Sigma_\epsilon\right]. \tag{23}$$

2.3 Discussion of the setup

An alternative modeling setup adopted in the literature – which we refer to as the "alternative" modeling approach – is to specify a loglinear SDF and linear dynamics for log dividend growth, and then obtain prices by imposing the Euler equation (e.g., see Brennan et al. (2004), Koijen et al. (2017), and Gupta and Van Nieuwerburgh (2021)). The main difference between this and our specification is that we model log prices (or, equivalently, excess log equity returns) as linear in factors F_t and shocks u_{t+1}, whereas the "alternative" specification models log dividend growth Δd_{t+1} as linear in F_t and u_{t+1}. Both assumptions are plausible and both are widely used in asset pricing. The "alternative" assumption of dividend growth linearity is primarily used in the macro-finance literature (e.g., in the context of Lucas economies). Our assumption of (log) equity price and return linearity is widespread in empirical asset pricing in the context of factor models, and in return-based models like the ICAPM. In this section we discuss in detail the relationship between the two approaches.

Returns, dividends and price-dividend ratios are related by the usual non-linear identity in (15). Given that we are aiming to price both dividend strips and stocks, there are two relevant Euler equations that need to be satisfied (for strips and stocks respectively):

$$\frac{P_t^{(n)}}{D_t} = E_t\left[M_{t+1}\frac{D_{t+1}}{D_t}\frac{P_{t+1}^{(n-1)}}{D_{t+1}}\right], \tag{24}$$

$$1 = E_t\left[M_{t+1}\frac{P_{t+1}}{P_t}\left(1+\frac{D_{t+1}}{P_{t+1}}\right)\right] \equiv E_t\left[M_{t+1}R_{t+1}\right], \tag{25}$$

where $P_t^{(n)}$ denotes the price to an n-year dividend claim.

As mentioned above, the alternative approach assumes that $\Delta d_{t+1} = \log\left(\frac{D_{t+1}}{D_t}\right)$ is linear in F_t

and u_{t+1}. If M_{t+1} is also loglinear in F_t and u_{t+1}, then it is immediate from (24) that the log price-dividend ratios of *dividend strips* of all portfolios, $pd_t^{(n)}$, will be linear in F_t. The stock's price-dividend ratio *cannot* be linear in F_t at the same time. To facilitate the estimation of the model for directly observable stock D/P ratios, the alternative approach typically relies on the Campbell-Shiller *approximation* (approximate log-linearization), under which pd_t is approximately linear in F_t.[5] Due to these approximations, the price of a stock in this approach is *not* exactly the same as the sum of the prices of the strips.

In our approach, from (25), if excess returns $r_{t+1} - r_{f,t}$ are linear in F_t and u_{t+1}, and M_{t+1} is loglinear, it follows that the *equity*—rather than strip—dividend-price ratios, $y_t = \log\left(1 + \frac{D_t}{P_t}\right)$, are linear in F_t. Equity returns are then linear, too, whereas, naturally, neither pd_t, $pd_t^{(n)}$, nor $\Delta d_{t+1} - r_{f,t}$ will be linear in F_t and u_{t+1}. Our set of assumptions has several advantages. First, our affine model, based on the observable state space, can be estimated *without* imposing the Campbell-Shiller approximation (instead imposing the exact identity linking prices, dividends and returns). Intuitively, this is so because we observe D/Ps of stocks rather than strips in the data, so making stocks' D/Ps (rather than dividend strips' D/Ps) linear in states and including them in the state vector makes the estimation process much easier. Second, *after* we have fully solved the affine model, we compute dividend strip yields and implied deflated dividend growth, allowing them to be exact nonlinear functions of F_t. As a result, we automatically satisfy the no-arbitrage condition that the price of a stock is exactly equal to the sum of the prices of all its dividends.

We also note that, if one *wants to* impose the Campbell-Shiller approximation, then our approach is exactly equivalent to the "alternative" approach. To see this, note that the Campbell-Shiller approximation of the returns equation (15) makes $\log\left(1 + e^{pd_t}\right)$ an approximately linear function of pd_t, that is:

$$\log\left(1 + e^{pd_t}\right) \simeq a + b \times pd_t \tag{26}$$

Under the Campbell-Shiller approximation, *both* y_t and pd_t are (approximately) linear in F_t, and so are dividend growth and returns. To see this, note that the approximation in (26) and the identity $y_t = \log\left(1 + \frac{D_t}{P_t}\right) = \log\left(1 + e^{pd_t}\right) - pd_t$ imply that y_t is affine in pd_t. Therefore, under either of the two setups discussed previously, both y_t and pd_t are linear in F_t: the two setups are equivalent.

To summarize, given the nonlinearity of the fundamental identity relating pd_t, r_{t+1}, and Δd_{t+1} in (15), if some of the three components is linear in some variables F_t, the others cannot all be also linear in F_t. The "alternative" approach assumes that Δd_{t+1} and $pd_t^{(n)}$ are linear (and so both pd_t and y_t are nonlinear), but the Euler equation cannot be solved analytically under this assumption alone; to close the model analytically, one needs to also assume the Campbell-Shiller approximation or approximately solve a system of nonlinear equations defining dividend-price ratios for each portfolio. Our approach assumes instead that $r_{t+1} - r_{f,t}$ and y_t are linear in F_t. These assumptions are sufficient to solve the model fully even without imposing the Campbell-Shiller approximate linearization. Only *after* our model is fully solved analytically, the remaining quantities (e.g., $\Delta d_{t+1} - r_{f,t}$, pd_t, $pd_t^{(n)}$)

[5] For example, see Koijen et al. (2017); Gupta and Van Nieuwerburgh (2021). Alternatively, rather than specifying an observable Gaussian state vector as we do in the paper, one could introduce a latent Gaussian state vector that is non-linearly related to D/P ratios and solve the alternative approach exactly.

can be expressed as exact nonlinear functions of the state vector F_t and shock u_{t+1}. Finally, if one is willing to impose the Campbell-Shiller approximation, the two setups are equivalent.

We conclude by noting that the idea of modeling $y_t = \log\left(1 + \frac{D_t}{P_t}\right)$, as opposed to $pd_t = \log\left(\frac{P_t}{D_t}\right)$, has appeared in the macro-finance literature. For example, Martin (2013) uses it to study the importance of cumulants in consumption-based models, and Gao and Martin (2021) derive a loglinear approximation of y_t instead of pd_t similar to that of Campbell-Shiller.

2.4 Identification and Estimation

While the model contains a large number of parameters, many of them are related through no-arbitrage restrictions or through some of the additional restrictions imposed above. In particular, it is useful to distinguish *reduced-form* parameters that directly enter moment conditions that can be estimated from the data, and *structural* parameters that in turn determine the reduced-form parameters.

Reduced-form parameters and moment conditions. The entire state vector F_t is fully observed. Equation (1), therefore, implies a first set of moment conditions, which depend on the parameters c and ρ. We use OLS moments, with additional restrictions that elements of ρ corresponding to loadings on lagged returns are all zero, that is, $\mathbb{E}\left(u_{t+1} \otimes [1, F_{y,t}]\right) = 0$, where $\otimes$ denotes a kronecker product. Note that for factor returns, which are part of F_t, the parameters $\beta_{f,0}, \beta_{f,1}$ and $\beta_{f,2}$ from (17) are directly a function of c and ρ (because these return equations are just the first p rows of (1)).

For test assets we have two sets of moment conditions: one relating contemporaneously their yields to the factors, (9), and another relating returns to lagged factors, (10). These equations form a set of moment conditions that depend on parameters β_0, β_1, β_2, b_0 and b_1 for all test assets. Again, we use OLS moments, with parameter restrictions, $\mathbb{E}\left(\varepsilon_{t+1} \otimes [1, F_{y,t}, u_{r,t+1}]\right) = 0$ and $\mathbb{E}\left(\epsilon_t \otimes [1, F_{y,t}]\right) = 0$.

To sum up, moment conditions from equations (1), (9), (10) depend on the parameters c and ρ, as well as β_0, β_1, β_2, b_0 and b_1 for all assets.

Structural parameters. We then have a set of structural parameters that are linked to the reduced-form parameters by identities and arbitrage restrictions. First, the parameters γ_0, γ_1 and γ_2 are related one-to-one to the reduced-form parameters ((11), (12), (13)); these parameters do not add any additional restriction to the system, and simply correspond to an alternative representation of (10), in terms of Δp_{t+1} instead of r_{t+1}.

More important are the other structural parameters λ and Λ. These parameters introduce restrictions on the reduced-form parameters of *all* portfolios, through the valuation equation in Equation (7).

Estimation and inference. We estimate the model using GMM. We use the moment conditions described above to estimate the reduced-form and structural parameters (i.e., imposing the valuation restriction (7) using all test assets. We use a prespecified, diagonal weighting matrix for GMM where the factor moments moments are weighted by 1, and the individual assets' moments are all normalized by the square root of the number of test assets, $\sqrt{n}$, to keep their contribution to the GMM objective

invariant to n.[6] This weighting matrix ensures a good balance between the two sets of moment conditions and yields reasonable estimates of the risk premia of all 102 portfolios—which is important in our context given the weaker factor structure of equities compared to other settings (e.g. bonds).

We estimate the dynamics at the annual horizon using monthly data (therefore, with overlapping yearly observations). We derive standard asymptotic GMM standard errors for all reduced-form and structural parameters. To account for overlapping data we use a spectral density covariance matrix of moments with 12 lags, following the approach in Hansen and Hodrick (1980).[7] Finally, we compute standard errors on any derived quantities—such as model-implied yields and returns on dividend strips, which are non-linear functions of structural parameters—using the Delta method.

Note that there are two sources of uncertainty for the model-implied moments of interest. First, there is uncertainty stemming from the fact that the model parameters have to be estimated, and so some time series are not observed but estimated (for example, the time series of the prices of dividend strips, which are not an input in our estimation, and therefore must be estimated). We refer to this part of uncertainty as *parameter uncertainty.* Second, there is the uncertainty stemming from the fact that some objects of interest are unconditional moments, and estimating those requires computing time-series averages. We refer to this as *sampling uncertainty.* For example, suppose that we are interested in the unconditional slope of the dividend strip term structure. The GMM estimator directly gives us the estimate for this moment, together with standard errors that incorporate *all* types of estimation uncertainty. However, we can separate the two by thinking of the estimation procedure as follows. First, an estimate of the slope at each point in time is obtained by estimating our model with GMM, then estimating the model-implied prices of long-term and short-term dividend strips at each point in time, and then computing the difference between the two in each period. Next, the unconditional average slope (the quantity of interest) is estimated by taking the time-series average of the slope estimated at each point in time. The two sources of uncertainty reflect the two steps in this procedure. Note that if dividend strip prices were directly observable, the *parameter uncertainty* would disappear, and the *sampling uncertainty* will account for the entirety of the estimation uncertainty.

The GMM estimation approach we described above automatically takes *both* sources of uncertainty into account, when estimating any moment. But comparing the two sources of uncertainty is useful to understand how much of our estimation uncertainty on objects of interest (like the unconditional slope of the dividend term structure) is due to the fact that we do not observe the dividend strip prices, and how much of that uncertainty would be there even if we did observe the prices, since it comes from the fact that we only have a finite sample to work with to estimate unconditional moments.

This estimation uncertainty decomposition allows us to understand a tradeoff that our way of estimating term structures faces, compared to working only with traded securities: our procedure adds *parameter uncertainty* to the estimation (since we need to estimate the model to generate prices

[6] We have also explored the use of efficient GMM, and find that it is numerically unstable in our setting (which features a large number of moment conditions and parameters). This result is similar to that in Campbell et al. (2013), who also use GMM to jointly estimate, in a lower-dimensional setting compared to ours, the dynamics of the model and the test asset moment conditions, and also find efficient GMM to be numerically unstable. Rather than imposing restrictions on the parameter space (e.g. bounds on risk prices) as in Campbell et al. (2013) to ensure convergence of the efficient GMM estimator, here we use a prespecified weighting matrix.

[7] We perform a non-parametric bootstrap exercise to validate our standard errors in Section 4.7.6.

of securities like dividend strips), but has the advantage of significantly lengthening the time series available for the study of unconditional moments, therefore reducing the *sampling uncertainty* (since traded dividend strips are available only for a short time period). In practice, we find that parameter uncertainty is small relative to data sampling uncertainty for most moments we consider. This observation suggests that our procedure leads to the overall reduction of standard errors compared to methods that use traded strips but rely on a shorter sample.

We provide more details on the estimation in Appendix B, and come back to the comparison of the sources of uncertainty when we present the empirical results.

2.5 Dividend strips

We next derive the prices and returns of theoretical dividend strips in the model. These can be computed *after* having estimated the entire model, because they are simple functions of the parameters we estimate via GMM. Consider first a fully-diversified portfolio with its dividend D_t and price P_t. Note first that since

$$P_{t+n} = P_t \exp\left[\sum_{i=1}^{n} \Delta p_{t+i}\right]$$

and

$$\frac{D_{t+n}}{P_{t+n}} = \left[\exp\left(y_{t+n}\right) - 1\right]$$

then

$$\frac{D_{t+n}}{P_t} = \frac{D_{t+n}}{P_{t+n}} \exp\left[\sum_{i=1}^{n} \Delta p_{t+i}\right] = \left[\exp\left(y_{t+n}\right) - 1\right] \exp\left(\sum_{i=1}^{n} \Delta p_{t+i}\right),$$

we then have that the price of a claim to D_{t+n}, $P_t^{(n)}$, as a fraction of the price of the portfolio, P_t, is:

$$\frac{P_t^{(n)}}{P_t} \equiv w_t^{(n)} = \mathbb{E}^Q\left\{\frac{D_{t+n}}{P_t} \exp\left[-\sum_{i=1}^{n} r_{f,t+i-1}\right]\right\} = \mathbb{E}^Q\left\{\left[\exp\left(y_{t+n}\right) - 1\right] \exp\left[\sum_{i=1}^{n} \left(\Delta p_{t+i} - r_{f,t+i-1}\right)\right]\right\} \tag{27}$$

$$= \mathbb{E}^Q\left\{\exp\left[y_{t+n} + \sum_{i=1}^{n} \left(\Delta p_{t+i} - r_{f,t+i-1}\right)\right]\right\} - \mathbb{E}^Q\left\{\exp\left[\sum_{i=1}^{n} \left(\Delta p_{t+i} - r_{f,t+i-1}\right)\right]\right\} \tag{28}$$

$$= \exp\left(a_{n,1} + d_{n,1} F_t\right) - \exp\left(a_{n,2} + d_{n,2} F_t\right), \tag{29}$$

The parameters in this equation can be shown to satisfy the following recursions:

$$a_{n,\cdot} = a_{n-1,\cdot} + \gamma_0^* + d_{n-1,\cdot} c^* + \frac{1}{2}\left(d_{n-1,\cdot} + \gamma_2\right) \Sigma \left(d_{n-1,\cdot} + \gamma_2\right)' + \frac{1}{2}\sigma_v^2 \tag{30}$$

$$d_{n,\cdot} = \gamma_1^* + d_{n-1,\cdot} \rho^*, \tag{31}$$

and initial values by $a_{0,1} = b_0 + \frac{1}{2}\left(\sigma_r^2 - \sigma_v^2\right)$, $d_{0,1} = b_1$, $a_{0,2} = 0$, $d_{0,2} = 0$, $\sigma_r^2 = \text{var}\left(\varepsilon_t\right)$, and $\sigma_v^2 =$

var (v_t). In these formulas, stars indicate risk-neutral parameters.[8] In our estimation, we find that the risk-neutral dynamics of the model are stationary, which is sufficient to guarantee that the price of the stock market, which is the infinite sum of the prices of the dividend strips, is finite. Note that the prices of dividend strips are obtained as a function of observable excess returns and dividend yields and their dynamics (captured by F_t and their dynamics). This is because specifying these dynamics implicitly fully specifies the dynamics of deflated dividends, which are sufficient to pin down dividend strip prices. Note also that $w_t^{(n)}$ can be interpreted as cap-weights of each dividend within a portfolio of all dividends of a stock (i.e., the stock itself).

Equity yields. We can also compute equity yields (for assets/portfolios with strictly positive dividends) at time t with maturity n, $e_{t,n}$, as defined in Van Binsbergen et al. (2013):

$$e_{t,n} = \frac{1}{n} \log \left(\frac{D_t}{P_t^{(n)}} \right) = \frac{1}{n} \left[\log \left(\exp \left(y_t \right) - 1 \right) - \log \left(\frac{P_t^{(n)}}{P_t} \right) \right]. \tag{32}$$

Forward equity can be easily computed using the (externally) observable bond yields:

$$e_{t,n}^f = \frac{1}{n} \log \left(\frac{D_t}{F_{t,n}} \right) = e_{t,n} - y_{t,n}^b, \tag{33}$$

where $y_{t,n}^b$ is the nominal government bond yield with no default risk, and where $F_{t,n}$ denotes the futures (or forward) price, which, under no arbitrage, is linked to the spot price by

$$F_{t,n} = P_{t,n} \exp \left(n y_{t,n}^b \right). \tag{34}$$

In Appendix A we provide additional definitions and derivations for returns and expected returns on dividend strips in our model. Additionally, we show that equity yields can be decomposed into log (expected) hold-to-maturity returns and log (expected) cumulative deflated dividend growth rates (see (A17) and (A19)).

3 Data

3.1 Stock data

We focus on a broad set of 51 stock-specific characteristics and long and short legs of portfolio sorts based on these characteristics.[9] We construct these portfolios as follows. We use the universe of CRSP and COMPUSTAT stocks and sort them into three value-weighted portfolios for each of the characteristics studied in Kozak et al. (2020) and Kozak and Santosh (2019) and listed in Table A.3, for a total of 51 characteristics.[10] Portfolios include all NYSE, AMEX, and NASDAQ firms; however, the

[8] Risk-neutral parameters are defined as: $\gamma_0^* = \gamma_0 - \gamma_2 \Sigma \lambda$, $\gamma_1^* = \gamma_1 - \gamma_2 \Sigma \Lambda$, $c^* = c - \Sigma \lambda$, $\rho^* = \rho - \Sigma \Lambda$.

[9] Our data are available at `https://www.serhiykozak.com/data`.

[10] We apply the following two automatic filters to characteristics as in Kozak et al. (2020) and Kozak and Santosh (2019) to arrive at 51 characteristics: (i) mark portfolio returns as missing if a portfolio contain fewer than 100 firms at any point in time, and (ii) remove characteristics for which more than 120 months of returns are missing, either in a

Table 1: Percentage of variance explained by anomaly PCs

Percentage of variance explained by each PC of the 51 long-short anomaly portfolio returns (top panel) and each PC of 102 returns on long and short legs of each characteristic sort (bottom panel). PC1 in the bottom panel refers to the market portfolio.

	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	PC9	PC10
	Long-Short (51 portfolios)									
% var. explained	24.9	19.4	10.1	5.7	4.3	3.9	3.1	2.9	2.5	2.2
Cumulative	24.9	44.3	54.4	60.1	64.4	68.3	71.4	74.3	76.8	79.0
	Long and short legs (102 portfolios)									
% var. explained	88.8	4.1	1.6	0.9	0.6	0.5	0.5	0.3	0.3	0.2
Cumulative	88.8	92.9	94.5	95.5	96.1	96.5	97.0	97.3	97.6	97.8

breakpoints use only NYSE firms as in Fama and French (2016). We obtain 102 portfolios, two for each anomaly (P1 & P3). Our sample consists of monthly returns from February 1973 to December 2020. For each portfolio we construct a corresponding measure of its valuation based on dividend-to-price ratios of the underlying stocks. We then construct the yield, y_t, as defined in Equation (6).[11]

We also construct long-short portfolios as differences between each anomaly's log return on portfolio 3 minus the log return on portfolio 1. Their valuation ratios are defined as the difference in yields y_t between the two legs. Most of these portfolio sorts exhibit a significant spread in average returns and CAPM alphas. This finding has been documented in the vast literature on the cross-section of returns and can be verified in Appendix Table A.3. In our sample, most anomalies show a strong pattern in average returns across tercile portfolios, consistent with prior research.

3.2 Choice of variables in the state vector

Kozak et al. (2018, 2020) show that a few dominant principal components (PCs) of a large cross-section of anomaly portfolio returns explain the cross-section of expected returns well. We use this insight to guide our choice of portfolios in the state space dynamics in Equation (1). In particular, we use excess log returns on the market and three PCs of 51 long-short portfolio returns based on the underlying long and short ends of each characteristic used in sorting, as our choice for $f_{r,t}$ in Equation (16). Formally, consider the eigenvalue decomposition of anomaly excess returns, $\text{cov}\left(r_{LS,t+1}\right) = \tilde{Q}\Lambda\tilde{Q}'$, where $\tilde{Q} = [\tilde{q_1}, ..., \tilde{q_{51}}]$ is the matrix of eigenvectors and Λ is the diagonal matrix of eigenvalues. The i-th PC portfolio is formed as $PC_{i,t+1} = \tilde{q}_i' r_{LS,t+1} = q_i'\left(r_{t+1} - r_{f,t}\right)$, where $\tilde{q}_i$ is the i-th column of $\tilde{Q}$, $q_i = [\tilde{q}_i, -\tilde{q}_i]'$, and $r_{t+1} - r_{f,t}$ are excess log returns on 102 portfolios underlying the long-short anomalies.

Table 1 shows that anomaly portfolio returns exhibit a moderately strong factor structure. The

short or long leg.

[11]The dividend-to-price ratio of a portfolio is defined as the sum of all dividends paid within the last year relative to the total market capitalization of all firms in that portfolio. Equivalently, it is the market-capitalization weighted average of individual stocks' D/P ratios.

top panel focuses on long-short portfolios with the market—which captures the vast majority of all co-movement across portfolios—effectively removed. It, therefore, focuses on explaining the remaining variation in portfolio returns once the market has been removed. The first PC accounts for one fourth of the total variation. The first three PCs explain more than 50% of the variation not captured by the market. The bottom panel extracts PCs of all 102 long and short legs of each sort, therefore not subtracting the market. It shows that the first four PCs (approximately the market and three cross-sectional PCs) capture more than 95% of variation in returns across 102 portfolios.

Haddad et al. (2020) further demonstrate that valuation ratios strongly and robustly predict expected returns on these PCs—and their risk prices in the SDF—in the time-series. They argue that the resulting time-variation in risk prices is critical for adequately capturing dynamic properties of the pricing kernel. Motivated by this evidence, we pick four yields corresponding to the market and three PC factors as our choice for y_t in Equation (16). Yields on PCs are constructed simply as a linear combination of log yields on underlying portfolios, with weights determined by returns' eigenvectors, $y_{pc_i} = q_i^{'} y_t$.

In summary, our state vector includes returns and yields on the market, three PCs of long-short log returns and their valuation ratios of all test assets. We further validate this choice by conducting an out-of-sample analysis in Section 4.7.4.

3.3 Bond data

Our model is designed to be orthogonal to the term structure of interest rates and does not require any bond data to estimate the model, other than the one-year risk-free rate $r_{f,t}$. To compute forward equity yields (for which bond yields are needed), we use data from Gurkaynak et al. (2006), who provide a long history of interpolated US bond yields.

After merging the equity and bond data, we obtain a full sample of annual (overlapping) observations at the monthly frequency from February 1973 to December 2020.

4 Results

In this section we report the empirical results of our estimation. We begin by reporting the estimates and the fit of the model; specifically, we show how the model fits the returns of the 51 anomaly portfolios, their price-dividend ratios, and their dividends.

Next, we evaluate the performance of the model against data that was *not* used in the estimation: the term structure of S&P 500 dividend strips and futures. We show that the dividend strips implied by our estimated model closely match the prices of the traded strips; we therefore replicate the main facts of van Binsbergen et al. (2012b), Van Binsbergen et al. (2013) and van Binsbergen and Koijen (2015) on the term structure of S&P 500 dividend strips and forwards. We view this exercise mainly as an out of sample validation of our empirical setup.

We then explore the novel empirical facts that emerge from our model, along both the cross-sectional and the time-series dimensions. In the cross-section, we show that our model produces a rich variety of shapes for the term structures for different portfolios. In the time series, we show that

Table 2: Estimates of the dynamics of the factors F_t

We report estimates of the dynamics of the factors F_t in Equation (1): c and $\rho_{\cdot,y}$ in $\rho = [0_{K\times p}, \rho_{\cdot,y}]$, and the R^2 for each of the eight equations of the dynamics of F_t. Hansen-Hodrick GMM standard errors using a spectral density matrix with 12 lags are reported in parentheses. The dynamics are estimated at an annual horizon using monthly overlapping observations in the February 1973 to December 2020 sample.

	c	$\rho_{\cdot,y}$				R^2
r_{mkt}	0.07	-0.16	1.35	3.21	0.50	10.44
	(0.068)	(4.8)	(1.7)	(1.5)	(0.9)	-
r_{pc1}	-0.05	-16.13	11.39	4.23	-2.57	22.09
	(0.2)	(15)	(4.8)	(3.2)	(2.1)	-
r_{pc2}	0.26	-7.54	0.34	1.12	0.06	7.19
	(0.12)	(3.8)	(2.2)	(2.9)	(1.6)	-
r_{pc3}	0.08	-0.64	1.46	3.88	1.40	11.62
	(0.081)	(5)	(1.7)	(1.5)	(1)	-
y_{mkt}	0.00	0.93	-0.00	0.02	-0.01	97.58
	(0.0015)	(0.086)	(0.035)	(0.038)	(0.034)	-
y_{pc1}	-0.00	0.57	0.72	0.02	-0.23	96.98
	(0.0046)	(0.25)	(0.091)	(0.083)	(0.051)	-
y_{pc2}	-0.01	-0.25	0.03	0.44	-0.10	71.19
	(0.0062)	(0.27)	(0.17)	(0.25)	(0.072)	-
y_{pc3}	-0.01	-0.58	0.31	0.08	0.01	20.32
	(0.0053)	(0.27)	(0.088)	(0.11)	(0.11)	-

the slope and shape of the term structure of different portfolios vary in interesting ways over time. Finally, we discuss various implications and applications of our results.

Taken together, our results show that the term structures of discount rates are heterogeneous across types of risks (captured by different portfolios) and vary significantly over time. The empirical patterns that we extract from the data in Section 4.4 (for example, the difference in the term structure behavior of value and growth stocks, or high- and low-profitability stocks) provide new moments that can help guide the construction and evaluation of asset pricing models.

4.1 Fit of the model

As discussed in section Section 2.4, we estimate our model in one step via GMM. The core of our model are the time-series dynamics of the dividend yields and returns of the four factors (the market plus three principal components from 51 anomalies). Table 2 reports the estimates of the dynamics of the factors F_t: c and ρ. As discussed in Section 2, we impose that the conditional mean of returns and yields is only a function of the lagged yields but not of the lagged returns, that, is, $\rho = [0_{K\times p}, \rho_{\cdot,y}]$. In Table 2 we, therefore, omit zeros and report only estimates and standard errors of $\rho_{\cdot,y}$, as well as c. In parentheses, we report GMM standard errors using a spectral density matrix with 12 lags.

For comparison, Table 3 reports the risk-neutral parameters c^Q and ρ^Q implied by the model. Note that all conditional loadings of returns on lagged yields are zero, $\rho_{r,y} = 0$, since under the risk-neutral

Table 3: Risk-neutral estimates of the dynamics of the factors F_t

We report risk-neutral estimates of the dynamics of the factors F_t in Equation (1): c_r^Q, c_y^Q,and $\rho_{y,y}^Q$ in $c^Q = \left[c_r^Q, c_y^Q\right]'$ and $\rho^Q = \left[0_{K\times p}, \rho_{\cdot,y}^Q\right]$, where $\rho_{\cdot,y}^Q = \left[0, \rho_{y,y}^Q\right]'$. Annual overlapping observations at monthly frequency from February 1973 to December 2020.

	c_r^Q	c_y^Q	$\rho_{y,y}^Q$			
MKT	-0.01	0.00	0.92	0.03	0.09	-0.00
PC1	0.03	-0.00	0.33	0.94	0.19	-0.26
PC2	0.03	-0.00	-0.24	-0.16	0.35	-0.06
PC3	-0.00	-0.00	-0.35	0.18	0.14	0.10

dynamics the expected excess return on any asset is equal to the risk-free rate; we, therefore, report only loadings of yields onto lagged yields, $\rho_{y,y}$, as well as the intercepts, c_r^Q and c_y^Q. The intercepts of the return regressions, c_r^Q, are non-zero and reflect a variance adjustment.

The last column of Table 2 also shows the R^2 for each of the eight equations of the dynamics of F_t. The first four regressions are effectively predictive regressions of yearly excess returns using the lagged dividend yields of the factor portfolios as predictors; the last four regressions are predictive regressions of dividend yields using lagged dividend yields as predictors, again for the four factor portfolios. The results are consistent with those in Haddad et al. (2020), who show that principal components of anomaly returns are strongly and robustly predictable by their valuation ratios, and that this predictability is essential for capturing dynamic properties of an SDF. That paper also shows that this predictability survives out-of-sample, suggesting that our analysis should also be robust to out-of-sample tests.

One of the main objectives of the model is to fit the cross-section of returns for the 51 anomalies (and not just the four linear combinations we use as factors). The average R^2 for the regression of portfolio returns in Equation (10) is 93.3%. The average R^2 for the regression of dividend yields onto the factor dividend yields in Equation (9) is 98.7%. Our model also fits well the cross-section of risk premia for the 51 anomalies, reaching a cross-sectional R^2 of 42.9% in the full sample from February 1973 to December 2020. This compares favorably with the R^2 of 73% one obtains on the much narrower cross-section of 25 size and book-to-market-sorted portfolios when using the three Fama-French factors (Fama-French factors explain the cross-section of 51 anomalies with an R^2 close to 0%).[12] This evidence demonstrates that our model performs well in fitting not only the dynamics of the PCs themselves, but also the time-series and cross-sectional properties of the 51 anomalies.

To get a sense of the factors' behavior in the data, Figure 1 and Figure 2 report the time-series of factor returns and factor yields. There are several interesting patterns to note. First, while the first factor (the market) has an extremely persistent dividend yield, as well known in the literature, the other factors' persistence is significantly lower. Second, the factors clearly capture different economic forces. For example, the fourth factor captures relatively high-frequency dynamics, whereas the third factor is most strongly associated with the financial crisis.

[12]Kozak et al. (2018) provide more detailed evidence on the relative performance of a PC-based model and the Fama-French model in explaining a wide cross-section of anomaly returns, which is consistent with our findings.

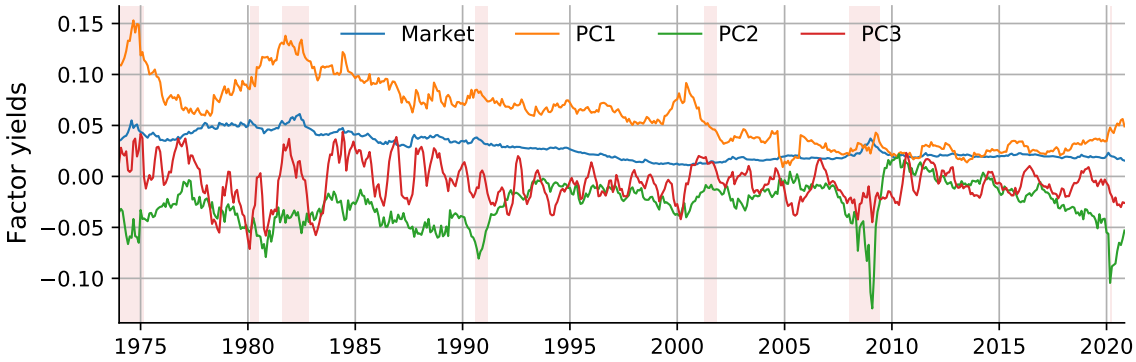

Figure 1: Time-series of factor yields. We show factor yields, $f_{y,t}$, for the aggregate market and three PCs of anomaly portfolios. Annual dividends, monthly overlapping observations.

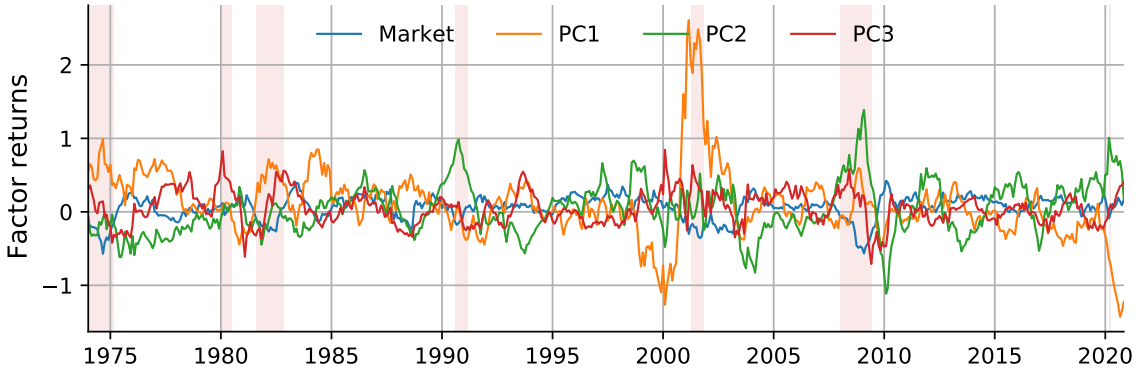

Figure 2: Time-series of factor returns. We show factor returns, $f_{r,t}$, for the aggregate market and three PCs of anomaly portfolios. Annual returns, monthly overlapping observations.

The fact that these factors display different dynamics is crucial for our identification. Haddad et al. (2020) argue that time-series predictability of these PC-based factors is critical to adequately capturing the dynamic properties of the pricing kernel — this is the key motivation for our choice of the state space vector. We can identify term structures of discount rates only because we can identify shocks to dividends and discount rates at different horizons — in other words, we need to be able to estimate shocks that give rise to different dynamic responses of the economy, to be able to estimate how investors price these shocks differently. Intuitively, by studying how prices of portfolios respond to short-lived shocks to their dividends we can learn about the discount rate at short horizons (controlling for discount rate movements that are correlated with the dividend shocks); by studying how prices respond to shocks that affect dividends far in the future we can learn about the long-term discount rate (again controlling for simultaneous discount rate changes). To identify different points along the term structure, we therefore need to observe different shocks at different horizons for dividends and discount rates. To validate empirically the idea that our model is able to capture both short-term and long-term shocks to portfolios, we regress returns on the long-short duration sorted portfolio onto the variables in the state vector, F_t, and find that the R-squared of this regression is 92% (see Table A.5

in the Internet Appendix). This result suggests that the factor innovations indeed capture news that drive the wedge between short-duration and long-duration portfolios.

4.2 Fit to traded S&P 500 dividend forwards

An important step forward in understanding the term structure of risk premia in the data and in the models has come from the study of traded claims to dividends of finite maturity (dividend strips and dividend futures), starting with the seminal work of van Binsbergen et al. (2012b). Recent papers have expanded the sample to include the prices of dividend forwards traded over the counter and, more recently, on exchanges (e.g. Bansal et al. (2017) and van Binsbergen and Koijen (2015)).

One of the main goals of our paper is to provide a framework to recover implied dividend strip and forward prices from data that only includes equity portfolios.[13] A simple and direct criterion to evaluate whether we can do so successfully is to verify whether our implied dividend forwards match those from the traded contracts (when the latter are available). In this section, we compare the time-series and cross-sectional moments of our implied dividend forwards against those reported in Bansal et al. (2017) and van Binsbergen and Koijen (2015) for the traded S&P 500 forwards. Given that the sample of traded dividend forwards used in these papers starts in 2004, we estimate our model in the full sample but focus in this section on the corresponding moments from the 2004–2020 sample, so that the sample moments are directly comparable.

Figure 3 plots the time series of forward equity yields implied by our model (defined in Equation (33)) against the most recent data from Bansal et al. (2017). The first four panels report the equity yields and standard errors for maturities 1, 2, 5, and 7 years, respectively. The figure shows that our model does an overall very good job in matching the traded forward equity yields. Individual model-implied yields stay close to the dividend strip yields we observe in the data, and most of the time they are within standard error bounds. That said, they are not an exact fit, just like we see in the bond literature.

Figure 4 summarizes the dynamics of the model-implied yields in this sample, and resembles closely Figure 1 in Bansal et al. (2017). As clear from this figure, the shape of the term structure varies over time. It is sometimes relatively flat (for example, between 2012 and 2019), sometimes upward sloping (as at the beginning of 2010), and sometimes steeply downward sloping (as during the financial crisis) – just like the traded forwards.

To be clear, this evidence should not be interpreted to imply that our synthetic dividend yields are precise estimates of true yields point by point, as the confidence bands reflecting parameter estimation uncertainty are fairly wide. Rather, by extending the effective length of the time-series of dividend strips, our model delivers more statistical power that can be used to more precisely estimate interesting moments (e.g., the average slope of the term structure, and so on), compared to using the shorter time-series of observable dividend strips.

Table 4 reports statistical patterns of empirical yields based on BMSY data ("Actual"), the model estimated in the same sample ("Fitted"), and the fitting error ("Actual–Fitted") for the level (top

[13] We make our recovered synthetic dividend strip yields data available at `https://www.serhiykozak.com/data`, together with the standard errors on the estimated prices.

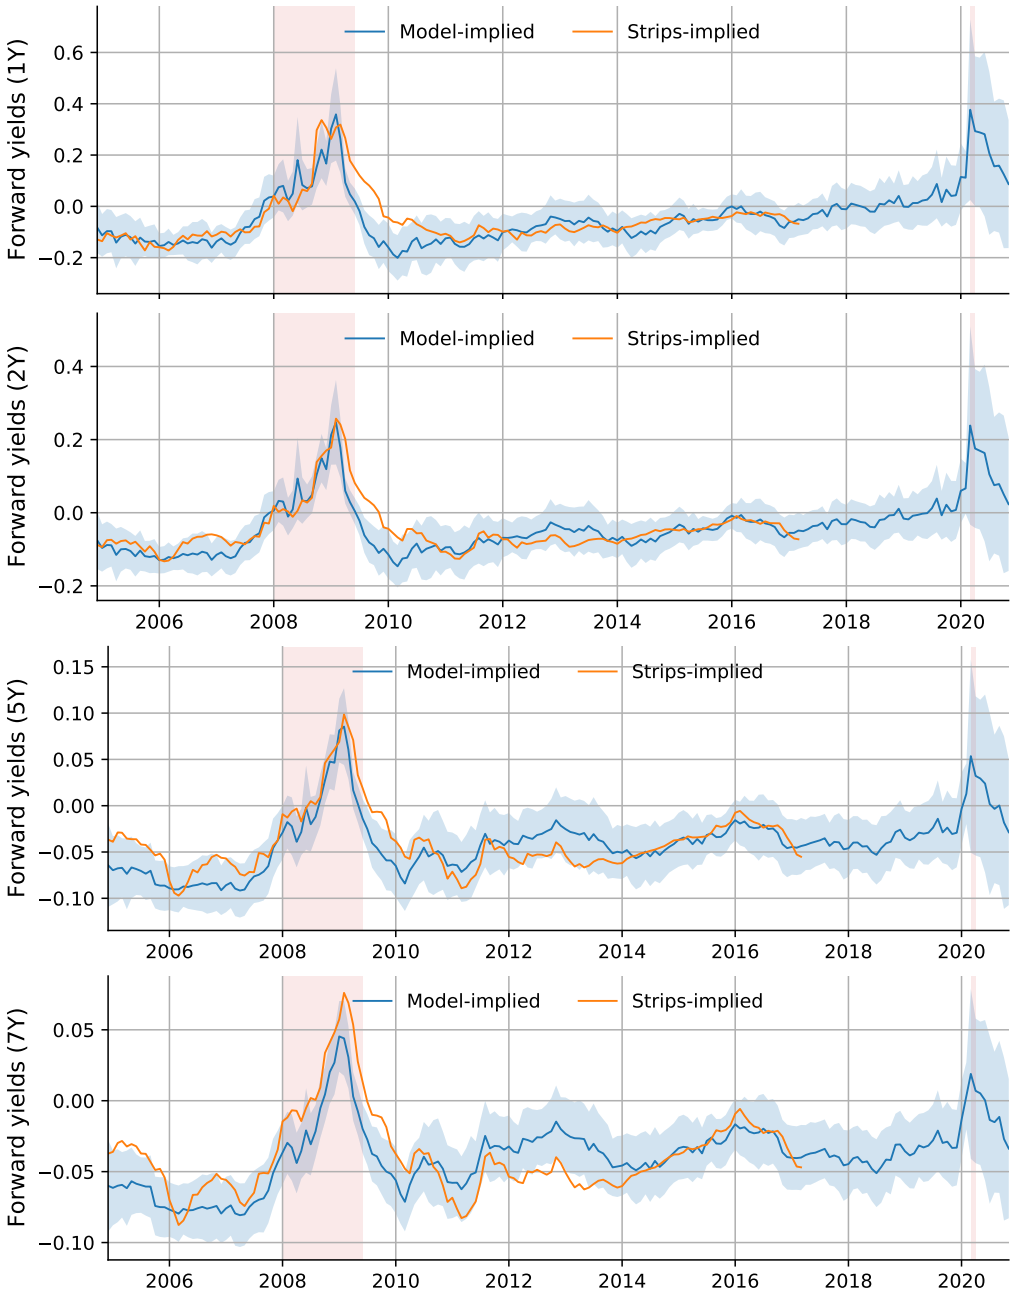

Figure 3: Model-implied forward equity yields vs. forward equity yield data. We compare our model-implied forward yields for maturities 1, 2, 5, and 7 years to their empirical counterparts in Bansal et al. (2017). Shaded areas depict two-standard-deviation bands around point estimates. Model parameters are estimated in full sample.

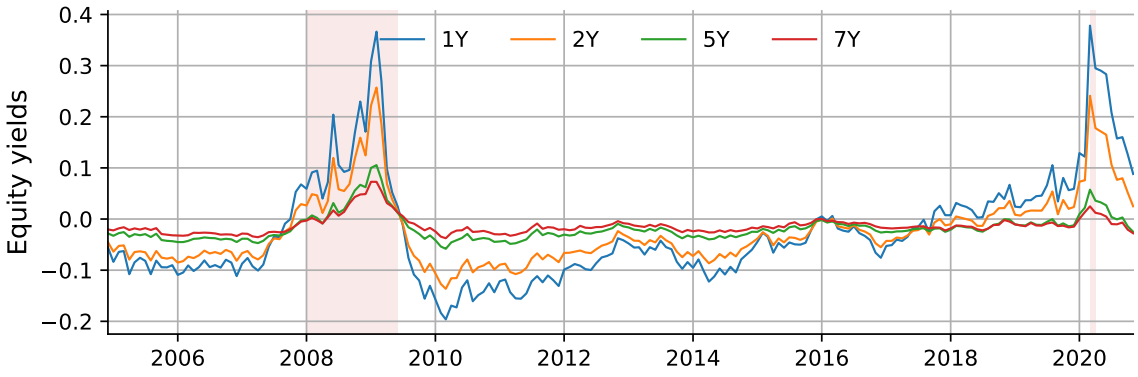

Figure 4: Dynamics of model-implied yields in the Bansal et al. (2017) **sample**. Equity yields are constructed using the trailing 12-month dividend. Model parameters are estimated in full sample, February 1973 to December 2020.

Table 4: Accuracy in fitted vs. actual dividend yields.

We present means and standard deviations (in %), as well as annual autocorrelation for observed BMSY forward yields ("Actual"), model-implied yields ("Fitted"), and the fitting errors ("Actual–Fitted") in levels (top panel) and slopes (bottom panel).

	1Y	2Y	5Y	7Y	1Y	2Y	5Y	7Y	1Y	2Y	5Y	7Y
						Level						
	Actual				Fitted				Actual-Fitted			
Mean	-5.09	-4.55	-3.88	-3.77	-6.80	-5.69	-4.45	-4.14	1.71	1.14	0.58	0.37
S.D.	9.85	6.71	3.36	2.94	9.19	6.44	3.16	2.39	5.58	3.54	1.96	1.89
AC	0.30	0.26	0.24	0.25	0.13	0.10	0.15	0.20	0.13	0.06	0.42	0.60
						Slope						
	Actual				Fitted				Actual-Fitted			
Mean	-	0.54	1.22	1.32	-	1.11	2.35	2.66	-	-0.57	-1.13	-1.33
S.D.	-	3.56	6.86	7.26	-	2.96	6.54	7.38	-	2.96	4.91	5.13
AC	-	0.31	0.30	0.29	-	0.25	0.21	0.20	-	0.31	0.18	0.14

panel) and slope (bottom panel) of the term structure. The table shows that the model fits yields well on average (first row). That said, the fit is not always precise in every time period, which can be seen formally in this table as the variability of pricing errors (second row). These findings carry over to slopes of dividend yields (the bottom panel).

Given these results, it should not be surprising that we also match well the unconditional moments. For example, we find that, consistent with van Binsbergen and Koijen (2015), the average term structure of forward premia is close to flat for the US after 2004, with a slight downward slope at the short end and a slight upward slope at the long end. The solid blue line in Figure 5 reports the estimated term structure of forward risk premia for this (post 2004) sample. The figure also shows the average term structure estimated using the full sample, which begins in 1973, along with two-standard-deviations bounds for both lines. We use model-implied GMM standard errors—which

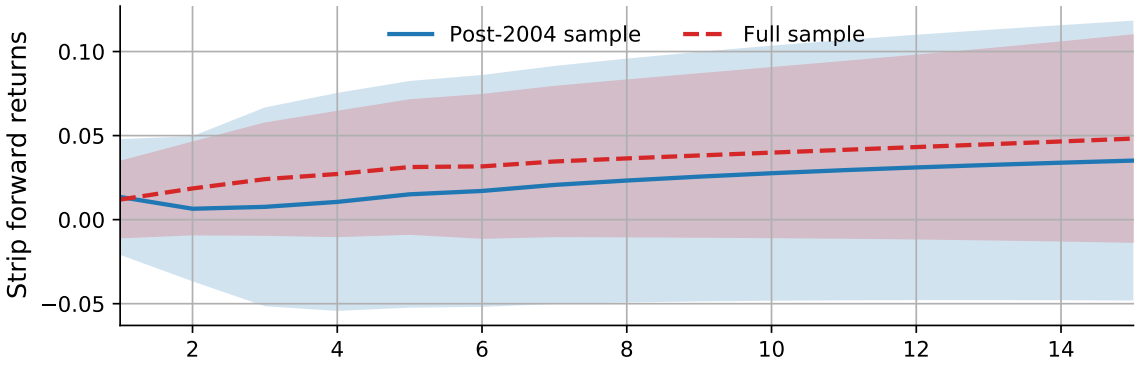

Figure 5: Estimated term structure of forward risk premia. Mean realized returns on dividend strips forward contracts in the 2005-2020 sample (solid blue) and full sample (dashed red). Shaded areas for each line depict two-standard-deviation error bands around point estimates. Model-implied GMM standard errors using a spectral density matrix with 12 lags for full-sample errors (narrow red bands; model and sampling uncertainty). In the 2005-2020 sample we show empirical HAC robust standard errors associated with the estimate of the mean of realized strip returns (sampling uncertainty; wide blue bands).

incorporate all model uncertainty—for the full-sample estimates, and empirical HAC robust standard errors associated with the estimate of the mean of realized strip returns—which only reflect sampling uncertainty—for the BMSY sample estimates. The plot shows that our standard errors are substantially narrower even though they incorporate all model estimation uncertainty – due to our ability to estimate means in the longer sample enabled by our use of stock returns data rather than dividend futures data.

We explore the two sources of uncertainty further in Appendix Figure A.18. Panels (a) and (b) focus on full and BMSY samples, respectively. In each figure, the blue shaded area shows GMM two-standard-error bounds which incorporate model parameter and sampling uncertainty, while the red area reflects only sampling uncertainty. The figure shows that effectively the entirety of the variation here comes from the sampling uncertainty, which would be there *even if prices were perfectly observable*. That is, while model uncertainty due to the fact that prices have to be estimated might be non-trivial, it is dominated and washed out by the sampling uncertainty of estimating mean returns or average yields. Going back to Figure 5, this explains why the standard errors over the 2004+ sample (in blue) are close to the ones implied by the estimates in Table 4 in Bansal et al. (2017), whereas the ones we obtain over the full sample (in red) are significantly tighter.

Another dimension in which our data matches the results in Van Binsbergen et al. (2013) is the decomposition of equity yields into cash flow and discount rate components. Van Binsbergen et al. (2013) show that expected cash-flow variation is, surprisingly, a major driver of the movement in the equity yields; Figure 7 in their paper shows that the sharp increase in the equity yield for the S&P 500 during the financial crisis can be almost entirely attributed to a sharp decline in expected dividend growth. We perform the same decomposition in our model: Figure 6 shows that the results looks extremely similar to those in Van Binsbergen et al. (2013).

The Appendix shows the similarity between our implied yields and the ones from Bansal et al. (2017) and van Binsbergen and Koijen (2015) along additional dimensions: the average market beta

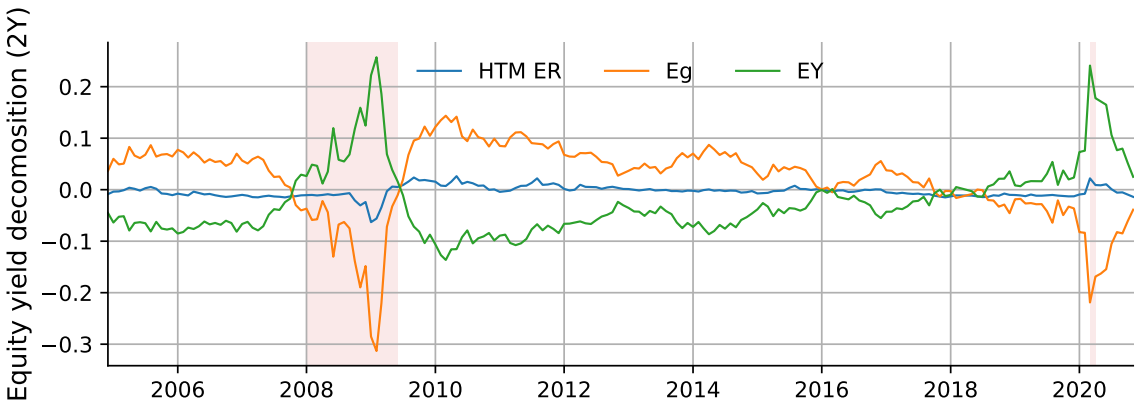

Figure 6: Decomposition of the 2-year equity yield in the Bansal et al. (2017) **sample**. The figure decomposes the 2-year equity yield on the market into log HTM risk premium and log dividend growth net of the risk-free rate.

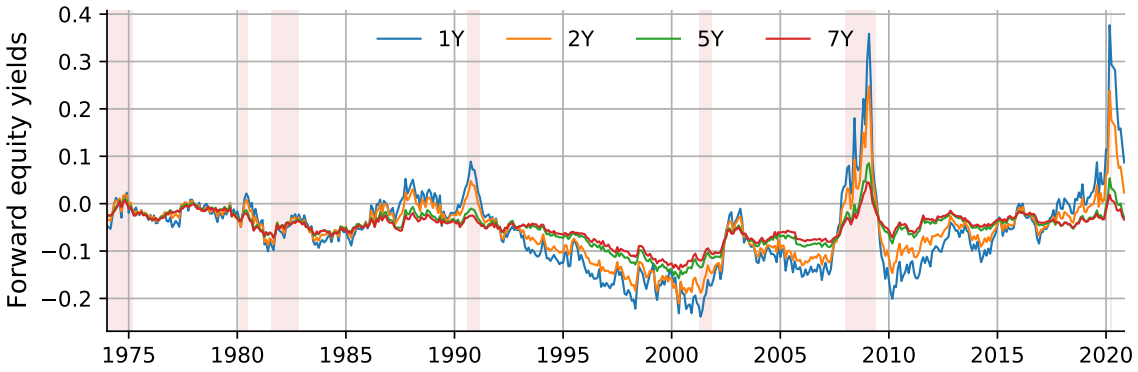

Figure 7: Dynamics of model-implied forward equity yields for the aggregate market for different maturities. The figure plots dynamics of model-implied forward equity yields of maturities 1, 2, 5, and 7 years. Equity yields are constructed using the trailing 12-month dividend. The sample is from February 1973 to December 2020.

of the strips of different maturity in Figure A.14 is similar to the one reported in van Binsbergen and Koijen (2015) (Figure 4); it also plots the one-year returns of the 1- and 2-year dividend forwards in the data and in our model in Figure A.15, showing that they match well.

To conclude: when we compute theoretical dividend forwards and strips from our model (estimated using only equity portfolios) and compare them with the prices of actually traded dividend strips, we see that they match well along several dimensions. This provides an external validation of the ability of our model to capture the dynamics of risk and cash flows and investor preferences for risks along the term structure.

Next, we use the model to explore the behavior of the equity term structure over the longer sample (beginning in 1975), and we then turn to the cross-section of term structures for different risks.

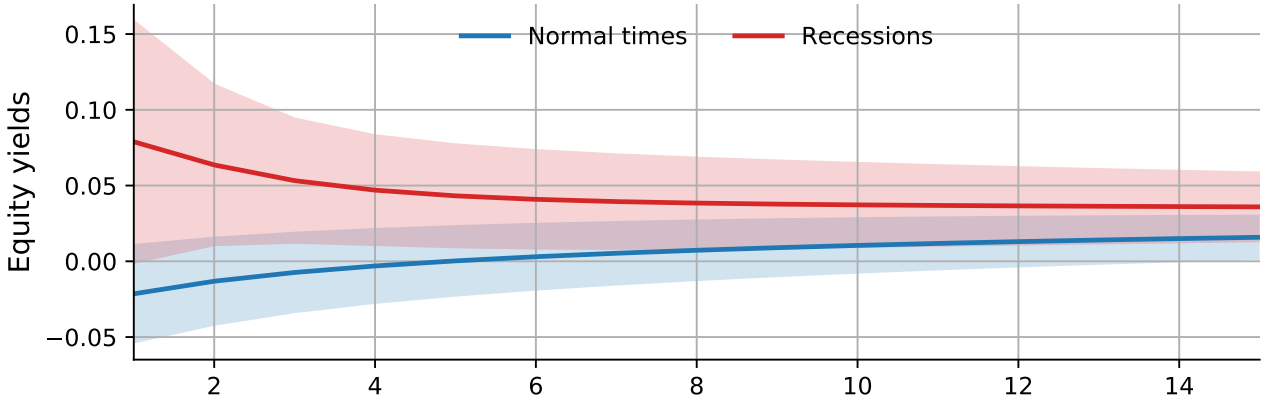

Maturity:	2	3	4	5	7	10	15
Level diff.	0.08	0.06	0.05	0.04	0.03	0.03	0.02
	(2.73)	(2.79)	(2.66)	(2.47)	(2.17)	(1.92)	(1.71)
Slope diff.	0.00	-0.02	-0.03	-0.03	-0.04	-0.05	-0.06
	-	(-1.71)	(-1.80)	(-1.87)	(-1.97)	(-2.08)	(-2.22)

Figure 8: Term structures of equity yields conditional on NBER recessions. The figure shows the term structures conditional on being (red) or not (blue) in an NBER recession by maturuty. Shaded areas depict two-standard-deviation error bands around point estimates. The table reports differences of equity yields ("Level diff.") and differences of equity yield slopes (relative to a 2-year equity yield; denotes as "Slope diff.") in recession minus expansion, and their associated t-statistics based on GMM standard errors with a spectral density matrix with 12 lags.

4.3 The time series of the equity term structure since the 1970s

Figure 7 shows the forward equity yields for the aggregate market for different maturities, as estimated from our model. The results confirm many of the patterns reported above for the post-2004 data: the term structure is generally close to flat, with periods of positive slope during booms, and periods of inversion during busts. Interestingly, the term structure appeared significantly more stable up to the early 1990s, remaining essentially flat for almost two decades. The brief recession of the 1990s induced an inversion of the curve, followed by a period in which the slope changed sign several times.

These changes in the slope of the term structure are strongly correlated with the macroeconomic cycle. To see this more clearly, Figure 8 shows the term structures conditional on being or not in an NBER recession. Outside recessions (blue line) the term structure is mildly upward sloping on average. In recessions, instead, the term structure is downward sloping.

The table under the plot computes the differences between these two lines and the t-statistics associated with them ("Level diff."). It also shows the difference in the slopes (difference between each maturity and maturity 2, "Slope diff"). Equity yields are statistically significantly higher in recessions than expansions for dividend claims of up to eight years. Equity yield slopes are significantly different at the 5% level at a maturities above 7 years (i.e., 7-2 slopes) and are significantly different at the 10% level for all maturities. Standard errors in this plot reflect both parameter and sampling uncertainty.

One of the advantages of our method is the ability to study the behavior of the term structure

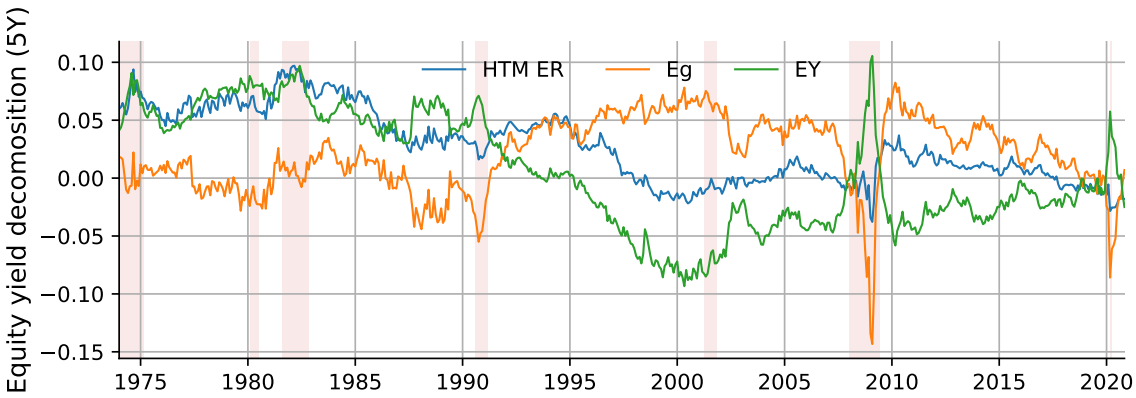

Figure 9: Decomposition of the 5-year equity yield. The figure decomposes the 5-year equity yield on the market into log HTM risk premium and expected log dividend growth net of the risk-free rate in full sample, February 1973 to December 2020.

over a much longer sample, that includes several recessions (contrary to the post-2004 sample where the only recession is the financial crisis). We can therefore ensure that the patterns we find are not entirely due to the specialness of the Great Recession. So for example, we see inversion in the yield curve at several other times in history, as confirmed by Figure 8 that averages across *all* recessions.

It is worth recalling that the slope of the term structure of equity yields has a direct interpretation in terms of term structure of discount rates for different horizons: as pointed out in Dew-Becker et al. (2017) and Backus et al. (2018),

$$\frac{1}{n}\,\mathbb{E}\log\left(R_{t,t+n}\right) - \mathbb{E}\log\left(R_{t,t+1}\right) = \mathbb{E}\,e_{t,n} - \mathbb{E}\,e_{t,1}.$$

This result highlights how the main facts about the conditional and unconditional term structure of equity yields from van Binsbergen et al. (2012b) extend to our full sample starting in the 1970s.

Finally, we report in Figure 9 the decomposition of the 5-year equity yield into expected annual dividend growth over 5 years and expected hold-to-maturity (HTM) excess returns. This figure confirms that a large fraction of the variation in equity yields is driven by expected dividend growth as opposed to discount rate variation.

Extending the equity term structure data to the 1970s yields several new insights. First, as mentioned above, the term structure appeared much less volatile before the 1990s. Second, it appeared to have been effectively flat for decades. Third, there had been times before the financial crisis in which markets strongly anticipated negative dividend growth: for example, during the recession of the early 1980s and early 1990s; these movements appear to have been reflected in the prices of equities and (implied) equity strips.

4.4 Cross-section of term structures of different risks

The most interesting advantage of our model is that it can generate term structures for different types of risks, as captured by different portfolios. For example, it can produce a term structure of discount

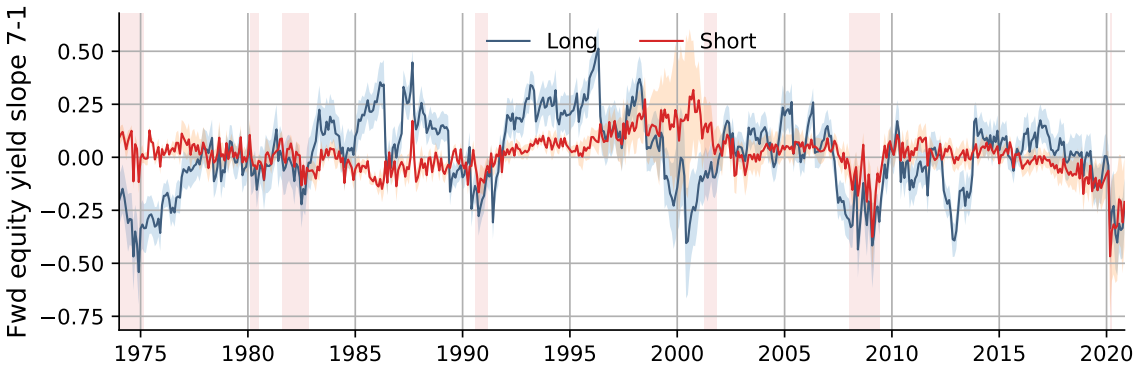

Figure 10: Slope 7-1 of forward equity yields for small and large stocks. The figure shows time-series of estimated slope of forward equity yields (7y - 1y) from our model for diversified portfolios of small (long) and large (short) stocks in full sample, February 1973 to December 2020.

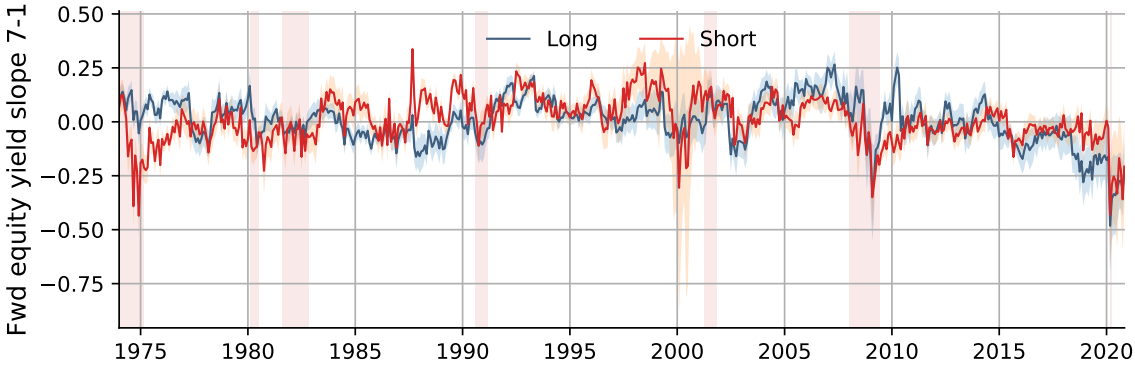

Figure 11: Slope 7-1 of forward equity yields for value and growth stocks. The figure shows time-series of estimated slope of forward equity yields (7y - 1y) from our model for diversified portfolios of value (long) and growth (short) stocks in full sample, February 1973 to December 2020.

rates for value firms and one for growth firms, one for small firms and one for large firms, and so on. In turn, these term structures can be used to test the implications of structural models that have cross-sectional implications.

As an example, different structural models have been proposed to explain the value premium. But these models have mainly confronted the *average risk premium* of value vs. growth portfolios. These models will also have implications for the *term structure* of discount rates on value and growth stocks; this will be an especially important moment for models in which the dynamics of shocks play a role in determining risk premia.

Figure 10 and Figure 11 show the time series of the estimated slope of forward equity yields (7y - 1y) from our model, for small stocks (long) and large stocks (short), and for value stocks (long) and growth stocks (short), respectively. Many interesting patterns emerge.

Large firms' and small firms' term structures moved in similar ways in several periods (e.g., during the financial crisis). But their equity yields went in opposite directions during the tech boom and bust. Whereas the equity term structure inverted for small firms during the tech boom, it didn't do so for large firms. During the financial crisis and the Covid episode, instead, both curves inverted. On the contrary, no such divergence in the behavior of the term structure can be seen for value and growth stocks in that period – the largest difference in that case occurred in the recovery from the financial crisis: after 2008, the slope of the term structure increased significantly for value but not for growth stocks.

Figure 12 depicts the term structure of risk premia for a number of long-short portfolios. Panel (a) shows the difference of risk premia between small and large stocks, value and growth stocks, and high vs. low momentum stocks, for different horizons. Shaded areas depict one-standard-deviation error bands around point estimates (for readability). Note that these long-short portfolios have positive risk premia in this sample. Yet, the term structures of these risk premia are quite heterogeneous. For example, the SMB portfolio appears to exhibit a downward-sloping term structure, while the HML portfolio has an upward-sloping term structure. Theoretical models aiming to explain the value and size spreads can make use of these estimates to help refine and calibrate the economic mechanisms (of course, taking into account the substantial estimation uncertainty associated with these unconditional moments).

It is important to note that our term structures correspond to term structures of actively managed portfolios that are rebalanced at an annual or monthly frequency. This choice of the basis assets is motivated by two observations. First, it directly parallels the construction of traded claims on dividend strips of aggregate indices, such as the S&P 500 in the US. Indeed, dividend strips on S&P 500 are claims to the dividends that future constituents of the S&P 500 index will pay in the future, rather than dividends on today's constituents of the index. Second, Keloharju et al. (2019) show that differences in expected returns across firms decay within five years. Their findings suggest that term structures of any portfolio sorted on today's characteristics should approximately converge to the term structure of the aggregate market at that horizon. We avoid such convergence in our results by focusing on term structures of actively managed portfolios. To the extent that firm characteristics capture economically relevant risk exposures (e.g., if book-to-market captures distress risk), it is meaningful

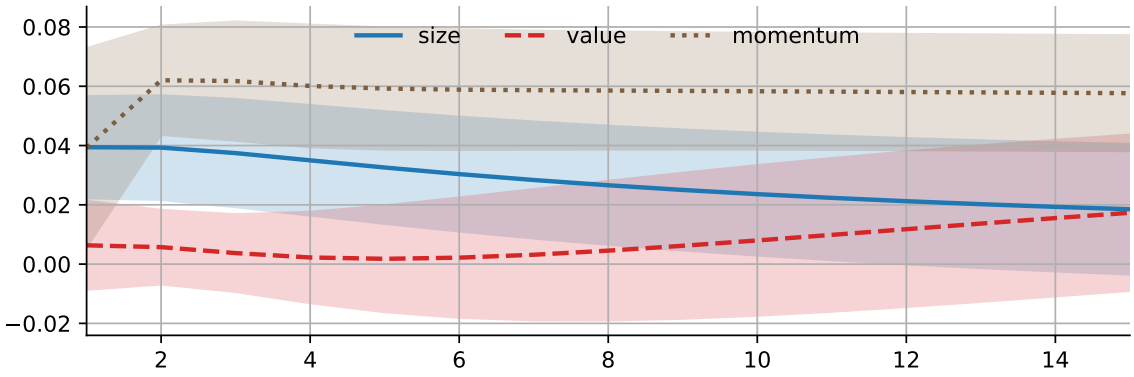

(a) SMB, HML, and MOM risk premia with standard error bounds

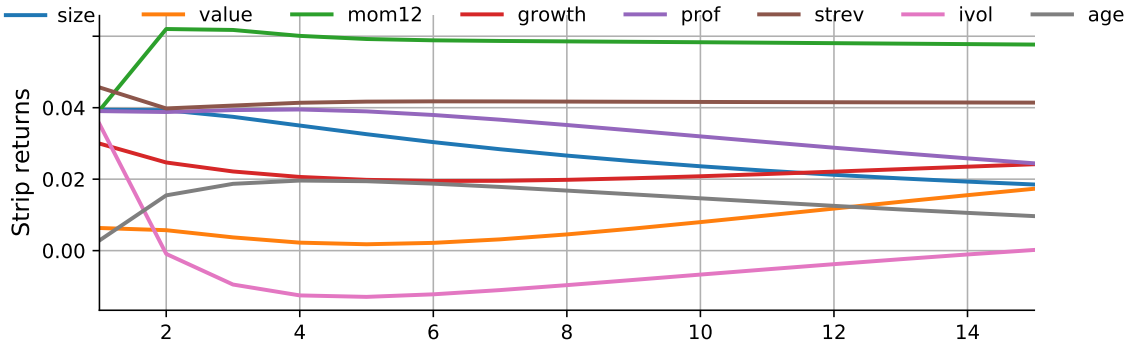

(b) Select long-short portfolios

Figure 12: Term structure of risk premia for long-short portfolios. Panel (a) shows the difference of realized risk premia between small and large stocks, value and growth stocks, and high vs. low momentum stocks. Shaded areas depict one-standard-deviation error bands around point estimates. Panel (b) shows the difference of realized risk premia across select anomaly portfolios: SMB, HML, MOM, INV, PROF, STREV, IVOL, and Age.

to try and understand the pricing of potential risks that affect stocks with those characteristics at various horizons; for example, the pricing of shocks that will affect distressed stocks in the future.[14]

Many other portfolios exhibit interesting term structures. Panel (b) of the figure shows the difference of realized risk premia across a few additional anomaly portfolios: SMB, HML, momentum, investment, profitability, short-term reversals, idiosyncratic volatility, and age. Some of them have strikingly different shapes. For example, the term structure of risk premia for the portfolio sorted on idiosyncratic volatility is sharply downward sloping up to around 5 years, then it becomes mildly upward sloping. The term structure of risk premia for the age portfolio is hump shaped.

We conclude that our model generates very rich predictions about the differential behavior of term structures across portfolios, both in the time series and on average. These results present new moments that structural asset pricing models can try to match (in addition to the term structure of

[14]That said, if one wanted to construct term structures of dividends of a fixed portfolio of stocks, they could be easily approximated by overlaying the term structures of managed portfolios we report (at short horizons) and that of the aggregate market (at longer horizons).

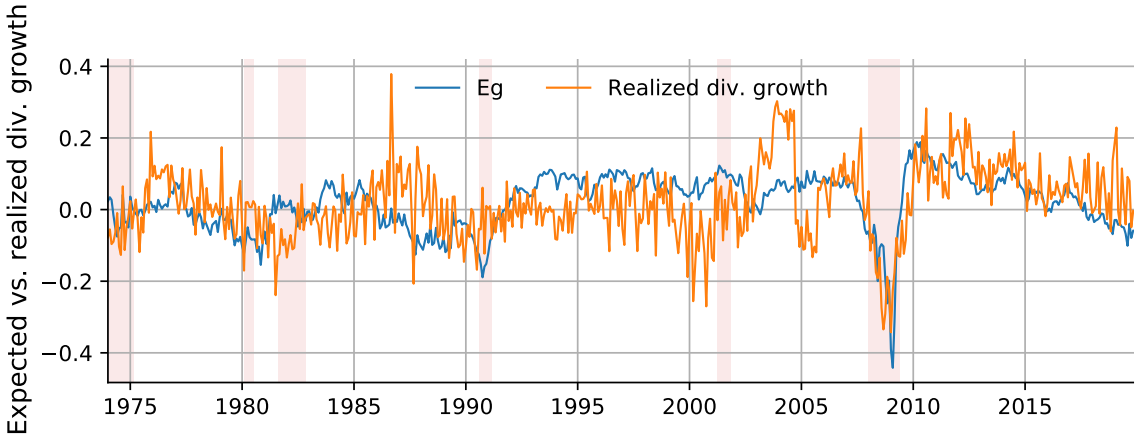

(a) Aggregate market index

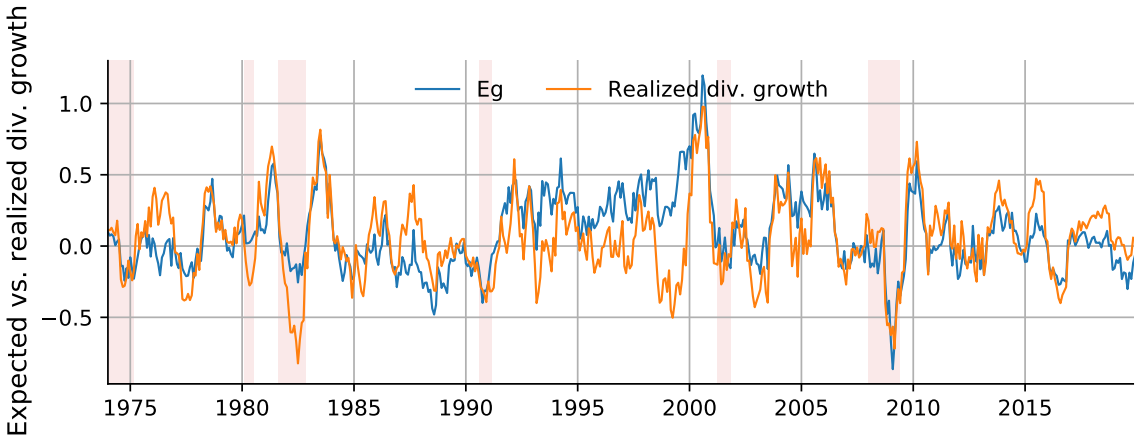

(b) Momentum portfolio

Figure 13: Expected and realized dividend growth for the market index and momentum stocks. Panel (a) shows time-series of the model implied expected (blue) and realized (orange) dividend growth for the aggregate market index, net of the risk-free rate. Panel (b) shows the same results for the momentum stock portfolio.

the aggregate S&P500 dividend that has already been studied using traded dividend forwards).

4.5 Implied dividend growth

In our model, deflated dividend growth (i.e., net of the risk-free rate) is pinned down exactly by Equation (15). Since we model both returns and dividend-price ratios as inputs, dividends are just a deterministic function of these observable variables and thus perfectly match the data, month by month. Note that we don't use log-linearizations in this expression; therefore, dividends are given by an exact (nonlinear) function of returns and equity dividend-price ratios. The dynamics of dividends are thus entirely pinned down by the dynamics of these variables.

We now come back to the discussion of dividend growth predictability in our model. We observe a relatively strong model-implied dividend predictability across the board, for many stock portfolios,

as well as the aggregate market (Panel (a) of Figure 13). Dividend growth predictability is strongest, however, for medium-to-high-frequency strategies, such as momentum, in which case dividend growth can be predicted with R^2 exceeding 80%. Panel (b) of Figure 13 shows time-series of the model implied expected (blue) and realized (orange) deflated dividend growth for the momentum stock portfolio. Similar to Figure 6, we find that most variation in equity yields on high momentum stock is driven by variation in expected dividend growth. This, in turn, implies that dividend growth on high momentum stocks is highly predictable by their equity yields, consistent with the findings of Van Binsbergen et al. (2013).

Dividends can thus be used to provide additional validation of the equity yields constructed using our model. To motivate this idea a little further, first note that log gross dividend growth at an n-year horizon, $\Delta d_{t,t+n}$ is given by

$$\Delta d_{t,t+n} = r_{t,t+n} - ne_{t,n}, \tag{35}$$

where $r_{t,t+n} = \log\left(\frac{D_{t+n}}{P_t^{(n)}}\right)$ is the log of hold-to-maturity return on an n-year dividend strip from t to $t+n$ and $e_{t,n}$ is the equity yield on an n-year dividend strip at time t as defined in Equation (32) (see Section A.3 for the derivation). This is an exact formula that does not require any approximations.

Second, notice that this equation holds in expectations, that is,

$$\mathbb{E}_t\left[\Delta d_{t,t+n}\right] = \mathbb{E}_t\left[r_{t,t+n}\right] - ne_{t,n}. \tag{36}$$

Third, recall from Figure 6 that most of the variation in short-term equity yields comes from expected dividend growth rather than discount rates.

Equity yields, therefore, are natural predictors of dividend growth rates at all horizons, and especially so at short horizons. We use this insight to investigate dividend growth time-series predictability using model-implied equity yields in the full sample. As a benchmark, we use dividend growth predictability based on the stock's own D/P ratio as is standard in the literature.[15]

In Table 5 we report means of R^2 computed in the cross-section of 102 portfolios from the regressions of n-year dividend growth on several variables (all univariate regressions). The top panel shows the results in the full sample, the bottom panel in the post-2004 sample. In the first row, we use as predictor the equity's own dividend-to-price ratio (D/P), and find limited predictability of dividends by this variable, consistent with findings in the literature.

In the second row, motivated by the analysis above, we use the model-implied dividend yield on an n-year dividend strip as a predictor. We observe substantially higher predictability – a statistically significant difference in predictability up to four years in maturity, as can be seen from the p-values of a one-sided t-test of the equality of the two means of R^2 distributions for D/P and EY (model) regressions (the last rown in the panel). In the third row, we include the model implied expected dividend growth net of the risk-free rate at each horizon, $\mathbb{E}[\Delta d - r_f]$, as a predictor. This is the same object we depicted in Figure 13 in blue. We see that this variable predicts realized dividends even better than the dividend yield.[16]

[15] We have also considered a specification with an additional predictor—lagged dividend growth—with similar results.

[16] We use the expected deflated dividend growth as opposed to the expected dividend growth because the former is

Table 5: Dividend growth predictability.

The top panel reports average R^2 in the cross-section of 102 portfolios from the regressions of an n-year dividend growth on: (i) the equity's own dividend-to-price ratio (D/P), and (ii) the model-implied dividend yield on an n-year dividend strip, (iii) BMSY dividend yield on an n-year dividend strip (bottom panel only), and (iv) model implied expected dividend growth net of the risk-free rate at each horizon, $\mathbb{E}[\Delta d - r_f]$. The last row in each panel shows p-values of a one-sided t-test of the equality of the two means of R^2 distributions for D/P and EY (model) regressions. The top panel show results for the full sample; the bottom panel focuses on BMSY sample.

Horizon:	1-year	2-year	3-year	4-year	5-year	6-year	7-year
	Full sample						
D/P	3.7	3.5	3.7	4.4	5.1	6.0	7.0
EY (model)	17.8	12.9	9.0	6.9	6.1	6.5	7.3
$\mathbb{E}[\Delta d - r_f]$	19.9	16.6	14.6	13.1	13.4	13.4	16.0
$p(R^2_{\text{D/P}} = R^2_{\text{EY}})$	0.00	0.00	0.00	0.00	0.10	0.28	0.35
	BMSY sample						
D/P	36.0	25.6	16.6	9.7	10.3	11.5	12.9
EY (model)	41.5	36.4	25.1	13.3	12.9	13.5	14.8
EY (BMSY)	20.7	18.5	-	-	7.9	-	9.1
$\mathbb{E}[\Delta d - r_f]$	42.6	44.9	36.6	21.5	21.5	23.5	25.4
$p(R^2_{\text{D/P}} = R^2_{\text{EY}})$	0.01	0.00	0.00	0.01	0.06	0.12	0.15

Lastly, we compare the predictability between our model-implied equity yields and their empirical counterparts based on dividend strip data from BMSY. To do that we limit our sample to theirs and report all the results in the second panel. We now include an additional row "EY (BMSY)" which includes R-squared averages of regressions that use BMSY equity yields as a predictor. Perhaps somewhat surprisingly, in spite of the visual similarity of BMSY and our equity yields depicted in Figure 3 in the paper, we find that our model-implied EY perform significantly better than the observed BMSY yields in predicting cumulated realized dividend growth in the data. Similarly to the full sample case, our model-implied expected dividend growth (row 4 of the second panel) predicts realized dividend growth even more robustly than model-implied equity yields, even though we do not directly target dividend predictability in our estimation.

We conclude that dividends up to four years are significantly more predictable using equity yields extracted from the model than when using stock's D/P. This fact serves as an another indirect validation of the model in that it is able to extract useful information embedded in equity yields which helps predict dividends on a wide cross-section of test asset portfolios. Figure A.21 in the appendix visualizes distributions of R^2 for both predictors across portfolios.

4.6 Applications

We briefly discuss here three potential applications of our findings.

an output of our model; since we don't model the process of the risk free rate, our model does not provide estimates for the expected non-deflated dividend growth at horizons beyond 1.

4.6.1 New moments for model evaluation

As mentioned in the introduction, we first of all view our methodology as a way to extract novel moments of the term structure of discount rates (its conditional and unconditional slope) for a variety of portfolios. Rather than working directly with the rich dynamics of the economy and the large cross-section of available equity portfolios, researchers can use the generated term structures directly, to calibrate and evaluate models and compute valuations of investments. We also produce standard errors on our estimates, which can be used together with the point estimates to properly take into account the estimation uncertainty.

One natural application of our term structures is in the evaluation and testing of asset pricing models. Many asset pricing models, for example the long-run risk, habit formation, and rare disaster models have strong predictions about the moments we estimate in the data. For example, van Binsbergen et al. (2012b) and van Binsbergen and Koijen (2015) show that the slope of the aggregate dividend term structure observed in the post-2004 sample is on average too low to be explained by the long-run risk and habit formation models – but their conclusions relied on a small sample which includes the financial crisis.

Using our method, we can evaluate these models against the unconditional slope of the forward yields curve estimated since 1975. Having 45 years of data, and being able to look at maturities beyond 7 years, gives our approach more power for these tests, and brings data from new economic cycles. In Appendix C, we illustrate this exercise by simulating different models and evaluating them against our estimated term structure moments: specifically, we look at the slope of the term structure of equity yields and forward equity yields, the risk premia associated with strips and forwards of different maturity, and the cyclicality of the slope of the term structure of equity yields. We study five models: the long-run risk models of Bansal and Yaron (2004) (BY) and Bansal et al. (2010) (BKY), the habit-formation model of Campbell and Cochrane (1999) (habit), the model of Lettau and Wachter (2007) (LW), and the rare-disaster model of Gabaix (2012) (disaster), and show that each model can match some, but not all, of the rich set of new moments we bring to the table. This exercise represents an illustration of the type of additional model evaluations that can be produced using our estimates.

In addition to extending the results on aggregate dividend strips, our paper can also be used to test directly the implications of models about the term structure of discount rates of specific portfolios. For example, Hansen et al. (2008) discuss the implied term structure of value and growth stocks in their model. Production-based models (for example, Belo (2010); Kogan and Papanikolaou (2013, 2014)) specify the dynamics of the stochastic discount factor and the risks of different types of firms; they therefore have direct implications about the term structure of discount rates for value firms, growth firms, high- and low-investment firms, and so on, for which we provide estimates. We leave the evaluation of the cross-sectional term structure implications of these models to future research.

4.6.2 Valuation

Second, our model can be used to value investments whose cash flows pay off at various horizons. Once the riskiness of any investment with a specified maturity is determined (that is, the relation between

the investment returns and our factors F_t and their innovations), either through economic theory or empirically, our model provides the appropriate term structure of discount rates for that investment. For example, it can be used to value private equity investments as in Gupta and Van Nieuwerburgh (2019), or to study long-term discounting for climate change mitigation investments (as in Giglio et al. (2015), who used the term structure of long-run discount rates on housing only).

4.6.3 Hedging portfolio construction

Our work identifies four priced factors and associated dynamics and pricing. The model can therefore be used for investors that would like to get exposure to specific type of risks. The main advantage relative to standard models that are specified only in terms of shocks is that the type of risks can be more easily identified. For example, the model can be used to build a portfolio that isolates cash flow risk or discount rate risk at any specified horizon.

In addition, as explored in Appendix D, the model can be used to explore the links between characteristics sorts (e.g., duration) and risk exposures. As an example, in that section we show that sorting portfolios by the level of the term structure of discount rates yields a portfolio that is principally exposed to the first PC (the second factor after the market), which commands the Sharpe ratio of 0.71 in our model.

4.7 Robustness

4.7.1 Counterfactual analysis: alternative models of the SDF

In this section we investigate whether simpler models (where the factors are not based on the principal components of a large set of anomalies) could also produce estimates for the equity strip prices that match the data from the traded contracts. We consider two benchmarks.

Our first benchmark is the conditional CAPM model. That is, we assume there is a single priced source of risk—the aggregate market—and its risk price time-variation is driven by the aggregate dividend-price ratio. We replicate the analysis in the paper for this choice of the state vector, which now contains only two variables.

Our second benchmark is the five-factor model of Fama and French (2016) supplemented by the momentum factor from Carhart (1997). Like our main specification, we consider the setup with four returns (the market and three PCs of the five cross-sectional factors), that is, our state vector contains eight variables.[17] We replicate the analysis in the paper for this choice of a state vector.

We report the results in Appendix Figure A.19. Panel (a) shows results for the CAPM benchmark. Panel (b) shows results for the Fama and French 5-factor model, supplemented with the momentum factor. Figure A.19 plots the time-series of yields in the Bansal et al. (2017) sample. It is evident from the plots that the dynamics of implied dividend strip yields for both benchmarks are rudimentary: all yields move almost synchronously with no yield curve inversions occurring in this sample for the

[17]This specification performs better than a specification based on the six unrotated factor returns. We, therefore, report the former specification.

CAPM model. The yields implied by the Fama and French 5-factor model look very different from dividend strips yields observed in the data.

We conclude that our rich specification is key for capturing the dynamics of equity strips in the data. This is because our specification is successful in explaining both the cross-section and the time-series of US equity returns. On the contrary, SDFs implied by stylized models, such as the CAPM or the FF 5-factor model, do not feature rich enough dynamics to match the empirically observed patterns in equity yields.

4.7.2 Additional out-of-sample analysis

As discussed in the previous sections, the main way the model is evaluated is in its ability to match the time series of the traded dividend strip prices. In this sense, our main analysis is already out of sample, because traded dividend strip prices are not used in the estimation. In this section we present two additional out of sample analyses.

First, we split our sample into two parts: prior to 2004, and 2005–present. We chose the split point in 2004 because this is when Bansal et al. (2017) data become available, which provides a natural out-of-sample test for us. We estimate the model using data in the first part of the sample only. We then fix the parameters and apply them to construct the estimates of equity yields in the 2005–present sample. Importantly, principal components and their associated eigenvectors are also estimated only in the pre-2004 sample.

We report RMSE of model-implied relative to empirical equity yields in panel (a) of Table 6 (second line; split sample). We also depict the dynamics of estimated equity yields in the 2005–present sample in Appendix Figure A.22. Overall, compared to the full-sample RMSE in the first line of panel (a) and to Figure 4, out-of-sample yields are very similar to their in-sample counterparts. Parameter estimates of the state dynamics are also similar (not reported). This evidence gives us confidence that our results are not driven by merely overfitting the data in sample.

As a second test, we also implement a full rolling estimation. The third line of panel (a) shows RMSE of the model which starts with the pre-2005 training sample and updates it annually on a rolling basis. Figure A.23 in the Appendix depicts the model-implied yields based on this rolling estimation procedure. Qualitatively, the patterns in yields are similar to those estimated both in full sample and in pre-2005 sample.

4.7.3 Alternative portfolio sorts

While our main specification focuses on a cross-section of 51 portfolios which were used in prior work (Kozak et al. (2020), Kozak (2020)), we consider here two alternative sets of portfolios to assess the robustness of our results.

The first set uses return predictive signals from Green et al. (2014) (GHZ) to form sorted portfolios, which are widely used in the literature.[18] This set contains 102 anomaly characteristics (therefore 204 between low- and high-tercile portfolios).

[18]Green et al. (2014) provide SAS code to generate their characteristics, which we have used to construct our own portfolio sorts.

Table 6: Robustness.

The table reports RMSE for forward equity yields of maturities of 1, 2, 5, and 7 years and the average RMSE (last column) across several alternative specifications: panel (a) reports RMSE of the main specification (51 Anomalies) in and out of sample; panel (b) explores alternative portfolio sets and reports RMSE in full sample; panel (c) reports OOS RMSE for models with varying number of PC factors; panel (d) adds bond factors; panel (e) adds volatility factors.

	1-year	2-year	5-year	7-year	Average
			(a) Main specification: 51 Anomalies		
Full sample	0.058	0.037	0.020	0.019	0.034
Split sample	0.057	0.037	0.022	0.020	0.034
Rolling estimation	0.058	0.038	0.021	0.020	0.034
			(b) Alternative portfolio sets		
GHZ	0.075	0.046	0.027	0.024	0.043
WRDS financial ratios	0.070	0.042	0.019	0.017	0.037
			(c) Varying number of PCs (OOS)		
MKT only	0.094	0.065	0.033	0.029	0.055
MKT + 1 PC	0.084	0.056	0.027	0.024	0.048
MKT + 2 PCs	0.064	0.045	0.027	0.024	0.040
MKT + 4 PCs	0.088	0.058	0.025	0.022	0.048
MKT + 5 PCs	0.088	0.061	0.033	0.028	0.052
MKT + 6 PCs	0.087	0.068	0.045	0.039	0.060
			(d) Bonds		
MKT+2PCs + 1 bond PC	0.073	0.053	0.032	0.028	0.047
MKT+3PCs + 1 bond PC	0.061	0.043	0.024	0.022	0.037
MKT+2PCs + 2 bond PCs	0.070	0.051	0.043	0.044	0.052
MKT+3PCs + 2 bond PCs	0.083	0.058	0.033	0.028	0.051
			(e) Volatility		
MKT+2PCs + vol (mkt)	0.073	0.053	0.033	0.029	0.047
MKT+3PCs + vol (mkt)	0.065	0.044	0.025	0.022	0.039
MKT+2PCs + vol (PC)	0.067	0.049	0.031	0.028	0.044
MKT+3PCs + vol (PC)	0.092	0.073	0.046	0.037	0.062

The second set uses WRDS financial ratios constructed by Wharton Research Data Services. This set was also used in Kozak et al. (2020). The set contains over 70 different financial ratios, categorized based on economic intuition into the following seven groups: (1) Capitalization: measures the debt component of a firm's total capital structure, e.g.: Capitalization Ratio, Total Debt-to-Invested Capital Ratio; (2) Efficiency: captures the effectiveness of firm's usage of assets and liability, e.g.: Asset Turnover, Inventory Turnover; (3) Financial Soundness/Solvency: captures the firm's ability to meet long-term obligations, e.g.: Total Debt to Equity Ratio, Interest Coverage Ratio; (4) Liquidity: measures a firm's ability to meet its short-term obligations, e.g.: Current Ratio, Quick Ratio; (5) Profitability: measures the ability of a firm to generate profit, e.g.: ROA, Gross Profit Margin; (6) Valuation: estimates the attractiveness of a firm's stock (overpriced or underpriced),

e.g.: P/E ratio, Shiller's CAPE ratio; (7) Others: Miscellaneous ratios, e.g.: R&D-to-Sales, Labor Expenses-to-Sales.

In Panel (b) of Table 6 we report full-sample RMSE for traded dividend strip yields of the model estimated using these alternative datasets. The table shows that using these alternative sets produce qualitatively similar results to the main specification (first line of Panel (a)), both in terms of general fit metrics and implications for dividend yields. Appendix Figure A.16 and Figure A.17 compare yields at each maturity to BMSY yields. Overall, the results indicate that our setup is not particularly sensitive to the choice of the test assets and that alternative choices deliver qualitatively comparable results.

4.7.4 Number of principal component factors

We explore the robustness of our results to picking a different number of PCs in the state space. There is an inherent trade-off between choosing a number that is too low or too high. Figure A.20 shows that in cases when the number of factors is one (market) or two (market + PC1), implied strip dividend yields exhibit extremely simplistic dynamics which are visibly incompatible with the prices of dividend futures we observe in the data. A model that has two cross-sectional PCs (and the market) performs only slightly worse than our primary specification. Increasing the number of PCs to four or five still performs reasonably well, but the out-of-sample (split sample) performance starts deteriorating as can be seen in Panel (c) of Table 6. The reason for this is that the number of parameters we need to estimate for the physical dynamics (VAR) grows as a square of the number of PCs. Predicting returns in the time-series is notoriously difficult, so this increased number of free parameters proves challenging for the estimation.

A model that includes the market and three PCs—our primary specification (second line in panel (a))—strikes a good balance between flexibility (bias) and overfitting (variance).

4.7.5 Additional factors and predictors of risk prices

The goal of this paper is to produce a parsimonious specification that can generate empirically plausible patterns of dividend strip yields. As shown above, we find that our specification that includes only valuation ratios (D/P) as predictors is sufficiently rich to achieve that objective. D/P is a somewhat special predictor: because it is a part of the return, it must be included in the state space to make the model solvable. Conveniently, Haddad et al. (2020) have shown that valuation ratios alone are powerful predictors of SDF risk prices. However, it is interesting to explore the potential of additional predictors to improve the fit of the model. Importantly, adding predictors does not necessarily improve the performance of the model out of sample (e.g. to fit the dividend strip prices), because it adds parameters to be estimated, which can lead to overfit. In this section we explore this question empirically.

Specifically, we explore modifications to our specification which add additional variables: (i) rolling 1-year volatility (of the market or the first PC of factor volatilities), or (ii) risk-free bond (the first two PCs of bond excess returns across maturities). For each of the new variables we supplement the state space by a corresponding factor return and a predictor. For volatility, we use a volatility-

mimicking portfolio return (as a return) and volatility itself (as a predictor).[19] For bond PCs we use a bond return PC (as a return) and its corresponding forward spread (as a predictor). We show these results in panels (d) and (e) of Table 6, respectively. Overall, adding these variables leads to a deterioration of fit, suggesting that benefits of more restrictive parsimonious specifications (without bond and volatility) generally outweigh the costs of the bias due to the exclusion of these variables.

4.7.6 Non-parametric bootstrap

We conduct a non-parametric bootstrap exercise to compute standard errors which we then compare to the GMM standard errors used in our main specification. We draw 48 years with replacement, each including 12 monthly observations, and include the 12 preceding monthly observations to maintain the time-series structure in estimating the dynamics in our state vector. We then use this bootstrapped sample to estimate the main parameters of the model $(c, \rho, b_0, b_1, \beta_2)$ in our GMM and compute the rest of quantities based on these estimates. Next, we use the actual time-series to compute and construct equity yields, returns, and other variables of interest. We repeat this entire process 500 times. Finally, we compute standard deviations of equity yields across all 500 bootstrapped estimates of forward equity yields and produce an analogue of Figure 3 in the paper in which standard errors are now based on this bootstrap procedure.

Figure A.24 in the appendix shows the results. Standard errors are remarkably close to our GMM-based standard errors with Hansen-Hodrick correction. To assess the relative magnitudes of standard errors across the two methods quantitatively, we also report time series averages of standard errors (in %) for forward equity yields of maturities 1-7 years in Table A.4 in the appendix. In the top panel we focus on the post-2004 (BMSY) sample which is used in Figure A.24. As in the Figure, we can see that bootstrap standard errors are very close to the model-implied GMM standard errors. Only at high maturities our standard errors tend to be somewhat smaller than the bootstrapped ones (up to 30% smaller for the 7-year yield). The bottom panel shows the same results for the entire sample. The conclusion remains unchanged.

5 Concluding Remarks

Our model effectively processes a rich information set (the time-series and cross-sectional behavior of 102 portfolios spanning a wide range of equity risks) to produce "stylized facts" – the time series and cross-sectional behavior of implied dividend term structures – that summarize a dimension of the data that is particularly informative about our economic models. Similarly in spirit to the way in which the introduction of vector autoregressions (VAR) by Sims (1980) provided new moments against which to evaluate structural macro models (the impulse-response functions that were generated by the VARs), the objective of this paper is to produce realistic term structure of discount rates for different portfolios that closely resemble the actual dividend claims we observe in the data, and that can be used by asset pricing models as a moment for evaluation and guidance.

[19] Volatility is partially spanned by the variables included in the state vector ($R^2 = 30\%$) and is negatively correlated with the slope of equity yields (-0.3).

Our approach uses only a cross-section of equity portfolios to produce new (implied) term structure data that expands the existing (observed) data along each of those dimension. The term structures we generate cover a large number of cross-sectional portfolios, in addition to the S&P 500: value, size, profitability, momentum, etc., for the total of 102 portfolios. They have a long time series, starting in 1975 and therefore covering several recessions and booms. They have all possible maturities, including the very short and the very long ends of the term structure.

The main building block of our specification is the carefully crafted state-space vector, which includes four excess returns (on the market and three PCs of anomaly portfolios), and four valuation ratios corresponding to these portfolios. This choice is motivated by recent empirical evidence in Kozak et al. (2020) and Haddad et al. (2020), who show that (i) an SDF constructed from dominant PCs of a large cross-section of characteristics-sorted portfolios explains the cross-section of expected returns well, and (ii) that SDF risk prices corresponding to these factors are highly predictable in the time-series by their valuation ratios and that this variation in risk prices is essential to adequately capturing dynamic properties of the SDF. It is this choice of the state space vector that represents the core of our paper: it allows us to produce term structures of discount rates that match well the observed ones – and gives us confidence in extending them over time, maturities, and portfolios.

We derive a variety of novel empirical results. First, we extend the study of the term structure of aggregate dividend claims (on the S&P 500, as in van Binsbergen and Koijen (2015)) over time (back to the 70s) and maturities. In the sample starting in 2004 that was used in van Binsbergen and Koijen (2015), we match the time series of dividend forward prices very closely, and, mechanically, we also match the term structure of discount rates. The term structure of discount rates appears in this sample slightly downward sloping, in contrast with the predictions of many models like the long-run risk model, that instead predict that it should be steeply upward-sloping. Extending the sample to the 1970s allows us to include several additional recessions to our sample; at the same time, the Great Recession carries less overall weight in the sample. It is interesting to see that many of the results of the post-2004 sample carry over to the longer sample. The term structure inverts in almost all of the additional recessions (for example, in the early 80s and 90s). The term structure of forward discount rates is still close to flat (it is mildly upward sloping on average, but it is not significantly so).

Most importantly, the model generates interesting differences in the average term structures *across portfolios*. For example, we show that some portfolios (e.g., size-sorted portfolios or idiosyncratic volatility-sorted portfolios) have downward sloping average term structures of risk premia, whereas others (e.g., book-to-market sorted portfolios) have upward-sloping term structures. These results give us new moments that can be used to evaluate structural asset pricing models that have direct implication about the risk premia of these portfolios (as well as any of the other 102 we include in our analysis).

There are also interesting patterns in the *time series* of slopes of the term structure of different portfolios. For example, the slopes of small and large stocks often move closely together; both term structures were upward sloping during the 1990s, and both were downward sloping during the Great Recession. Yet, only the term structure for small stocks inverted during the late 90s stock cycle, marking an important divergence between the two portfolios that lasted several years. On the contrary,

no such divergence in the shape of the term structure can be seen for value and growth stocks in that period – instead, the largest difference in that case occurred in the recovery from the financial crisis: after 2008, the term structure of value stock expected returns increased significantly, whereas this didn't happen for growth stocks.

Overall, this set of stylized facts provide new conditional moments that asset pricing and structural macro models should seek to match. While in the paper we have explored some of the term structure implications of a small number of models, more work remains to be done to fully explore the implications of our cross-section of equity yield term structures.

Finally, we would like to emphasize that our estimates of equity yields, discount rates, and returns on specific dividend claims are subject to estimation uncertainty. Compared to the traditional approach of using the observed prices of traded dividend strips to learn about the term structure of risky claims, our method has important advantages and disadvantages. On the one hand, it requires estimating the full model to obtain estimated prices for the dividend strips. This introduces "parameter uncertainty". On the other hand, it allows us to dramatically expand the time series available, which reduces the "sampling uncertainty" that is common to both approaches. This tradeoff, and, more generally, the overall uncertainty, need to be taken into account when bringing empirical moment from these estimation methods to the theoretical models. For this reason, we provide full data on standard errors of our synthetic equity yields along with their point estimates on our website.[20]

[20] `https://www.serhiykozak.com/data`.

References

Adrian, T., R. K. Crump, and E. Moench (2015). Regression-based estimation of dynamic asset pricing models. *Journal of Financial Economics 118*(2), 211–244.

Ang, A. and M. Ulrich (2012). Nominal bonds, real bonds, and equity.

Backus, D., N. Boyarchenko, and M. Chernov (2018). Term structures of asset prices and returns. *Journal of Financial Economics 129*(1), 1–23.

Bansal, R., R. F. Dittmar, and C. T. Lundblad (2005). Consumption, dividends, and the cross-section of equity returns. *Journal of Finance 60(4)*, 1639–1672.

Bansal, R., D. Kiku, and A. Yaron (2010). An empirical evaluation of the long-run risks model for asset prices. Working paper.

Bansal, R., S. Miller, and A. Yaron (2017). Is the term structure of equity risk premia upward sloping. *Duke University and University of Pennsylvania Mimeo*.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance 59*(4), 1481–1509.

Beeler, J. and J. Y. Campbell (2012). The long-run risks model and aggregate asset prices: An empirical assessment. *Critical Finance Review 1(1)*, 141–182. Working paper.

Belo, F. (2010). Production-based measures of risk for asset pricing. *Journal of Monetary Economics 57*(2), 146–163.

Brennan, M. J., A. W. Wang, and Y. Xia (2004). Estimation and test of a simple model of intertemporal capital asset pricing. *The Journal of Finance 59*(4), 1743–1776.

Campbell, J. Y. (1991). A variance decomposition for stock returns. *The Economic Journal 101(405)*, 157–179.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), 205–251.

Campbell, J. Y., S. Giglio, and C. Polk (2013). Hard times. *The Review of Asset Pricing Studies 3*(1), 95–132.

Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance 52*(1), 57–82.

Chernov, M., L. A. Lochstoer, and S. R. Lundeby (2018). Conditional dynamics and the multi-horizon risk-return trade-off. Technical report, National Bureau of Economic Research.

Cochrane, J. H. and M. Piazzesi (2005). Bond risk premia. *American Economic Review 95*(1), 138–160.

Cochrane, J. H. and M. Piazzesi (2008). Decomposing the yield curve. *Graduate School of Business, University of Chicago, Working Paper*.

Croce, M. M., M. Lettau, and S. C. Ludvigson (2014). Investor information, long-run risk, and the term structure of equity. *The Review of Financial Studies 28*(3), 706–742.

Da, Z. (2009). Cash flow, consumption risk, and the cross-section of stock returns. *The Journal of Finance*, 923–956.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2017). The price of variance risk. *Journal of Financial Economics 123*(2), 225–250.

Fama, E. F. and K. R. French (2016). Dissecting anomalies with a five-factor model. *Review of Financial Studies 29*(1), 69–103.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics 127(2)*, 645–700.

Gao, C. and I. W. R. Martin (2021). Volatility, valuation ratios, and bubbles: An empirical measure of market sentiment. *The Journal of Finance 76*(6), 3211–3254.

Giglio, S., M. Maggiori, and J. Stroebel (2014). Very long-run discount rates. *The Quarterly Journal of Economics 130*(1), 1–53.

Giglio, S., M. Maggiori, J. Stroebel, and A. Weber (2015). Climate change and long-run discount rates: Evidence from real estate. Technical report, National Bureau of Economic Research.

Giglio, S. and D. Xiu (2021). Asset pricing with omitted factors. *Journal of Political Economy*.

Gomes, L. and R. Ribeiro (2019). Term structure (s) of equity risk premia.

Gonçalves, A. (2019). The short duration premium. *Available at SSRN 3385579*.

Gonçalves, A. S. (2021). Reinvestment risk and the equity term structure. *The Journal of Finance 76*(5), 2153–2197.

Gormsen, N. J. (2018). Time variation of the equity term structure. *Available at SSRN 2989695*.

Gormsen, N. J. and E. Lazarus (2019). Duration-driven returns. *Available at SSRN*.

Green, J., J. Hand, and F. Zhang (2014). The remarkable multidimensionality in the cross-section of expected us stock returns. *Available at SSRN 2262374*.

Gupta, A. and S. Van Nieuwerburgh (2019). Valuing private equity investments strip by strip.

Gupta, A. and S. Van Nieuwerburgh (2021). Valuing private equity investments strip by strip. *The Journal of Finance 76*(6), 3255–3307.

Gurkaynak, R. S., B. Sack, and J. H. Wright (2006). The u.s. treasury yield curve: 1961 to the present. Federal Reserve Board Finance and Economics Discussion Series paper 2006-28.

Haddad, V., S. Kozak, and S. Santosh (2020, 02). Factor Timing. *The Review of Financial Studies 33*(5), 1980–2018.

Hansen, L. and R. Hodrick (1980). Forward exchange rates as optimal predictors of future spot rates: An economic analysis. *Journal of Political Economy 88*, 829–854.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the econometric society*, 1029–1054.

Hansen, L. P., J. C. Heaton, and N. Li (2008). Consumption strikes back? measuring long-run risk. *Journal of Political Economy 116(2)*, 260–302.

Keloharju, M., J. T. Linnainmaa, and P. Nyberg (2019). Long-term discount rates do not vary across firms. Technical report, National Bureau of Economic Research.

Kogan, L. and D. Papanikolaou (2013). Firm characteristics and stock returns: The role of investment-specific shocks. *The Review of Financial Studies 26*(11), 2718–2759.

Kogan, L. and D. Papanikolaou (2014). Growth opportunities, technology shocks, and asset prices. *The journal of finance 69*(2), 675–718.

Koijen, R. S., H. Lustig, and S. Van Nieuwerburgh (2017). The cross-section and time series of stock and bond returns. *Journal of Monetary Economics 88*, 50–69.

Kozak, S. (2020). Kernel trick for the cross section. *Available at SSRN 3307895*.

Kozak, S., S. Nagel, and S. Santosh (2018). Interpreting factor models. *The Journal of Finance 73*(3), 1183–1223.

Kozak, S., S. Nagel, and S. Santosh (2020). Shrinking the cross-section. *Journal of Financial Economics 135*(2), 271–292.

Kozak, S. and S. Santosh (2019). Why do discount rates vary? *Available at SSRN 2037521*.

Kragt, J., F. De Jong, and J. Driessen (2014). The dividend term structure. *Journal of Financial and Quantitative Analysis*, 1–72.

Lettau, M. and J. A. Wachter (2007). Why is long-horizon equity less risky? a duration-based explanation of the value premium. *Journal of Finance 62*(1), 55–92.

Lettau, M. and J. A. Wachter (2011). The term structures of equity and interest rates. *Journal of Financial Economics 101*(1), 90–113.

Martin, I. W. (2013). Consumption-based asset pricing with higher cumulants. *Review of Economic Studies 80*(2), 745–773.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica: journal of the Econometric Society*, 1–48.

van Binsbergen, J., M. Brandt, and R. Koijen (2012a). On the timing and pricing of dividends. *The American Economic Review 102(4)*, 1596–1618.

Van Binsbergen, J., W. Hueskes, R. Koijen, and E. Vrugt (2013). Equity yields. *Journal of Financial Economics 110*(3), 503–519.

van Binsbergen, J. H., M. W. Brandt, and R. Koijen (2012b). On the timing and pricing of dividends. *American Economic Review 102*, 1596–1618.

van Binsbergen, J. H. and R. S. Koijen (2015). The term structure of returns: Facts and theory. Working paper.

Weber, M. (2018). Cash flow duration and the term structure of equity returns. *Journal of Financial Economics 128*(3), 486–503.

Yan, W. (2015). *Essays on the term structures of bonds and equities*. Ph. D. thesis, The London School of Economics and Political Science (LSE).

Appendix

A The Model

A.1 Dividend yields

We start with the Euler equation in (7) and guess the solution of the form in (9). Plugging in the expressions (1)–(3), (8), and (9) into (7), we obtain:

$$0 = -r_{f,t} - \frac{1}{2}\lambda_t'\Sigma_t\lambda_t + (r_{f,t} + \gamma_0 + \gamma_1 F_t + b_0 + b_1 F_{t+1}) + \frac{1}{2}\operatorname{var}_t\left[-\lambda_t' u_{t+1} + (b_1+\gamma_2)u_{t+1} + \epsilon_{r,t+1}\right] \quad \text{(A1)}$$

$$0 = (\gamma_0 + \gamma_1 F_t) + [b_0 + b_1(c + \rho F_t)] - (b_1+\gamma_2)\Sigma\lambda_t + \frac{1}{2}\operatorname{diag}[\Omega], \quad \text{(A2)}$$

where $\operatorname{diag}[\Omega] = \operatorname{diag}[(b_1+\gamma_2)\Sigma(b_1+\gamma_2)' + \Sigma_\epsilon]$. Matching coefficients on the constant and F_t, we have:

$$0 = (\gamma_0 - \gamma_2\Sigma\lambda) + b_0 + b_1(c - \Sigma\lambda) + \frac{1}{2}\operatorname{diag}[\Omega]. \quad \text{(A3)}$$

$$0 = (\gamma_1 - \gamma_2\Sigma\Lambda) + b_1(\rho - \Sigma\Lambda). \quad \text{(A4)}$$

Given the estimated risk price parameters λ and Λ, price dynamics parameters γ_0, γ_1, γ_2, state space dynamics parameters c, ρ, and variances Σ, Σ_ϵ, this gives us two equations to solve for two parameters determining dividend yield parameters b_0 and b_1.

A.2 Expected returns on dividend strips

Excess (level) returns on dividend strips can be computed using

$$R_t^{(n)} = \frac{P_{t+1}^{(n-1)}/P_{t+1}}{P_t^{(n)}/P_t}\frac{P_{t+1}}{P_t} \times \exp(-r_{f,t}) = \frac{P_{t+1}^{(n-1)}/P_{t+1}}{P_t^{(n)}/P_t}\exp\left[(r_{t+1} - r_{f,t}) - y_{t+1}\right]. \quad \text{(A5)}$$

Note that

$$\frac{P_t^{(n)}}{P_t} = \mathbb{E}_t^{\mathbb{Q}}\left[\frac{P_{t+1}^{(n-1)}}{P_{t+1}}\frac{P_{t+1}}{P_t}e^{-r_{f,t}}\right] \quad \text{(A6)}$$

$$= \mathbb{E}_t^{\mathbb{Q}}\left[\left(\exp(a_{n-1,1} + d_{n-1,1}F_{t+1}) - \exp(a_{n-1,2} + d_{n-1,2}F_{t+1})\right)e^{\Delta p_{t+1} - r_{f,t}}\right] \quad \text{(A7)}$$

$$= \exp(a_{n,1} + d_{n,1}F_t) - \exp(a_{n,2} + d_{n,2}F_t), \quad \text{(A8)}$$

where $a_{n,\cdot}$, $d_{n,\cdot}$ are given by iterations Equation (30)-(31).

We can now compute

$$\mathbb{E}_t\left[\frac{P_{t+1}^{(n-1)}}{P_t}R_{f,t}^{-1}\right] = \mathbb{E}_t\left[\frac{P_{t+1}^{(n-1)}}{P_{t+1}}\frac{P_{t+1}}{P_t}e^{-r_{f,t}}\right] \quad \text{(A9)}$$

$$= \mathbb{E}_t\left[\left(\exp(a_{n-1,1} + d_{n-1,1}F_{t+1}) - \exp(a_{n-1,2} + d_{n-1,2}F_{t+1})\right)e^{\Delta p_{t+1} - r_{f,t}}\right] \quad \text{(A10)}$$

$$= \exp\left(\tilde{a}_{n,1} + \tilde{d}_{n,1}F_t\right) - \exp\left(\tilde{a}_{n,2} + \tilde{d}_{n,2}F_t\right), \quad \text{(A11)}$$

which is exactly the same equation as above, except that expectation is taken under physical dynamics. The solution for $\tilde{a}_{n,\cdot}$, $\tilde{d}_{n,\cdot}$, therefore, is given by the recursion Equation (30)-(31), where we simply replace γ_0^*, γ_1^*, c^*, and ρ^* by their physical counterparts.

Log expected excess returns on the strip are

$$\log\left(\mathbb{E}_t\left[R_{t+1}^{(n)}\right]\right) - r_{f,t} = \log \mathbb{E}_t\left[\frac{P_{t+1}^{(n-1)}}{P_t} R_{f,t}^{-1}\right] - \log\left[\frac{P_t^{(n)}}{P_t}\right] \tag{A12}$$

$$= \log\left[\exp\left(\tilde{a}_{n,1} + \tilde{d}_{n,1} F_t\right) - \exp\left(\tilde{a}_{n,2} + \tilde{d}_{n,2} F_t\right)\right] \tag{A13}$$

$$- \log\left[\exp\left(a_{n,1} + d_{n,1} F_t\right) - \exp\left(a_{n,2} + d_{n,2} F_t\right)\right]. \tag{A14}$$

A.3 Equity yield decomposition into HTM returns and expected growth rates

Let $R_{t,t+n} = \frac{D_{t+n}}{P_t^{(n)}}$ be the hold-to-maturity (HTM) return from t to $t+n$. Excess return is then:

$$\frac{R_{t,t+n}}{R_{f,t,t+n}} = \frac{D_{t+n}/P_t}{P_t^{(n)}/P_t} R_{f,t,t+n}^{-1}. \tag{A15}$$

Take expectations:

$$\mathbb{E}_t\left[\frac{R_{t,t+n}}{R_{f,t,t+n}}\right] = \frac{\mathbb{E}\left(R_{f,t,t+n}^{-1} D_{t+n}/P_t\right)}{P_t^{(n)}/P_t}, \tag{A16}$$

where the numerator is just Equation (27) computed using physical parameters, and the denominator is the same expression computed under risk-neutral dynamics.

Then, taking logs of the above and dividing through by n, gives an annualized log expected HTM return.

Further, note that we can write

$$R_{t,t+n} = \frac{D_{t+n}}{P_t^{(n)}} = \frac{D_t}{P_t^{(n)}} G_{t,t+n}, \tag{A17}$$

where $G_{t,t+n}$ is the cumulative growth rate. We then get that

$$\mathbb{E}_t\left[\frac{R_{t,t+n}}{R_{f,t,t+n}}\right] = \frac{D_t}{P_t^{(n)}} \times \mathbb{E}_t\left[\frac{G_{t,t+n}}{R_{f,t,t+n}}\right], \tag{A18}$$

or

$$\frac{1}{n}\log \mathbb{E}_t\left[\frac{R_{t,t+n}}{R_{f,t,t+n}}\right] = \frac{1}{n}\log\left(\frac{D_t}{P_t^{(n)}}\right) + \frac{1}{n}\log \mathbb{E}_t\left[\frac{G_{t,t+n}}{R_{f,t,t+n}}\right], \tag{A19}$$

that is, dividend yield can be decomposed into annualized log expected HTM return and log cumulative growth rate (both excess of the risk-free rate).

Note that log returns on dividend strips can be expressed as

$$\frac{1}{n}\log\left(R_{t,t+n}\right) = \frac{1}{n}\log\left(\frac{D_{t+n}}{P_t^{(n)}}\right) = \frac{1}{n}\log\left(\frac{D_t}{P_t^{(n)}} G_{t,t+n}\right) \tag{A20}$$

$$= e_{t,n} + \frac{1}{n}\log\left(G_{t,t+n}\right) \tag{A21}$$

As in Backus et al. (2018), taking unconditional expectations of this equation for strips of maturities n and 1, and subtracting them, gives:

$$\frac{1}{n}\mathbb{E}\log\left(R_{t,t+n}\right) - \mathbb{E}\log\left(R_{t,t+1}\right) = \mathbb{E}\, e_{t,n} - \mathbb{E}\, e_{t,1},$$

that is, the unconditional slope of the dividend yield curve is equal to the slope of the expected log strip return curve.

B Estimation

The system of equations consisting of the state space dynamics (1), test asset returns (10), and yields (9) is estimated jointly to obtain the estimates of the following vector of parameters:

$$\theta = [c, \rho_{r,y}, \rho_{y,y}, \beta_{2,u_r}, b_0, b_{1,y}]',$$

where we restrict β_2 to load only on shock to returns, that is, $\beta_2 = [\beta_{2,u_r}, \mathbf{0}_{p\times p}]$, and $b_1 = [\mathbf{0}_{p\times p}, b_{1,y}]$ is restricted to load only on valuation ratios. $\rho_{r,y}$ and $\rho_{y,y}$ denotes sub-blocks of the matrix ρ which correspond to return and D/P loadings on y, respectively.

The system is solved using an iterative Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm. At each iteration we perform the following steps, which we repeat until convergence:

1. Given current values of $c, \rho_{r,y}, \rho_{y,y}$, the transition matrix ρ is reconstructed, shocks u_{t+1} in (1) are computed, and Σ is estimated.
2. Given current values of β_{2,u_r} and $b_{1,y}$, variables β_2 and b_1 are constructed for test assets, respectively, by imposing restrictions that β_2 loads only on shock to returns and b_1 loads only on valuation ratios.
3. For state variables (market and PCs), βs are given by Equations (18)-(20).
4. For given parameter values at any step, we compute prices of risk parameters λ and Λ in Equation (3) via a simple solution to the system of $p = \frac{1}{2}k$ equations for variables in the state space given by Equations (A4) and (A3). Note that parameters λ and Λ inherit restrictions on physical dynamics. Specifically, only the first p elements of λ are non-zero, since only shocks to PC returns are priced; similarly, Λ takes the following form:
$$\Lambda = \begin{bmatrix} 0_{p\times p} & \tilde{\Lambda} \\ 0_{p\times p} & 0_{p\times p} \end{bmatrix},$$
where $\tilde{\Lambda}$ is an $p \times p$ matrix of risk price loadings on D/Ps, which are all fully pinned down by the physical dynamics as well.
5. Given the estimates of parameters λ and Λ we can now compute β_0, β_1 parameters for all test assets using an analogue of Equations (A4) and (A3) applied to test assets.
6. Lastly, we stack state and test asset return and D/P equations, as well as all parameters. We then use them to estimate the shocks, instruments, and moment conditions. The moment conditions are as follows:
   (a) Using the state vector dynamics in Equation (1), we construct shocks u_{t+1} and interact them with instruments which contain a vector of ones and valuation ratios F_t. This corresponds to standard OLS moments with additional restrictions that elements of ρ corresponding to loadings on lagged returns are all zero,
   $$\mathbb{E}\left(u_{t+1} \otimes [1, F_{y,t}]\right) = 0,$$
   where $\otimes$ denotes a kronecker product.
   (b) Using return equations (5), we construct shocks to returns, $\epsilon_{r,t+1}$, and interact them with instruments which contain a vector of ones, valuation ratios in F_t, as well as contemporaneous return shocks in u_{t+1},
   $$\mathbb{E}\left(\epsilon_{r,t+1} \otimes [1, F_{y,t}, u_{r,t+1}]\right) = 0.$$
   (c) Using the dividend price equation (9) we construct residuals $\epsilon_{y,t}$ and interact them with instruments which contain a vector of ones and contemporaneous D/P ratios in the state vector,
   $$\mathbb{E}\left(\epsilon_{y,t} \otimes [1, F_{y,t}]\right) = 0.$$

We use a prespecified GMM weighting matrix in the GMM objective, where the time-series moments have weight 1 and individual assets' moments are all weighted by the inverse of the square root of the number of test assets, $\sqrt{n}$, to keep their contribution to GMM objective invariant of n.

$$\hat{\theta} = \arg\min_{\theta} g_T(\theta)' W g_T(\theta). \tag{A22}$$

The following parameters are estimated by GMM:

1. State space vector variables:
   (a) Intercept c: k parameters
   (b) Loadings ρ: k loadings onto D/P ratios, $k \times p$ parameters
2. Asset-specific parameters:
   (a) Intercepts of D/P equations b_0: n parameters
   (b) Loadings of assets' D/Ps onto state-space D/Ps: $n \times p$ (loadings on returns are restricted to zeros)
   (c) Loadings of assets' returns onto state-space shocks to returns: $n \times p$ (loadings on D/Ps are restricted to zeros). All other loadings of test assets' returns are pinned down by no arbitrage and SDF risk prices.

The spectral density covariance matrix of moments uses 12 lags and follows the approach in Hansen (1982).

Lastly, to compute standard errors on means of yields or returns (risk premia) that account for both model parameter uncertainty and data sampling uncertainty we expand the set of moments to include moments corresponding to a regression of these variables on a constant. We then use the structure of our GMM estimator to obtain correct standard errors on the intercept, which is the estimate of interest.

Internet Appendix

C Model calibrations and evaluation

In this section we illustrate how our synthetic term structures can be used to evaluate asset pricing models. We propose two different tests of the models. The first test follows the standard procedure in the literature (e.g., Beeler and Campbell (2012)): it simulates a model many times (specifically, 10,000 times) generating samples as long as the actual data sample X. Each simulation of the model produces a simulated dataset $\tilde{X}$. A particular moment $g()$ is computed using the actual data sample, $g(X)$, and using the various simulated datasets, producing a distribution of statistics $g(\tilde{X})$ across simulated samples. The first test then computes how likely the observed moment $g(X)$ is to be generated by the model. For example, we generate the average slope of the term structure of discount rates in the Bansal and Yaron (2004) model in many samples, and ask how often it is as low as we measure it in the data using our reduced-form estimate. The test takes the entire model as a null, and therefore uses the calibrated model parameters (and assumptions about shocks distributions) when generating the simulated data. Note that this test does take into account the time-series component of uncertainty, that is the fact that a time-series average is used to estimate an unconditional expectation of the realized moments in the sample (through the resampling of the data). However, this procedure does *not* take into account the fact that the realized moment we measure in the data (e.g., the term structure of discount rates at each point in time) is not actually *observed* from traded prices, but rather it is estimated using our reduced-form model, which, as discussed in Section 2.4, adds another component of uncertainty.

We therefore propose an alternative way to test asset pricing models using our estimated moments, that explicitly takes into account both the time-series uncertainty and the uncertainty coming from estimating our reduced-form model. In this test, we compute for each model the *population* value of each moment, and then use our point estimate and GMM standard errors (which account for both types of uncertainty) to test the null that the moment of interest is what is implied by the model in population.

We calibrate and test 5 models: the Bansal and Yaron (2004) (BY) and Bansal et al. (2010) (BKY) models, the habit formation model of Campbell and Cochrane (1999), the model of Lettau and Wachter (2007), and the rare disaster model of Gabaix (2012). Model calibrations and simulations are implemented as follows. For the solution and calibration of the BY and BKY models, we follow Beeler and Campbell (2012). We build equity yields (EY) and forward equity yields (FEY) as in van Binsbergen et al. (2012a) and bond yields as in Beeler and Campbell (2012). For the estimation and calibration of Campbell and Cochrane (1999) model, we follow Gonçalves (2021), who also computes equity yields, basing our code on the one he kindly shares; in this model, the term structure of bond yields is flat and constant, so that the EY and FEY curves are the same. For the estimation and calibration of Lettau and Wachter (2007), we follow Gormsen (2018), who also computes equity yields, basing our code on the one he kindly shares; in this model, the term structure of bond yields is flat and constant, so that the EY and FEY curves are the same. For the rare disaster model in Gabaix (2012), we use the closed-form expressions (for equity yields and bonds) provided in the paper, and simulate the model based on the procedure detailed in the appendix of that paper.

Table A.1 reports results for two different time periods. In the top panel, it considers only the period 2004-2020. In the bottom panel, all results are reported for the full sample (1976-2020). Each panel is divided horizontally in two parts. The left part performs the first test (simulating the model and comparing the results with our point estimates of each moment). The right part performs the second test (using our GMM standard errors to test the null imposed by the population values of each model). Each of the two parts reports the results for all five models mentioned above.

We consider a variety of moments (one in each row), motivated by the recent literature. First, we consider the average slope of the term structure of EY and FEY. Note that we use the 2 year maturity as the short end when computing the slope, because as discussed by van Binsbergen et al. (2012a) the exact assumptions about the payment of dividends within the year matter for the return of the strips (especially the 1-year strips whose return is directly determined by the dividend paid out); in addition, the models are calibrated at different frequencies (some monthly, some quarterly), and therefore imply different timing of the dividend payments of a 1-year strip. In previous versions of this paper, we had used the 1-year strip as the short end of the curve, and all results were qualitatively and quantitatively similar. Second, we study the risk premia (average excess returns for a 1-year holding period) of strips and forward strips at different maturities, 2, 7 and 15 years. Finally, we test the coefficient of a regression of the slope of the equity yield term structures on the log price-dividend ratio, capturing the cyclicality of the term structure of equity yields.

Table A.1: Model-implied term structures.

We simulate five models and compare the model simulations to the data moments: the long-run risk models of Bansal and Yaron (2004) (BY) and Bansal et al. (2010) (BKY), the habit-formation model of Campbell and Cochrane (1999) (habit), the model of Lettau and Wachter (2007) (LW), and the rare-disaster model of Gabaix (2012) (disaster). Each row corresponds to a different moment. The numbers in the left side of the table correspond to the fraction of simulated samples (from the models) in which the moments are lower than the one we estimate. The numbers in the right side of the table are the one-sided p-value for the hypothesis that the true population moment is higher than the population moment implied by the models; the test uses our point estimate and GMM standard errors. So both tests reports numbers close to 0 when the estimated moment is below the model-implied one, and close to 1 when the estimated moment is above the model-implied one. The table reports results for two time periods: post-2004 (top panel) and full sample (bottom panel).

Panel A: 2004+	Test 1: via model simulation					Test 2: using SE from data				
Moments	BY	BKY	LW	Habit	Disaster	BY	BKY	LW	Habit	Disaster
EY slope, 7-2	0.16	0.54	0.91	0.96	1.00	0.26	0.46	1.00	0.61	0.76
EY slope, 15-2	0.31	0.60	0.95	0.52	1.00	0.36	0.55	1.00	0.43	0.90
FEY slope, 7-2	0.02	0.00	0.85	0.00	0.49	0.04	0.11	0.99	0.35	0.51
FEY slope, 15-2	0.03	0.00	0.90	0.00	0.71	0.04	0.08	1.00	0.15	0.67
Avg. ret. EY, 2	0.34	0.51	0.02	0.49	0.02	0.29	0.19	0.00	0.48	0.00
Avg. ret. EY, 7	0.50	0.60	0.12	0.93	0.19	0.48	0.50	0.04	0.87	0.13
Avg. ret. EY, 15	0.70	0.66	0.62	0.98	0.51	0.72	0.70	0.62	0.93	0.46
Avg. ret. FEY, 2	0.22	0.45	0.02	0.42	0.01	0.15	0.10	0.00	0.41	0.00
Avg. ret. FEY, 7	0.13	0.41	0.03	0.61	0.03	0.08	0.09	0.01	0.57	0.03
Avg. ret. FEY, 15	0.14	0.06	0.19	0.62	0.09	0.10	0.07	0.18	0.54	0.11
7-2 EY slope on PD	0.99	0.92	0.83	1.00	1.00	0.67	0.94	0.62	0.95	0.89
15-2 EY slope on PD	0.87	0.56	0.35	1.00	1.00	0.59	0.97	0.48	0.99	0.89
7-2 FEY slope on PD	0.97	0.85	0.98	1.00	1.00	0.67	0.98	0.73	0.98	0.94
15-2 FEY slope on PD	0.82	0.51	0.78	1.00	1.00	0.59	0.99	0.65	1.00	0.96

Panel B: Full sample	Test 1: via model simulation					Test 2: using SE from data				
Moments	BY	BKY	LW	Habit	Disaster	BY	BKY	LW	Habit	Disaster
EY slope, 7-2	0.24	0.63	0.98	0.99	1.00	0.26	0.69	1.00	0.90	0.99
EY slope, 15-2	0.23	0.59	0.99	0.46	1.00	0.23	0.60	1.00	0.35	1.00
FEY slope, 7-2	0.01	0.28	0.97	0.86	0.98	0.01	0.08	1.00	0.55	0.82
FEY slope, 15-2	0.01	0.00	0.97	0.00	0.94	0.00	0.01	1.00	0.05	0.89
Avg. ret. EY, 2	0.51	0.53	0.00	0.80	0.01	0.46	0.20	0.00	0.83	0.00
Avg. ret. EY, 7	0.60	0.60	0.01	1.00	0.12	0.58	0.62	0.00	0.99	0.05
Avg. ret. EY, 15	0.74	0.62	0.61	0.99	0.38	0.74	0.71	0.57	0.99	0.32
Avg. ret. FEY, 2	0.25	0.45	0.00	0.67	0.00	0.13	0.05	0.00	0.68	0.00
Avg. ret. FEY, 7	0.08	0.41	0.00	0.92	0.01	0.03	0.04	0.00	0.83	0.00
Avg. ret. FEY, 15	0.07	0.33	0.08	0.87	0.03	0.07	0.04	0.17	0.72	0.07
7-2 EY slope on PD	0.00	0.40	0.00	1.00	0.75	0.06	0.94	0.03	0.95	0.70
15-2 EY slope on PD	0.00	0.43	0.00	1.00	0.64	0.03	0.97	0.01	1.00	0.69
7-2 FEY slope on PD	0.00	0.33	0.00	1.00	0.65	0.01	0.94	0.02	0.93	0.65
15-2 FEY slope on PD	0.00	0.34	0.00	1.00	0.50	0.00	0.97	0.00	1.00	0.62

On the left side of the table, the numbers report the fraction of model simulated samples for which the moment is *lower* than the one we estimate in the data. On the right side, the table reports the one-sided p-value for the hypothesis that, given our point estimates and the uncertainty in our estimates, the true population moment is equal to (or higher than) the one implied by the model. Therefore, in both cases, when the number reported is close to 0, the model is rejected because we estimate an empirical moment *lower* than what the model implies, and when the number is close to 1 the model is rejected because we estimate a moment *higher* than what the model implies. For example, for the BY model, the slope of the term structure of FEY is too steep: the p-values reported in the table are close to 0 (for the 15-2 slope, .03 and .04 with the two tests respectively).

Looking across the left and right parts of the table, we note that the two tests give similar evaluations of the models, even if the right side (that incorporates the uncertainty from the estimation of the moments) has, as expected, less power in most cases.

The table presents a rich evaluation of the models. Overall, it appears that each model succeeds in some dimensions and fails in others. While this exercise is illustrative, it gives an example of how these additional moments can help the calibration and evaluation of the economic forces at play in the models. For example, the table shows that the BY model does a better job explaining the slope of the EY curve than that of the FEY curve. Given that the difference in the two curves depends on the term structure of interest rates, this reinforces the point made by Beeler and Campbell (2012) that the BY model has counterfactual implications about the term structure of interest rates.

More specifically, the BY model tends to generate term structures of FEY that are too steep, and, as pointed out by Beeler and Campbell (2012), it also generates term structures of interest rates that are downward sloping (and therefore not steep enough compared to the data). Given that the EY are the sum of the FEY and the bond yields, the two errors in matching FEY and bond yields cancel out, and the model actually generates an EY term structure much more consistent with the data. This illustrates that looking jointly at different moments can be quite informative about the various mechanisms at play in the model.

When we look at the individual risk premia on strips (and forwards) of different maturity, we find that the risk premia for long maturity forwards are too high in the BY and BKY model. Interestingly, the individually estimated short-maturity risk premia are not rejected statistically (partly because of low power, but partly because they do not differ dramatically from those predicted by the model). While we don't have direct evidence on the very short end of the curve, when we look at the risk premia of the 2-year strips we do not find them to be exceedingly high, which makes the results similar to the predictions of BY and BKY. Finally, we estimate a procyclical slope of the EY and FEY term structures (the slope is high when the PD ratio is high). This is in line with the prediction of the BY and BKY model in the post-2004 sample. In the full sample, we actually estimate a much lower degree of procyclicality, lower than what we find after 2004 and than what the models predict. The cyclicality of the term structure of equity yields and risk premia is an interesting moment to study, as it reveals interesting features of the dynamics of the variables driving the slope of the term structure. Our analysis allows us to expand the sample significantly, including many more business cycle compared to the analysis that uses traded strips only and that starts in 2004.

The LW model is known to generate downward sloping term structures. Given that our estimated term structures are somewhat upward sloping, it should not be surprising that the LW model is rejected as having *too low* a slope in our sample. The LW model generates risk premia on EY and FEY that are too high at the short term, both in the full sample and the post-2004 sample. This again highlights the fact that estimated short-term risk premia are relatively low in our estimate (and therefore not consistent with models where the term structure of risk premia is steeply downward sloping as in LW). Finally, the LW model also generates procyclical term structures, about in line with our post-2004 estimates, and higher than our full-sample estimates (which are much closer to zero).

The habit formation model generates EY slopes that are broadly in line with our estimates and FEY slopes that are too steep. The model cannot be rejected statistically on many of the individual maturity risk premia (partly because risk premia are estimated with less power than term structure slopes). Interestingly, the model predicts that the slope of the EY and FEY term structure is countercyclical, which is at odds with the data.

Finally, the rare disaster model implies a flat term structure of discount rates for EY and FEY, so it tends to predict a slope that is too low compared to the data. The rare disaster model is the model that best matches the slope of the FEY term structure. The calibration of the model gets the risk premia off by a few percentage points (in other words, it predicts the level of risk premia higher than we find in the data), so it is rejected on the level of risk premia dimension. Finally, it gets the procyclicality about right, at least in the full sample.

To sum up, the table shows that different models match some aspects of the data well and others less so. Overall, the rare disaster model appears to be the one that matches the data the best, except for the level of the risk premium which is calibrated to be too high. Beyond these illustrative examples, the moments we provide in this paper can be useful to help guide and refine the calibrations of future models (taking into account, of course, that there is substantial estimation error in our estimated term structures).

Table A.2: Correlations of returns and Sharpe Ratios of sorted portfolios

The table reports selected columns of the correlation matrix and the annualized Sharpe ratios (last column) of portfolios sorted by Macaulay duration (MD), level and slope of the discount rate curve ($\mathbb{E}(r)$), and level and slope of the expected dividends curve ($\mathbb{E}(g)$).

	Correlations							Sharpe
	MD	$\mathbb{E}[r_l]$	$\mathbb{E}[r_s]$	$\mathbb{E}[\Delta g_l]$	$\mathbb{E}[\Delta g_s]$	β_1	β_2	
MD	-	0.37	0.26	-0.50	0.51	-0.15	0.75	0.37
$\mathbb{E}[r_l]$	0.37	-	0.55	-0.49	0.58	0.13	0.55	0.65
$\mathbb{E}[r_s]$	0.26	0.55	-	-0.27	0.42	0.12	0.30	0.01
$\mathbb{E}[\Delta g_l]$	-0.50	-0.49	-0.27	-	-0.90	0.20	-0.48	0.29
$\mathbb{E}[\Delta g_s]$	0.51	0.58	0.42	-0.90	-	-0.03	0.64	0.03
β_1	-0.15	0.13	0.12	0.20	-0.03	-	0.08	1.27
β_2	0.75	0.55	0.30	-0.48	0.64	0.08	-	0.71
β_3	-0.01	0.36	0.05	-0.16	0.34	0.10	0.61	0.31
β_4	-0.27	-0.02	-0.12	0.28	-0.21	0.14	-0.19	0.81
α_l	0.42	0.94	0.52	-0.55	0.66	0.10	0.65	0.74
α_s	0.42	0.68	0.87	-0.47	0.62	0.08	0.54	0.19
D/P	0.84	0.56	0.29	-0.81	0.78	-0.11	0.68	0.33

D Duration

Following the work of van Binsbergen et al. (2012b), documenting the declining term structure of discount rates in equity dividend derivatives, subsequent work has studied the term structure of discount rates in equities in reduced form, by sorting firms by measures of duration and studying the cross-sectional patterns that emerge. For example, Weber (2018), Gonçalves (2019), and Gormsen and Lazarus (2019) all show that low-duration portfolios appear to command higher risk premia (or CAPM alphas) compared to high-duration portfolios, obtaining therefore results that are consistent with the analysis of dividend derivatives.

Given that our model is estimated from equities, but also matches the prices of dividend strips, we can look at duration sorts through the lens of this model, to gain a better understanding of what duration-sorted portfolios are capturing.

The measure of duration typically used in this literature (Macaulay duration) is derived from the bond literature. In that context, it is clear that duration captures a specific risk exposure: exposure to shocks to the level of the yield curve. Applying this notion of duration to equities, however, raises an important issue: equities are exposed to a multiplicity of different shocks, to both dividends and discount rates, both long term and short term. For a general risky asset, Macaulay duration *may* line up with exposure to long-term discount rate shocks (similarly to the case of bonds), but it may also line up with exposure to long-term dividend (cash flow) shocks, or, for example, to a combination of the two.

For risky assets, therefore, duration is economically less informative than looking separately at exposures to the different shocks, which is something that our estimated model can deliver for any asset and for any type of risk. For example, if one is interested in understanding the risk premium associated with long-term discount rate shocks, one can use a model to build a portfolio that is directly exposed to that shock.

To illustrate the issues related to the interpretation of duration-sorted portfolio, we build, using our model, portfolios sorted monthly by duration and by exposure to different types of shocks: level and slope of the term structure of discount rates, and level and slope of the term structure of cash flow.[21] We then check empirically which exposures the duration-sorted portfolios line up with.

Table A.2 reports the correlation between the returns of portfolios sorted on different variables: duration

[21] After estimating the model, we approximate duration for each stock by taking the average of durations of all portfolios this given stock directly enters.

(MD), level and slope of the discount rate curve ($\mathbb{E}(r_l)$ and $\mathbb{E}(r_s)$ respectively), level and slope of the expected dividends curve ($\mathbb{E}(\Delta g_l)$ and $\mathbb{E}(\Delta g_s)$ respectively), exposures to shocks u_{t+1} (denoted as $\beta_1 - \beta_4$), CAPM alphas (α_l and α_s), and on D/P (a value sort).

Two interesting results emerge from this table. First, Macaulay duration is strongly correlated with both the expected return and the dividend term structures. Sorts on Macaulay duration are also highly correlated with sorts on measures of value (D/P) and betas with respect to the first principal component of long-short anomalies (the dominant risk priced in the cross-section – β_2), with correlation coefficients of 0.84 and 0.75, respectively. Therefore, Macaulay duration cannot be clearly interpreted as pure exposure to discount rate shocks, as in the case of bonds.[22]

Second, while (consistent with the existing literature) Macaulay duration is associated with a positive Sharpe ratio (0.37), sorts based on the other variables also produce similar (or larger) Sharpe ratios. In particular, sorts by the level of the term structure of discount rates command a Sharpe ratio of 0.65, and sorts by the level of the expected dividend term structure command a Sharpe ratio of 0.29. Sharpe ratio of the sort on betas with respect to the first PC of anomalies, β_2, is 0.71. Interestingly, Sharpe ratio of the portfolio weighted by the stocks' loadings on the market orthogonalized with respect to the cross-sectional factors yields the highest Sharpe ratio of 1.27.

These results have interesting implications for investors. An investor who wants to tailor her exposure to long-term discount-rate or cash-flow shocks can do better than sorting on Macaulay duration. The alternative sorts presented in this section offer more targeted and interpretable exposures (to discount rates vs. cash-flow shocks), and achieve higher Sharpe ratios.

In addition, note that the portfolio sorted by the level of expected returns is highly correlated (0.55) with the portfolio sorted by exposure to the second factor in the model (β_2), which carries the Sharpe ratio of 0.71. So these results suggest that an investor who wants to obtain exposure to the second factor, can do so by buying assets sorted on the level of the term structure of discount rates. In fact, one can do even better by sorting on the level of CAPM alphas, α_l, because this sort has an even higher correlation with the second PC, 0.65, and delivers effectively the same Sharpe ratio as the second factor itself, 0.74. Whereas these were obtained within the model, using a long time series, going forward an investor can simply look at the term structure from dividend strips to form portfolios exposed to this highly priced risk factor.

Finally, portfolios like the ones we present in this section can also be used as a moment to evaluate and test asset pricing models: compared to duration-based portfolios, they provide a more powerful test, because they allow the researcher to distinguish between the different types of shocks (long-term vs. short-term, cash flow shocks vs. discount rate shocks).

E Additional results

E.1 Properties of anomaly portfolio returns

Table A.3 shows annualized mean excess returns on the fifty anomaly long-short portfolios as well as the underlying characteristic-sorted tercile portfolios.

[22] There are other interesting correlations in this table: for example, portfolios sorted on level and slope shocks are positively correlated with correlation of 0.55, whereas portfolios sorted level and slope dividend growth shocks are highly negatively correlated, -0.9. These patterns reflect a combination of the dynamics of the underlying shocks to dividends and returns, as well as the cross-sectional correlation of exposures to the different shocks.

Table A.3: Anomaly portfolios mean excess returns, %, annualized

Columns Long and Short show mean annualized returns (in %) on each anomaly portfolio long (P3) and short ends (P1) of a sort, respectively, net of risk-free rate. The column L-S lists mean returns on the strategy which is long portfolio 3 and short portfolio 1. Portfolios include all NYSE, AMEX, and NASDAQ firms; however, the breakpoints use only NYSE firms. Monthly data from February 1973 to December 2020.

	Short	Long	L-S
Accruals	4.3	7.4	3.0
Asset Growth	5.3	8.4	3.1
Asset Turnover	5.0	6.9	1.9
Beta Arbitrage	3.9	7.2	3.3
Cash Flows/Price	5.2	8.1	2.9
Composite Issuance	4.2	8.0	3.8
Debt Issuance	5.5	7.0	1.5
Dividend Growth	6.2	5.8	-0.4
Dividend/Price	5.1	6.7	1.6
Duration	5.6	7.8	2.1
Earnings/Price	4.5	7.8	3.3
F-score	5.1	6.6	1.4
Firm's age	5.8	5.4	-0.5
Gross Margins	5.6	5.8	0.1
Gross Profitability	4.5	6.8	2.3
Growth in LTNOA	5.9	6.6	0.7
Idiosyncratic Volatility	3.7	6.5	2.7
Ind. Mom-Reversals	4.0	9.0	4.9
Industry Momentum	4.0	6.6	2.6
Industry Rel. Rev. (L.V.)	3.2	10.9	7.7
Industry Rel. Reversals	2.6	9.8	7.2
Investment Growth	5.4	7.4	2.0
Investment/Assets	5.2	7.2	2.0
Investment/Capital	5.6	7.0	1.5
Leverage	5.4	6.3	0.9
Long Run Reversals	6.0	7.3	1.2
Momentum (12m)	2.2	8.6	6.4
Momentum (6m)	6.7	6.6	-0.1
Momentum-Reversals	5.6	7.1	1.5
Net Issuance (A)	5.0	7.9	2.9
Net Issuance (M)	5.2	7.6	2.3
Net Operating Assets	4.3	7.2	2.9
Price	5.0	5.8	0.9
Return on Assets (A)	4.9	6.1	1.3
Return on Assets (Q)	3.0	6.5	3.5
Return on Book Equity (A)	5.1	6.2	1.1
Return on Book Equity (Q)	3.1	6.7	3.6
Return on Market Equity	2.4	9.5	7.1
Sales Growth	5.9	6.7	0.8
Sales/Price	5.0	8.8	3.8
Seasonality	4.0	7.9	3.9
Share Repurchases	5.4	6.9	1.5
Share Volume	4.8	5.8	1.0
Short Interest	3.5	6.8	3.2
Short-Term Reversals	4.0	7.4	3.4
Size	5.7	6.5	0.8
Value (A)	5.5	7.4	2.0
Value (M)	5.3	7.3	2.0
Value-Momentum	5.7	7.3	1.5
Value-Momentum-Prof.	5.9	9.1	3.3
Value-Profitablity	4.3	9.1	4.8

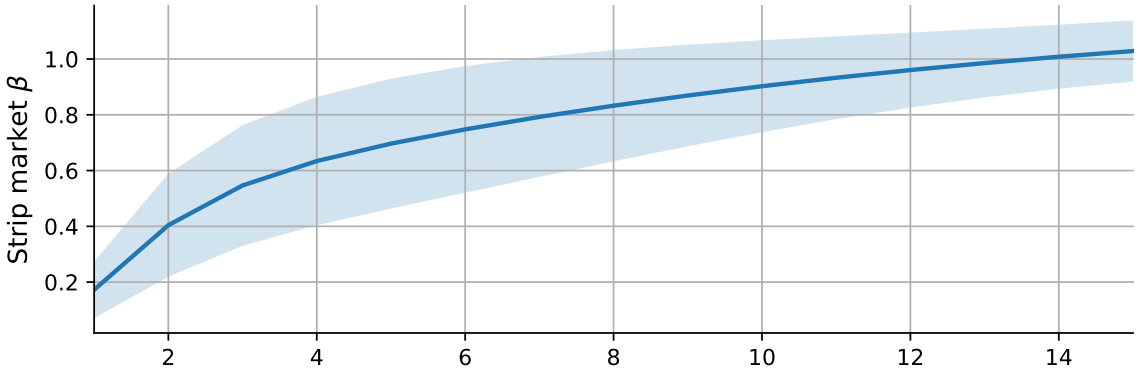

Figure A.14: Average market beta of dividend strips by maturity. We plot average betas between returns on a dividend strip of any given maturity and the aggregate market index. Shaded areas depict two-standard-deviation bands around point estimates.

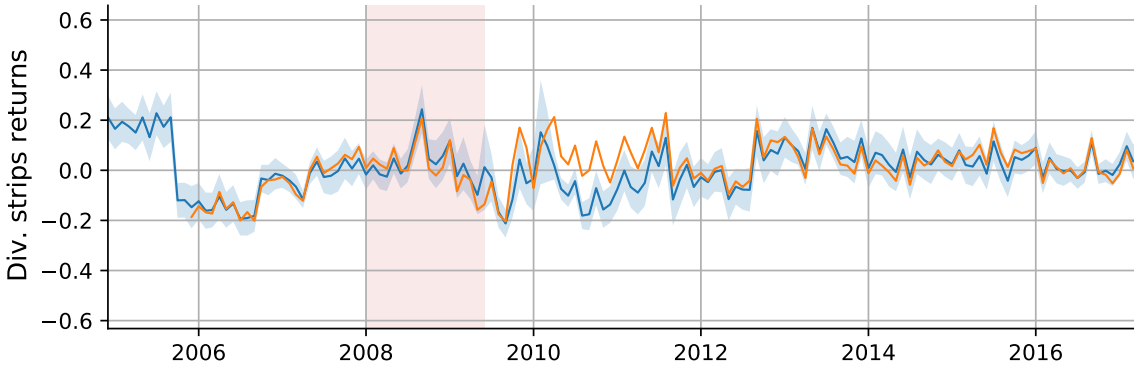

(a) 1-year

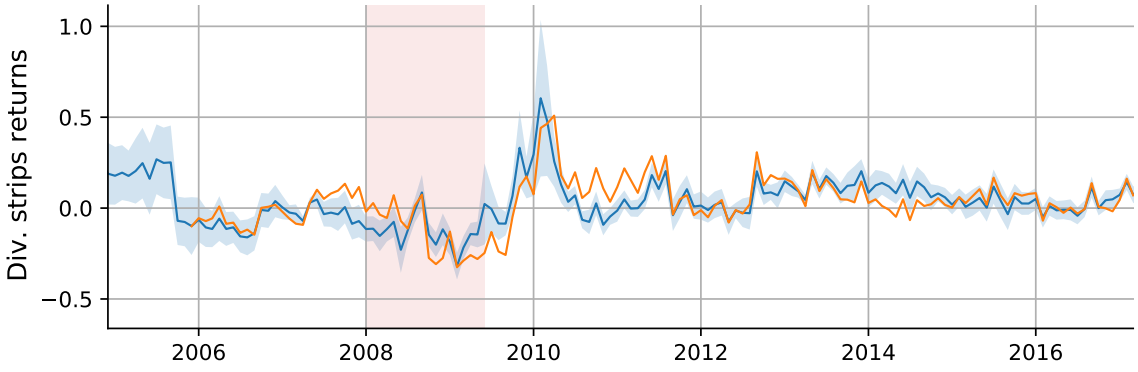

(b) 2-year

Figure A.15: One-year returns of the 1- and 2-year dividend strips in the data and in our model. We compare our model-implied returns on dividend strips to their empirical counterparts based on Bansal et al. (2017). Shaded areas depict two-standard-deviation bands around point estimates. Model parameters are estimated in full sample.

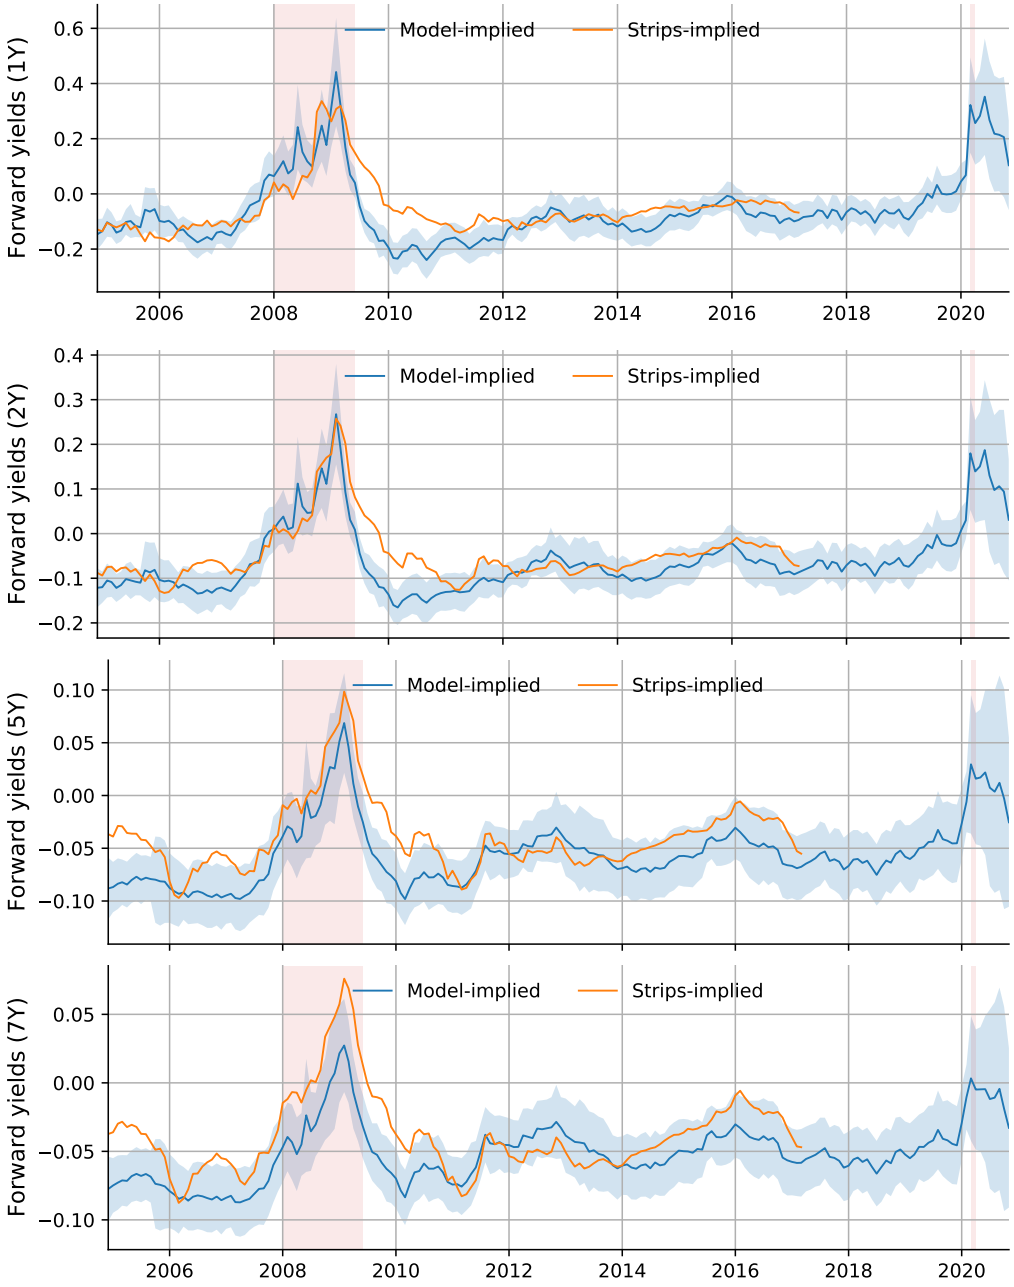

Figure A.16: Model-implied forward equity yields vs. forward equity yield data (GHZ RPS). We compare our model-implied forward yields to their empirical counterparts in Bansal et al. (2017). Shaded areas depict two-standard-deviation bands around point estimates. Model parameters are estimated in full sample.

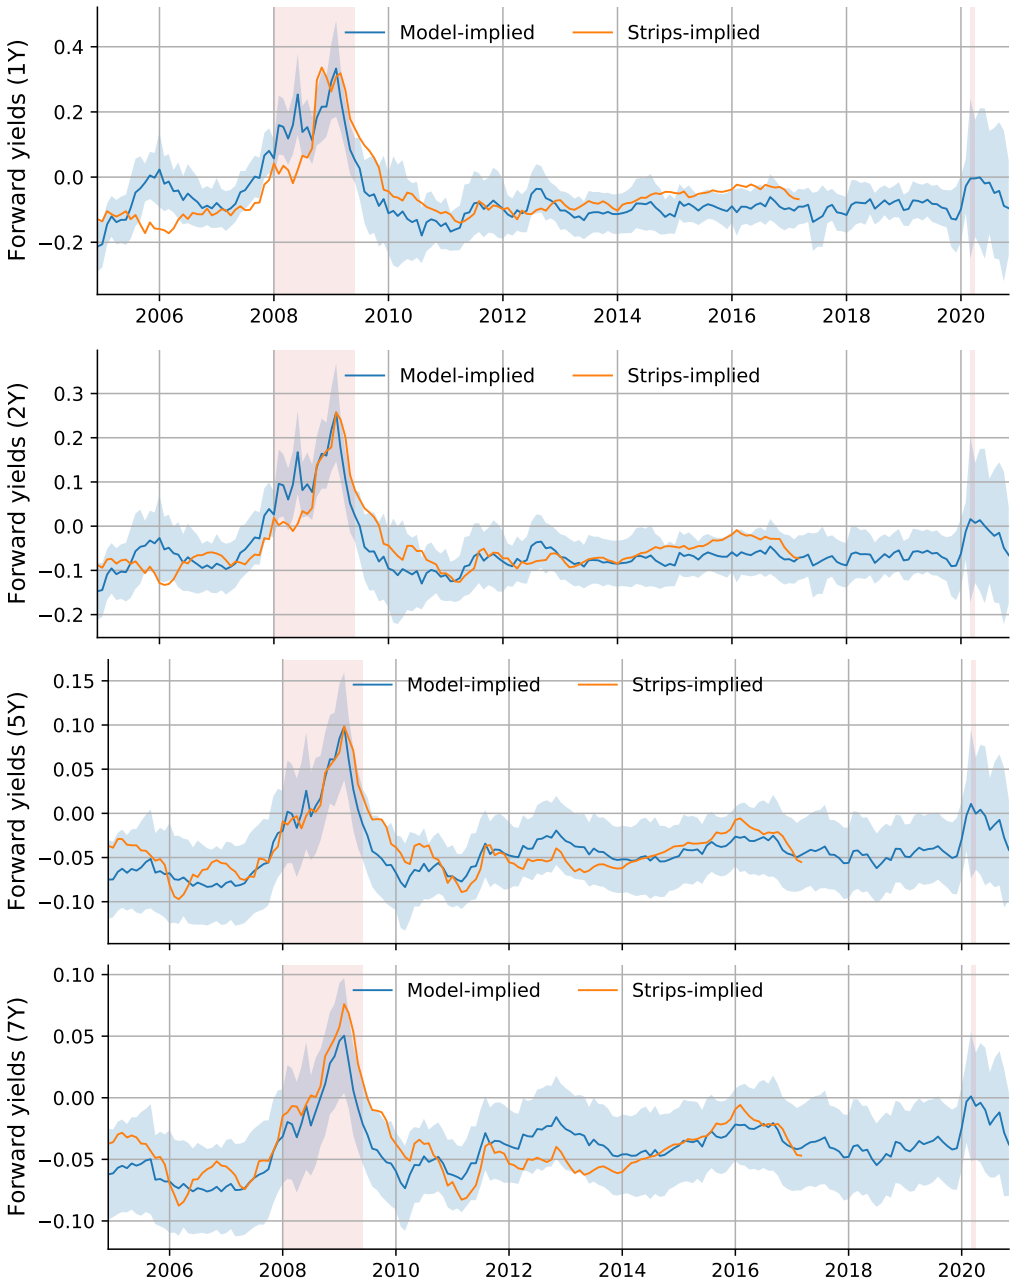

Figure A.17: Model-implied forward equity yields vs. forward equity yield data (WRDS financial ratios). We compare our model-implied forward yields to their empirical counterparts in Bansal et al. (2017). Shaded areas depict two-standard-deviation bands around point estimates. Model parameters are estimated in full sample.

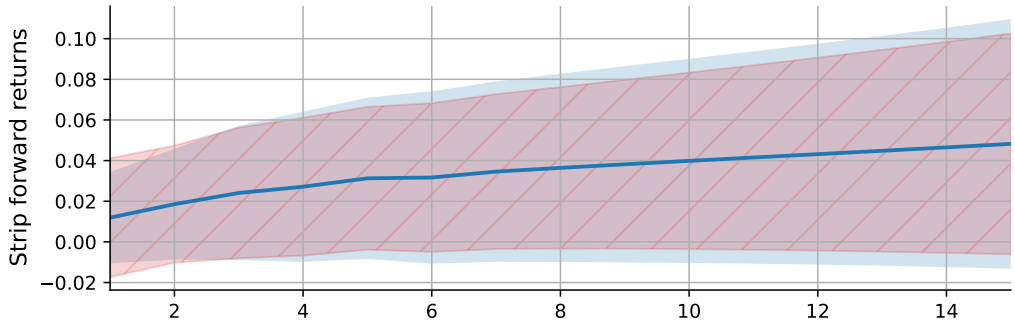

(a) Forward risk premium (full sample)

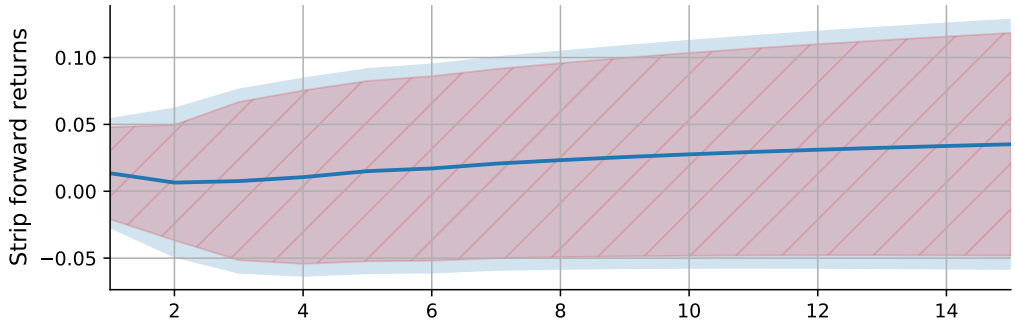

(b) Forward risk premium (BMSY sample)

Figure A.18: Comparison of standard errors for forward risk premia by maturity. The figure compares model-implied GMM two-standard-error bounds which incorporate model parameter and sampling uncertainty (blue) to HAC robust standard error bounds reflecting only sampling uncertainty (red).

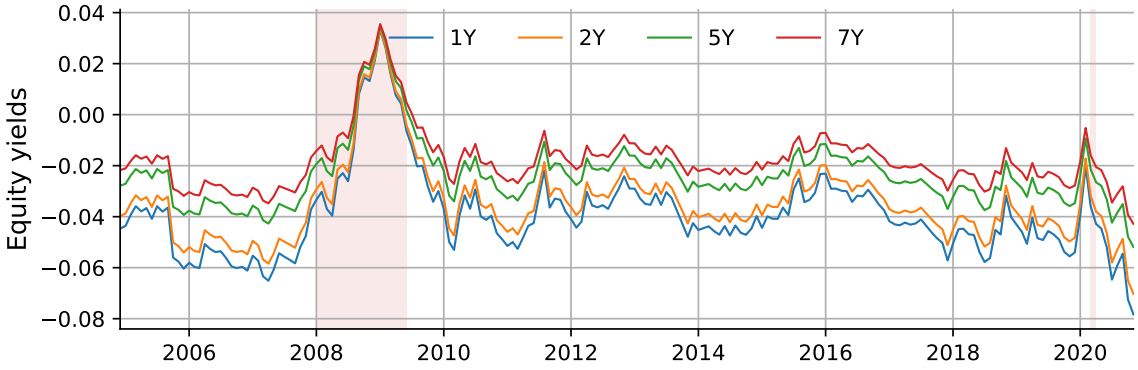

(a) CAPM

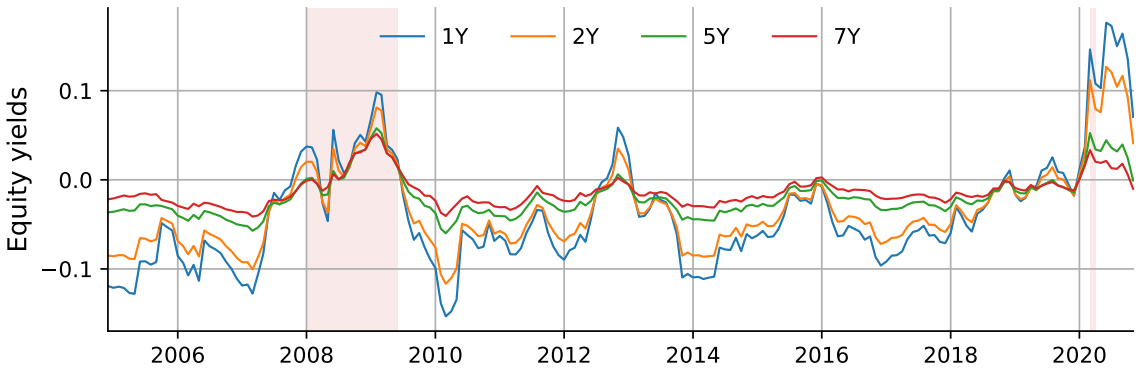

(b) FF 5-factor + Momentum

Figure A.19: Dynamics of benchmark-implied equity strip yields for the aggregate market for different maturities. The figure plots dynamics of yields implied by the CAPM (Panel a) and Fama-French 5-factor model supplemented with the momentum factor (Panel b) for maturities 1, 2, 5, and 7 years.

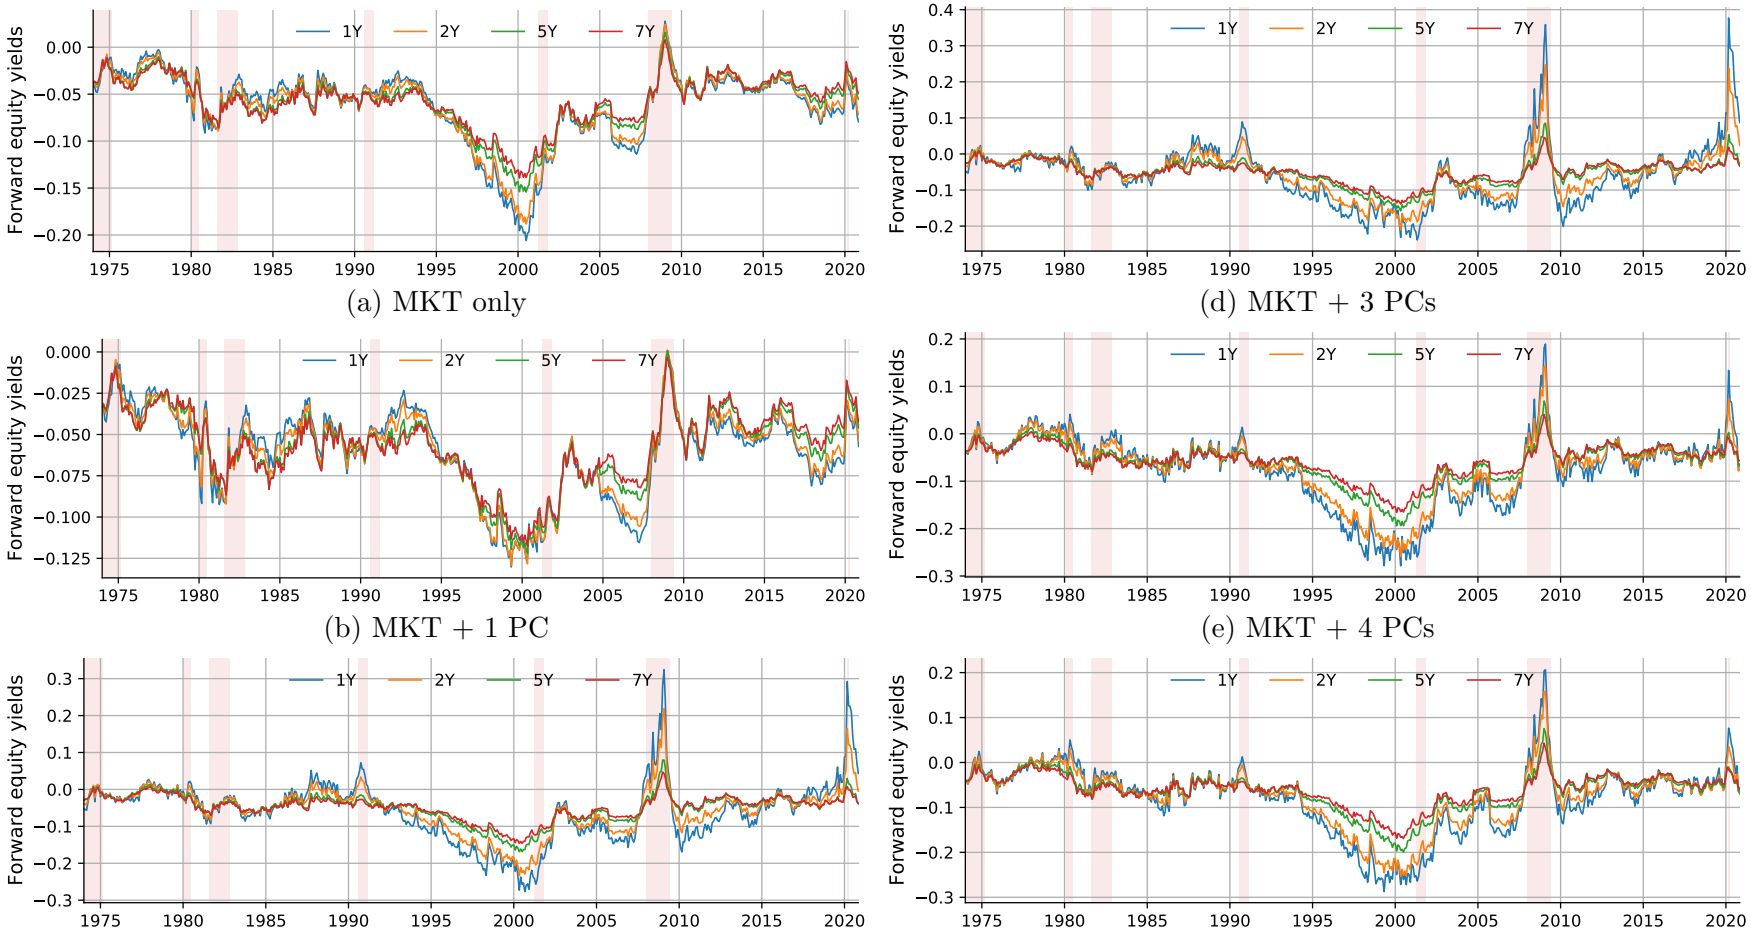

Figure A.20: Model-implied forward equity yields for models with varying number of PCs. We compare forward equity yields for models fitted using: (a) the market only, and (b)-(f) the market plus 1-5 PCs of long-short anomaly returns.

Table A.4: Bootstrap vs. model-implied standard errors

The table reports time series averages of standard errors (in %) for forward equity yields of maturities 1-7 years. The top row of each panel uses model-implied GMM standard errors. The bottom row reports standard errors based on the non-parametric bootstrap method. Results for the post-2004 (BMSY) sample are reported in the top panel and for the full sample in the bottom panel.

Horizon:	1-year	2-year	5-year	7-year
	Post-2004 sample			
model	4.37	3.17	1.67	1.21
bootstrap	4.24	3.28	2.08	1.70
	Full sample			
model	4.21	3.18	1.88	1.47
bootstrap	4.22	3.50	2.50	2.11

Table A.5: SDF spanning of duration portfolio returns.

We regress returns on the long-short duration-sorted portfolio onto the return factors included in the SDF.

	Dependent variable:
const	0.026^{***}
	(0.005)
MKT	-0.042^{***}
	(0.014)
DP_{MKT}	-1.449^{***}
	(0.235)
PC1	0.218^{***}
	(0.005)
DP_{PC1}	0.234^{**}
	(0.091)
PC2	-0.217^{***}
	(0.006)
DP_{PC2}	-0.114
	(0.129)
PC3	0.078^{***}
	(0.008)
DP_{PC3}	-0.180
	(0.110)
Observations	563
R^2	0.920
Adjusted R^2	0.918
Residual Std. Error	0.040(df = 554)
F Statistic	791.398^{***} (df = 8.0; 554.0)
Note:	*p<0.1; **p<0.05; ***p<0.01

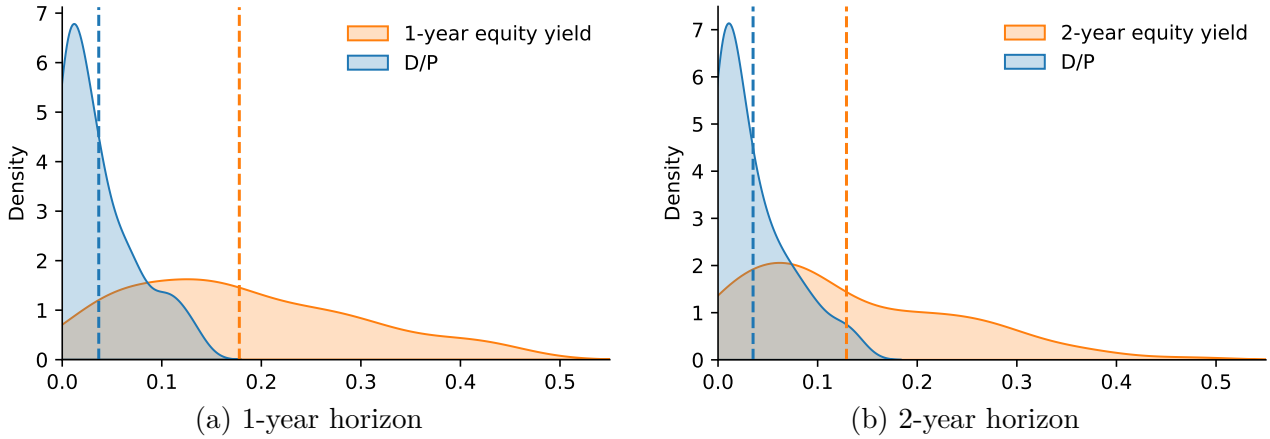

(a) 1-year horizon

(b) 2-year horizon

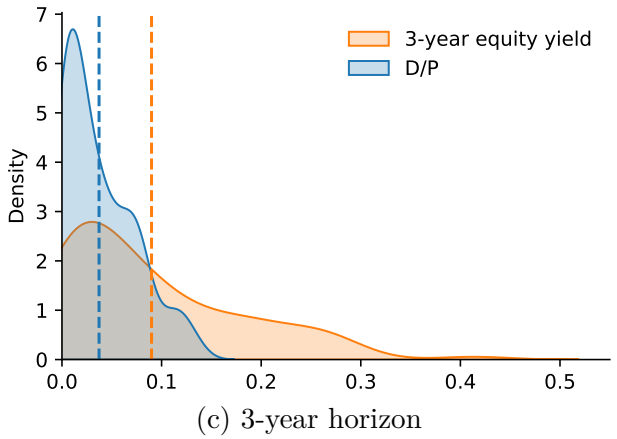

(c) 3-year horizon

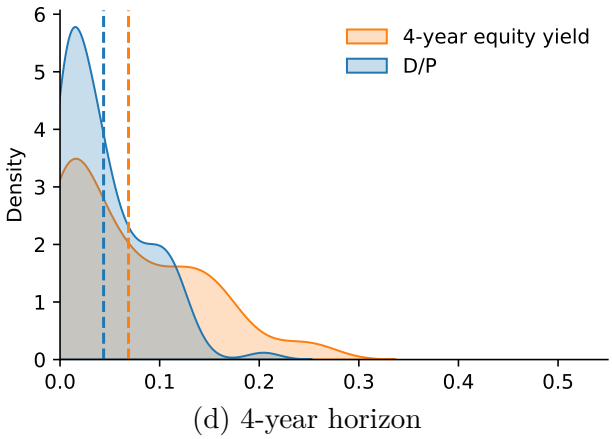

(d) 4-year horizon

Horizon:	1-year	2-year	3-year	4-year	5-year	6-year	7-year
p-value	0.00	0.00	0.00	0.00	0.10	0.28	0.35

Figure A.21: Distributions of dividend growth predictability R^2. We show cross-sectional distribution empirical densities of R^2 of 102 portfolios from two regressions of an n-year dividend growth on: (i) the equity's own dividend-to-price ratio (D/P), and (ii) the model-implied dividend yield on an n-year dividend strip. Panels (a)-(d) correspond to $n = 1, 2, 3, 4$, respectively. Dashed lines depict means of their corresponding distributions. The table under the graph reports p-values of a one-sided t-test of the equality of the means of the R^2 distributions.

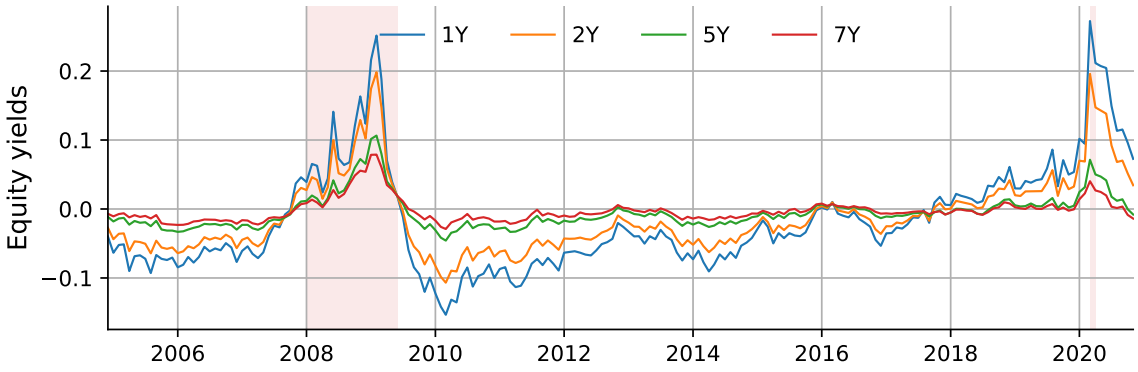

Figure A.22: Out-of-sample dynamics of model-implied yields in the Bansal et al. (2017) **sample**. Equity yields are constructed using the trailing 12-month dividend. Model parameters are estimated in the 1975–2004 sample and held constant throughout the rest of the sample.

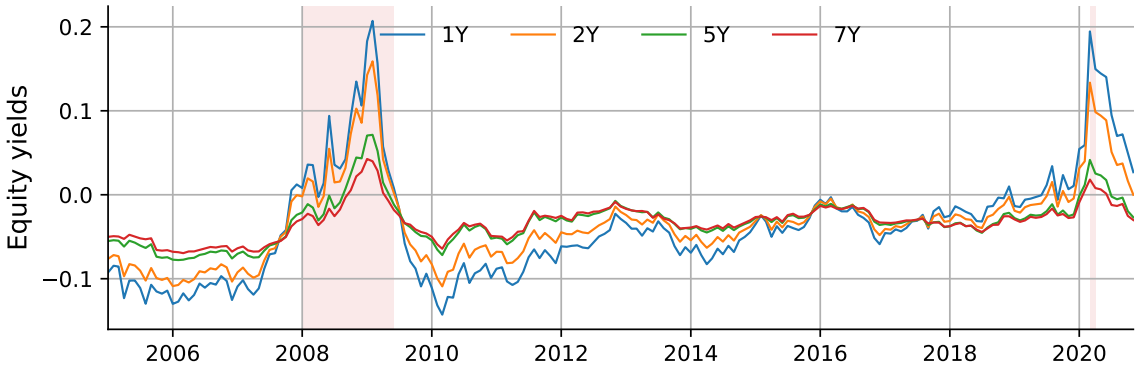

Figure A.23: Dynamics of model-implied forward equity yields for the aggregate market for different maturities (rolling). The figure plots dynamics of model-implied forward equity yields of maturities 1, 2, 5, and 7 years. Equity yields are constructed using the trailing 12-month dividend.

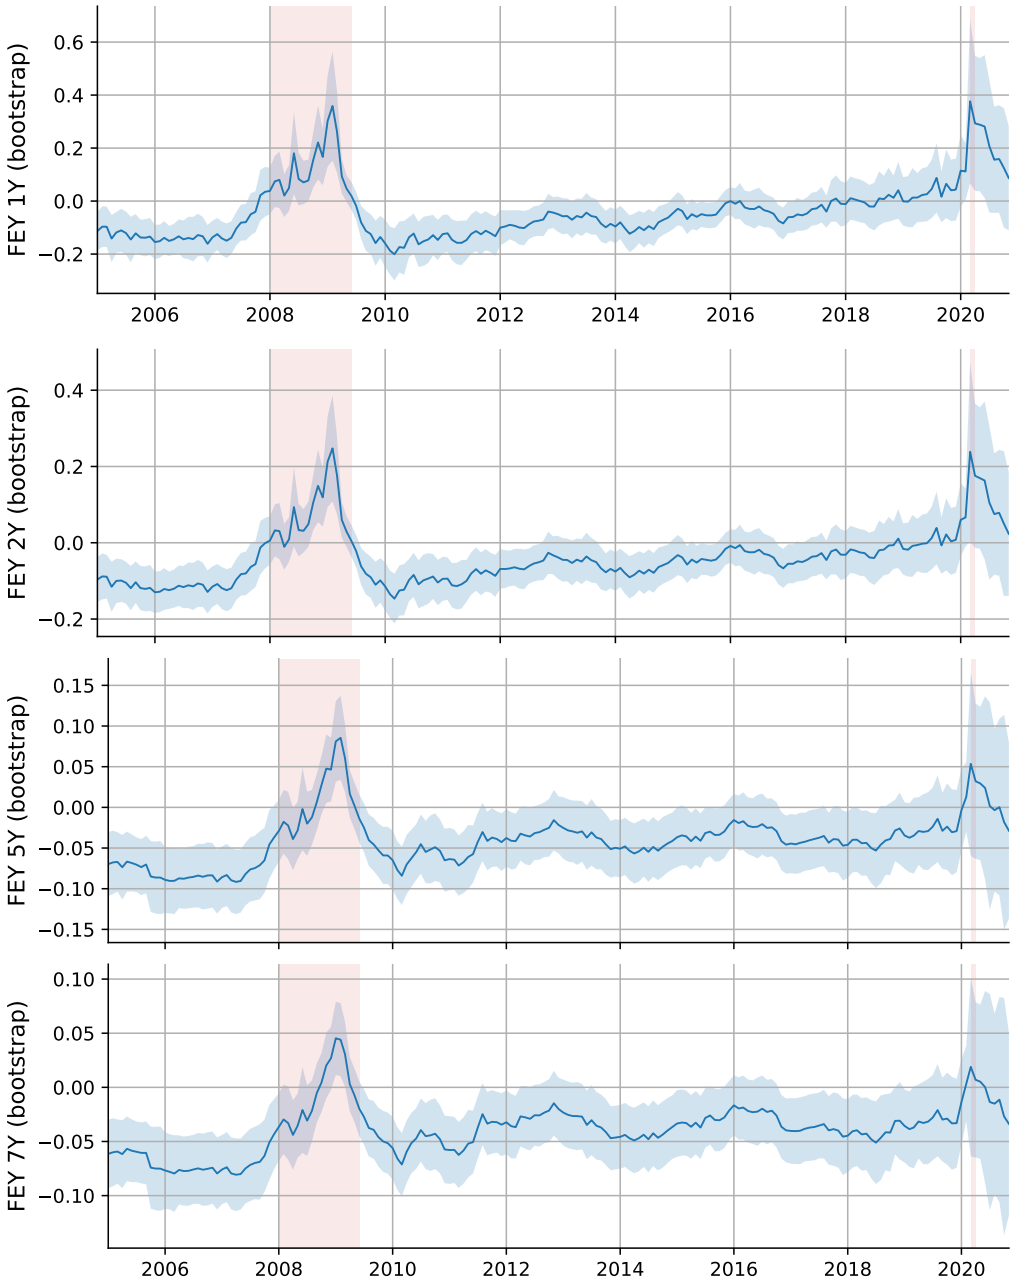

Figure A.24: Model-implied forward equity yields with bootstrapped standard errors. The figure replicates Figure 3 in the main text but uses non-parametric bootstrap based standard errors instead of model-implied GMM standard errors.