STAFF

REPORTS

NO. 1203

SEPTEMBER 2026

The Implied Equity Term Structure

Lieven Baele | Joost Driessen | Tomas Jankauskas

FEDERAL RESERVE BANK *of* NEW YORK

The Implied Equity Term Structure
Lieven Baele, Joost Driessen, and Tomas Jankauskas
Federal Reserve Bank of New York Staff Reports, no. 1203
September 2026
https://doi.org/10.59576/sr.1203

Abstract

We propose a new methodology to estimate the equity term structure. Instead of using realized returns of dividend assets, we generalize the implied cost of capital approach and imply the term structure of ex-ante expected returns from the cross-section of observed stock prices and projected firm-level cash flows. Using US data for 1980-2024, we find an unconditionally upward sloping term structure of risk premia with rich cross-sectional patterns in the size, value and credit risk dimensions. Strikingly, value firms and speculative-grade firms have flat or even downward-sloping term structures. We also detect that the term structure flattens out in recessions.

JEL classification: G10, G12
Key words: cross-section, equity term structure, implied cost of capital

---

Jankauskas: Federal Reserve Bank of New York (email: tomas.jankauskas@ny.frb.org). Baele, Driessen: Tilburg University (emails: lieven.baele@tilburguniversity.edu, J.J.A.G.driessen@tilburguniversity.edu).
The authors thank Frank de Jong, Stefano Cassella, Stefano Giglio, Benjamin Golez, Christian Julliard, Shane Miller, Michael Weber, Isabel Garrigues, and other seminar participants at the 2025 European Finance Association annual meeting, 5th Asset Pricing Conference by LTI@UniTorino, 29th Finance Forum at University of Santiago de Compostela and Tilburg University. They thank Henry Dyer for research assistance.

1 Introduction

An active body of literature studies and measures the expected returns on equity dividend assets with different maturities. The properties of the term structure of equity risk premia have far-reaching implications for at least three key reasons. To begin with, it provides a direct test of competing theoretical asset pricing models, generally implying a flat or positive relationship between an asset's maturity and its expected return. Second, the risk premium term structure is vital for understanding the aggregate cost of capital and net present value (NPV) calculations, which may drive companies' and investors' decisions at the micro level. Third, there is evidence suggesting that many documented anomalies (Cochrane (2011) 'factor zoo'), such as the value premium, may be captured or tamed by the differences in cash flow duration across firms (Dechow, Sloan & Soliman (2004), Weber (2018), Gormsen & Lazarus (2023)). In light of these arguments, it is crucial to obtain a precise measurement of the equity term structure.

In our paper, we propose a new methodology that tackles the inherent empirical challenges of measuring equity risk premia across different maturities. The idea behind the procedure is straightforward: we project expected cash flows and estimate discount rates that bring the model price of a firm's equity as close as possible to its observed market value. In the implied cost of capital literature this is typically done by identifying one discount rate specific to the firm (e.g. Pástor, Sinha & Swaminathan (2008), Li, Ng & Swaminathan (2013)). This is equivalent to assuming no commonality across observations, which could further improve the precision of the estimation procedure. In addition, this approach leaves no degrees of freedom to recover maturity-specific rates. For these reasons, we diverge from the implied cost of capital literature by imposing some similarity across firms based on observable characteristics (e.g., size, book-to-market, credit risk) and leveraging the cross-sectional variation in cash flows and prices to identify maturity-specific risk premia. Intuitively, because different firms have different cash flow patterns across maturities we can identify the term structure of risk premia from the cross-section of firms. As a starting point, we assume that the term structure of risk premia is the same across all firms. This allows us to focus on pricing differences that are purely driven by the timing and size of cash flows. However, we subsequently allow these term structures to vary across firms and over time by flexibly modeling them as functions of time-varying firm characteristics and business cycle variables.

The key value of our method is that it delivers equity term structure estimates over the last 45 years that are neither based on (noisy) realized returns nor assume a specific asset pricing model. We achieve this by combining insights from two strands of the literature: cross-sectional studies on the realized returns of portfolios with different cash flow durations (Dechow et al. (2004), Weber (2018)) and the implied cost of capital literature (Pástor et al. (2008), Li et al. (2013)). The key advantage of the first branch of literature is that the U.S. cross-section of stock prices and fundamentals spans a much longer sample (1980-2024 in our case) compared to the option contracts and dividend strip futures data used in the seminal work of van Binsbergen, Brandt & Koijen (2012) and van Binsbergen, Hueskes, Koijen & Vrugt (2013). The longer sample is particularly valuable for studying the dynamics of the equity term structure over the business cycle. Moreover, the liquidity of the dividend derivative markets may pose problems, especially in short samples (Bansal, Miller, Song & Yaron (2021), Boguth, Carlson, Fisher & Simutin (2023)).[1] The sparse and only recent coverage of firm-level dividend strips also limits thorough cross-sectional estimation of the equity term structure.[2]

However, realized equity returns used in cross-sectional studies face the critique of non-stationarity and large degrees of noise (Fama & French (2002)). Additionally, firm duration may be correlated with other determinants of expected returns, complicating the identification of pure duration effects. For these reasons, rather than averaging realized returns,

[1] Golez & Jackwerth (2024) build on van Binsbergen et al. (2012) by using a longer options sample, longer holding periods for options and option-implied interest rates. They report a flat term structure of aggregate risk premia on average, with countercyclical variation.

[2] Gormsen & Lazarus (2023) are amongst the first to study the interesting market of single-stock dividend futures. Their sample starts, however, only in 2010 and covers only 180 different stocks.

we draw on the implied cost of capital literature and imply the expected returns using constructed cash flow forecasts and observed stock prices. In other words, by assuming a flexible specification for the term structure of risk premiums, we can combine the best of two worlds without explicitly assuming a specific asset pricing model (as in Giglio, Kelly & Kozak (2024)). As shown later, our approach uncovers novel cross-sectional patterns.

We now describe our methodology in greater detail. We augment the mean-reverting cash flow model of Dechow et al. (2004) with IBES analyst forecasts data on sales, return on equity (ROE) and earnings-per-share (EPS). In this manner, we are able to include a richer investors' information set and capture more flexible cross-maturity cash flow patterns than standard mean-reversion models alone. Inserting these forecasts in the present value equation, we are able to estimate the term structure of discount rates by fitting these rates to observed equity values. Here we impose a parsimonious Nelson-Siegel type functional form for the discount rate term structure, commonly applied in the fixed income literature (Nelson & Siegel (1987)), which accommodates both monotonically increasing and decreasing term structures. We also carefully model the structure of the pricing error terms to obtain standard errors for our term structure estimates at each maturity.

Our first main finding is that, when we apply our method to the full sample of stocks, the aggregate term structure of risk premia is upward sloping. However, we also find substantial cross-sectional differences in risk premium term structures. Focusing first on the univariate effects of firm-specific variables, we see that the slope becomes less positive (and even negative in some cases) when firms are smaller, have a higher book-to-market ratio, or have a lower credit rating. Our approach also allows for a joint, multivariate estimation of the effects of firm-specific variables on the equity term structure. When performing such a multivariate analysis, we again find that for strong value firms and firms with speculative-grade credit ratings the term structure is flat or even downward sloping, while the effect of firm size is not very robust. Interestingly, we also see that risk premiums decrease monotonically with the creditworthiness of the issuer, a result which is intuitive, but often undetected in the distress risk literature using realized returns (Campbell, Hilscher & Szilagyi (2008)). In addition, when we estimate the term structure separately for each industry, we find that all industries exhibit upward sloping term structures. Turning to the statistical significance, we find that our slope and long-term risk premium estimates are almost always significant across specifications, while the very short-maturity risk premia (for the first few years) are estimated least precisely. The latter result is not surprising given that the duration of most firms is arguably much longer than just a few years.

These novel results illuminate the role of cross-sectional characteristics not only for the level of expected returns (Fama & French (1993)), but also for the shape of the term structure. This is a powerful insight because the empirical success of asset pricing models may depend on the sample of firms one is looking at. On the one hand, aggregate market risk premia are increasing with maturity as canonical habit (Campbell & Cochrane (1999)) and long run risk models (Bansal & Yaron (2004)) predict. On the other hand, to our knowledge, there is no equilibrium model which could match the cross-sectional patterns that we observe (as well as Giglio et al. (2024) document in their recent work). Thus, our findings reveal a new dimension to analyze and explain.

Our empirical analysis also shows that there is substantial time series variation in the equity term structure slope. In recessions, the term structure remains upward sloping but does become flatter, with short-term discount rates rising and long-term discount rates falling slightly. This is consistent with the results of Bansal et al. (2021) and Giglio et al. (2024), who also detect countercylicality of short term risk premia, at least at the annual frequency.[3] On the other hand, our findings challenge Gormsen (2021) who, at the market level, finds a countercyclical term premium implying that most of the variation is driven by changes in long term rates. In contrast, our estimates reveal that most of the time variation comes from the short and intermediate maturity risk premia. Moreover, we do not find evidence that in

[3]The equity risk premia here is measured using hold-to-maturity returns.

recessions the variation is so large that it can create an inversion of the term structure slope at the aggregate market level.

Our approach strongly relies on projections of firm's future cash flows. We therefore perform a range of robustness checks on these cash flow projections and show that our qualitative findings persist in a wide array of robustness tests and alternative specifications. First, using results of van Binsbergen, Han & Lopez-Lira (2022) and Zhang, Zhu & Linnainmaa (2024), we implement estimates of analyst forecast biases to correct for documented optimism in IBES consensus estimates. Second, we consider time-varying long-term cash flow growth rates proxied by the Survey of Professional Forecasters' (SPF) predictions on long-term inflation and real GDP growth rates. Third, we average cash flow forecasts across different models, cap them, and exclude the most extreme size and book-to-market portfolios to alleviate the possible impact of outliers. None of these extensions affect the upward sloping pattern of the aggregate market equity risk premia. Fourth, we also perform the estimation with alternative values for cash flow mean reversion parameters and different forecasting horizons. We see that for moderate to high levels of mean reversion, which includes both our mean reversion estimates and those of Dechow et al. (2004) and Weber (2018), the recovered term structure is upward sloping. Only with very slow mean reversion, much slower than what we see in the data, the slope can flip sign.

Our paper contributes to the empirical asset pricing literature aiming to measure the shape of the equity term structure. Though we find empirical support for a negative slope among small value firms and firms with low credit ratings, on aggregate we find upward-sloping term structures, consistent with the evidence in Bansal et al. (2021). This carries substantial ramifications for theoretical research since it finds suggestive evidence supporting the classical theoretical asset pricing models of Campbell & Cochrane (1999) and Bansal & Yaron (2004), at least for the aggregate market. On the other hand, our results challenge the current frameworks in the cross-sectional dimension as, to our knowledge, there is no microfounded model that can generate different equity term structures across firms sorted on book-to-market, size and credit rating.

The closest study to our paper is Giglio et al. (2024). They also estimate term structures in cross-sectional portfolios, though they focus mostly on equity yield term structures, which are a combination of risk premiums and expected dividend growth rates. At the stock index level they also document an upward sloping term structure of risk premia. They find that the SMB portfolio has a downward sloping term structure of risk premia, which is consistent with our results. For the HML portfolio they report an upward sloping term structure, in contrast to our results for value firms. Though our findings thus are consistent with their results in most cases, their paper is based on a fundamentally different approach. Giglio et al. (2024) assume a reduced form model for the pricing kernel (in the light of Lettau & Wachter (2007)) and fit its parameters to the principal components of realized returns. In contrast, we make assumptions about cash flows and imply ex-ante expected returns without assuming a specific asset pricing model, using a flexible empirical specification for the discount rate term structures instead. In this way, our methodology provides new degrees of freedom and a perspective that contributes to the ongoing debate in the equity term structure literature.

Another related paper is Ulrich, Florig & Seehuber (2023). They use survey forecasts for dividends and option-implied dividends to construct an ex-ante term structure of risk premia and report a flat term structure on average, while it becomes negative during contractions. Their focus is on the S&P 500 index level and relatively short maturities, up to three years. We complement this work by studying the cross-section of term structures and focusing on longer-term discount rates.

Our paper is also different from the studies that rely on the realized returns of duration portfolios (Weber (2018), Gormsen (2021), Gonçalves (2021)), given that we focus on ex-ante returns. Other related work includes Callen & Lyle (2020) who combine option prices and the CAPM to estimate discount rates at different maturities at the firm level. Consistent with our results, they report largely upward sloping term structures. Schröder (2024) shows that, if one assumes that the term structure of risk premia is flat at the firm level, the aggregate term structure may be upward sloping due to aggregation effects. We find empirically that firm-level term structures can be upward or downward sloping depending on firm characteristics.

The structure of the paper is as follows. Section 2 describes the identification strategy and methodology. Section 3 presents the data, whereas section 4 elaborates on the construction of expected cash flows. Section 5 outlines the results and key comparisons with the literature. Section 6 contains robustness checks, and section 7 contains a summary and final remarks.

2 Implying Term Structures of Discount Rates: A New Methodology

The starting point of our estimation approach is that, at time t, the value of firm i's equity $P_{i,t}$ can be described by the discounted cash flow equation:

$$\sum_{m=1}^{\infty} \frac{\mathbb{E}_t[CF_{i,t+m}]}{(1+r_{i,t,m})^m} = P_{i,t} \tag{1}$$

where $CF_{i,t+m}$ is the free cash flow accruing to the company's shareholders, $P_{i,t}$ is the market value of equity, and $r_{i,t,m}$ is the maturity-specific and firm-specific (zero-coupon) discount rate. We assume that after T years the firm reaches its steady-state where expected cash flows grow at the common growth rate of g_t, and $r_{i,t,m}$ converges to $r_{i,t,T}$. In other words, both the term structure of cash flow growth rates and discount rates are flat beyond this maturity. Given this assumption, one can simplify this infinite sum to:[4]

$$\sum_{m=1}^{T} \frac{\mathbb{E}_t[CF_{i,t+m}]}{(1+r_{i,t,m})^m} + \frac{\mathbb{E}_t[CF_{i,t+T+1}]}{(1+r_{i,t,T})^T (r_{i,t,T} - g_t)} = P_{i,t} \tag{2}$$

In the implied cost of capital literature one uses the observed equity market values $P_{i,t}$, combined with estimates of expected cash flows, to obtain a firm-specific discount rate $r_{i,t}$ from equation (2). However, this leaves no degrees of freedom to identify maturity-specific discount rates. We address this limitation by imposing a key identifying assumption: the discount rate curve $\{r_{i,t,m}\}$ is a parametric function of maturity m and observable characteristics $\boldsymbol{Z_{i,t}}$ that can vary over time and/or across firms:

$$r_{i,t,m}(\boldsymbol{\beta}) = r^f_{t,m} + f(\boldsymbol{\beta}, \boldsymbol{Z_{i,t}}, m), \tag{3}$$

where $\boldsymbol{\beta}$ is a vector of parameters, and $r^f_{t,m}$ captures the term structure of risk-free rates. By parametrizing discount rates and imposing some common structure across firms, we extend the implied cost of capital approach along both the maturity and cross-sectional dimensions. Now we can estimate model parameters $\boldsymbol{\beta}$ jointly by minimizing the squared relative pricing error between the firm's total market capitalization and model valuation. Specifically, let us denote the relative pricing error for firm i at time t as:

$$\epsilon_{i,t}(\boldsymbol{\beta}) = \sum_{m=1}^{T} \frac{\mathbb{E}_t[CF_{i,t,t+m}]/P_{i,t}}{(1+r_{i,t,m}(\boldsymbol{\beta}))^m} + \frac{\mathbb{E}_t[CF_{i,t,t+T+1}]/P_{i,t}}{(1+r_{i,t,T}(\boldsymbol{\beta}))^T (r_{i,t,T}(\boldsymbol{\beta}) - g_t)} - 1 \tag{4}$$

Then, we estimate $\boldsymbol{\beta}$ by minimizing the sum of the squared pricing errors over the panel of firm-level observations across firms and time:

$$\underset{\boldsymbol{\beta}}{argmin} \sum_t \sum_i w_{i,t} \epsilon^2_{i,t}(\boldsymbol{\beta}) \tag{5}$$

[4] More formal derivations are included in Appendix E.

where $w_{i,t}$ is a weight assigned to firm i in year t.

The main estimation criterion is how well the chosen risk premia parameters fit the panel of stock values. This estimation procedure can be seen as a joint calibration of model values to market prices where observations with the largest weight or highest pricing deviations attain more importance. Intuitively, the identification of discount rates of different maturities relies on the cross-maturity variation of cash flows within the cross-section of firms. Some firms have cash flows that are concentrated at shorter maturities while other firms have a longer duration of cash flows. This variation helps to identify the slope of the risk premium term structure. By including firm characteristics we control for other forms of discount rate variation across firms and time.

We now discuss the functional form of the discount rate term structure. As we see in (3), we take into account the term structure of risk-free rates, $r^f_{t,m}$, which allows us to focus directly on the level of risk premia rather than total discount rates. It also implies that total discount rates will naturally vary over time as interest rates change. To accommodate various possible shapes, the functional form of the risk premia term structure, $f(\boldsymbol{\beta}, \boldsymbol{Z_{i,t}}, m)$, is specified flexibly. However, with a highly non-linear optimization we want to avoid an excessive number of parameters. Since the m period cash flow is not expected to be much different from the $m+1$ period cash flow, the discount rates for similar maturities will be strongly correlated. For this reason, we choose the Nelson-Siegel (NS) functional form (Nelson & Siegel (1987)), often used to model fixed income yield curves, as it is both sufficiently rich in capturing different discount rate curve patterns and parsimonious in the number of parameters:[5]

$$f(\boldsymbol{\beta}, \boldsymbol{Z_{i,t}}, m) = \underbrace{(\beta_1 + \boldsymbol{\beta_2'} \boldsymbol{Z}_{i,t})}_{\text{level}} - \underbrace{(\beta_3 + \boldsymbol{\beta_4'} \boldsymbol{Z}_{i,t})}_{\text{slope}} \varphi(m), \tag{6}$$

[5] In this version, we do not include the hump-term as it is known to be highly correlated with other parameters. Thus, the term structure of risk premiums is either monotonically increasing or decreasing - an assumption which can be relaxed in future work.

where $\varphi(m) = \frac{1-e^{-m/\lambda}}{m/\lambda}$ is a decay function. In equation (10) β_1 and β_3 are scalars, while $\boldsymbol{\beta_2}$ and $\boldsymbol{\beta_4}$ are vectors with length equal to the number of firm characteristics, and $\boldsymbol{\beta}$ is the vector containing all these parameters. Most importantly, the specification allows both the level and the slope of the term structure to vary with the firm-specific variables $\boldsymbol{Z}_{i,t}$.

As discussed above, for maturities $m > T$, the total discount rates are assumed to be constant, i.e. for all $m > T$ we have $r_{i,t,m} = r_{i,t,T}$. Finally, the parameter λ in $\varphi(m)$ governs the curvature of the equity term structure and is known to be difficult to estimate in interest rate applications. For this reason, we fix λ to economically plausible values throughout the analysis. In our robustness tests, we show that the curvature does not affect the qualitative sign or statistical significance of the slope.

The estimation procedure described in equation (5) is equivalent to weighted non-linear least squares. Therefore, the standard errors can be computed numerically following Greene (2000). We correct these standard errors for heteroscedasticity and take into account time-series and cross-sectional correlation in the pricing errors. Subsequently, we use the delta method to compute the standard errors for each maturity point. More details are provided in Appendix D. To assess the sensitivity of our results to the expected cash flows we perform various robustness checks on the way these cash flow forecasts are constructed in Section 6.

Finally, to relate our framework to the literature on dividend strips, we can apply equation (1) to a single cash flow, $CF_{i,t+m}$. We then see that $r_{i,t,m}$ equals the expected return from holding this strip asset to its maturity date, $t + m$. Existing work on realized returns on dividend strips typically focuses on returns over short holding periods. As we show later, as long as the variation in discount rates is not too large, there is a tight relation between our hold-to-maturity risk premia and risk premia over short holding periods (additional details are provided in Appendix A).

3 Data

The analysis is performed using US data over the 1980-2024 period. Even though all required variables are observable as of 1976, we start in 1980 since from that year onwards we have a cross-section of more than 500 firms per year. This restriction is important, as our approach relies on substantial cross-sectional variation, particularly given that we allow cross-sectional characteristics to affect the term structure of risk premia. The balance sheet information comes at the yearly frequency. Stock price data and balance sheet data, as well as historical credit ratings and S&P500 index composition, are obtained from Compustat-CRSP and FISD Mergent.[6] Supplementary book equity data is collected from Kenneth French's website. The median analyst forecasts are from Thomson Reuters I/B/E/S summary database, whereas historical constant maturity (zero-coupon) yields are from Gürkaynak, Sack & Wright (2007). Historical GDP and NBER recession periods can be accessed from St. Louis FED, and the long-term inflation and real GDP growth expectations come from the survey of professional forecasters collected by the Philadelphia FED.

For our estimation approach, we use an annual panel of equity values with June as valuation month. We construct book equity measures following Davis, Fama & French (2000). We measure book equity (and other historical accounting variables) using the most recent financial year results prior to the valuation month (June). Similarly, the market cap is calculated as price times the number of shares outstanding at the end of the valuation month. Both prices and shares outstanding come from CRSP. June is chosen as the valuation month because it guarantees that most of the financial statement information has already arrived to the market, and analyst forecasts are mostly covering the upcoming financial years. To reduce potential extremities, we keep only those I/B/E/S median forecasts that were provided by more than 1 analyst. Following Weber (2018), utilities (4900$<=$ SIC$<$5000) and

[6] We use CRSP-provided ratings up to 2017, and from 2018 onward, we use issuer-level credit ratings from FISD Mergent, proxied by the median corporate bond rating assigned by S&P, Moody's, and Fitch (in that order of priority).

financials (6000<=SIC<7000) are dropped from the sample, as well as firms that are not listed on one of the major stock exchanges (NYSE, AMEX and Nasdaq). Finally, since we value the total equity size of each firm, we exclude all mismatches larger than 5% between Compustat and CRSP database market caps.[7] Further, we winsorize the sample at 2.5% and 97.5% based on the model implied internal rate of return of a stock. This step effectively excludes firms with extremely high or low valuation ratios (for instance, companies in severe financial distress) that are challenging to value with standard fundamental value models. In this way, without much loss of generality (most stocks excluded are smaller with extreme ROE and sales growth rates), we reduce the impact of outliers. In the robustness section we show that we obtain similar results with less winsorization. The filters leave us with 97,032 Compustat-CRSP firm-year observations over the 1950-2024 period. After merging with I/B/E/S, the final sample consists of 56,126 firm-year observations over the period 1980-2024. More details on the filters and sample construction are provided in the Appendix B.

In terms of summary statistics, the final sample is representative of the broader Compustat-CRSP dataset. As reflected in Table I, the estimation sample consists of larger firms (with median market cap of USD 776 mln. as compared to USD 209 mln. in Compustat-CRSP) as larger firms tend to be more frequently covered by analysts. Nevertheless, financial ratios, such as ROE or dividend payouts, have almost identical distributions in both samples. The only exception is the BM ratio which is slightly lower in the final sample (median of 0.41 vs 0.48) because larger firms tend to have higher valuation ratios. Since most financial ratios have large outliers due to small denominators, we winsorize ROE (including sales growth) at the 5% and 95% level on a yearly basis.

[7] Differences in end of the year market caps between Compustat and CRSP records may lead to a severe mispricing in our model. For this reason, we exclude all observations with large discrepancies that potentially indicate multiple classes of traded or non-traded shares.

Table I: Summary statistics

The table compares the Compustat-CRSP-I/B/E/S sample used for the analysis with a broader Compustat sample obtained after standard filtering procedures (as in Weber (2018)). Sales, book and market equity are presented in millions of USD. ROE, sales growth, book-to-market, earnings-per-share and dividend payout ratio are all presented in units. The dividend payout ratio is calculated similarly as in Li et al. (2013) - the total dividend paid divided by net income (in case it is positive). The ROE is defined as current earnings before extraordinary items divided by lagged book equity. To alleviate outliers, ROE and sales growth are winsorized at 5 and 95% levels.

Compustat 1980-2024						
	p5	p50	p95	mean	sd	N
Sales	3.7	180	8,931	2,365	12,255	97,932
Book-Equity	4.2	96	4,017	1,194	6,704	97,932
Market-Equity	8.7	209	12,215	3,892	33,801	97,932
Sales-Growth	-0.24	0.09	0.78	0.15	0.30	90,149
ROE	-0.62	0.08	0.37	0.02	0.28	89,592
B/M	0.08	0.48	1.83	0.68	0.89	97,932
E/P	-0.45	0.03	0.15	-0.06	0.80	97,932
Payout	0.00	0.00	0.72	0.13	0.25	97,932
IBES-Compustat 1980-2024						
Sales	17	594	15,301	3,982	15,982	56,126
Book-Equity	24	310	7,503	2,027	8,758	56,126
Market-Equity	59	776	23,780	6,688	44,441	56,126
Sales-Growth	-0.22	0.10	0.77	0.16	0.29	54,010
ROE	-0.57	0.11	0.39	0.06	0.27	53,717
B/M	0.08	0.41	1.39	0.55	0.63	56,126
E/P	-0.26	0.04	0.13	-0.02	0.47	56,126
Payout	0.00	0.00	0.76	0.15	0.25	56,126

4 Cash flow forecasts

In this section we describe how we construct expected cash flows across maturities for all firms. These are key inputs to our estimation methodology. Following Dechow et al. (2004) and others, we deduce the dynamics of the cash flow $CF_{i,t}$ of firm i in year t from the identity:[8]

$$CF_{i,t} = \underbrace{E_{i,t}}_{\text{earnings}} - \underbrace{(BE_{i,t} - BE_{i,t-1})}_{\text{reinvestments}} = BE_{i,t-1} \times (ROE_{i,t} - \%BE_{i,t}) \tag{7}$$

where BE stands for the book value of equity, $\%BE$ denotes its growth rate, and ROE is the return on equity. Since Nissim & Penman (2001) document that the growth rate of sales is a better predictor for book equity growth than most other accounting variables, we follow Weber (2018) and proxy $\%BE$ by the growth rate of sales.

In order to forecast ROE and sales growth, our starting point is the common approach in the literature of assuming mean-reverting processes and use those projections to forecast cash flows (Dechow et al. (2004), Weber (2018), Chen & Li (2018)). It is a parsimonious but restrictive approach as all variation in cash flows is attributed to initial differences across firms which dissipate at the rate of reversion parameters. As discussed below, we therefore also incorporate analyst forecasts to the first years of cash flow predictions (similarly to Claus & Thomas (2001), Gebhardt, Lee & Swaminathan (2001), Li et al. (2013)). This allows for more flexible patterns arising from different firm life cycles, and for more accurate depiction of investor beliefs. Needless to say, the analyst projections also alleviate the concern that our results are highly correlated with the initial book-to-market ratios (as in Dechow et al. (2004) and Weber (2018)) and provide additional cross-sectional variation in the estimation of long- versus short-term discount rates. Of course, analyst forecasts may carry some biases or additional noise too, and for this reason we later investigate specifications where cash flow forecasts are averaged across different approaches and also explicitly correct for analyst biases. We find that the main results remain intact.

We also pay special attention to the estimation of the mean reversion parameters of the cash flow model. As discussed below in more detail, we show that matching the longer-

[8]To be precise, this identify is true when clean surplus accounting holds. To a large extent, this is empirically supported (Dechow et al. (2004)).

term dynamics of cash flows, instead of the short-term dynamics, leads to higher persistence parameters. This is relevant because for valuing firms the long-term dynamics of cash flows are most important.

In sum, based on the arguments mentioned above, we model cash flows in two stages. First, for the 1-5 year horizon, depending on the data availability, we extract information from the median analyst forecasts. Second, starting from the last year of analyst data availability until period T, we assume the same mean-reverting process as in Dechow et al. (2004) and Weber (2018), calibrated to long-term cash flow dynamics. The key modeling choices of these two stages are discussed below.

Stage one: Analyst forecasts

The estimation procedure benefits from the short- and medium-horizon forecasts available from IBES by introducing additional cross-firm variation. However, a key empirical challenge is that IBES forecasts are primarily available for earnings per share (EPS), while forecasts for return on equity (ROE) are comparatively scarce.[9] To mitigate these data limitations, we construct implied ROE forecasts based on the median EPS forecasts, which almost triples our sample and data range. Assuming that the number of shares remains fixed, we imply ROE forecasts by dividing median EPS forecasts by forecasted book equity (BE) per share. In a similar manner, we derive revenue growth projections from the levels of available revenue forecasts. The complete set of projection procedures is detailed in Appendix B.

This procedure results in a wide heterogeneity across firms in their ROE and sales growth patterns (Table II). This is particularly valuable in the context of ROE forecasts, as ROE empirically exhibits slower mean reversion than sales growth, implying that much of the cross-sectional variation in cash flows can be attributed to differences in ROE. A potential concern is that the implied forecasts may not accurately reflect investors' actual expectations.

[9] The availability of ROE forecasts is limited to 2001-2014 period with relatively few instances in the early years. The level of revenue forecasts is accessible starting from 1992 with much larger coverage across firms. Yet, EPS forecasts date back to 1976 with the largest coverage of over 60,000 observations.

However, the correlation between implied and actual ROE forecasts is high, ranging from 70% to 80% at shorter horizons, and remains around 55% even at the five-year maturity. Moreover, when comparing the distribution of ROE forecasts with implied ones (Table II), one can notice similar means and medians with slightly larger standard deviations and tails of the distribution. Overall, it seems that the procedure delivers a good approximation, however, tighter filters and forecast averaging across alternative models may help to mitigate some noise and reduce the influence of outliers (see Section 6).

Table II: Analyst forecasts

The table below summarizes the implied sales growth, ROE and actual ROE forecasts in our merged Compustat-CRSP-I/B/E/S sample over 1980-2024. The implied sales growth forecasts are calculated as the growth rate in median I/B/E/S sales forecasts. The implied ROE forecasts are constructed by dividing median EPS forecasts by projected book value of equity (assuming that it grows at a sales growth rate). Horizon specifies the future financial year for which the forecasts are recorded. To alleviate outliers, all forecasts are winsorized at a 5% and 95% level.

	p5	p50	p95	mean	sd	N
Implied sales growth forecasts						
1y	0.01	0.10	0.66	0.16	0.17	34,061
2y	-0.01	0.09	0.70	0.17	0.20	17,935
3y	0.00	0.12	1.27	0.27	0.34	8,244
4y	0.00	0.10	1.80	0.33	0.50	6,258
ROE forecasts						
1y	-0.09	0.15	0.43	0.17	0.13	14,018
2y	-0.02	0.16	0.43	0.18	0.11	13,802
3y	-0.01	0.18	0.55	0.21	0.13	8,262
4y	0.07	0.18	0.46	0.20	0.10	695
5y	0.08	0.18	0.49	0.21	0.11	408
Implied ROE forecasts						
1y	-0.49	0.11	0.55	0.10	0.22	56,111
2y	-0.25	0.13	0.66	0.15	0.19	52,397
3y	-0.03	0.15	0.82	0.20	0.20	39,115
4y	0.01	0.16	1.00	0.24	0.24	34,147
5y	0.02	0.18	1.10	0.26	0.26	32,779

Stage two: Mean-reverting process

Starting from year 6 or the last available horizon of IBES forecasts, we assume similar mean-reverting processes for ROE and sales growth as in Dechow et al. (2004) and Weber (2018):

$$\begin{aligned} ROE_{i,t} &= 0.12 + 0.81(ROE_{i,t-1} - 0.12) \\ BE_{i,t} &= 0.06 + 0.40(BE_{i,t-1} - 0.06) \end{aligned} \tag{8}$$

Both AR(1) coefficients (0.40 for BE and 0.81 for ROE) are chosen based on a pooled OLS regressions using 3-year lags in the full Compustat sample over 1964-2024 (the fourth row of Table III). The equations state that ROE converges to a long-run level of 12%; $\%BE$ is proxied by sales growth, which also reverts to the assumed long-term nominal GDP growth rate of 6% (g). We follow Dechow et al. (2004) and Weber (2018) when choosing the levels of 12% and 6%.

The key difference here with previous research is that we employ different regressions that lead to much higher reversion coefficients (in other words, much slower mean reversion). Dechow et al. (2004) find a coefficient for ROE equal to 0.57, and 0.24 for sales growth; Weber (2018) obtains 0.41 for ROE and 0.24 for sales growth. Both studies estimate the annual AR(1) coefficients using annual pooled OLS regressions over the full sample starting in 1964. Yet, these estimates are sensitive to filters applied and, as illustrated in Figure C.I for the ROE and sales growth quintile portfolios, they appear biased downwards due to the very quick initial mean-reversion and slow decay afterwards. In fact, similar pooled regressions where we use more distant lags of 3 to 5 years to estimate the AR(1) coefficient or tighter winsorization of the sample deliver slower mean reversion estimates, in the range of 0.73-0.85 (0.40-0.54) for ROE (sales growth) (Table III).[10] Since the long term behavior of firm's cash flows is most important for valuing firms, and with short-term patterns mostly

[10] We also explored specifications with multiple lags and moving averages, however, the additional terms often were not statistically significant. For this reason, we keep the parsimonious AR(1) model as our base specification.

captured by the analyst forecasts, we conclude that slower mean reversion is most plausible for our application.

Table III: AR coefficient estimates

The table presents the estimated AR coefficients from pooled OLS regressions using the full Compustat sample over 1964-2024 (first five rows) and merged Compustat-CRSP-I/B/E/S sample over 1980-2024 (last three rows). The second and third rows describe regressions where both ROE and sales growth are winsorized at 5% and 95% levels and ROE is capped at -100% (some firms incur losses exceeding their book value of equity). Rows four, five, seven and eight present specifications where the AR coefficient is estimated using more distant lags. All AR coefficients are standardized to the annual level. The standard errors (SE) are clustered by year. For longer lag regressions, the delta method is applied to compute standard errors.

	ROE				Sales growth			
	AR(1)	SE	R^2	N	AR(1)	SE	R^2	N
Full Compustat sample	0.01	0.00	0.00	237,810	-0.00	0.00	-0.00	247,049
Winsorize (5% and 95%)	0.74	0.00	0.49	237,810	0.22	0.00	0.05	247,049
Cap ROE at -100%	0.73	0.00	0.50	237,810				
3 years lag	0.81	0.01	0.25	193,938	0.40	0.02	0.00	203,787
5 years lag	0.85	0.01	0.17	161,597	0.54	0.02	0.00	170,447
Merged sample	0.73	0.00	0.52	42,425	0.25	0.00	0.07	42,635
3 years lag	0.81	0.02	0.28	33,165	0.50	0.03	0.02	33,458
5 years lag	0.85	0.01	0.19	27,064	0.59	0.03	0.01	27,351

In the robustness section we also consider alternative autoregressive coefficient magnitudes. While small perturbations do not affect the results much, larger shifts in mean reversion can indeed alter the qualitative conclusions.

5 Empirical Results

5.1 Benchmark setting

The general model presented in Section 2 necessitates the specification of several key parameters. First, following Dechow et al. (2004), Li et al. (2013), and Weber (2018), in our benchmark model we set $T = 15$ years, which implies that the term structures of growth rates and risk premia are flat beyond this maturity point. We also test the robustness of our results to this assumption, varying T from 10 to 50 years. Second, we normalize the equity term structure slope as the difference between risk premia at 15-year and 5-year maturities. We do this because shorter-term risk premia (maturities of 1 to 4 years) are typically more difficult to estimate with precision given that the duration of firms' cash flows typically varies between 5 years and 40 years. We thus rewrite, without loss of generality, the term structure specification as follows:

$$f(\beta, Z_{i,t}, m) = (\beta_1 + \boldsymbol{\beta_2'} \boldsymbol{Z}_{i,t}) - (\beta_3 + \boldsymbol{\beta_4'} \boldsymbol{Z}_{i,t}) \frac{\varphi(m) - \varphi(15)}{\varphi(5) - \varphi(15)}, \tag{9}$$

This specification is normalized in such a way that β_1 reflects the 15-year risk premium and β_3 represents the 15- minus 5-year risk premium (if $\boldsymbol{Z_{i,t}} = \mathbf{0}$). Third, we follow Dechow et al. (2004) and Weber (2018) and set g_t to 6% in equation (2), which equals the long-term average nominal GDP growth rate. As shown in Robustness Section 6.2, allowing for time variation in g does not qualitatively affect the results. Fourth, we adopt a value-weighting scheme for the estimation in equation (5), in order to focus the estimates on economically significant firms and to mitigate the influence of large pricing errors from smaller firms on the estimation procedure. At the same time, we also impose equal weighting across all observation years, thereby ensuring that later years - with typically higher market values - do not unduly influence the results. We do test the robustness of our results to alternative weighting schemes below. Finally, for the parameter λ that governs the curvature of the term structure function, we choose $\lambda = 4$ as this value provides sufficient curvature to ensure a smooth convergence to the long-run levels. In Appendix C, we verify that our results remain similar when we use other values for λ (Figure C.II).

5.2 Market-wide estimates

Given that many earlier studies of the equity term structure have focused on the market level, we begin by estimating our model on the full firm-level panel covering the entire sample period (1980-2024). We allow for a time-varying risk-free rate term structure but assume (for now) the same equity risk premium term structure for all firm-year observations.

Figure I presents the implied term structure alongside its 95% confidence interval. We find strong evidence for an upward sloping equity term structure, representing the first key result in this paper. The long-term risk premium flattens out and approaches a constant level of 8.3% at a 15-year horizon. Figure I also highlights that our benchmark specification implies a positive slope of 7.5% per year, which is economically and statistically significant. As shown below, the qualitative finding of a positive slope is robust across many alternative specifications, suggesting that, all else equal, investors have a preference for shorter maturity cash flows.

The result carries similar implications for the term structure of one-year ahead expected returns. Over long samples, average holding-to-maturity returns converge to average one-year-ahead returns on different maturity assets. We illustrate this idea in Appendix A where we demonstrate that time-invariant discount rates imply that the one-period risk premium term structure equals the forward risk premium term structure derived from our holding-to-maturity premiums.[11] In our sample, as expected, the forward rate term structure is steeply upward sloping, driven mostly by the low estimated risk premia at short maturities. Thus, our baseline findings suggest similar conclusions for the term structure of risk premia over shorter holding periods.

At first glance, our finding contradicts the downward-sloping term structure reported in van Binsbergen et al. (2012). However, their analysis is restricted to the very short end of the equity term structure, as their measure is based on options with maturities of up to two years only. In contrast, our approach is better suited for estimating risk premia at intermediate and

[11] One can show that in settings with fully time-varying discount rates, additional assumptions on the dynamics of risk premia are required. However, over long samples, most of these transitory effects average out or become second order effects. In other words, it takes very volatile term structures and substantially different covariances of short- and long-term cash flows with discount rates to construct examples in which the one-year-ahead term structure is downward sloping while the hold-to-maturity term structure is upward sloping.

long-term horizons, which—due to their longer average duration—are likely a more dominant driver of firm valuation. In fact, our estimates of short-term risk premia are negative and associated with large confidence bands, mainly because they are mostly extrapolated (using the monotonic Nelson–Siegel specification) from the intermediate maturities. In principle, our findings do not rule out an U-shaped term structure: one that features declining premia at very short maturities and increasing premia beyond the intermediate horizon. Note, however, that some of the more recent studies also focusing on short-horizon term structures found either flat (Ulrich et al. (2023)) or upward sloping (Bansal et al. (2021), Callen & Lyle (2020)) term structures. Finally, we also rule out that our contradicting results are due to a different sample and time period: when we restrict the sample to the 1996-2008 period and to constituents of the S&P 500 index, which is what van Binsbergen et al. (2012) focus on, we still find an upward sloping term structure (Appendix Figure C.IV).[12]

Figure I: Full-sample term structure of equity risk premia

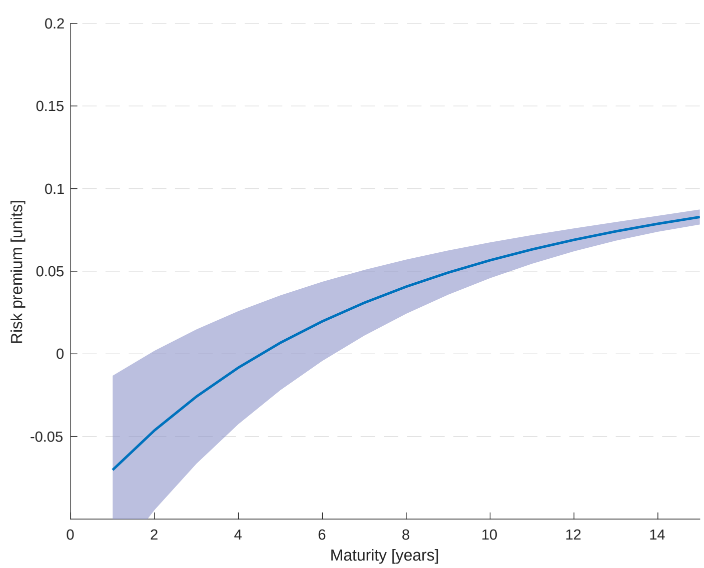

The graph depicts the estimated equity risk premium term structure over the full sample for the benchmark value-weighted case. Compustat-CRSP-I/B/E/S sample: 1980-2024.

[12]The match between samples is not perfect since, on average, we have only 250-300 S&P500 firms in our sample. The main reason is that we exclude financials and utilities, which constitute roughly 150 firms of the index.

As discussed, our estimates of risk premia for the first few annual maturity points are the least precise and are negative in our benchmark estimation. However, these estimates are not very influential for equity valuation, since most cash flows occur at intermediate and long maturities. To illustrate this, we perform two restricted estimations of our benchmark specification. First, we impose that the term structure of risk premiums is flat for the first five years. In Figure C.III we see that in this case the risk premium is very close to zero for the first five years, and then largely follows the pattern of the unrestricted estimates for longer maturities. In the second alternative estimation approach we constrain risk premia to be nonnegative at all maturity points. In this case we again find an upward sloping term structure, starting at 0% and increasing to a level at 15-year maturity that is just below the benchmark estimate, to compensate for the higher premia at shorter maturities. In sum, while the negative risk premia at short maturities might seem unusual, they are economically not very important for our analysis and restricting these premia to be nonnegative does not alter our conclusions.

Lastly, while the level of the long-term risk premium might appear high, it is tightly linked to the assumed long-term cash flow growth rate of 6%. Robustness analysis in Section 6 shows that the level parameter of term structure moves approximately one-to-one with the long-term growth rate, without substantially affecting the slope. For example, in Table VII, we obtain a long-term risk premium of 5% when the long-term growth rate is set to 3%. The slope becomes slightly flatter, and short-term risk premia turn slightly positive when expected growth rates are allowed to vary over time—measured using historical moving-average GDP growth rates or professional forecasters' data on real GDP and inflation. This evidence pushes us to investigate time-series patterns over the business cycle in greater depth.

5.3 Cyclical Variation

We test for cyclical variation in the risk premium term structure by making the long-term risk premium and slope a linear function of a bad times dummy B_t^z:

$$r_{t,m} = rf_{t,m} + (\beta_1 + \beta_2 B_t^z) - (\beta_3 + \beta_4 B_t^z)\frac{\varphi(m) - \varphi(15)}{\varphi(5) - \varphi(15)} \tag{10}$$

Our first proxy for bad times is a yearly dummy that equals one when at least three months within that year are classified as an NBER recession. The estimates in column 2 of Table IV show that the slope flattens significantly in bad times (from 7.5% in good to 3.8% in bad times). At the same time we see a moderate drop in long-term discount rates, from 8.3% to 6.3%. As illustrated in Figure II, this pattern implies that, during recessions, short- and medium-term discount rates increase markedly, whereas the long-term risk premium declines, though not to a degree that would invert the term structure. That is, even in downturns the equity term structure remains upward sloping. Columns 3 and 4 show that these patters are robust to using two alternative bad times indicators, namely a dummy equalling one when the US equity market return over that year was in the bottom quartile of its historical distribution (column 3 and Figure II) or when at least 25% of days within the year belong to a drawdown period, as identified by Baele & Bekaert (2025) (column 4).

Table IV: Risk premia parameter estimates

The table presents level and slope parameter estimates based on the modified Nelson-Siegel form:

$$r_{t,m} = rf_{t,m} + (\beta_1 + \beta_2 B_t^z) - (\beta_3 + \beta_4 B_t^z)\frac{\varphi(m)-\varphi(15)}{\varphi(5)-\varphi(15)},$$

where B_t^z is a recession dummy, and $\varphi(m) = \frac{1-e^{-m/\lambda}}{m/\lambda}$ is a decay function normalized in a way that β_1 reflects the 15-year risk premium and β_3 reflects the 15 minus 5-year risk premium. Bad Times represent either (i) years with at least three recession months classified by NBER, (ii) years where US Equity Market return was in the bottom quartile of historical distribution (EBQ), or (iii) years where at least 25% of days are in a drawdown period as identified in Baele & Bekaert (2025) (DRAW). Numbers in the brackets represent heteroscedasticity-robust standard errors, and stars denote statistical significance: * - 0.10; ** - 0.05. Sample period: 1980-2024.

	Unconditional	Bad Times (NBER)	Bad Times (EBQ)	Bad Times (DRAW)
Level (β_1)	0.083**	0.084**	0.085**	0.085**
	(0.002)	(0.003)	(0.003)	(0.003)
$Level_{Bad}$		-0.021**	-0.012**	-0.011**
		(0.002)	(0.002)	(0.002)
Slope (β_2)	0.075**	0.076**	0.082**	0.084**
	(0.014)	(0.016)	(0.016)	(0.016)
$Slope_{Bad}$		-0.048**	-0.040**	-0.045**
		(0.023)	(0.019)	(0.019)
N	50265	50265	50265	50265

Direct comparisons across studies are challenging due to differences in identification methods, the maturities examined, and the precise definitions of equity risk premia. Nonetheless, our results broadly align with those of Bansal et al. (2021), who also document countercyclicality in short-term risk premia. Unlike Bansal et al. (2021), however, we find that the observed rise in short-term equity premia is not large enough to invert the term structure. Giglio et al. (2024) show that equity yields, of which equity risk premia are an important part - are slightly upward sloping during good times and downward sloping during NBER recessions. As in our paper, changes in the slope are nearly entirely driven by increasing yields over short maturities. Ulrich et al. (2023) construct an ex-ante term structure of risk premia using survey forecasts for dividends and option-implied dividends and find the slope to be flat or positive in times of economic expansions and negative during contractions. They also find short-term ex-ante dividend premiums to be more volatile than their longer-term counterparts. Note, however, that because of the option data, they focus on relatively short maturities, up to three years, which is much shorter than in our paper. Our findings - and those in the papers just described - do challenge Gormsen (2021) who, at the market level, finds a countercyclical term premium implying that most of the variation is driven by changes in long term rates. In contrast, our estimates reveal that most of the time variation comes from the short and intermediate maturity risk premia. Moreover, we do not find evidence that in recessions the variation is so large that it can create an inversion of the term structure slope at the aggregate market level.

Figure II: Good and bad times

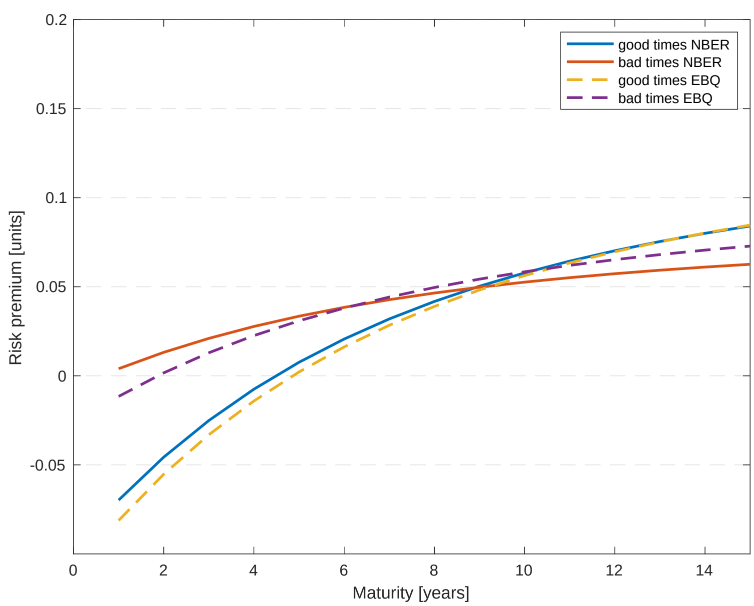

The graph depicts the estimated equity risk premium term structure in recessions and normal times based on two different cyclical measures (NBER - dashed line, EBQ - solid line). The year is classified as an NBER recession if at least 3 months in that year are considered to be a recession by NBER. Sample: 1980-2024.

5.4 The Cross-Section of Equity Term Structures

A notable strength of our model is its ability to generate firm-specific term structures by flexibly modeling them as functions of time-varying firm characteristics. As shown by equation (10) in Section 2, this is achieved by interacting the level and slope coefficients of the

Nelson-Siegel specification with firm-specific indicators that may vary over time. In our analysis, we focus on three characteristics that have been widely studied in the literature: firm size, book-to-market ratio, and credit rating. Section 5.4.1 reports results from specifications that include one characteristic at a time, while Section 5.4.2 extends the analysis to specifications that incorporate multiple characteristics simultaneously.

5.4.1 Univariate Analysis: Size, Value, Credit Risk

Size and Value Table V reports results from specifications in which the level and slope coefficients are modeled as functions of size or value dummies. These dummies take the value of one one if a firm falls within the smallest 50% (or 25%) by size in a given year, or the highest 50% (or 25%) by book-to-market ratio (i.e., value firms). Focusing first on size, the first two columns of Table V show that the slope is strongly positive for larger firms (7.7%), declines to 3.8% for below-median firms, and is essentially zero (0.2%) for the smallest quartile. This pattern indicates a steep term structure for large firms and a nearly flat term structure for the smallest firms. In Table C.I, where we employ size-quintile dummies, we find that the slope increases monotonically across the quintiles, from -0.4% for the smallest quintile to 8.2% for the largest quintile. The level estimates, which represent the 15-year risk premium, are only marginally higher for small relative to large firms (0.2% higher for below median sized firm, 0.7% higher for 25% smallest firms). Our results are thus consistent with the extensive literature that documents a moderate size effect in realized returns[13]. Based on our estimates, the size premium results from a combination of higher premia associated with short- and intermediate-maturity cash flows and a slightly higher long-term risk premium. Our results are also consistent with a downward sloping term structure for the Small-Minus-Big (SMB) factor, a finding we share with Giglio et al. (2024).

It thus seems that investors have a preference for short-maturity cash flows generated by

[13] Smith & Timmermann (2024) report a drop in the size premium over time, from 4% per year prior to 1972 to 1.9% after 1981, followed by a further reduction to 0.7% after the GFC.

large firms. Both rational and behavorial mechanisms could be at work here. For example, this finding could be consistent with the intuition that the majority of small cap stocks are young firms that are exposed to short-term fluctuations priced by investors. Yet, if the firm survives long enough, it matures and evolves similarly as large firms whose dynamics are largely explained by the long-run risks.

Table V: Cross-section of risk premium term structures

The table presents level and slope parameter estimates based on the modified Nelson-Siegel form:

$$r_{t,m} = rf_{t,m} + (\beta_1 + \boldsymbol{\beta_2'} \boldsymbol{Z}_{i,t}) - (\beta_3 + \boldsymbol{\beta_4'} \boldsymbol{Z}_{i,t}) \frac{\varphi(m) - \varphi(15)}{\varphi(5) - \varphi(15)},$$

where $\boldsymbol{Z}_{i,t}$ are characteristic dummies, and $\varphi(m) = \frac{1-e^{-m/\lambda}}{m/\lambda}$ is a decay function normalized so that β_1 reflects the 15-year risk premium and β_3 represents the 15 minus 5-year risk premium. *Small* stands for firms in the bottom 50% or 25% size percentile (on a yearly basis), *Value* denotes the top 50% or 25% in terms of book-to-market ratios, and *Speculative* represents firms with a speculative credit rating. Numbers in parentheses are heteroscedasticity-robust standard errors; stars denote statistical significance: * $p < 0.10$, ** $p < 0.05$. Sample period: 1980-2024.

	Size		Value		Speculative	Combined	
	50%	25%	50%	25%		50%	25%
Level (β_1)	0.083**	0.083**	0.078**	0.082**	0.084**	0.079**	0.082**
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
$Level_{Small}$	0.002	0.007**				0.004*	0.008
	(0.002)	(0.002)				(0.002)	(0.007)
$Level_{Value}$			0.006**	0.004*		0.005**	0.004*
			(0.002)	(0.002)		(0.002)	(0.002)
$Level_{Speculative}$					0.000	-0.000	-0.002
					(0.002)	(0.002)	(0.002)
Slope (β_3)	0.077**	0.076**	0.090**	0.085**	0.077**	0.090**	0.084**
	(0.015)	(0.014)	(0.015)	(0.014)	(0.014)	(0.015)	(0.014)
$Slope_{Small}$	-0.039**	-0.075**				0.073**	0.063
	(0.015)	(0.022)				(0.025)	(0.049)
$Slope_{Value}$			-0.087**	-0.127**		-0.097**	-0.153**
			(0.016)	(0.025)		(0.019)	(0.032)
$Slope_{Speculative}$					-0.113**	-0.101**	-0.098**
					(0.020)	(0.022)	(0.020)
N	50265	50265	50265	50265	17149	17149	17149

We now turn to the comparison of value and growth stocks. The literature provides mixed evidence on their behavior, with considerable debate surrounding the size of the premium, whether the value premium reflects compensation for risk and, if so, what the nature of these risks is. Recent studies suggest that the value anomaly may be explained by differences in the duration of firms' cash flows, which, when combined with a downward-sloping term structure, could account for the observed value premium (Lettau & Wachter (2007); Weber (2018); Gormsen & Lazarus (2023)). We contribute to this debate by allowing the level and slope parameters to depend on a value dummy, which equals one when a firm's book-to-market (BTM) ratio falls within the top 50% or top 25% in a given year. The main conclusion from columns 3 and 4 in Table V is that the term structure is upward sloping for growth (low BTM) firms but close to flat or even downward sloping for value (high BTM) firms. More specifically, we find that the term structure is only moderately upward sloping (0.3%) for firms with above-median BTM ratios but steeply upward sloping (9%) for below-median BTM firms. The slope even becomes negative (-4.2%) for top 25% BTM firms.[14] We also see that the long-term risk premia are slightly higher for value firms. Our estimates thus imply a value premium for stocks, which results from a combination of substantially higher risk premia associated with short- and intermediate-maturity cash flows and slightly higher long-term premia. The implied term structures for growth and value firms further suggest that the HML strategy - long value and short growth - displays a downward-sloping term structure. While many of our empirical findings align with those of Giglio et al. (2024), this result contrasts with their evidence of an upward-sloping term structure for the HML factor, highlighting how our alternative methodology that does not impose an asset pricing model but instead makes assumptions on future cash flows can lead to different conclusions.

We also estimate a version of our model that uses continuous values for size and book-to-market, instead of dummies. We standardize (the log) size and book-to-market variables to have mean zero and unit standard deviation. The advantage of this approach is that the term structure is potentially different for every firm (but fully determined by its characteristics);

[14] In Table C.I we find that the slope is decreasing from 11.8% for the lowest quintile BTM firms (Growth Stocks) to -5.7% for the highest quintile firms (Value Stocks). Slopes in between the two extreme cases are in the more moderate range of 0.0% to 6.7%.

a disadvantage that results are more vulnerable to potential outliers in characteristics. We therefore cap both variables at 3 standard deviations. As shown in Table C.II, the term structure becomes progressively less steep—and ultimately inverts—as the book-to-market ratio increases, confirming our previous analysis. The turning point from positive to negative slope is at a standardized BTM value of 0.32. Contrary to our specifications with dummy variables, we do not find a significant interaction between slope and (standardized) size. Given the very skewed nature of firm size, even after taking logs, this shows that using dummies may give more robust results than using the continuous variable.

Credit Risk Many articles have analyzed how realized equity returns differ across firms with different credit risk. Often, it has been difficult to detect a significant relationship between creditworthiness and realized returns, see e.g. Campbell et al. (2008). We revisit this question by making the level and slope coefficients a function of credit rating dummies. Column 5 of Table V shows results for a specification that includes a dummy that equals one when the firm has a speculative credit rating, that is, a rating below BBB-, and zero otherwise.[15] We find the slope to be positive and steep (7.7%) for investment-grade firms but downward sloping (7.7%-11.3% = -3.6%) for high credit risk firms. Figure III deepens the analysis by depicting the implied term structures for issuers in three credit rating categories: above BBB+, between BBB+ and BBB-, and below BBB-. The first two groups correspond to investment-grade ratings, while the last group is classified as speculative. The figure shows that investment-grade firms exhibit an upward-sloping term structure, whereas speculative-grade firms display a slightly downward-sloping curve. As one would expect, firms belonging to the highest rating category face a slightly lower risk premium compared to firms belonging to the intermediate (but still investment grade category). Interestingly, as maturity increases, the risk premium across all rating categories converges to the same long-term level, indicating that differences in credit quality are mainly priced for short to

[15] The credit rating we use in our analysis are S&P issuer's long-term credit rating. Because ratings are only available for a selection of firms, our sample size drops from 50,258 to 17,148 firm-year observations.

intermediate term cash flows.

Figure III: Credit ratings

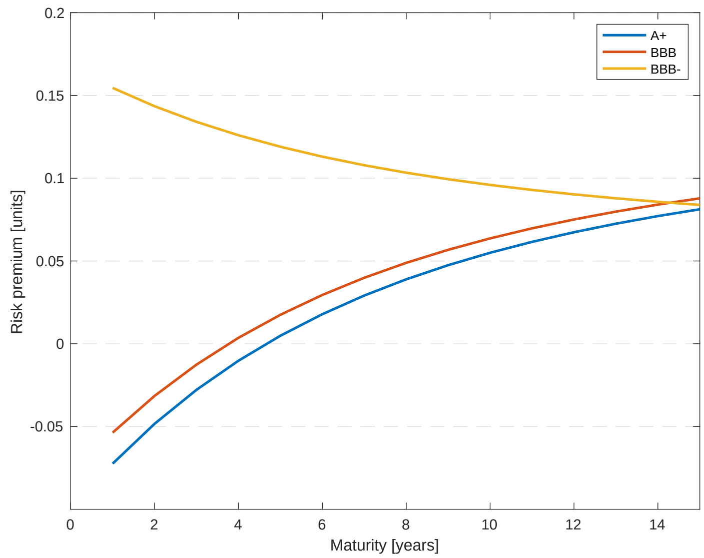

The figure presents the estimated equity risk premium term structure in different S&P credit rating portfolios. A+ group contains all credit ratings above and including an $A-$ rating. BBB group includes all $BBB+$, BBB and $BBB-$ rated firms. Finally, BBB- includes all non-investment grade ratings. Subsample: 1985-2024.

5.4.2 Multivariate Analysis

The flexible way in which we model the equity term structure in equation (10) easily accommodates multiple firm characteristics at the same time. The final two columns of Table V present results for specifications that jointly include small size, value, and speculative grade dummies. As before, these dummies take the value of one if a firm falls within the smallest 50% (25%) by size, ranks in the highest 50% (25%) by book-to-market ratio, or holds a speculative credit rating, respectively. The multivariate analysis confirms that a high book-to-market ratio and/or speculative credit rating both contribute to a reduction in the equity term structure's slope, confirming the univariate findings. As before, the reduction in slope is more extreme for the top-25% BTM firms (-15.3%) than for the above-median BTM firm (-9.7%). In contrast, we no longer find evidence that the term structure becomes less steep for small firms, probably because many small firms are value firms and/or firms with a speculative credit rating. In fact, once we control for value and credit rating, the effect of size on the slope is positive. When using continuous size and BTM ratios instead of dummies in a multivariate setting, we find similar results for value and credit risk (Table C.II in the Appendix). For firm size the effect differs from the specification with dummies, a phenonemon we also observed in the univariate analysis. In sum, we observe robust effects of book-to-market and credit risk on the equity term structure, while the effect of firm size is less robust across specifications.[16]

6 Robustness

In this section we perform various robustness checks. In particular, we focus on different approaches to construct the cash flow forecasts that are an input to our estimation methodology.

6.1 Biases in cash flow expectations

The literature documents that analyst forecasts entail biases which can, in turn, affect our estimates of risk premia because analyst forecasts are one ingredient of our procedure to obtain cash flow forecasts. Recent empirical work highlights that the biases tend to increase with maturity (van Binsbergen et al. (2022), Cassella, Golez, Gulen & Kelly (2023), Zhang et al. (2024)), at least for the first few years. Moreover, there is some evidence that analysts overreact/underreact to information contributing to known cross-sectional anomalies (Bordalo, Gennaioli, Porta & Shleifer (2019)). All these errors can distort our estimates if the investors have different expectations from the analyst (or introduce a wedge between subjective and objective risk premia). In order to take this into account, we use the esti-

[16] Table C.III in the Appendix provides additional detail by reporting level and slope estimates for portfolios constructed using a 3×3 double sort on size and book-to-market terciles. Within each size tercile, the slope declines markedly as we move from the lowest to the highest BTM tercile; a significantly negative slope is found only for the smallest value firms. Within each BTM tercile, the effect of firm size on the slope is smaller and not always monotonic.

mates of analyst biases based on the methodology of van Binsbergen et al. (2022) and Zhang et al. (2024) in our cash flow forecasts. Specifically, we correct all ROE forecasts for the biases in EPS forecasts.[17] However, since van Binsbergen et al. (2022) and Zhang et al. (2024) focus only on the first 2 years of forecasts, we make a simplifying assumption that for years 2-5 the bias stays constant. Afterward, it dwindles at the speed of cash flow reversion. The summary statistics of analyst biases are presented in Table C.IV, where each EPS bias is normalized by the firm's stock price. Consistent with van Binsbergen et al. (2022) and Zhang et al. (2024) we see that analysts are, on average, optimistic: the bias is positive and increasing with maturity. One-year cash flows display a bias of roughly 1.9%, whereas the two-year bias stands at close to 3.8%.

The results with corrected cash flow forecasts are displayed in Figure C.V. We see that correcting for analyst biases has a small effect on our estimates. In particular, the correction does not alter the main conclusion much: the slope remains positive, though the short-maturity risk premia become slightly larger. At first sight, one may expect a decrease in discount rates. On average, the analyst bias correction lowers our cash flow forecasts and hence lower discount rates would be needed to fit observed equity values. However, our methodology also uses the cross-sectional variation of cash flow forecasts. As Table C.IV shows, the analyst biases vary substantially across firms (with a standard deviation of 22-23%) and this explains why discount rate estimates can increase when correcting for analyst biases, though the effect is quantitatively small.

6.2 Time-varying g

The cash flow modeling hinges on the assumption that cash flows converge to some steady state g over time. As outlined in the general version of the model, g_t can be made time-varying to account for changes in growth expectations and inflation - two key macroeconomic

[17] In that sense, we investigate an extreme effect of analyst biases. If there was the same-direction bias in sales growth, the bias would (largely) cancel out. We can see this from equation (7).

variables that may capture phenomena such as structural shifts or secular stagnation. One possibility is to follow Li et al. (2013) who proxy g_t with the average sample nominal GDP rate up to the valuation year. As shown in Appendix C, the benchmark result of an upward sloping aggregate term structure is highly robust to using this time-varying g.

We also exploit market expectations about the long-term inflation and real GDP trends. Specifically, we incorporate the Federal Reserve survey of professional forecasters data which dates back to 1982 for long-term inflation and to 1992 for real GDP. Mixing inflation forecasts with the average realized real GDP growth rate or real GDP inflation forecasts does not affect the full sample estimates: we again find an upward sloping term structure. In addition, we find non-negative risk premia for all maturities, suggesting that negative baseline point estimates may be overcompensating for unaccounted business-cycle and other structural variation in cash flow growth rates. Overall, it appears that differences in the steady state growth rate do not generate shifts in the term structure as much as changes in cross-sectional firm characteristics do.

6.3 Cash flow forecast averaging

Considering that our findings may be driven by particular filters or noisy cash flow forecasts, we investigate alternative procedures. The generic insight is that, regardless of additional refinements, we always find upward sloping aggregate term structure estimates with in most cases an even more positive slope compared to the benchmark results (Table C.V).

First, we average cash flow forecasts between our model and the simple mean-reversion approach, as in Dechow et al. (2004) and Weber (2018), which should reduce the impact of outliers. The findings again show a strong positive slope at the aggregate market level. We also introduce a more intricate weighting of these two different models where each forecast i is assigned a score based on the likelihood of such cash flow to materialize:

$$w_i = \frac{\phi(forecast_i; \mu, \sigma)}{\sum_{k=1}^{2} \phi(forecast_k; \mu, \sigma)} \tag{11}$$

In order to determine the likelihood, we assume that cash flows are normally distributed with the mean of 1.9% and standard deviation of 25.3% (empirical sample statistics). For instance, if one of the forecasts is very far from the prior mean of 1.9%, then all the weight is attributed to the more conservative forecast. On top of that, we cap the maximum (minimum) cash flow at 100% (-100%) of firm's value. All these steps should smooth out the extremities. As we see from the second row, this averaging leads to stronger positive slope compared to our baseline results. Next, if we winsorize the sample at 5% and 95% percentiles of the IRR we continue to find a strong positive slope, which becomes even more positive if we combine this with cash flow averaging (the third, fourth, fifth rows). Finally, we exclude the most extreme size and book-to-market portfolios, and again find a more strongly positively sloped term structure.

It is important to note that in all these robustness checks, we reduce the cross-sectional variation in cash flow patterns across firms, particularly in the term structure of cash flows. Since the identification hinges on differences in the maturity patterns of cash flows across firms, it is not surprising that these robustness checks yield somewhat more extreme estimates for the term structure of risk premia. By definition, large cross-sectional pricing differences must be driven either by large differences in risk premia or by cross-maturity variation in cash flows. Consequently, in samples with less cross-maturity variation, the slope may be more sensitive to relatively small cross-sectional differences.

6.4 Persistence of cash flows

Concerns about the methodology may also point at the cash-flow reversion process itself. Given the estimates we employ, most of the firms reach their steady state growth rate in 6-10 years. In order to test if the reversion is driving the qualitative results, we rerun our base estimation with perturbed ROE and sales growth reversion parameters (Table VI). The general tendency is that for small and moderately high mean reversion coefficients the term structure slope is positive. However, very large coefficients, much larger than our empirical estimates, can lead to a downward sloping curve. However, in the range of our benchmark mean reversion estimates (0.41 for sales growth and 0.81 for ROE) the estimated term structure is upward sloping.

The persistence of cash flow reversion matters in two significant ways. When the firm's transition to the steady state is very short, all firms behave similarly irrespective of their initial differences in growth rates and profitability. As a result, the pricing differences across firms would be largely affected by the initial book-to-market ratios across firms, as in Dechow et al. (2004) and Weber (2018). Thus, the identification of short- versus long-run risk premia would be weaker and estimates are naturally much more noisy and extreme. On the other hand, the persistence can affect which firms earn a higher expected return. Largely profitable firms tend to have most of their value in the immediate and intermediate future since all firms later converge to the same steady state. A decrease in the speed of cash flow mean reversion (higher AR coefficient) would unilaterally increase their expected return. On the other hand, the unprofitable firms (the flip side of profitable firms) may experience a decrease in expected return due to much more persistent losses (the relationship is not entirely clear if there are non-monotonicities). This can potentially shed some light on why the term structure ultimately depends on the persistence parameters.

To conduct a similar experiment, we also alter the forecasting horizon to values from 10 to 50 years, and assume different steady state growth rates for g. As outlined in Table VII, this again generates two insights. Firstly, it appears that as long as the forecasting period is long enough for the convergence to occur, the slope sign remains unchanged. Although the magnitude of estimates does depend on the forecasting horizon (as more and more value of the firm is attributed to the terminal horizon), the slope of the equity term structure does not. Second, the steady state growth rate only mildly affects the slope estimates, except for the high growth and short horizon combination. Yet, a long-term cash flow growth rate of 9% is both rather exceptional and convergence at 10 years horizon is unlikely to fully occur. Overall, the qualitative results seem quite robust to the selection of g and T.

Table VI: Sensitivity of NS estimates

The table presents the sensitivity of slope and level estimates to changes in ROE and sales growth reversion parameters. The equilibrium cash flow growth rate g and forecasting horizon T are set to 6% and 15%, respectively. Numbers in the brackets represent heteroscedascity-robust standard errors, and stars denote statistical significance: * - 0.10; ** - 0.05. The negative slope estimates are shaded in dark gray. Sample period: 1980-2024.

		Slope: AR(1) ROE					
		0.4	0.5	0.6	0.7	0.8	0.9
AR(1) sales	0.2	0.19**	0.18**	0.16**	0.14**	0.09**	0.00
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
	0.3	0.19**	0.18**	0.16**	0.14**	0.09**	-0.00
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
	0.4	0.18**	0.17**	0.16**	0.13**	0.08**	-0.01
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
	0.5	0.17**	0.16**	0.15**	0.12**	0.07**	-0.03
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
	0.6	0.15**	0.14**	0.13**	0.10**	0.05**	-0.05**
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
		Level: AR(1) ROE					
		0.4	0.5	0.6	0.7	0.8	0.9
AR(1) sales	0.2	0.08**	0.08**	0.08**	0.08**	0.08**	0.09**
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	0.3	0.08**	0.08**	0.08**	0.08**	0.08**	0.09**
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	0.4	0.08**	0.08**	0.08**	0.08**	0.08**	0.09**
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	0.5	0.08**	0.08**	0.08**	0.08**	0.08**	0.09**
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	0.6	0.08**	0.08**	0.08**	0.08**	0.08**	0.09**
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Table VII: Sensitivity of NS estimates

The table presents the sensitivity of slope and level estimates to changes the equilibrium cash flow growth rate g and forecasting horizon T. In all specifications below, the ROE and sales growth mean reversion parameters are set to 0.81 and 0.40, respectively. For brevity, the level estimates are excluded from the table. The negative slope estimates are shaded in dark gray. Numbers in the brackets represent heteroscedascity-robust standard errors, and stars denote statistical significance: * - 0.10; ** - 0.05. The negative slope estimates are shaded in dark gray. Sample period: 1980-2024.

		Slope: T				
		10	15	20	30	50
	0.03	0.05**	0.07**	0.07**	0.08**	0.08**
		(0.02)	(0.02)	(0.01)	(0.01)	(0.01)
g	0.06	0.02	0.08**	0.09**	0.10**	0.10**
		(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
	0.09	-0.03	0.08**	0.10**	0.11**	0.12**
		(0.02)	(0.02)	(0.01)	(0.01)	(0.02)
		Level: T				
		10	15	20	30	50
	0.03	0.07**	0.05**	0.05**	0.05**	0.05**
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
g	0.06	0.09**	0.08**	0.08**	0.09**	0.09**
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	0.09	0.12**	0.11**	0.12**	0.12**	0.13**
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

6.5 Industry portfolios

As a final robustness check we perform another cross-sectional sort, where we sort stocks into the 15 Fama-French industries. Because, arguably, the technology sector has changed most substantially over our sample period, we separate technology firms into a separate category following Barron, Byard, Kile & Riedl (2002).

In Table C.VI the term structure estimates are shown. We always find positively sloped term structures. In addition, the level estimates (the 15-year risk premium) are fairly stable across industries. While all slope estimates are positive, we do see substantial variation across industries in the size of the slope estimate. It seems that consumer-oriented industries have the steepest term structures. One risk-based interpretation would be that these industries are less exposed to short-lived economic risks, but further work is needed to understand these patterns in more detail.

7 Conclusion

The empirical literature documenting the properties of the term structure of expected equity returns is divided. The prediction of the seminal asset pricing models that longer maturity assets contain more systematic risk was put to test by a series of papers by van Binsbergen et al. (2012), van Binsbergen et al. (2013), van Binsbergen & Koijen (2017). Specifically, the price data in options and dividend strips markets seem to embed a downward-sloping equity term structure meaning that most established models are not in line with empirics. Nevertheless, the evidence is not without controversy due to relatively short historical samples, containing two extreme crashes in 2001 and 2007-2008, which makes the realized returns noisy proxies of the unconditional average expected returns. The challenge is further exacerbated by liquidity effects in derivative markets (Bansal et al. (2021)).

This paper proposes a new methodology to estimate the term structure of equity risk premia based on ex-ante returns. We do so by exploiting the cross-sectional heterogeneity across firm cash flow forecasts and combining them with the implied cost of capital approach. Rather than sorting into duration portfolios and looking at noisy realized returns, we estimate the risk premium term structure by minimizing errors between the discounted cash flow model price of the firm's equity (in the spirit of Dechow et al. (2004)) and the observed market prices. To incorporate as much identifying variation as possible, we augment Dechow et al. (2004) with the information extracted from the IBES analyst forecasts. Our approach does not require prices of dividend assets and can thus be implemented for longer sample periods, 1980 – 2024 in our estimation for the U.S. market. Also, our approach allows us to investigate cross-sectional and time-series variation in risk premium term structures.

Our baseline results show that the equity term structure is upward sloping at the aggregate level, consistent with predictions of classical asset pricing theories such as Campbell & Cochrane (1999) and Bansal & Yaron (2004). However, we also find substantial undocumented heterogeneity in the size, value, and credit risk dimensions. In particular, our estimates imply that value firms and speculative-grade firms exhibit a downward sloping term structure.

We do find evidence that the term structure becomes flatter during recessions, but it remains upward sloping also in these periods. Adding to that, in this paper we cast some light on how credit risk is priced in firms. The implied discount rates are monotonically increasing with credit risk, proxied by the credit ratings, and are particularly high for speculative grade firms. All in all, our ex-ante measure of returns does not support the credit risk anomaly detected in realized returns (Campbell et al. (2008)).

In sum, our key contribution is to develop and implement a coherent framework that merges the cross-section of stock prices and implied cost of capital approach to recover the term structure of equity risk premia. For the modeling cash flows we closely follow the original framework of Dechow et al. (2004) and Weber (2018), and amend it with the information extracted from analysts' forecasts. In doing so, we incorporate more heterogeneity in cash flow expectations, necessary to identify separate horizon discount rates, while at the same time keeping the parsimonious structure of the expected cash flow process. As shown, our methodology serves as a powerful tool to flexibly measure ex-ante returns and inspect the cross-sectional and time-varying properties of the equity term structure.

References

Baele, L. & Bekaert, G. (2025), Equity market cycles. Working paper, Available on SSRN.

Bansal, R., Miller, S., Song, D. & Yaron, A. (2021), 'The term structure of equity risk premia', *The Journal of Financial Economics* **142**(3), 1209–1228.

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

Barron, O. E., Byard, D., Kile, C. & Riedl, E. J. (2002), 'High-technology intangibles and analysts' forecasts', *Journal of Accounting Research* **40**(2), 289–312.

Boguth, O., Carlson, M., Fisher, A. & Simutin, M. (2023), 'The term structure of equity risk premia: Levered noise and new estimates', *Review of Finance* **27**(4), 1155–1182.

Bordalo, P., Gennaioli, N., Porta, R. L. & Shleifer, A. (2019), 'Diagnostic expectations and stock returns', *The Journal of Finance* **74**(6), 2839–2874.

Callen, J. L. & Lyle, M. R. (2020), 'The term structure of implied costs of equity capital', *Review of Accounting Studies* **25**(1), 342–404.

Campbell, J. Y. & Cochrane, J. H. (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., Hilscher, J. & Szilagyi, J. (2008), 'In search of distress risk', *Journal of Finance* **63**(6), 2899–2939.

Cassella, S., Golez, B., Gulen, H. & Kelly, P. (2023), 'Horizon bias and the term structure of equity returns', *The Review of Financial Studies* **36**(3), 1253–1288.

Chen, S. & Li, T. (2018), 'A unified duration-based explanation of the value, profitability, and investment anomalies', *Journal of Empirical Finance* **84**.

Claus, J. & Thomas, J. (2001), 'Equity premia as low as three percent? evidence from analysts' earnings forecasts for domestic and international stock markets', *The Journal of Finance* **56**(5), 1629–1666.

Cochrane, J. H. (2011), 'Presidential address: Discount rates', *Journal of Finance* **66**(4), 1047–1108.

Davis, J. L., Fama, E. F. & French, K. R. (2000), 'Characteristics, covariances, and average returns: 1929 to 1997', *The Journal of Finance* **55**(1), 389–406.

Dechow, M. P., Sloan, R. G. & Soliman, M. T. (2004), 'Implied equity duration: A new measure of equity risk', *Review of Accounting Studies* **9**(2-3), 197–228.

Fama, E. F. & French, K. R. (1993), 'Common risk factors in the returns on stocks and bonds', *Journal of Financial Economics* **33**(1), 3–56.

Fama, E. F. & French, K. R. (2002), 'The equity premium', *The Journal of Finance* **57**(2), 637–659.

Gebhardt, W. R., Lee, C. M. & Swaminathan, B. (2001), 'Toward an implied cost of capital', *Journal of accounting research* **39**(1), 135–176.

Giglio, S., Kelly, B. T. & Kozak, S. (2024), 'Equity term structures without dividend strips data', *The Journal of Finance* **79**(6), 4143–4196.

Golez, B. & Jackwerth, J. (2024), 'Holding period effects in dividend strip returns', *The Review of Financial Studies* **37**(10), 3188–3215.

Gonçalves, A. S. (2021), 'The short duration premium', *Journal of Financial Economics* **141**(3), 919–945.

Gormsen, N. J. (2021), 'Time variation of the equity term structure', *The Journal of Finance* **76**(4), 1959–1999.

Gormsen, N. J. & Lazarus, E. (2023), 'Duration-driven returns', *Journal of Finance* **78**(3), 1393–1447.

Greene, W. H. (2000), *Econometric analysis 4th edition*.

Gürkaynak, R. S., Sack, B. & Wright, J. H. (2007), 'Treasury yield curve: 1961 to the present', *Journal of Monetary Economics* **54**(8), 2291–2304.

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

Li, Y., Ng, D. T. & Swaminathan, B. (2013), 'Predicting market returns using aggregate implied cost of capital', *Journal of Financial Economics* **110**(2), 419–436.

Nelson, C. R. & Siegel, A. F. (1987), 'Parsimonious modeling of yield curves', *Journal of business* pp. 473–489.

Nissim, D. & Penman, S. H. (2001), 'Ratio analysis and equity valuation: From research to practice', *Review of accounting studies* **6**(1), 109–154.

Pástor, L., Sinha, M. & Swaminathan, B. (2008), 'Estimating the intertemporal risk–return tradeoff using the implied cost of capital', *The Journal of Finance* **63**(6), 2859–2897.

Schröder, D. (2024), 'The term structure of equity yields—a bottom-up approach', *Review of Finance* **28**(2), 661–697.

Smith, S. & Timmermann, A. (2024), 'Have risk premia vanished?', *Journal of Financial Economics* **145**, 553–576.

Ulrich, M., Florig, S. & Seehuber, R. (2023), 'A model-free term structure of u.s. dividend premiums', *The Review of Financial Studies* **36**(3), 1289–1318.

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

van Binsbergen, J. H., Han, X. & Lopez-Lira, A. (2022), 'Man vs. machine learning: The term structure of earnings expectations and conditional biases', *The Review of Financial Studies* .

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

van Binsbergen, J. & Koijen, R. (2017), 'The term structure of returns: Facts and theory', *Journal of Financial Economics* **124**(1), 1–21.

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

Zhang, Y., Zhu, Y. & Linnainmaa, J. T. (2024), 'Man versus machine learning revisited', *Review of Financial Studies, forthcoming* .

Appendix A Connection to One-Period Returns

Let's start with a simplified scenario of time-invariant discount rates. Let $P_{t,m}$ be the price of a dividend strip, paying out at time $t+m$. Then:

$$P_{t,m} = \frac{\mathbb{E}_t(CF_{t+m})}{R(m)^m},$$

with $R(m)$ the holding-to-maturity total discount rate. The one-period expected return is then:

$$R_{t:t+1}(m) = \mathbb{E}_t\left(\frac{P_{t+1,m-1}}{P_{t,m}}\right) = \frac{R(m)^m}{R(m-1)^{m-1}} = R(m)\left(\frac{R(m)}{R(m-1)}\right)^{m-1}$$

The expression above is identical to the forward rate from $m-1$ to m. Thus, the shape of the one-period return term structure is analogous to the forward rate curve, linking two expected return measures - $R(m)$ and $R_{t:t+1}(m)$.

In many cases $R(m)$ slope alone is sufficient to summarize the slope of $R_{t:t+1}(m)$. To see that, we can express discount rates and expected returns in logs:

$$r_{t:t+1}(m) = r(m) + (m-1)\Delta r(m), \qquad \text{with } \Delta r(m) \equiv r(m) - r(m-1).$$

If we calculate the difference between the one-period expected return of the $(m+1)$-strip minus the m-strip, we get:

$$\Delta r_{t:t+1}(m+1) = (m+1)\Delta r(m+1) - (m-1)\Delta r(m).$$

Equivalently,

$$\Delta r_{t:t+1}(m+1) = 2\underbrace{\Delta r(m+1)}_{slope} + (m-1)\Big(\underbrace{r(m+1) + r(m-1) - 2r(m)}_{convexity}\Big).$$

For moderate levels of concavity, the first term always dominates. That means if the discount rate term structure $r(m)$ is strictly upward sloping, the short-holding-period expected return term structure is also upward sloping. For strongly concave holding-to-maturity term structures, one-period expected returns can initially be upward sloping and then moderately downward sloping at the longest maturities.[18] These nonlinearities become less important when we measure the term-structure slope as the difference between more distant short- and long-maturity points on the term structure. Thus, the proper comparison between the slope of $r_{t:t+1}(m)$ and $r(m)$ involve either computing the forward curve or avoiding intermediate maturities when calculating slopes.

Figure A.I displays the implied forward yield curve based on our baseline estimates. As expected, the forward rate term structure is steeply upward sloping, driven mostly by the low short-maturity risk premium. We also do not observe any hump-shaped patterns in the forward curve, implying that the concavity effects are small, at least for the maturities we investigate.

[18]By definition, the reverse is always true - upward sloping forward rates imply upward sloping holding-to-maturity returns.

Figure A.I: Implied Forward Rates

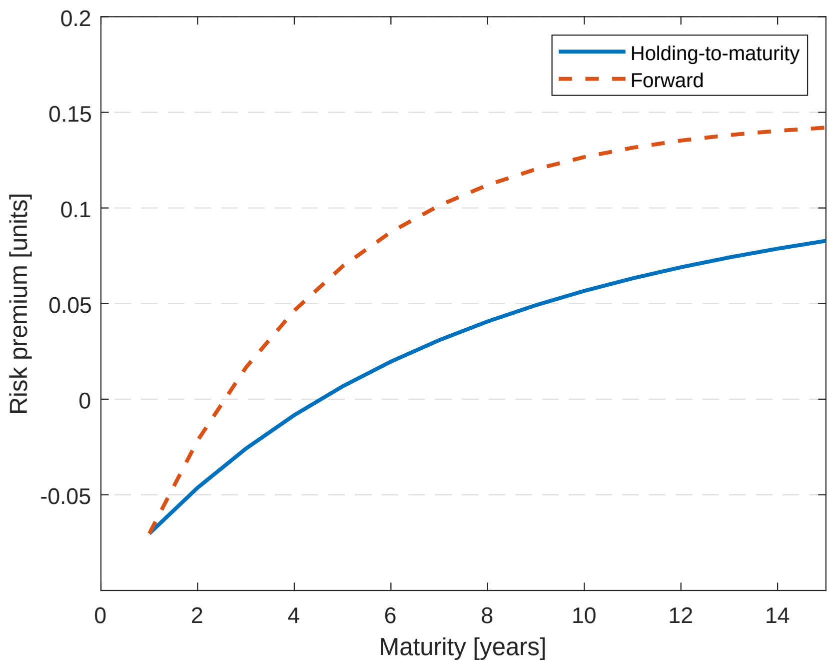

The figure displays the baseline implied forward risk premia curve, which under the constant discount rate assumption coincides with the one-year-ahead term structure of risk premia.

Appendix B Cash Flow Forecasts

1. Integrating analyst forecasts

We employ and modify the analyst data in the following way:

1. as in Li et al. (2013), we construct the constant maturity median forecasts for revenues, ROE and EPS as a weighted average between year $t+m$ and $t+m+1$ forecasts. This step accounts for the differences in financial years across firms creating discrepancies in timing of cash flows and their discounting.

2. if a firm has revenue forecasts, we used them for the 1-5 year growth rate of book equity projections (depending on the availability). Namely, we calculate the m'th year sales growth rates as $forecastSALE_{t+m}/forecastSALE_{t+m-1}-1$. If a firm has a non-missing IBES long-term growth rate forecast for sales, we extrapolate all remaining missing values for the 2-5 year horizon using this rate. All resulting sales growth rates are additionally winsorized at 5% and 95% level.

3. if a firm has ROE forecasts, we used them for years 1-5 (depending on the availability). If the forecasts are not available, we proceed with step 4. All ROE forecasts are winsorized at 5% and 95% level.

4. if EPS forecasts are available when ROE are not, we perform the following exercise. For years 1-4, we projected the book value of equity using the mean-reverting sales growth process with an AR coefficient of 0.40 and long-term mean of 6% (à la Dechow et al. (2004) and Weber (2018)). Again, if a firm has a non-missing long-term growth rate forecast for EPS and the last forecast value was positive, we extrapolated all remaining EPS missing values for the 2-5 year horizon using this rate. Then, we implied the ROE forecasts for maturity m using this formula: $EPS_{t+m} * \#shares_t/BE_{t+m-1}$.[19]. All the implied forecasts are winsorized at 5% and 95% level.

[19] The historical number of shares outstanding is assumed to remain fixed in the future.

2. Constant maturity forecast

Forecasts are interpolated using this principle:

$$12monthForecast = w * FY1 + (1 - w) * FY2 \tag{12}$$

where $FY1$ and $FY2$ are the current and subsequent financial year I/B/E/S forecasts and w stands for the number of months left until the realization of current financial year cash flow divided by 12. Since we have forecasts for up to 5 years, we repeat the same procedure for 24, 36, and 48 month forecasts.

Following Li et al. (2013), for the last available year of forecasts we apply this rule to extrapolate the constant maturity forecast:

$$24monthForecast = 12monthForecast * g \tag{13}$$

where g is the implicit growth between $FY1$ and $FY2$ forecasts. To reduce the effect of outliers, we winsorize this g at 5% and 95% level. If $FY1$ and $FY12$ switch signs, the growth rates become ill-defined. Therefore, we do not apply this rule for those cases.

The underlying assumption is that firm cash flows realize at the end of the firm's financial year. Because most firms have their financial year ending on Dec 31, we choose June as a valuation month.

3. Data Filters

There are 3 key dates in the I/B/E/S summary database:

- *statpers* - the date on which the forecasts are summarized in this database (e.g. 18 Dec 2014)
- *fpedats* - the end date of the forecasting period (e.g. 31 Dec 2014)

- *anndatsact* - the date when the actuals, i.e. financial statements, for that forecasting period were announced (e.g. 30 Jan 2015)

Filters:

1. The valuation date is June. Thus, only forecasts available in this month are used.
2. We account for potential mistakes by: i) dropping all observations where the forecast record day (*statpers*) is later than the announcement day of actuals (*anndatsact*); ii) dropping all observations where the forecast record day (*statpers*) exceeds the corresponding forecasting period end date (*fpedats*) by more than 4 months (122 days).[20] The idea is that forecasts announced after the actual financial statements are no longer forecasts.
3. We keep firms that have forecasts denominated in US dollars. This filter deletes double entries and eliminates concerns that the reporting currency in I/B/E/S does not match the balance sheet data in Compustat (which is important for measures like EPS, but not for ROE though).
4. In case more than one forecast measure is used, we follow the conservative approach and i) always choose the latest announcement date of actuals across measures, ii) delete all observations where the forecasting period end dates do not match across measures.

[20]The filter is also applied to all mismatches between 1-5 year actual and implied forecasting period end dates calculated as $statpers + forecastinghorizon$.

Appendix C Additional Tables and Figures

Figure C.I: ROE and Sales growth quintile portfolios

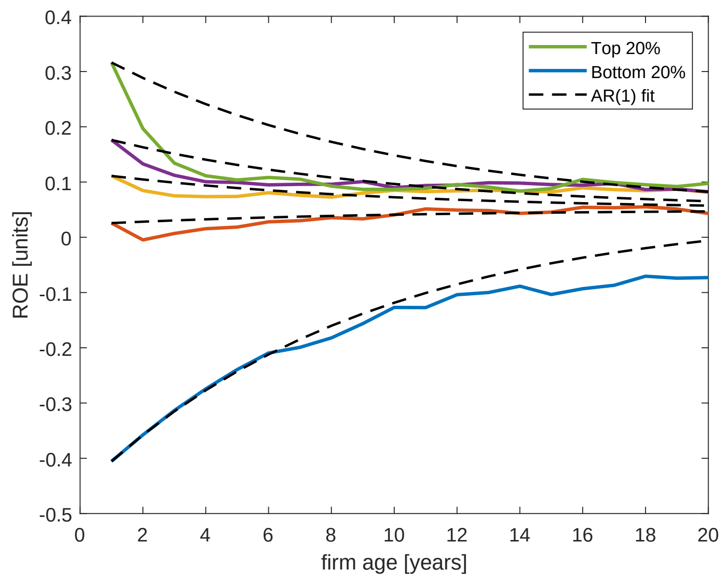

(a) ROE quintiles

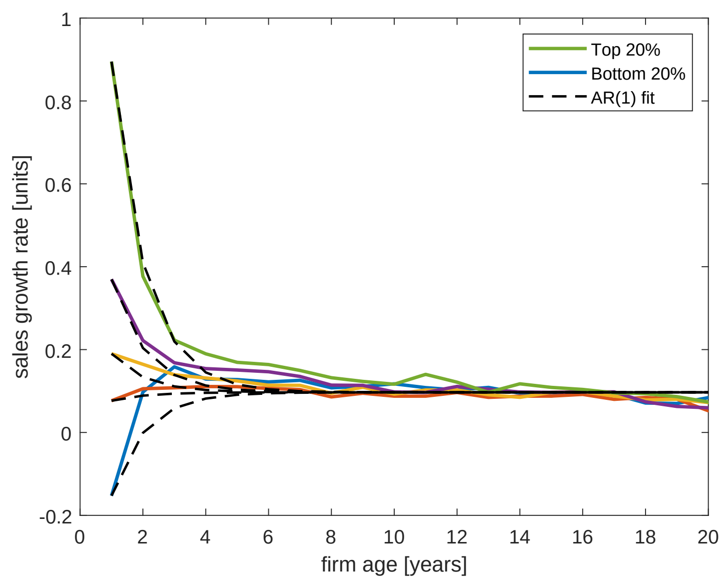

(b) Sales growth quintiles

The graph shows the term structure of firm's ROE and sales growth across firm age. Each solid line represents an average of the ROE (sales growth) quintile formed when the age equals 1. The age is proxied by the first occurrence of the firm in the Compustat database. Each dashed line represents the least squares fit of an AR(1) model, where AR(1) for ROE and sales growth portfolios equal to 0.81 and 0.40 respectively. Sample range: 1964-2021.

Figure C.II: Sensitivity to λ

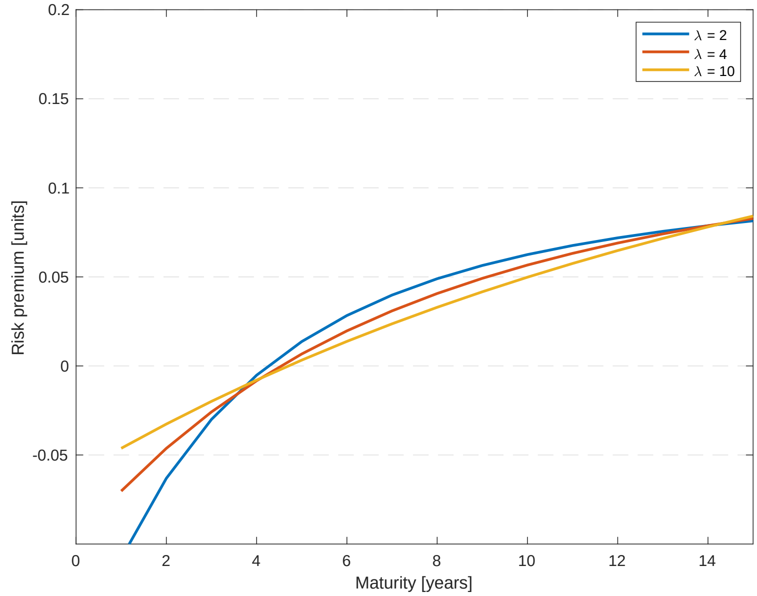

The figures illustrates the sensitivity of baseline estimates to changes in the decay parameter λ. Compustat-CRSP-I/B/E/S subsample: 1980-2024.

Figure C.III: Robustness of baseline results

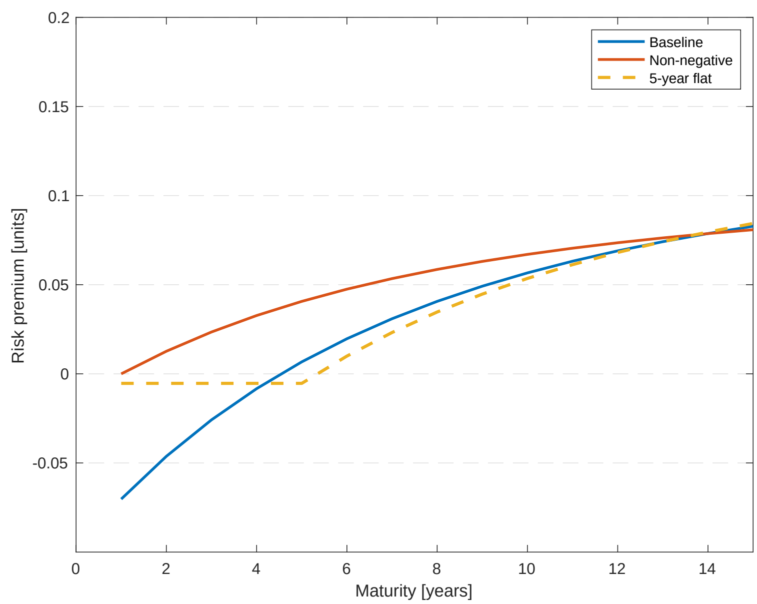

The graph presents alternative equity risk premium term estimates over the full sample with . Compustat-CRSP-I/B/E/S sample: 1980-2024.

Figure C.IV: van Binsbergen, Brandt & Koijen (2012) sample

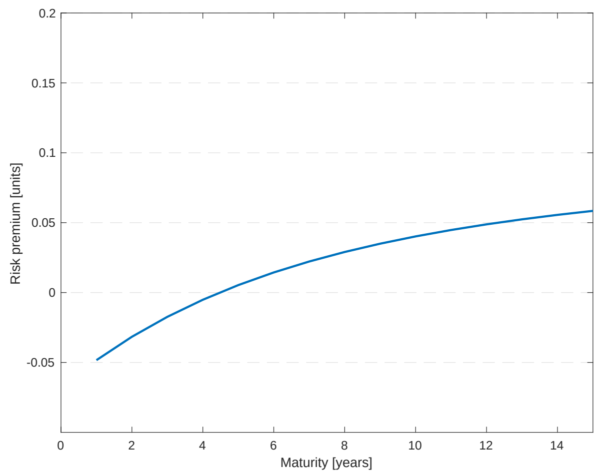

(a) Top size quintile

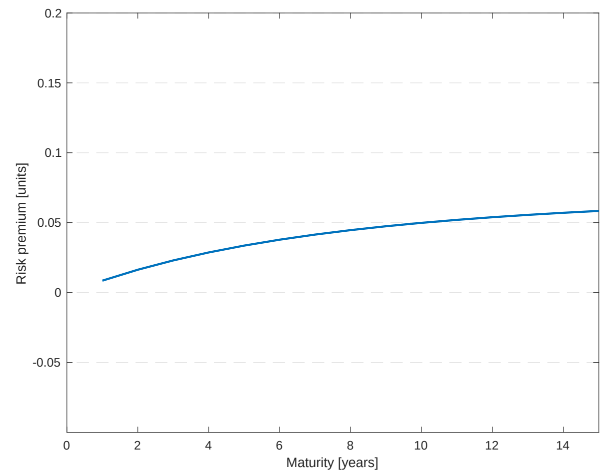

(b) S&P500 firms

The figures depict the estimated equity risk premia term structure in the van Binsbergen et al. (2012) sample proxied by the top size quintile firms (left graph) or the $S\&P500$ index firms in our sample (right graph). Compustat-CRSP-I/B/E/S subsample: 1996-2008.

Figure C.V: Cash flows with corrected analyst bias

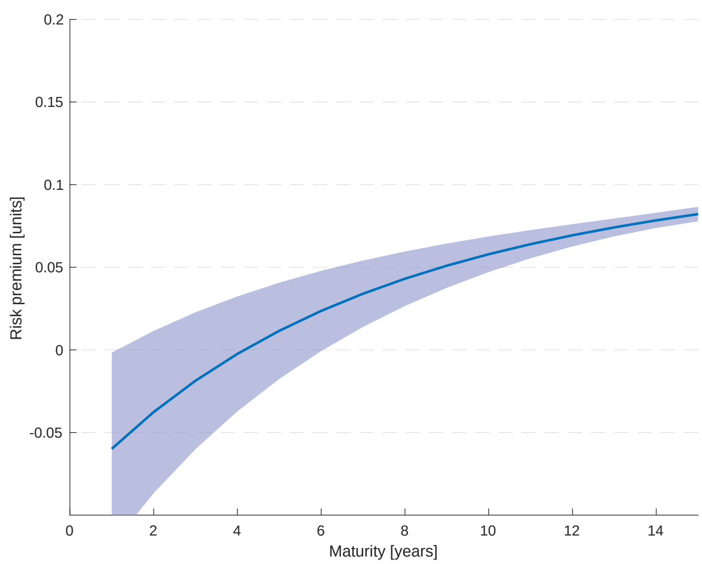

The figure presents the estimated equity risk premium term structure with corrected analyst forecasts. The correction is made by using estimates based on van Binsbergen et al. (2022) and Zhang et al. (2024) (the procedure is described more thoroughly in Section 6.1). Compustat-CRSP-I/B/E/S sample: 1980-2024.

Figure C.VI: Time-varying g

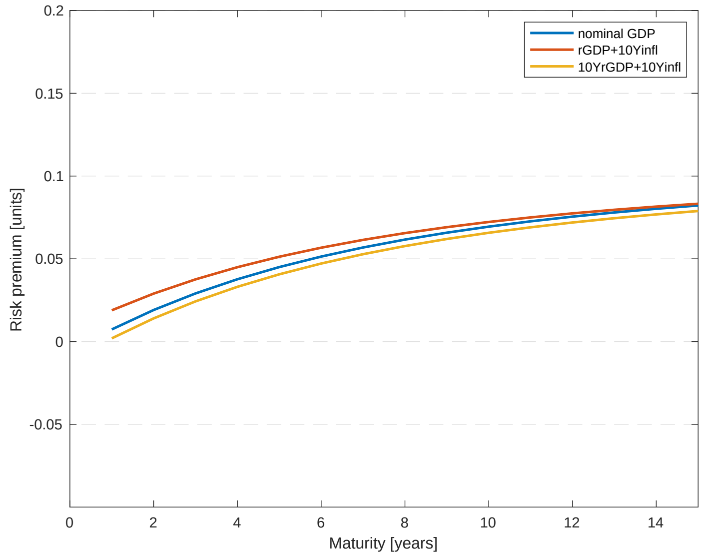

The figures illustrates equity term structure estimates using alternative proxies for the long-term cash flow growth rate g. The blue line represents estimates where g equals the average nominal GDP growth rate up to year t, as in Li et al. (2013) (Compustat-CRSP-I/B/E/S sample: 1980-2024). The red dashed line approximates g with the average real GDP growth rate up to year t plus the long-term (10Y) inflation forecasts (Compustat-CRSP-I/B/E/S subsample: 1982-2024). The yellow line measures g as a sum of the long-term (10Y) real GDP and long-term (10Y) inflation forecasts (Compustat-CRSP-I/B/E/S subsample: 1992-2024).

Table C.I: Size and Book-to-Market portfolios

The table shows the estimation results within size and book-to-market quintiles formed on a yearly basis. The first row always represents the level and the second row - slope estimates. Size is defined as price times the number of shares outstanding at the end of valuation month (June). Numbers in the brackets represent heteroscedascity-robust standard errors, and stars denote statistical significance: * - 0.10; ** - 0.05. Sample period: 1980-2024.

	Size		BM	
	$15Y$(Level)	$15Y-5Y$(Slope)	$15Y$(Level)	$15Y-5Y$(Slope)
Low	0.090**	-0.004	0.074**	0.118**
	(0.002)	(0.017)	(0.009)	(0.037)
2	0.085**	0.038**	0.076**	0.048**
	(0.002)	(0.014)	(0.004)	(0.024)
3	0.081**	0.063**	0.084**	0.068**
	(0.002)	(0.013)	(0.003)	(0.019)
4	0.078**	0.071**	0.081**	0.001
	(0.002)	(0.012)	(0.002)	(0.020)
High	0.084**	0.082**	0.087**	-0.056*
	(0.002)	(0.011)	(0.002)	(0.032)
High-Low	-0.006**	0.086**	0.013	-0.173**
	(0.002)	(0.020)	(0.009)	(0.049)

Table C.II: Cross-sectional risk premia parameter estimates

The table presents level and slope parameter estimates based on the modified Nelson-Siegel form:

$$r_{i,t,m} = rf_{t,m} + (\beta_1 + \boldsymbol{\theta_1'} \boldsymbol{Z}_{1,i,t}) + (\beta_2 + \boldsymbol{\theta_2'} \boldsymbol{Z}_{1,i,t}) \frac{\varphi(m) - \varphi(15)}{\varphi(5) - \varphi(15)},$$

where $\boldsymbol{Z}_{1,t}$ are cross-sectionally standardized characteristic, and $\varphi(m) = \frac{1-e^{-m/\lambda}}{m/\lambda}$ is a decay function normalized in a way that β_1 reflects the 15-year risk premium and β_3 reflects the 15 minus 5-year risk premium. *Size* stands for cross-sectionally standardized logarithm of size (on a yearly basis), B/M denotes the standardized book-to-market ratios, and *Speculative* represents firms with a speculative credit rating. Both *Size* and B/M variables are capped at 3 standard deviations. Numbers in parentheses are heteroscedasticity-robust standard errors; stars denote statistical significance: * $p < 0.10$, ** $p < 0.05$. Sample period: 1980-2024.

	Size	Value	Credit Risk	Combined
Level (β_1)	0.084**	0.083**	0.084**	0.088**
	(0.002)	(0.002)	(0.002)	(0.003)
$Level_{Size}$	-0.000			-0.002
	(0.001)			(0.001)
$Level_{Speculative}$			0.000	-0.003
			(0.002)	(0.002)
$Level_{B/M}$		0.001		0.002
		(0.001)		(0.002)
Slope (β_2)	0.086**	0.036**	0.077**	0.166**
	(0.014)	(0.019)	(0.014)	(0.014)
$Slope_{Size}$	-0.006			-0.066**
	(0.007)			(0.010)
$Slope_{Speculative}$			-0.113**	-0.155**
			(0.020)	(0.018)
$Slope_{B/M}$		-0.112**		-0.152**
		(0.014)		(0.015)
N	50265	50265	17149	17149

Table C.III: BM-Size double sorts

The table shows the estimation results within 3 book-to-market (BM) and size portfolios formed over the full sample. The data is first sorted on BM ratios and then on size. The first row always represents the level (i.e., 15-year risk premium) and the second row - slope estimates (15 minus 5-year risk premium spread). Compustat-CRSP-I/B/E/S sample: 1980-2024. Risk premia are measured in units.

		Size							
		Low		2		High		High-Low	
		Level	Slope	Level	Slope	Level	Slope	Level	Slope
BM	Low	0.107**	0.110**	0.081**	0.132**	0.082**	0.121**	-0.025**	0.011
		(0.004)	(0.013)	(0.003)	(0.013)	(0.006)	(0.022)	(0.007)	(0.025)
	2	0.084**	0.055**	0.070**	0.058**	0.084**	0.073**	0.000	0.018
		(0.002)	(0.012)	(0.002)	(0.015)	(0.003)	(0.015)	(0.003)	(0.019)
	High	0.086**	-0.049**	0.081**	-0.006	0.085**	-0.025	-0.001	0.024
		(0.002)	(0.024)	(0.002)	(0.018)	(0.002)	(0.018)	(0.002)	(0.030)
	High-Low	-0.021**	-0.159**	0.001	-0.138**	0.003	-0.146**		
		(0.005)	(0.027)	(0.004)	(0.022)	(0.006)	(0.028)		

Table C.IV: Analyst forecast bias in EPS

The table presents the average bias, its standard deviation and the number of observations of analyst EPS forecast biases obtained from Zhang et al. (2024), based on the approach of van Binsbergen et al. (2022). Compustat-CRSP-I/B/E/S sample: 1980-2024.

	mean	std	N
1y	0.020	0.217	44,290
2y	0.038	0.227	40,128

Table C.V: Additional filters and model averaging

The table presents the sensitivity of level and slope estimates to alternative filters and cash flow models. The first and the second row stand for scenarios where cash flow forecasts are averaged across two approaches (the mean-reversion model and the augmented version with I/B/E/S forecasts) using the equal-weighting scheme or weights based on the normal distribution, where the mean and the standard deviation are based on the empirical values. The third row represents the case where the estimation sample is trimmed at the higher level of 5th and 95th percentile of the sample IRR. The forth and the fifth rows describe scenarios where different procedures are combined. The sixth row summarizes the results when the estimation is performed excluding the most extreme book-to-market portfolios within each size quintile. The final row presents estimates when using the approach of Li et al. (2013) to forecast cash flows. Compustat-CRSP-I/B/E/S subsample: 1980-2024.

	level	slope
1) cash flows equally averaged	0.08^{**}	0.18^{**}
	(0.00)	(0.01)
2) cash flows averaged using normal dist	0.09^{**}	0.21^{**}
	(0.00)	(0.02)
3) trimmed at 5% and 95% sample IRR	0.08^{**}	0.11^{**}
	(0.00)	(0.01)
combination of 1) & 3)	0.08^{**}	0.19^{**}
	(0.00)	(0.01)
combination of 2) & 3)	0.09^{**}	0.23^{**}
	(0.00)	(0.01)
excluding extreme BM portfolios	0.09^{**}	0.24^{**}
	(0.01)	(0.02)

Table C.VI: Industry analysis

The table presents the estimation results within Fama-French 15 industries. The technology industry classification is constructed separately following Barron et al. (2002). Weight denotes the market cap weight of the industry-specific observations in the total sample. All Fama-French industries are sorted based on the slope estimates. Compustat-CRSP-I/B/E/S subsample: 1980-2024.

	level	slope	weight
Transportation	0.08	0.01	0.05
Oil and Petroleum Products	0.08	0.04	0.09
Steel Works Etc	0.08	0.04	0.01
Mining and Minerals	0.08	0.06	0.01
Other	0.08	0.06	0.32
Automobiles	0.10	0.08	0.02
Machinery and Business Equipment	0.09	0.08	0.12
Textiles, Apparel and Footware	0.07	0.09	0.00
Chemicals	0.08	0.10	0.02
Construction and Construction Materials	0.08	0.11	0.01
Fabricated Products	0.07	0.13	0.00
Retail Stores	0.08	0.14	0.05
Food	0.08	0.14	0.04
Drugs, Soap, Perfumes, Tobacco	0.10	0.16	0.10
Consumer Durables	0.08	0.18	0.01
Technology	0.09	0.10	0.37

Appendix D Standard Errors

The estimation procedure is equivalent to non-linear least squares (NLS). For each firm i and time t, we can use equation (4) to express the value of the firm normalized by its stock price as:

$$1 = g(x_{i,t}; \boldsymbol{\beta}) + \epsilon_{i,t} \tag{14}$$

Here $\epsilon_{i,t}$ is a pricing error and $g(x_{i,t}; \boldsymbol{\beta})$ is a (scaled) model price of the firm. We can stack all observations and rewrite the optimization in a vector form:

$$\underset{\boldsymbol{\beta}}{\operatorname{argmin}}\, \boldsymbol{\epsilon}^T \boldsymbol{W} \boldsymbol{\epsilon} \tag{15}$$

where $\boldsymbol{W}$ is a diagonal matrix of optimization weights, $diag(\boldsymbol{W}) = \boldsymbol{w}$. Following Greene (2000), the NLS weighted heteroscedasticity-robust standard errors can be computed in the following way:

$$\text{Cov}(\hat{\boldsymbol{\beta}}) = (\boldsymbol{J}^T \boldsymbol{W} \boldsymbol{J})^{-1} \boldsymbol{J}^T \boldsymbol{W} \boldsymbol{\Omega} \boldsymbol{W} \boldsymbol{J} (\boldsymbol{J}^T \boldsymbol{W} \boldsymbol{J})^{-1} \tag{16}$$

where $\boldsymbol{J}$ is the Jacobian of the model with respect to the parameters, evaluated at the estimated parameters, $\hat{\boldsymbol{\beta}}$:

$$\boldsymbol{J} = \left. \frac{\partial g(\boldsymbol{X}; \boldsymbol{\beta})}{\partial \boldsymbol{\beta}} \right|_{\boldsymbol{\beta} = \hat{\boldsymbol{\beta}}} \tag{17}$$

and $\boldsymbol{\Omega}$ is the variance-covariance matrix of pricing errors. Given that most firms are correlated and do not have many time-series observations, estimating this huge variance-covariance matrix can be highly inefficient and unstable over time. For this reason, we parametrize $\boldsymbol{\Omega}$ as a function of two key parameters that govern the time-series, ρ_t, and cross-sectional correlation, ρ_c, of pricing errors:

$$\boldsymbol{\Omega} = \sigma_\epsilon^2 \boldsymbol{\Sigma}, \qquad \text{where } \boldsymbol{\Sigma}_{j,k} = \begin{cases} \rho_c & \text{if } \boldsymbol{t}_j = \boldsymbol{t}_k \ \& \ \boldsymbol{\iota}_j \neq \boldsymbol{\iota}_k, \\ \rho_t^{|\boldsymbol{t}_j - \boldsymbol{t}_k|} & \text{if } \boldsymbol{\iota}_j = \boldsymbol{\iota}_k, \\ 0 & \text{otherwise.} \end{cases} \tag{18}$$

Here $\boldsymbol{t}$ and $\boldsymbol{\iota}$ are vectors of all time indices and cross-sectional identifiers, and σ_ϵ^2 is the variance of pricing errors. We can represent this correlation matrix $\boldsymbol{\Sigma}$ as a sparse matrix with non-missing block-diagonal elements and non-zero contemporaneous correlations. An example is shown below:

$$\boldsymbol{\Sigma} = \begin{pmatrix} \begin{matrix} 1 & \rho_t \\ \rho_t & 1 \end{matrix} & \mathbf{0} & \cdots & \begin{matrix} \rho_c & 0 \\ 0 & 0 \end{matrix} \\ \mathbf{0} & \begin{matrix} 1 & \rho_t^2 \\ \rho_t^2 & 1 \end{matrix} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \mathbf{0} \\ \begin{matrix} \rho_c & 0 \\ 0 & 0 \end{matrix} & \mathbf{0} & \mathbf{0} & \begin{matrix} 1 & \rho_t \\ \rho_t & 1 \end{matrix} \end{pmatrix}$$

Given that discount rate parameters, $\boldsymbol{\beta}$, enter the pricing equation in a highly non-linear fashion, the elements of Jacobian are estimated numerically:[21]

$$\left.\frac{\partial f(\boldsymbol{X}; \beta_1, \beta_2)}{\partial \beta_1}\right|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}} \approx \left.\frac{f(\boldsymbol{X}; \beta_1 + \Delta, \beta_2) - f(\boldsymbol{X}; \beta_1 - \Delta, \beta_2)}{2\Delta}\right|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}}. \tag{19}$$

[21] Δ is arbitrarily set to 10^{-4}, but the results are insensitive to other choices of this parameter, i.e. $\Delta = 10^{-6}, 10^{-5}, 10^{-3}, 10^{-2}$.

Appendix D.1 Assumptions on Correlation Structure

To highlight the significance of $\mathbf{\Omega}$ in standard error calculations, Table D.I presents baseline estimation results under six correlation matrix assumptions:

Table D.I: Standard errors under different Σ assumptions

The table summarizes standard errors of baseline estimate under six different specifications: (i) column presents IID error terms, (ii) - error terms under the high time-series (0.8) correlation scenario, (iii) - error terms under the high cross-sectional (0.8) correlation scenario, (iv) - error terms under the block-diagonal matrix with at least 2 observations to calculate pairwise correlations, (v) - error terms under the full correlation matrix, (vi) - error terms under the block-diagonal matrix with at least 5 observations to calculate pairwise correlations.

	$\hat{\beta}$ $[\rho_t, \rho_c] =$	IID $[0, 0]$	High ρ_t $[0.8, 0]$	High ρ_c $[0, 0.8]$	Block-Diag $[0.20, 0.16]$	Full $[0.20, 0.16]$	Block-Diag, $n \geq 5$ $[0.29, 0.22]$
Level	0.083	0.001	0.002	0.004	0.002	0.002	0.002
		(100.28)	(43.00)	(20.55)	(40.81)	(35.88)	(35.96)
Slope	0.075	0.006	0.014	0.023	0.013	0.014	0.014
		(11.71)	(5.34)	(3.24)	(5.99)	(5.31)	(5.33)

- The first approach assumes IID pricing errors, which simplifies the variance-covariance matrix to $\mathbf{\Omega} = \sigma_\epsilon^2 I$. Given the large number of observations (50,265), the standard errors are tiny for the level estimates and moderate for the slope.
- The next two columns highlight the importance of time-series and cross-sectional dependence while keeping the block-diagonal nature of $\mathbf{\Omega}$, as defined in equation (18). Since most of the identification comes from the cross-sectional dimension (the sample contains 45 years of data and, on average, 1,100 firms), the standard errors barely change with high ρ_t but experience a staggering increase once ρ_c goes up to 0.8. Note that even in such extreme circumstances, the slope estimates are still statistically significant.
- The fifth column presents estimates based on the average sample time-series (0.20) and pairwise cross-sectional (0.16) correlation. The standard errors more than double in this instance. However, these firm-level correlation estimates are noisy and potentially

downward-biased because there are plenty of companies with just a few years of data.

- The six column utilizes a more flexible specification of correlation matrix where we allow for a non-contemporaneous correlation across firms:

$$\tilde{\mathbf{\Sigma}}_{j,k} = \begin{cases} \rho_c \rho_t^{|\boldsymbol{t}_j - \boldsymbol{t}_k|} & \text{if } \boldsymbol{\iota}_j \neq \boldsymbol{\iota}_k, \\ \rho_t^{|\boldsymbol{t}_j - \boldsymbol{t}_k|} & \text{if } \boldsymbol{\iota}_j = \boldsymbol{\iota}_k. \end{cases} \tag{20}$$

  Since the empirical correlations are fairly small, the adjustment is not substantial and delivers slightly more conservative standard errors.

- Finally, the last column lists estimates under the original form of $\mathbf{\Sigma}$ defined in equation (18), with more conservatively estimated ρ_t (0.29) and ρ_c (0.22). Here we use at least 5 common observations to compute firm-level time and cross-sectional correlations. This approach yields higher standard errors, which are fairly close to those in column six. As it strikes a good balance between computational simplicity and robustness, we adopt it for the remainder of the analysis.