The Cross-Section and Time Series of Stock and Bond Returns

Ralph S.J. Koijen*
NYU, NBER, and CEPR

Hanno Lustig†
Stanford and NBER

Stijn Van Nieuwerburgh‡
NYU, NBER, and CEPR

April 24, 2017

Abstract

We show that bond factors, which predict future U.S. economic activity at business cycle horizons, are priced in the cross-section of U.S. stock returns. High book-to-market stocks have larger exposures to these bond factors than low book-to-market stocks, because their cash flows are more sensitive to the business cycle. Because of this new nexus between stock and bond markets, a parsimonious three-factor dynamic no-arbitrage model can be used to jointly price book-to-market-sorted portfolios of stocks and maturity-sorted bond portfolios, while reproducing the time-series variation in expected bond returns. The business cycle itself is a priced state variable in stock and bond markets. JEL: G12

*Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; rkoijen@stern.nyu.edu; http://www.koijen.net.

†Department of Finance, Stanford GSB, 655 Knight Way, Stanford CA 94305; hlustig@stanford.edu; https://www.gsb.stanford.edu/faculty-research/faculty/hanno-lustig.

‡Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; svnieuwe@stern.nyu.edu; http://www.stern.nyu.edu/~ svnieuwe.

This is a preprint version of the article. The final version may be found at < https://doi.org/10.1016/j.jmoneco.2017.05.006 >.

Value investors buy stocks that have low prices relative to measures of fundamentals such as dividends or book assets, and sell stocks that have high prices relative to fundamentals. These strategies earn high returns that appear anomalous relative to standard models such as the CAPM (e.g., Basu, 1977; Fama and French, 1992). The profession has hotly debated whether these superior returns reflect a behavioral bias or a compensation for systematic risk. Under the behavioral hypothesis, extrapolative investors push up the price of growth ("glamour") stocks that performed well in the recent past, allowing contrarian investors to profit from their over-optimism by investing in out-of-favor value stocks and/or shorting the growth stocks (De Bondt and Thaler, 1985). Leading risk-based explanations of the value premium rely on differences in the riskiness of assets in place relative to growth options (Zhang, 2005) or differences in the duration of cash flows of value and growth stocks (Lettau and Wachter, 2007).

Early attempts to connect the cash flows of value and growth firms to macro-economic sources of risk were unsuccessful (Lakonishok, Schleifer, and Vishny, 1994). Our paper provides new evidence that links the excess returns on high minus low book-to-market stock portfolios to cash flow and output risk at business cycle frequencies. We study a much longer sample with more adverse macroeconomic events than previously examined (1926-2012 compared to 1968-1989 in Lakonishok, Schleifer, and Vishny (1994), or 15 recessions compared to 4). We develop and apply a new methodology to study macroeconomic events as well.

The connection between the value spread and the macro-economy is easiest to detect in the bond market. We study several linear combinations of bond yields that forecast future economic activity at business cycle horizons: the Cochrane-Piazzesi factor (CP, Cochrane and Piazzesi (2005)), the slope of the term structure, and the best yield-based linear predictor of economic activity at the one-year horizon. We show that innovations in these bond market factors strongly co-move with the returns on value-minus-growth. Since the bond market variables isolate a component of expected growth that is not persistent, our findings assign a central role to the business cycle as a priced state variable.

Our paper makes three contributions. The first is to document that value portfolio returns have a higher covariance with innovations in the bond factors that predict future economic activity at business cycle horizons than growth portfolio returns. This pattern of exposures is consistent with a value premium provided that bond factor innovations carry a positive risk price. Since these innovations represent good news about future output growth and thus lower the marginal utility of wealth for an average forward-looking investor, it is natural that investors assign them a positive risk price.

The second contribution is to attribute these different bond exposures to differences in the underlying cash flow dynamics. We find that value stocks experience negative cash-flow shocks in economic downturns. There are large differences in the behavior of cash-flow growth on value and growth over the macro-economic cycle. We also show that periods in which the bond factors are low are periods of significantly lower future dividend growth rates on the market portfolio and on the value-minus-growth portfolio.

One useful way to highlight the macro-economic risk in value strategies is to select periods during which value stocks and the value-minus-growth strategy experience exceptionally low returns, which we label "low-value events." Such low-value events are not only associated with low contemporaneous CP realizations, but also with low future economic activity and lower future dividend growth on value-minus-growth, consistent with a risk-based explanation. This event-based approach allows us to detect the link between prices, cash-flows, and macroeconomic aggregates in high marginal utility states of the world that matter most for pricing. The approach could prove fruitful for investigating other return anomalies and their link to the macro-economy.

Our third contribution is to build on this evidence linking the value spread to the bond factors to develop a parsimonious three-factor model that prices the cross-section of stock and bond returns. Our first pricing factor consists of innovations to the CP factor (results are similar for the other bond factors that predict future economic activity): differential exposure of the five book-to-market portfolios accounts for the average value spread in the data. Second, differential exposure to shocks to the level of the term structure accounts for the difference in excess returns on five maturity-sorted government bond portfolios, consistent with Cochrane and Piazzesi (2008). Third, exposure to the market return accounts for the aggregate equity premium. This three-factor model reduces mean absolute pricing errors on our test assets from 5.04% per year in a risk-neutral benchmark economy to 0.49% per year. By having the price of level risk depend on the lagged bond factor, the model also captures the predictability of bond returns by the CP factor. All of the estimated risk prices have the expected sign, and are collectively significantly different from zero.

We explore the robustness of the results for different sub-samples and for different sets of test assets. The model prices a set of corporate bond portfolios sorted by credit rating, jointly with equity and government bond portfolios. While it prices several other equity portfolio sorts, our model cannot explain the cross-section of momentum or return-to-equity portfolios.

What results is a coherent picture of value-minus-growth returns, the bond yield factor, macroeconomic activity, and dividend growth on value-minus-growth that is potentially consistent with a risk-based resolution of the value premium puzzle. Our parsimonious stochastic discount factor model makes progress towards a unified pricing model of stock and bond markets.

Furthering this connection, in the last part of the paper, we set up and solve a simple dynamic asset pricing model that can account for the empirical facts we document and that lends a structural interpretation to the three priced sources of risk. Shocks to the bond risk premium are shocks to a leading indicator of the business cycle, level shocks are expected inflation shocks, and market shocks reflect compensation for dividend growth risk.

The rest of the paper is organized as follows. Section 1 discusses the related literature. Section 2 reports the main results documenting the link between CP and the macro-economy, while Section 3 contains the main asset pricing results. In Section 4, we document the robustness of our results to other test assets and estimation methods. Section 5 contains a simple asset pricing model that formalizes the connections between the bond risk premium, the value premium, and the macro-economy. Section 6 concludes.

1 Related Literature

Researchers working in a small but growing literature model stock and bond returns jointly, most often in affine settings like ours. They have mostly examined the relation between the aggregate stock and bond markets,[1] with the exception of Lettau and Wachter (2007, 2011) and Gabaix (2012), who also study the cross-section of stock returns. The former is a model with common shocks to the risk premium in stock and bond markets, while the latter is a time-varying rare disasters model.

[1]Examples are Bakshi and Chen (2005) and Bekaert, Engstrom, and Xing (2009) in a Gaussian setting and Campbell, Sunderam, and Viceira (2013) in a linear-quadratic model. Lustig, Van Nieuwerburgh, and Verdelhan (2013) price both nominal bond yields and the aggregate stock market return in a no-arbitrage model in order to measure the wealth-consumption ratio in the data; they do not study the cross-section of bond nor stock returns.

In addition, work in production-based asset pricing has linked the investment behavior of value and growth firms during recessions to the value premium (Zhang, 2005). This literature has focused on explaining the cross-section of stock returns. Other notable contributions linking the cross-section of stock returns to firm characteristics in the production-based asset pricing literature are Berk, Green, and Naik (1999), Gomes, Kogan, and Zhang (2003), Papanikolaou (2007), Liu, Whited, and Zhang (2009), Kogan and Papanikolaou (2014), and Hou, Xue, and Zhang (2015). We are unaware of work in this area that also accounts for the cross-section of Treasury and corporate bond returns.

The business cycle itself plays a secondary role in modern dynamic asset pricing theory.[2] We uncover new evidence that the business cycle in output and consumption growth is itself a priced state variable in stock markets. Value stock returns are more sensitive than growth stock returns to innovations in bond market factors such as CP. Therefore, they are more exposed to cyclical news about the economy's future cash flow growth, because their subsequent cash flow growth is more sensitive to output growth. Value stocks earn a premium as a result. Relative to existing dynamic asset pricing models, our work uncovers the cyclical component in expected output growth as a new priced state variable, distinct from the low frequency state variables in long-run risk of Bansal and Yaron (2004) and external habit models of Campbell and Cochrane (1999). These models are designed to match the lower frequency variation in the market dividend yield.[3] Whether the market price assigned to transitory business cycle risk in existing dynamic asset pricing models is large enough to match equity market, value, and bond risk premia with reasonable parameter choices is an open question.

Our paper advances the empirical ICAPM literature, starting with the seminal work of Chen, Roll, and Ross (1986). These authors use term structure factors either as a predictor of the aggregate return on the stock market or as a conditioning variable in an estimation of a conditional beta model of the cross-section of stock returns. Ferson and Harvey (1991) study stock and bond returns' sensitivity to aggregate state variables, one of which is the slope of the yield curve. They conclude that time variation in equity risk premia is important for understanding the cross-sectional variation in size

[2]A related literature studies the temporal composition of risk in asset prices, (e.g., Cochrane and Hansen, 1992; Kazemi, 1992; Bansal and Lehman, 1997; Hansen, Heaton, and Li, 2008).

[3]These models are successful in accounting for many of the features of both stocks and bonds. For the external habit model, the implications for bonds were studied by Wachter (2006) and the implications for the cross-section of stocks were studied by Menzly, Santos, and Veronesi (2004) and Santos and Veronesi (2010). Likewise, the implications of the long-run risk model for the term structure of interest rates were studied by Piazzesi and Schneider (2006), Kung (2015), and Bansal and Shaliastovich (2013), while Hansen, Heaton, and Li (2008) study the implications for the cross-section of equity portfolios.

and industry equity portfolios, and that time variation in interest rate risk premia are important for understanding the cross-sectional variation in bond return portfolios. Brennan, Wang, and Xia (2004) analyze an ICAPM model in which the real rate, expected inflation, and the Sharpe ratio dynamically change investment opportunity set and show that this model prices the cross-section of stocks. Similarly, Petkova (2006) studies the connection between the Fama-French factors and innovations in state variables such as the default spread, the dividend-price ratio, the yield spread, and the short rate. Using a VAR model, Campbell and Vuolteenaho (2004) and Campbell, Polk, and Vuolteenaho (2010) argue that common variation in book-to-market portfolio returns can be attributed to news about future cash flow growth on the market. In their approach, the cash flow innovations are highly persistent. In contrast to this literature, our focus is on the joint pricing of stock and bond returns, business cycle shocks, and the link with dividend growth on stock portfolios. Baker and Wurgler (2012) show that government bonds co-move most strongly with "bond-like stocks," which are stocks of large, mature, low-volatility, profitable, dividend-paying firms that are neither high growth nor distressed. They propose a common sentiment indicator that drives stock and bond returns.

2 Measuring Business Cycle Risk in Value Stocks

We start by documenting that value stocks suffer from bad cash-flow shocks at times when a representative investor experiences high marginal utility growth. Because dividends adjust to bad shocks with a lag, it is natural to look for early indicators of poor future economic performance. Researchers have traditionally looked at bond markets for expectations about future economic activity. We follow that tradition and document the predictive ability of several linear combinations of bond yields. These bond market variables are strong predictors of both future aggregate economic activity, future aggregate dividend growth, and future dividend growth on value-minus-growth stock portfolios. To bolster the macro-economic risk explanation, in the last part of this section, we examine periods where realizations on both the value and the value-minus-growth portfolios are exceptionally low, and finds that these are periods characterized by bad news about future aggregate economic activity. The main text focuses on the CP factor as the bond factor, we explore the robustness to other yield-curve variables in Appendix A.

2.1 Cash-Flow Risk in Value-Growth and the Business Cycle

We use monthly data from the Center for Research on Securities Prices (CRSP) on dividends and inflation from July 1926 until December 2015. Inflation is measured as the change in the Consumer Price Index from the Bureau of Labor Statistics. We use the return on the value-weighted NYSE-AMEX-NASDAQ index from CRSP as the market return. Dividends on book-to-market-sorted quintile portfolios are calculated from cum-dividend and ex-dividend returns available from Kenneth French's data library. To eliminate seasonality in dividends, we construct annualized dividends by adding the current month's dividends to the dividends of the past 11 months.[4] We form log real dividends by subtracting the log of the consumer price index from the log of nominal dividends. Our focus is on cash dividends.[5] It is important to note that all quintile portfolios, including the growth portfolio 1, distribute substantial amounts of dividends. The average annual dividend yield varies only modestly across book-to-market quintile portfolios: 2.5% (portfolio 1), 3.4% (2), 3.8% (3), 3.9% (4), and 3.0% (5). The market portfolio has an average dividend yield of 3.3%. Additional summary statistics for variables in this section are reported in Appendix A.

In the left panel of Figure 1, we plot log real dividends on book-to-market quintile portfolios 1 (G for growth), 5 (V for value), and the market portfolio (M) against the NBER recession dates defined by the NBER's Business Cycle Dating committee. For consistency with the asset pricing results that are to follow, we focus on the post-1952.7 sample. The figure shows strong evidence that the dividends on value stocks fall substantially more in recessions than those of growth stocks. Value stocks' cash flows show strong cyclical fluctuations whereas dividends on growth stocks are, at best, a-cyclical. The picture for the pre-1952 period, reported in Appendix A, is consistent with this behavior. The two starkest examples of the differential cash-flow behavior of value and growth are the Great Depression (September 1929 - March 1933) and the Great Recession (December 2007 - June 2009), but the same pattern holds during most post-war recessions (e.g., 1973, 1982, 1991, 2001).

[4] Investing dividends at the risk-free rate yields similar results. Binsbergen and Koijen (2010) show that reinvesting monthly dividends at the market return severely contaminates the properties of dividend growth.

[5] Cash dividends are the right measure in the context of a present-value model that follows a certain portfolio strategy, such as value or growth (Hansen, Heaton, and Li, 2008). An alternative is to include share repurchases to cash dividends, but this would correspond to a different dynamic strategy (Larrain and Yogo, 2007). However, in the most recent recession, which is the largest downturn in cash dividends during the period in which repurchases became more popular, share repurchases also declined substantially. This suggests that during the episodes that we are most interested in, cash dividends and share repurchases comove positively and are exposed to the same aggregate risks.

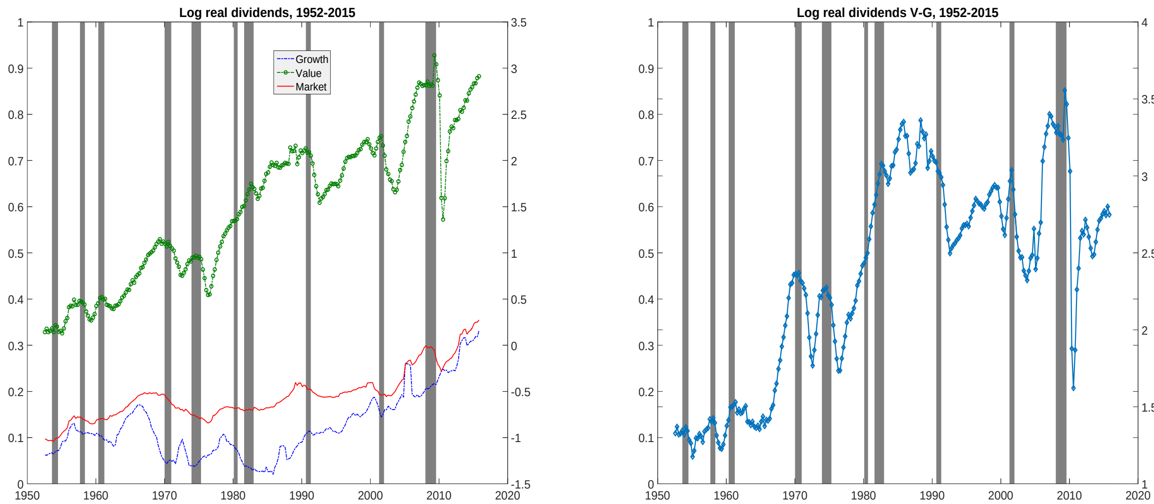

Figure 1: Dividends on value, growth, and market portfolios.

The left panel shows the log real dividend on book-to-market quintile portfolios 1 (growth, dashed line with squares) and 5 (value, dotted line with circles) and on the CRSP value-weighted market portfolio, plotted against the right axis. The right panel shows the log real dividend on book-to-market quintile portfolios 5 (value) minus the log real dividend on the book-to-market portfolio 1 (growth), plotted against the right axis. The grey bars indicate official NBER recession dates. Dividends are constructed from the difference between cum- and ex-dividend returns on these portfolios, multiplied by the previous month's ex-dividend price. The ex-dividend price is normalized to 1 for each portfolio in 1926.06. Monthly dividends are annualized by summing dividends received during the year. We take logs and subtract the log of the CPI price level (normalized to 100 in 1983-84) to obtain log real dividends. The data are monthly from July 1952 until December 2015 and are sampled every three months in the figure.

Strictly adhering to the NBER recession dates understates the change in dividends from the highest to their lowest point over the cycle. The right panel of Figure 1 shows the log difference between value and growth portfolios (right axis) as well as NBER recessions (bars). The figure illustrates not only large declines in dividends on value-minus-growth around recessions, as well a lag in the declines when compared to the NBER peak. This may reflect the downward stickiness in dividend adjustments that is well understood in the literature on firms' dividend payment behavior.[6]

2.2 Bond Factors and the Business Cycle

Having shown that dividends on value-minus-growth fall during and after recessions, this section shows that bond yield factors predict the incidence of recessions. Here, we show that the CP factor forecasts aggregate economic activity, aggregate dividend growth, and dividend growth on value-minus-

[6]For example, Yoon and Starks (1995) present evidence that firms cut their dividends much less frequently than they increase them, but when they cut them, they cut them at a rate that is five times larger than when they increase them. See also Chen (2009) for aggregate evidence on dividend smoothing.

growth stock portfolios. We follow Cochrane and Piazzesi (2005) in constructing the CP factor as a linear combination of 2- through 5-year government bond yields that bests forecasts future excess bond returns.[7] Appendix A shows that these results extend to two alternative linear combinations: the slope of the yield curve and the linear combination of bond yields that best forecasts future economic activity. Our findings contribute to the recent literature that links bond market variables to macroeconomic activity.[8]

We consider the following predictive regression in which we forecast future economic activity, measured by the Chicago Fed National Activity Index ($CFNAI$),[9] using the current CP factor:

$$CFNAI_{t+k} = c_k + \beta_k CP_t + \varepsilon_{t+k}, \tag{1}$$

where k is the forecast horizon expressed in months. The regressions are estimated by OLS and we calculate Newey-West standard errors with $k-1$ lags. The sample runs from March 1967 until December 2015, dictated by data availability. The left panels in Figure 2 show the coefficient β_k in the top panel, its t-statistic in the middle panel, and the regression R-squared in the bottom panel. The forecast horizon k is displayed on the horizontal axis and runs from 1 to 36 months. CP is strongly and significantly positively associated with future economic activity. All three statistics display a hump-shaped pattern, gradually increasing until approximately 12-24 months and then gradually declining afterwards.

[7] We use monthly Fama-Bliss zero-coupon yield data, available from June 1952 until December 2015, on nominal government bonds with maturities of one- through five-years to construct one- through five-year forward rates. We then regress the equally-weighted average of the one-year excess return on bonds of maturities of two, three, four, and five years on a constant, the one-year yield, and the two- through five-year forward rates. The yields are one-year lagged relative to the return on the left-hand side. The CP factor is the fitted value of this predictive regression. The R^2 of this regression in our sample of monthly data is 18.1%, roughly twice the 10.7% R^2 of the five-year minus one-year yield spread, another well-known bond return predictor.

[8] Brooks (2011) shows that the CP factor has a 35% contemporaneous correlation with news about unemployment, measured as deviations of realized unemployment from the consensus forecast. Gilchrist and Zakrajsek (2012) shows that a credit spread, and in particular a component related to the bond risk premium, forecasts economic activity. A related literature examines the predictability of macro-economic factors for future bond returns. Cooper and Priestley (2008) show that trend deviations in industrial production forecast future bond returns; Joslin, Priebsch, and Singleton (2014) incorporate this finding in an affine term structure model. Ludvigson and Ng (2009) shows that a principal component extracted from many macroeconomic series also forecasts future bond returns. While macro-economic series do not fully incorporate the variation in bond risk premia, there clearly is an economically meaningful link between them.

[9] The $CFNAI$ is a weighted average of 85 existing monthly indicators of national economic activity. $CFNAI$ peaks at the peak of the business cycle and bottoms out at the trough. Since economic activity tends toward trend growth over time, a positive index reading corresponds to growth above trend and a negative index reading corresponds to growth below trend. $CFNAI$ is normalized to have mean zero and standard deviation one.

The maximum slope is 24, with a t-statistic of 4.2 and an R^2 of 13.2%. This maximum predictability is for $CFNAI$ 21 months later. From Figure 2 we infer that a high CP factor precedes higher economic activity about 12 to 24 months later. At the 24-month horizon, CP is close to the best predictor in the class of linear combinations of 1- through 5-year bond yields. The predictability is statistically significant for horizons from 1 month to 31 months. Appendix A shows similar results when forecasting GDP growth rather than $CFNAI$.

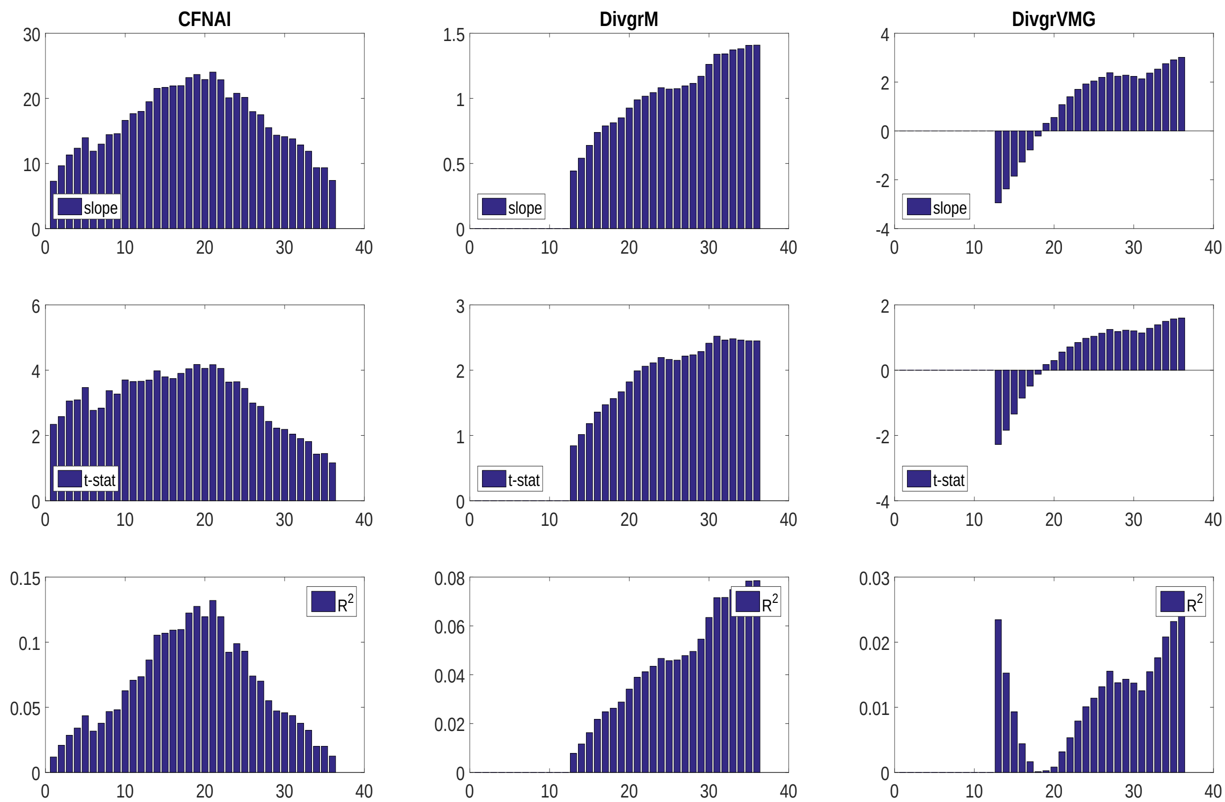

Figure 2: Economic activity predicted by bond factors.

We consider a regression of future values of $CFNAI$, which we normalize to have mean zero and standard deviation one, on the current CP factor: $CFNAI_{t+k} = c_k + \beta_k CP_t + \varepsilon_{t+k}$, where k is the forecast horizon expressed in months. The regressions are estimated by OLS and we calculate Newey-West standard errors with $k-1$ lags. The top panel displays the predictive coefficient β_k, the middle panel the t-statistic, and the bottom panel the corresponding R^2. We consider $k = 1, \ldots, 36$ months of lags, displayed on the horizontal axis in each panel, and the t-statistics are computed using Newey-West standard errors with $k-1$ lags. In all three columns, the predictor is the CP factor. In the left column, $CFNAI_{t+k}$ is the dependent variable. In the middle column, the aggregate dividend growth rate Δd_{t+k} is the dependent variable. In the last column, the dividend growth rate on value minus growth $\Delta d^V_{t+k} - \Delta d^G_{t+k}$ is the dependent variable. The sample is March 1967 until December 2015.

Having shown earlier that both aggregate dividend growth and dividend growth on value-minus-growth stocks declines around recessions, we now ask whether the bond yield factor (CP) predicts aggregate dividend growth and dividend growth on value-minus-growth stocks. We employ linear regressions like equation (1). Since dividend growth is constructed using 12 months of data, we only consider horizons $k \geq 12$. The predictive coefficients, t-statistics, and R-squared values for the aggregate dividend growth on the market (value minus growth) are summarized in the middle (right) column of Figure 2. CP strongly predicts aggregate dividend growth, especially 2-3 years out. The right column shows that our bond market variable also linearly predict dividend growth on value-minus-growth, although the statistical significance is weaker. The predictability of CP is concentrated at longer horizons of 33-36 months ahead. Table A.II in the Appendix contains the point estimates. This regression evidence implies that the bond market contains useful information about future cash flow growth in the aggregate and about differential cash-flow prospects for value and growth firms.

2.3 A Macro-Event Study of Value

In this section, we further explore the connection between value and growth returns, CP, and the macro-economy.

2.3.1 Low-CP Events

While the bond yield variables clearly lead the cycle, their exact timing vis-a-vis the official NBER recession dating is fragile because the lead-lag pattern may fluctuate from one recession to the next (see Figure A.2 in the Appendix). Thus, it may be informative to isolate periods in which CP is low and then to ask how the level of economic activity behaves around such events.

In each quarter since 1952.Q3 we compute quarterly CP as the CP factor value in the last month of that quarter, and we select the 25% of quarters with the lowest quarterly CP readings. Figure 3 shows how several series of interest behave six quarters before (labeled with a minus sign) until ten quarters after (labeled with a plus sign) the low-CP event, averaged across such events. The quarter labeled '0' in Figure 3 is the event quarter with the lowest CP reading. The top right panel shows the dynamics of CP itself, which naturally falls from a positive value in the preceding quarters to a highly negative value in the event quarter, after which it recovers.

The top left panel of Figure 3 shows quarterly returns on value-minus-growth. The value spread is demeaned over the full sample. The evidence presented in the introduction suggests a link between *innovations* in CP and returns on value-minus-growth. This panel is consistent with that evidence. Between quarters -2 and -1 and -1 and 0, the CP factor falls sharply while between quarter 0, +1, and +2, CP rises sharply. This figure shows that realized returns on the value-minus-growth strategy are negative in quarter -1 and 0 and but rise in quarters +1 and +2 (at which point they are slightly positive once we add back in the mean). This is consistent with the higher exposure of value stocks to CP innovations than the exposure of growth stocks. The top left panel of Figure 3 provides evidence against the interpretation of the CP shock as a discount rate shock (instead of, or in addition to, a shock to expected cash flows on value-minus-growth). Indeed, for CP shocks and *realized* value-minus-growth returns to be positively contemporaneously correlated, *expected* future returns on value-minus-growth would have to be particularly high upon a negative CP shock. This is belied by the low average value-minus-growth return in the quarters following the low CP event. We return to the relationship between value-minus-growth returns and the CP factor in detail in Section 3.[10]

The bottom left panel of Figure 3 shows annual dividend growth on value-minus-growth (fifth-minus-first book-to-market portfolio) over the CP cycle. The dividend growth differential is demeaned over the full sample, so as to take out the trend in the dividend growth rate differential. Dividend growth on value-minus-growth is high when CP is at its nadir and starts falling afterwards. This decline in value-minus-growth dividend growth is persistent and economically large. Comparing the bottom two panels, we see that dividend growth lags economic activity by several quarters. This lagged reaction arises in part because firms are reluctant to cut dividends, and only do so after a bad shock (like a low-CP event). In other part, the lag arises from the construction of the dividend growth measure. Since dividend growth is computed using the past twelve months of dividends, it is not until the end of quarter +4 that all dividends, used in the measured growth rate, are realized after the time-0 shock. In sum, low CP realizations predict low future dividend growth rates on value-minus-growth, but with a considerable lag. This evidence confirms the formal regression evidence discussed above.

Finally, the bottom right panel shows the economic activity index $CFNAI$ over this CP cycle.

[10] An adequate description of dividend dynamics contains at least two shocks: one shock that equally affects dividend growth rates on all portfolios, and a second shock (to the CP factor) that affects value dividends relative to growth dividends. The Appendix discusses the evidence against a one-factor model.

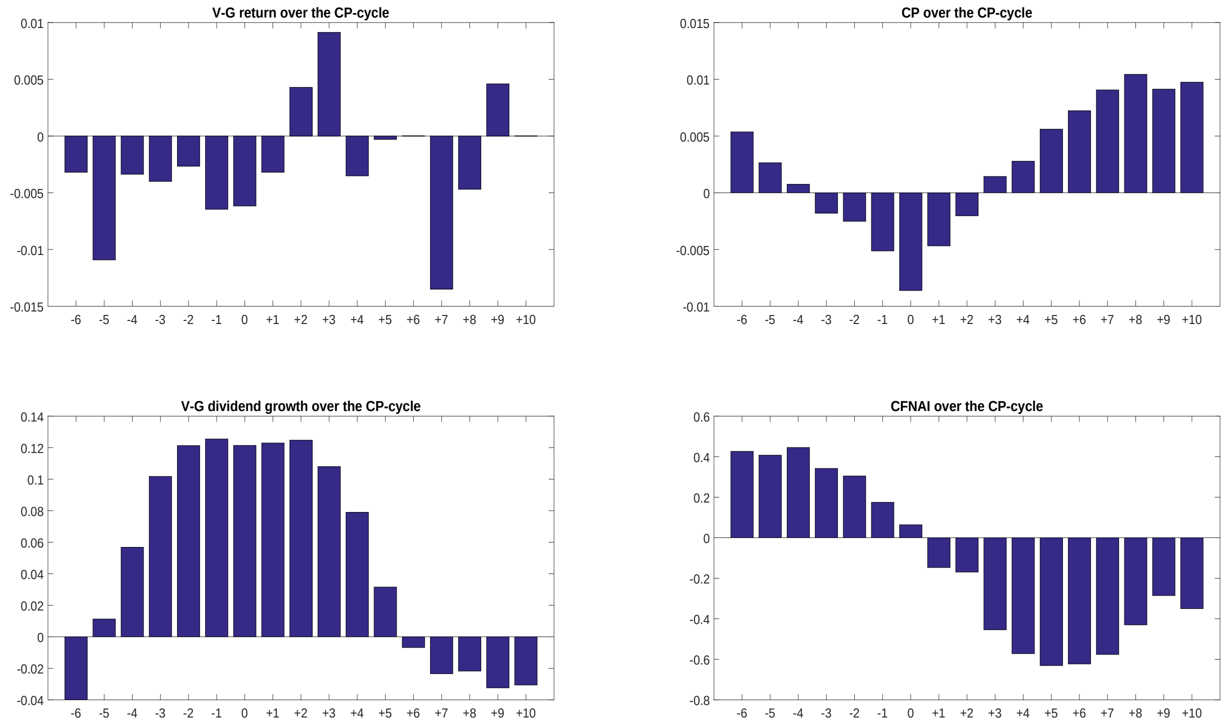

Figure 3: Low CP events

The figure shows four quarterly series in event time. The event is defined as a quarter in which the quarterly CP factor in its respective lowest 25% of observations. This selection leads to 63 events out of 254 quarters. The sample runs from 1953.Q3 until 2015.Q4. In each panel, the period labeled '0' is the quarter in which the event takes place. The labels -1, -2, -3, etc refer to one, two, three, etc quarters before the event whereas the labels +1, +2, +3, etc. refer to one, two, three, etc quarters after the event. The top left panel plots the realization of the quarterly log return on value-minus-growth. The bottom left panel reports annual log dividend growth on value-minus-growth. The top right panel plots the CP factor. The bottom right panel plots the $CFNAI$ index of economic activity. The latter is available only from 1967.Q2 onwards. Formally, the graph reports $c_k + \beta_k$ from a regression $X_{t+k} = c_k + \beta_k \mathcal{I}_{CP_t<LB} + \epsilon_{t+k}$, for various k, where $\mathcal{I}$ is an indicator variable, LB is the 25th percentile of CP, and X is the dependent variable which differs in each panel. Value-minus-growth returns and value-minus-growth dividend growth have been demeaned over the full sample; $CFNAI$ is also mean zero by construction.

There is a clear pattern in economic activity in the quarters surrounding the low-CP event. When CP is at its lowest point, economic activity is about average ($CFNAI$ is close to zero). $CFNAI$ then turns negative for the next ten quarters, bottoming out five to seven quarters after the CP event. This lead-lag pattern is consistent with the predictability evidence shown above. The change in $CFNAI$ from four quarters before until four quarters after is economically large, representing one standard deviation of $CFNAI$. The appendix shows similarly strong dynamics in real GDP growth around low-CP events.

2.3.2 Low-value events

Alternatively, we can isolate periods in which value stocks do particularly poorly. Around such periods, we should find evidence of the poor performance of cash-flows and the macroeconomy. To investigate this possibility, we select quarters in which both the realized log real return on value (the fifth book-to-market portfolio) and the realized log return on value-minus-growth (fifth minus first book-to-market portfolio) are in their respective lowest 30% of observations. These "low-value events" are periods in which value does poorly in absolute terms as well as in relative terms. The double criterion rules out periods in which value returns are average, but value-minus-growth returns are low because growth returns are high. This intersection leads to 41 events out of 254 quarters (or about 16% of the sample). The top left panel of Figure 4 shows the quarterly log returns on value-minus-growth around the event quarter. The value-minus-growth returns are again demeaned over the full sample. By construction, value-minus-growth returns are low in period 0. They are on average around 8% below the quarterly mean.

The top right panel of Figure 4 shows that the level of CP falls in the two quarters leading up to the low value-minus-growth return, bottoms out in the quarter of the value-minus-growth return, and increases in the following two quarters. There is a positive *contemporaneous* relationship between value-minus-growth returns and *changes* in the CP factor. This suggests that innovations in the CP factor capture the risk associated with low value-minus-growth returns.

The bottom left panel shows that dividend growth on value-minus-growth falls considerably in the aftermath of the low-value return event. Between the end of quarters 1 and 10, cumulative dividend growth on value-minus-growth is -52.2%, on average across low-value events. This finding dovetails nicely with the fall in dividends on value-minus-growth over the course of recessions, shown above. Indeed, many of the low-value events occur just prior to the official start of NBER recessions.

We see the same decline in macroeconomic activity following a low-value return event. The bottom right panel of Figure 4 shows the level of $CFNAI$. In the event quarter, the level of economic activity is 0.4 standard deviations below average and it stays below average for the ensuing eight quarters. The change in economic activity from two quarters before to two quarters after the event is 0.5 standard deviations of $CFNAI$. The Appendix shows a similar effect using real GDP growth as a measure

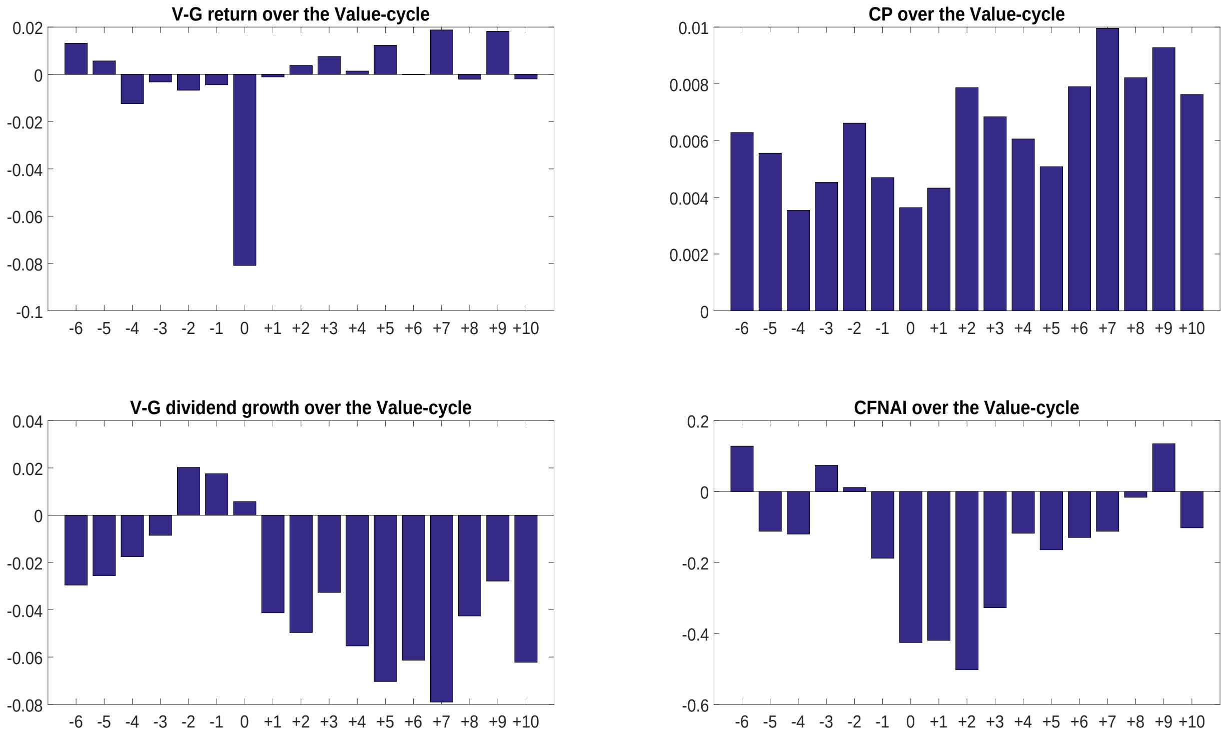

Figure 4: Low-value events

The figure shows four quarterly series in event time. An event is defined as a quarter in which both the realized log real return on the fifth book-to-market portfolio (value) and the realized log return on value-minus-growth (first book-to-market portfolio) are in their respective lowest 30% of observations. The sample runs from 1953.Q3 until 2015.Q4. In each panel, the period labeled '0' is the quarter in which the event takes place. The labels -1, -2, -3, etc refer to one, two, three, etc quarters before the event whereas the labels +1, +2, +3, etc. refer to one, two, three, etc quarters after the event. The top left panel plots the realization of the quarterly log return on value-minus-growth. The bottom left panel reports annual log dividend growth on value-minus-growth. The top right panel plots the CP factor. The bottom right panel plots the $CFNAI$ index of economic activity. The latter is available only from 1967.Q2 onwards. Formally, the graph reports $c_k + \beta_{1k} + \beta_{2k}$ from a regression $X_{t+k} = c_k + \beta_{1k}\mathcal{I}_{excret_V<LB_V} + \beta_{2k}\mathcal{I}_{excret_{V,t}-excret_{G,t}<LB_S} + \epsilon_{t+k}$, for various k, where $\mathcal{I}$ is an indicator variable, LB_V is the 30th percentile of excess returns on the value portfolio, LB_V is the 30th percentile of excess returns on the value-minus-growth portfolio, and X is the dependent variable which differs in each of the four panels. Value-minus growth returns and value-minus-growth dividend growth have been demeaned over the full sample; $CFNAI$ is also mean zero by construction.

of economic activity. The delayed adjustment in dividends vis-a-vis that of macroeconomic activity is consistent with that found in the low-CP event analysis. The evidence in the bottom two panels suggests that firms only cut dividends (and those in the value more than those in the growth portfolio) after a prolonged period of below-average levels of economic activity.

Methodologically, the advantage of the event-time approach is that it focuses on those periods where the investment strategy performs poorly. By looking at windows around these low value return events, the relationships between returns, cash flows, and macroeconomic activity become more transparent and therefore easier to detect. While the low value-minus-growth return events are clearly associated with recessions, the exact timing vis-a-vis the official NBER recession dates varies from recession to recession. This makes it hard to detect clear relationships between value returns and NBER recessions.

3 A Factor Model for Stocks and Bonds

Based on the evidence on the link between the value spread and the CP factor, we provide an asset pricing model for the cross-section of book-to-market equity portfolios, the equity market portfolio, and the cross-section of maturity-sorted bond portfolios. In a second pass, we also include corporate bond portfolios, sorted by credit rating. Our model is parsimonious in that only three pricing factors are needed to capture the bulk of the cross-sectional return differences. As a reduced-form stochastic discount factor model, it imposes little structure beyond the absence of arbitrage opportunities between these equity and bond portfolios. Section 5 presents a structural asset pricing model, which starts from cash flow growth rather than returns, and formalizes the intuition for the empirical connection between dividends and prices of stocks, bond prices, and the business cycle.

3.1 Setup

Let P_t be the price of a risky asset, D_{t+1} its dividend, and R_{t+1} the cum-dividend return. Then the nominal stochastic discount factor (SDF) implies $E_t[M^{\$}_{t+1}R_{t+1}] = 1$. Lowercase letters denote natural logarithms: $m^{\$}_t = \log\left(M^{\$}_t\right)$. We propose a reduced-form SDF, akin to that in the empirical term structure literature (Duffie and Kan, 1996):

$$-m^{\$}_{t+1} \quad = \quad y^{\$}_t + \frac{1}{2}\Lambda'_t\Sigma\Lambda_t + \Lambda'_t\varepsilon_{t+1}, \tag{2}$$

where $y^{\$}_t$ is the nominal short-term interest rate, ε_{t+1} is a $N \times 1$ vector of shocks to the $N \times 1$ vector of demeaned state variables X_t, and where Λ_t is the $N \times 1$ vector of market prices of risk associated with these shocks at time t. The state vector follows a first-order vector-autoregression with companion matrix Γ and conditionally normally, i.i.d. distributed innovations $\varepsilon_t \sim \mathcal{N}(0, \Sigma)$:

$$X_{t+1} = \Gamma X_t + \varepsilon_{t+1}, \tag{3}$$

$$\Lambda_t = \lambda_0 + \lambda_1 CP_t^\star, \tag{4}$$

where $CP_t^\star$ denotes the demeaned CP factor. The market prices of risk are affine in the state vector, where λ_0 is an $N \times 1$ vector of constants and λ_1 is an $N \times 1$ vector that governs the time variation in the prices of risk. In Appendix D, we provide a formal specification analysis to determine the order of the VAR. We also show that our asset pricing results are robust to changing the model to a second- or third-order VAR.

Log returns on an asset j can be stated as the sum of expected and unexpected returns: $r_{t+1}^j = E_t[r_{t+1}^j] + \eta_{t+1}^j$. Unexpected log returns η_{t+1}^j are assumed to be normally distributed and homoscedastic. We denote the covariance matrix between shocks to returns and shocks to the state variables by Σ_{Xj}. We define log excess returns including a Jensen adjustment:

$$rx_{t+1}^j \equiv r_{t+1}^j - y_t^\$(1) + \frac{1}{2} V[\eta_{t+1}^j].$$

The no-arbitrage condition then implies:

$$E_t\left[rx_{t+1}^j\right] = Cov_t\left[rx_{t+1}^j, -m_{t+1}^\$\right] = Cov\left[\eta_{t+1}^j, \varepsilon_{t+1}'\right]\Lambda_t \equiv \Sigma_{Xj}\left(\lambda_0 + \lambda_1 CP_t^\star\right). \tag{5}$$

Unconditional expected excess returns are computed by taking the unconditional expectation of (5) to generate:

$$E\left[rx_{t+1}^j\right] = \Sigma_{Xj}\lambda_0. \tag{6}$$

The main object of interest, λ_0, is estimated below. Equation (6) suggests an interpretation of our model as a simple factor model, where the factor innovations ε are the priced sources of risk. Alternatively, we can rewrite (6) to a beta representation where expected returns are decomposed in betas multiplied by factor risk premia. To focus on the pricing of individual shocks, we prefer to estimate the pricing model in (6). Lastly, we assume in (6) that second moment of returns and pricing innovations are constant. In Appendix E, we show that our results are robust to allowing for time-varying second moments.

3.2 Data and Implementation

In our main asset pricing result, we explain the average excess returns on the five value-weighted quintile portfolios sorted on their book-to-market ratio from Fama and French (1992), the value-weighted stock market return from CRSP (NYSE, AMEX, and NASDAQ), and five zero-coupon nominal government bond portfolios with maturities of 1, 2, 5, 7, and 10 years from CRSP. The return data are monthly from July 1952 until December 2015 (762 observations). In our second exercise, we add corporate bond returns. We use data from Citibank's Yield Book for four investment-grade portfolios: AAA, AA, A, and BBB. Return data for these portfolios are available monthly from January 1980 until December 2015, which restricts our estimation to this sample (432 observations). Section 4 examines other sets of test assets for robustness. We include three asset pricing factors in the state vector X_t. The first factor is the bond factor CP, which forecasts future macro-economic activity as discussed in Section 2. The second asset pricing factor measures the level of the term structure of interest rates, LVL. It is constructed as the first principal component of the one- through five-year Fama-Bliss forward rates following Cochrane and Piazzesi (2008). The third factor, MKT, is the value-weighted stock market return from CRSP.

We construct the unexpected bond returns in η as the residuals from a regression of each bond portfolio's log excess return on the lagged CP factor. Similarly, we assume that stock returns are also predictable by the lagged CP factor, and construct the unexpected stock returns in η as the residual from a regression of each stock portfolio's log excess return on the lagged CP factor.

We impose an autoregressive structure on the state vector X_t. We estimate a monthly VAR(1) with the CP, LVL, and MKT factors. Innovations to the state vector ε follow from equation-by-equation OLS estimation of the VAR model in (3). The innovation correlations between our three factors are close to zero. We find correlations of 0.05 between CP and LVL, 0.04 between CP and MKT, and -0.10 between LVL and MKT.

The first column of Table 1 shows the full sample average excess returns, expressed in percent per year for the 11 test assets. They are the pricing errors resulting from a model where all prices of risk in λ_0 are zero, that is, from a risk-neutral SDF model (*RN SDF*). Average excess returns on bonds are between 0.9% and 2.0% per year and tend to increase in maturity. The aggregate excess stock market return is 6.8%, the excess returns on the book-to-market portfolios range from 6.5% (BM1, growth stocks) to 10.6% (BM5, value stocks), implying a value premium of 4.1% per year.

The first column of Table 2 shows the average excess returns for the shorter 1980-2015 sample. Average excess returns on long-dated government bonds are substantially higher in this sample, for example 3.7% per year for the 10-year bond. The equity risk premium is also slightly higher at 7.3% while the value risk premium is slightly lower at 3.5% per year. The rating-sorted corporate bond portfolios have average excess returns between 3.5% per year for the highest-rated portfolio (AAA) and 4.3% for the lowest-rated portfolio (BBB).

We estimate the three risk price parameters in $\hat{\lambda}_0$ by minimizing the sum of squared pricing errors on the $J = 11$ test assets in Table 1. Formally, we define the GMM moments, conditional on the second moment matrix Σ_{Xj}, as:

$$g_T(\Lambda_0) = E_T\left[rx^j_{t+1}\right] - \Sigma_{Xj}\lambda_0, \tag{7}$$

where $E_T[\cdot]$ denotes the sample average. We estimate λ_0 as:

$$\widehat{\lambda}_0 = argmin_{\lambda_0}\ g_T(\lambda_0)'g_T(\lambda_0), \tag{8}$$

which is equivalent to regressing the $J \times 1$ average excess returns on the $J \times 3$ covariances in Σ_{XJ}. We use the same objective function in all models that we estimate. We estimate the risk prices of all the other models in the same way.[11]

Having estimated the constant market prices of risk, λ_0, we turn to the estimation of the vector λ_1, which governs the time variation in the prices of risk. The vector λ_1 is chosen to exactly match the observed predictability of the stock market and the average bond return by the CP factor as

[11] As is the case for two-pass regressions, the risk price may deviate from the in-sample mean of traded factors, such as the market factor. To impose this additional constraint, one could include the factor as a test asset and use the inverse of the covariance matrix of the pricing errors, instead of the identity matrix as we do, as the weighting matrix in (8). However, as we wish to compare the same cross-section of test assets in all of our tests, which do not include, for instance, the Fama and French factors, we do not impose this constraint in our estimation.

documented by Cochrane and Piazzesi (2005).[12] Specifically, we allow the price of level risk $\lambda_{1(2)}$ and the price of market risk $\lambda_{1(3)}$ to depend on the Z factor, where Z is either CP (benchmark case), the slope of the yield curve or one of the other bond factors that we consider in the appendix. We do not find strong evidence that the value-minus-growth portfolio returns are predicted by the CP factor, and we therefore set $\lambda_{1(1)} = 0$.

We use two predictive regressions to pin down this variation in risk prices. We regress excess returns on a constant and lagged Z:

$$rx^j_{t+1} = a_j + b_j Z_t + \eta^j_{t+1},$$

where we use either excess returns on the stock market portfolio or an equally-weighted portfolio of all bond returns used in estimation. Using equation (5), it then follows:

$$\begin{pmatrix} \lambda_{1(2)} \\ \lambda_{1(3)} \end{pmatrix} = \begin{pmatrix} \Sigma_{X,Market(2:3)} \\ \Sigma_{X,Bonds(2:3)} \end{pmatrix}^{-1} \times \begin{pmatrix} b_{Market} \\ b_{Bonds} \end{pmatrix}.$$

3.3 Estimation Results

The results from our model are in the second column of Table 1 (CP SDF). Panel A shows the pricing errors. Our model succeeds in reducing the mean absolute pricing errors (MAPE) on the 11 stock and bond portfolios to a mere 49 basis points (bps) per year. The model largely eliminates the value spread: The spread between the fifth and the first book-to-market quintile portfolios is 99 bps per year. We also match the market equity risk premium and the average bond risk premium. Pricing errors on the stock and bond portfolios are an order of magnitude lower than in the first column and substantially below those in several benchmark models we discuss below.

Panel B of Table 1 shows the point estimates for λ_0. To obtain risk premia, the risk prices need

[12] Time variation in the market prices of risk drives time variation in expected returns, thereby affecting the unexpected returns η^j_{t+1} and the unconditional asset pricing model in equation (6). Cochrane and Piazzesi (2005) provide evidence of predictability of the aggregate market return by the lagged CP factor. In addition, we could include the aggregate dividend-price ratio as a predictor of the stock market. Given the low R^2 of these predictive regressions, the resulting unexpected returns are similar whether we assume predictability by CP, the dividend-price ratio, both, or no predictability at all. But whatever predictability there is via CP, we match it through the estimate of λ_1.

Table 1: SDF Model for Stocks and Bonds - Pricing Errors

Panel A of this table reports pricing errors on five book-to-market sorted quintile stock portfolios, the value-weighted market portfolio, and five bond portfolios of maturities 1, 2, 5, 7, and 10 years. They are expressed in percent per year (monthly numbers multiplied by 1200). Each column corresponds to a different stochastic discount factor (SDF) model. The first column contains the risk-neutral SDF and therefore reports the average pricing errors to be explained. The second column presents our CP SDF model with three priced risk factors (CP, LVL, MKT). The third column presents the results for a bond pricing model, where only the level factor is priced (LVL). The fourth column (LVL-only bonds) only uses the bond returns as test assets to estimate the same SDF as in the third column. The SDF model of the fifth column has the market return as the only factor (MKT). The sixth column allows for both the prices of LVL and MKT risk to be non-zero. The seventh column refers to a model with the MKT, SMB, and HML factors of Fama and French (1992). In the final column, we use the same SDF as in (2), but we replace the CP innovations with their factor-mimicking portfolio return. To construct the factor-mimicking portfolio, we regress CP innovations on a set of excess returns, R_t^e, $\epsilon_t^{CP} = \nu_0 + \nu_1' R_t^e + u_t$. We use a small value, a small growth, a large value, and a large growth portfolio from the standard 25 size- and book-to-market-sorted portfolios. For instance, to construct the small value portfolio, we take the bottom quintile in terms of size and average the two portfolios with the lowest book-to-market portfolios. We follow the same procedure for the other three portfolios. The last row of Panel A reports the mean absolute pricing error across all 11 test assets (MAPE). Panel B reports the estimates of the prices of risk λ_0. In the seventh column, the pricing factors are the innovations in the excess market return (MKT), in the size factor (SMB), and in the value factor (HML), where innovations are computed as the residuals of a regression of these factors on the lagged dividend-price ratio on the market. Panel C reports asymptotic p-values of chi-squared tests of (i) the null hypothesis that all market prices of risk in λ_0 are jointly zero ($\lambda_0 = 0$), and (ii) of the null hypothesis that all pricing errors are jointly zero (Pr. err. = 0). The data are monthly from June 1952 through December 2015.

	Panel A: Pricing Errors (in % per year)						
	(1) RN SDF	(2) CP SDF	(3) LVL	(4) LVL only bonds	(5) MKT	(6) LVL, MKT	(7) MKT, SMB, HML
10-yr	1.74	0.20	-3.77	-0.41	1.34	-0.40	0.68
7-yr	2.03	0.41	-2.86	0.12	1.74	0.20	1.16
5-yr	1.69	-0.22	-2.44	0.08	1.49	0.18	1.01
2-yr	1.19	-0.73	-0.83	0.40	1.04	0.40	0.82
1-yr	0.94	-0.49	-0.10	0.53	0.83	0.51	0.72
Market	6.83	-0.79	5.66	6.37	-1.33	-1.29	-0.12
BM1	6.49	-0.20	5.27	6.01	-2.02	-1.98	0.43
BM2	7.20	-0.19	5.90	6.69	-0.77	-0.78	-0.44
BM3	8.24	0.90	6.86	7.70	0.96	0.89	-0.34
BM4	8.43	-0.45	7.14	7.92	0.85	0.83	-1.07
BM5	10.64	0.79	9.92	10.36	2.30	2.49	1.25
MAPE	5.04	0.49	4.61	4.24	1.33	0.90	0.73
	Panel B: Prices of Risk Estimates λ_0						
MKT		2.60			3.71	3.52	5.43
LVL/SMB		-18.90	-32.95	-12.88		-10.52	-8.02
CP/HML		85.63					5.97
	Panel C: Test on Risk Prices and Pricing Errors						
$H_0 : \lambda_0 = 0$, p-value (%)		0.11	0.00		0.02	0.01	0.01
H_0 : Pr. error = 0, p-value (%)		0.69	0.00		0.00	0.00	0.01

to be multiplied with the covariance of return and pricing factor innovations, see (6). We estimate a positive price of CP risk, while the price of LVL risk is negative and that of MKT risk is positive. The signs on these risk prices are as expected. As explained in Section 2, the positive price of CP risk arises because positive shocks to CP are good news for future economic activity, which implies a negative innovation to the SDF or equivalently low marginal utility of wealth states for the representative investor. A positive shock to the level factor leads to a drop in bond prices and negative bond returns. A negative shock to bond returns increases the SDF and, hence, carries a negative risk price. A positive shock to the market factor increases stock returns and lowers the SDF, and should carry a positive risk price. We also compute asymptotic standard errors on the λ_0 estimates using GMM with the identity weighting matrix. The standard errors are 32.24 for the CP factor price (point estimate of 85.63), 8.15 for the LVL factor price (-18.90), and 1.17 for the MKT factor price (2.60). Hence, all risk prices are significant at the 5% level.

The first row of Panel C in Table 1 tests the null hypothesis that the market price of risk parameters are jointly zero. This null hypothesis is strongly rejected, with p-value of 0.11%. The last row reports the p-value for the chi-squared test that all pricing errors are jointly zero. All models considered in Table 1 are rejected at the 1% level. However, our three-factor pricing model is able to account for the bulk of the cross-sectional variation in stock and bond returns with a single set of market price of risk estimates.

Following the separate estimation procedure for λ_1 in the full sample, we find $\hat{\lambda}_{1(2)} = -995$ and $\hat{\lambda}_{1(3)} = 42$ when $Z = CP$. This implies that equity and bond risk premia are high when CP is high, consistent with the findings of Cochrane and Piazzesi (2005). We find similar results with $\hat{\lambda}_{1(2)} = -793$ and $\hat{\lambda}_{1(3)} = 120$ when the predictor is the yield spread ($Z = YSP$).

How does our three-factor model manage to price the cross-section of returns on these test assets? Figure 5 decomposes each asset's risk premium into its three components: risk compensation for exposure to the CP factor, the LVL factor, and the MKT factor. The top panel shows risk premia for the five bond portfolios, organized from shortest maturity on the left (1-year) to longest maturity on the right (10-year). The bottom panel shows the decomposition for the book-to-market quintile portfolios, ordered from growth to value from left to right, as well as for the market portfolio (most right bar).

The top panel of Figure 5 shows the risk premium decomposition for the five bond portfolios. Risk premia are positive and increasing in maturity due to their exposure to LVL risk. The exposure to level shocks is negative and the price of level risk is negative, resulting in a positive contribution to the risk premium. This is the standard duration effect. But bonds also have a negative exposure to CP shocks for longer-maturity bonds. CP being a measure of the risk premium in bond markets, positive shocks to CP lower bond prices and realized returns. This effect is larger the longer the maturity of the bond. Given the positive price of CP risk, this exposure translates into an increasingly negative contribution to the risk premium. Because exposure of bond returns to the equity market shocks MKT is positive but near-zero, the sum of the level and CP contributions delivers the observed pattern of bond risk premia that increase in maturity.

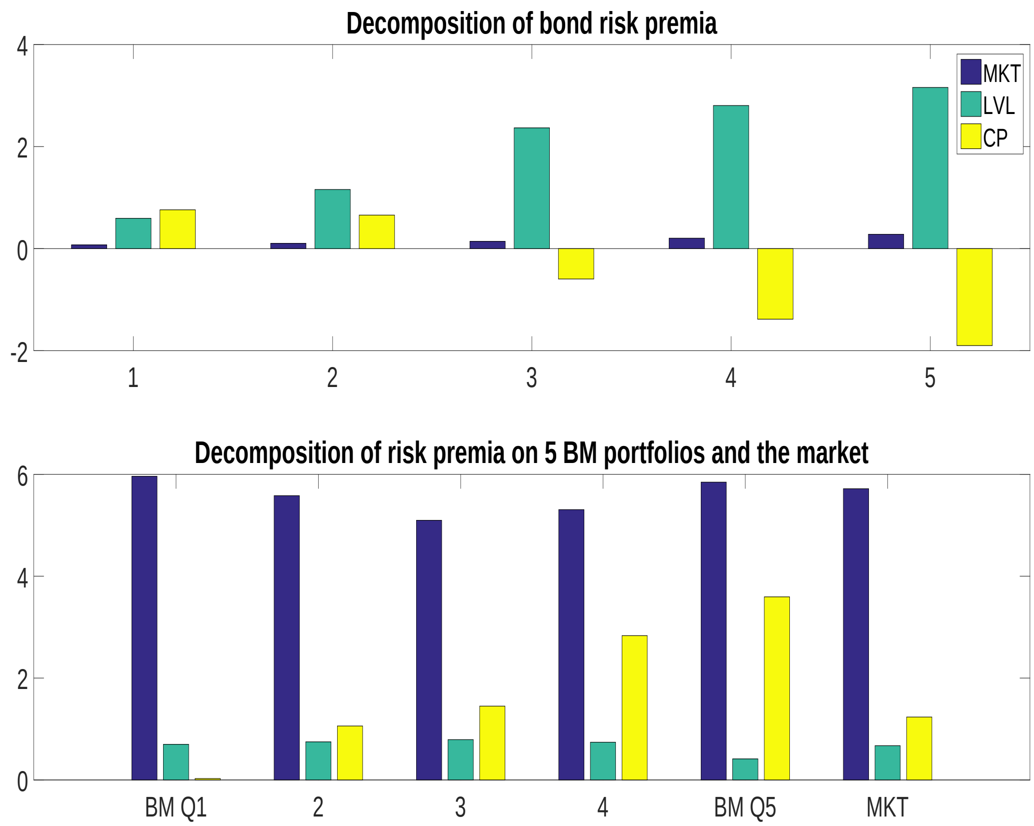

Figure 5: Decomposition of annualized excess returns in data.

The figure plots the risk premium (expected excess return) decomposition into risk compensation for exposure to the MKT, LVL, and CP factors. Risk premia, plotted against the left axis, are expressed in percent per year. The top panel is for the five bond portfolios: one-, two-, five-, seven-, and ten-year maturities from left to right, respectively. The bottom panel is for the book-to-market decile quintile portfolios, from growth (G) to value (V), and for the market portfolio (M). The three bars for each asset are computed as $\Sigma'_{XR}\lambda_0$, the risk exposures times the risk price for each of the three factors. The data are monthly from June 1952 until December 2015.

One might be tempted to conclude that any model with three priced risk factors can always account for the three salient patterns in our test assets. To highlight that such a conjecture is false and to highlight the challenge in jointly pricing stocks and bonds, Appendix C provides a simple example where (1) the CP factor is a perfect univariate pricing factor for the book-to-market portfolios (it absorbs all cross-sectional variation), (2) the LVL factor is a perfect univariate pricing factor for the bond portfolios, and (3) the CP and the LVL factors are uncorrelated. It shows that such a model generally fails to price the stock and bond portfolios jointly. This failure arises because the bond portfolios are exposed to the CP factor, while the stock portfolios are not exposed to the LVL factor. The example in Appendix C underscores the challenges in finding a model with consistent risk prices across stocks and bonds, or put differently, the challenge of going from a single asset class to multiple asset classes. In this setting, consistent risk pricing across stocks and bonds only works if the exposures of maturity-sorted bond portfolios to CP are linear in maturity. This linearity is what allows the model to jointly price stocks and bonds, but it is not a foregone conclusion.

The bottom panel of Figure 5 shows that all book-to-market portfolios have about equal exposure to both MKT and LVL shocks. If anything, growth stocks (G) have slightly higher MKT betas than value stocks (V), but the difference is small. Similarly, there is little differential exposure to LVL shocks across book-to-market portfolios. The spread between value and growth risk premia entirely reflects differential compensation for CP risk. Value stocks have a large and positive exposure to CP shocks while growth stocks have a low exposure. The differential exposure between the fifth and first book-to-market portfolio is statistically different from zero. Multiplying the spread in exposures by the market price of CP risk delivers a value premium of 30 bps per month or 3.6% per year. That is, the CP factor's contribution to the risk premia accounts for most of the 4.1% value premium. Given the monotonically increasing pattern in exposures of the book-to-market portfolios to CP shocks, a positive price of CP risk estimate is what allows the model to match the value premium.

To further quantify the separate roles of each of the three risk factors in accounting for the risk premia on these stock and bond portfolios, we return to columns (3)-(6) of Table 1. Column (3) of Table 1 minimizes the pricing errors on the same 11 test assets but only allows for a non-zero price of level risk. This is the bond pricing model proposed by Cochrane and Piazzesi (2008). They show that the cross-section of average bond returns is well described by differences in exposure to the level factor. Long-horizon bonds have returns that are more sensitive to interest rate shocks than short-horizon bonds; a familiar duration argument. However, this bond SDF is unable to jointly explain the cross-section of stock and bond returns; the MAPE is 4.6%. All pricing errors on the stock portfolios are large and positive, there is a 4.7% value spread, and all pricing errors on the bond portfolios are large and negative. Clearly, exposure to the level factor alone does not account for the high equity risk premium nor the value risk premium. Value and growth stocks have similar exposure to the level factor, that is, a similar "bond duration." The reason that this model does not do a better job at pricing the bond portfolios is that the estimation concentrates its efforts on reducing the pricing errors of stocks, whose excess returns are larger than those of the bonds.

To illustrate that this bond SDF is able to price the cross-section of bonds, we estimate the same model by minimizing only the bond pricing errors (the first five moments in Table 1). Column (4) of Table 1 confirms that the bond pricing errors fall substantially: The mean absolute bond pricing error goes from 200 bps in column (3) to 31 bps in column (4). However, the overall MAPE remains high at 4.24%. The canonical bond pricing model offers one important ingredient for the joint pricing of stocks and bonds, bonds' heterogeneous exposure to the level factor, but this ingredient does not help to account for equity returns.

Another benchmark is the market model where the only non-zero price of risk is the one corresponding to the MKT factor. Column (5) of Table 1 reports the corresponding pricing errors. Not surprisingly, this model is unable to jointly price stock and bond returns. The MAPE is 1.33%. One valuable feature is that the aggregate market portfolio is priced reasonably well and the pricing errors of book-to-market portfolio returns are centered around zero. Our estimation procedure does not impose that the MKT factor is priced exactly, explaining the pricing error on the market portfolio itself of -1.33%. So, while the LVL factor helps to explain the cross-sectional variation in average bond returns and the MKT factor helps to explain the level of equity risk premia, neither factor is able to explain why value stocks have much higher risk premia than growth stocks. Column (6) of Table 1 indeed shows that having both the level and market factor priced does not materially improve the pricing errors and leaves the value premium puzzle in tact. The MAPE is 90 bps, which highlights the need for the CP factor as a third priced factor.

Column (7) in Table 1 reports results for a three-factor model that includes the market, SMB, and HML factors (Fama and French, 1992), which offers a better-performing alternative to a model with the MKT only to price the cross-section of stocks. The model's MAPE is 73 bps. The worse fit than that of the CP SDF model is due to higher pricing errors on the bond portfolios. Tests of the null hypothesis that all pricing errors are jointly zero are rejected at conventional levels.[13]

Table A.IV in Appendix B shows that our model prices the bond portfolios and the market portfolio alongside the size deciles, the earnings-price deciles, and the 5-by-5 size and value portfolios using data from Ken French from June 1952 to December 2015. The results are qualitatively similar.

3.4 Adding Corporate Bond Portfolios

One asset class that deserves particular attention is corporate bonds. Stocks and corporate bonds are both claims on the firm's cash flows albeit with a different priority structure. In standard credit models, corporate bonds are a combination of risk-free bonds and equity. Lower-rated bonds, which are closer to default, have a larger equity component, while highly-rated bonds have a smaller equity component. Based on this logic, we ask whether our SDF model is able to price portfolios of corporate bonds sorted by ratings class. Fama and French (1993) also include a set of corporate bond portfolios in their analysis. They conclude that a separate credit risk factor is needed to price these portfolios. In contrast, we find that the same three factors we used so far also do a good job pricing the cross-section of corporate bond portfolios, while providing an economic interpretation to the pricing factors.

The sample of corporate bond data starts only in 1980; the excess returns to be explained in this sample are listed in the first column of Table 2. We start by re-estimating our main results on this subsample. Column (2) shows that the MAPE on the 11 tests assets we considered in Section 3.3 is 54 bps, nearly identical to the 49 bps in the full sample. In terms of risk prices, we find a similar price of market risk, a more negative price of LVL risk, and a smaller price of CP risk. However, the risk price estimates are not statistically different from their full sample values. The null hypothesis that all risk price estimates are zero is strongly rejected for both models. As before, we reject the null that all pricing errors are jointly zero.

[13] We find that the price of SMB risk is negative, which arises because this risk price cannot be estimated precisely from the five book-to-market portfolios, five bond portfolios, and the equity market portfolio. In Appendix B, we show that once we include portfolios sorted on market capitalization, the price of SMB risk turns positive.

The third column adds the the credit portfolios. We do not re-estimate the market prices of risk, but use those from column (2). The model does a good job pricing the corporate bonds: mean absolute pricing errors on the credit portfolios are 51 bps per year, compared to excess returns of almost 4% per year under risk-neutral pricing. The mean absolute pricing error among all fifteen test assets is 53 bps per year in column (3), which is virtually the same as without the corporate bond portfolios (54 bps per year).

Table 2: SDF Model for Stocks, Treasuries, and Corporate Bonds

The table is similar to Table 1 except that the sample is January 1980 until December 2015. The table adds four corporate bond portfolios sorted by S&P credit rating: AAA (Credit1), AA, A, and BBB (Credit4). Their returns are expressed in percent per year. Column 2 excludes the credit portfolios in the estimation. Column 3 uses the market price of risk estimates of Column 2, and evaluates all pricing errors including those on the corporate bond portfolios. Column 4 includes the credit portfolios when estimating the risk prices.

	(1) RN SDF	(2) *CP* SDF	(3) *CP* SDF not re-estim.	(4) *CP* SDF	(5) *MKT, SMB, HML*
Panel A: Pricing Errors (% per year)					
10-yr	3.69	0.18	0.18	0.44	0.24
7-yr	3.58	0.18	0.18	0.44	0.80
5-yr	2.91	-0.06	-0.06	0.17	0.89
2-yr	1.75	-0.27	-0.27	-0.11	0.85
1-yr	1.18	-0.05	-0.05	0.04	0.77
Market	7.33	-1.06	-1.06	-1.08	0.44
BM1	7.47	-0.17	-0.17	-0.28	0.49
BM2	8.57	0.56	0.56	0.53	-0.59
BM3	8.37	0.66	0.66	0.66	-2.02
BM4	7.84	-1.39	-1.39	-1.26	-2.17
BM5	10.97	1.36	1.36	1.46	2.78
Credit1	3.50		-0.76	-0.46	0.98
Credit2	3.64		-0.63	-0.34	0.94
Credit3	3.92		-0.50	-0.22	1.31
Credit4	4.31		-0.16	0.10	1.99
MAPE	5.27	0.54	0.53	0.51	1.15
Panel B: Prices of Risk Estimates λ_0					
MKT		2.73	2.73	2.80	6.42
LVL/SMB		-22.20	-22.20	-20.35	-19.68
CP/HML		45.64	45.64	41.70	2.88
Panel C: Test on Risk Prices and Pricing Errors					
$H_0 : \lambda_0 = 0$, p-value (%)		0.55		0.68	0.41
H_0 : Pr. error = 0, p-value (%)		2.66		2.26	1.37

Equally interesting is to re-estimate the market price of risk parameters of the SDF model when the corporate bond portfolios are included in the estimation. Column (4) shows that the corporate bond pricing errors now go through zero. For the CP SDF, the MAPE on the credit portfolios is 28 bps per year and the overall MAPE on all 15 assets is 51 bps per year, 3 bps below the MAPE when corporate bonds are not considered, and 2 basis points less than when the corporate bonds were not included in the estimation. Higher exposures to CP innovations and MKT innovations both contribute to higher average returns on the lowest-rated credit portfolio.

The last column of Table 2 reports results for the three-factor model with the MKT, SMB, and HML factors. Its pricing errors are higher than in our three-factor model; the MAPE is 1.15%. Average pricing errors on the corporate bond portfolios are 1.31% per year. The model severely underprices the BBB-rated portfolio (Credit4).

3.5 Other Yield Curve Factors

The CP factor is a specific linear combination of one- through five-year bond yields that predicts economic activity and whose innovations have a monotonic covariance pattern with returns on the book-to-market portfolios. Other linear combinations of these yields may better predict economic activity. Similarly, other linear combinations of yields may do a better job pricing the cross-section of stock and bond returns. We consider three natural alternatives to CP. The first one is the slope of the yield curve, YSP, measured as the difference between the 5-year and the 1-year bond yields. The second one, YGR, is the linear combination of bond yields that best forecasts economic activity levels 12 months ahead. The third one, YAP, is the linear combination of bond yields that best prices the 11 test assets over the full sample. The CP factor has a correlation of 77% with YSP, 56% with YGR, and 76% with YAP. For ease of comparison, we rescale these three factors so they have the same standard deviation as CP. The predictability of CP for future economic activity, discussed in Section 2, extends to YSP and YGR as detailed in Appendix A, in particular Tables A.II and A.III.

Next, we revisit the main asset pricing exercise with three alternative bond yield factors in lieu of the CP factor. Detailed results are in Table 3. The model with the yield spread factor produces results broadly consistent with those for CP. It leads to a larger MAPE of 62 bps per year in the full sample, and leaves more of the value risk premium and the difference between long- and short-term bonds unexplained than the model with CP as a factor. The signs and approximate magnitude of the market prices of risk of YSP and CP are the same. Over the 1980-2015 sample, the MAPE is 68 bps (column 2). When we add the corporate bond portfolios, the MAPE falls further to 58 bps and we fail to reject the null hypothesis that all pricing errors are jointly zero at the 5% level.

Table 3: Alternative Yield Curve Factors

This table reports pricing errors on five book-to-market sorted quintile stock portfolios, the value-weighted market portfolio, five bond portfolios of maturities 1, 2, 5, 7, and 10 years, and four credit-sorted portfolios. They are expressed in percent per year (monthly numbers multiplied by 1200). We also report the mean absolute pricing error across all securities (MAPE) and the estimates of the prices of risk. The first three columns correspond to the YSP SDF, the middle three columns to the YGR SDF, while the last three columns refer to the YAP SDF model. YSP is the slope of the yield curve, measured as the difference between the 5-year bond yield and the one-year bond yield. YGR is the fitted value of a regression of macro-economic activity $CFNAI_{t+12}$ on the one- through five-year yields at time t. YAP is the linear combination of one- through five-year yields which best prices the 11 test assets. The first, fourth, and seventh columns are for the full 1952-2015 sample, while the other columns are for the 1980-2015 sub-sample in which we observe corporate bond returns.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
		YSP SDF			YGR SDF			YAP SDF	
Panel A: Pricing Errors (% per year)									
10-yr	0.50	0.69	0.91	0.26	0.52	0.84	-0.17	0.32	0.76
7-yr	0.29	0.18	0.47	0.28	0.05	0.63	0.32	0.11	0.63
5-yr	-0.36	-0.42	-0.11	-0.13	-0.09	0.44	-0.05	0.00	0.45
2-yr	-0.96	-0.80	-0.54	-0.53	-0.75	-0.09	-0.03	-0.81	-0.30
1-yr	-0.68	-0.45	-0.28	-0.73	-0.55	-0.04	-0.04	-0.23	0.05
Market	-0.86	-1.14	-1.16	-0.81	-1.14	-1.21	-0.73	-1.13	-1.17
BM1	-0.60	-0.72	-0.81	-0.03	0.15	-0.50	-0.07	0.31	-0.15
BM2	-0.49	0.37	0.34	-1.29	0.00	0.05	0.00	0.05	0.07
BM3	0.99	0.78	0.76	0.63	0.21	0.34	0.93	-0.19	0.01
BM4	0.27	-0.59	-0.51	0.40	-0.43	-0.26	-1.02	0.00	0.01
BM5	0.82	1.36	1.51	1.14	1.18	1.77	0.90	0.95	1.39
Credit1			-0.42			-0.66			-0.59
Credit2			-0.44			-0.53			-0.57
Credit3			-0.34			-0.63			-0.48
Credit4			0.18			-0.08			-0.16
MAPE	0.62	0.68	0.58	0.57	0.46	0.54	0.39	0.37	0.45
Panel B: Prices of Risk Estimates λ_0									
MKT	2.42	3.16	3.21	1.48	2.68	2.99	2.93	2.84	3.00
LVL	-6.61	-16.52	-15.03	-3.32	-11.89	-11.07	-14.62	-30.26	-24.17
YSP/YGR/YAP	93.20	62.78	55.21	115.95	83.50	50.45	60.83	129.58	96.12
Panel C: Test on Risk Prices and Pricing Errors									
$H_0 : \lambda_0 = 0$, p-value (%)	0.02	0.23	0.27	1.51	2.97	1.27	0.44	9.66	5.65
H_0 : Pr. error = 0, p-value (%)	0.72	4.35	5.74	18.67	34.05	20.24	4.16	23.18	7.36

In the next three columns of Table 3 we use YGR alongside the MKT and LVL factors in our asset pricing exercise. This model generates a MAPE of 57 bps for the full sample, 46 bps for the post-1980 sample, and 54 bps when we include the credit portfolios. In all three exercises, we cannot reject the null hypothesis that all pricing errors are zero. The price of risk estimate for YGR in the main exercise is similar in magnitude and not statistically different from that of CP. These pricing results indicate that there is a lot of information about future economic growth in the term structure that is useful for pricing stocks and bonds. They also confirm that there is nothing special about CP

for asset pricing beyond its ability to forecast economic growth.

The last three columns show that we can lower MAPE to a mere 39 bps per year in the full sample by finding the best-pricing linear combination of 1- through 5-year bond yields. Using that same linear combination YAP, pricing errors are 37 bps for the post-1980 sample, and 45 bps for the same sample with credit portfolios. The 76% correlation of CP with YAP helps to explain why our main pricing results are strong.

All these results are consistent with the view that there is an component in expected economic growth, as measured from the term structure of interest rates, that prices the joint cross-section of stock and bond returns.

4 Robustness

In this section, we confirm the robustness of our asset pricing results. First, we consider other cross-sections of stock returns to be priced alongside bonds. We also verify robustness to using the Fama-MacBeth methodology.

4.1 Other Test Assets

In addition to the credit portfolios discussed above, we consider several other cross-sections of stock returns. We focus this exercise on the sample January 1967 until December 2015, which is the longest sample for which all factor returns are available. The columns labeled RN in Table 4 report the average returns that need to be explained in this sample. In columns (1)-(5), we consider the 5 bond portfolios, the aggregate stock market, and 10 book-to-market portfolios. In columns (6)-(10), we add 10 portfolios sorted on market capitalization (size). In column (11)-(15), we also add 10 momentum portfolios constructed by Hou, Xue, and Zhang (2015).

It is well known that the 3-factor Fama and French model cannot explain momentum portfolios, and the same is true for our model. We therefore augment both models with a momentum factor (UMD). We refer to the extended 3-factor Fama and French model as the Carhart model (Carhart, 1997). In our comparison of factor models, we also consider the recent 4-factor model of Hou, Xue, and Zhang (2015) (HXZ4) and the 5-factor model of Fama and French (2015) (FF5).

The second column of Table 4 shows that our model successfully reduces virtually all of the value spread and explains the cross-section of bond returns. The MAPE for this sample is 48 bps. This shows that our main results hold in a more recent sample and for both book-to-market deciles and quintiles. The Carhart and HXZ4 models also eliminate most of the value spread, while the FF5 model leaves 1.2% unexplained. All three models, however, struggle to simultaneously price the cross-sections of stocks and bonds. The MAPE is 93 bps for the HXZ4, 74 bps for the Carhart, and 86 bps for the FF5 model.

If we add size portfolios in columns (6)-(10), then our model reduces the size spreads to 1.6% without sacrificing much of the explanatory ability for the other 16 test assets. The MAPE equals 51 bps. All other models are able to explain the cross-section of size and book-to-market portfolios. Their MAPE are 66 bps for the HXZ4, 61 bps for the Carhart, and 66 bps for the FF5 models. While the MAPE fall for each of these models, the pricing errors of the bond portfolios remain equally large, or increase, for all three alternative models.

In columns (11)-(15) of Table 4, we add the momentum decile portfolios. The momentum spread is large and equal to 10.4% per year. Our model leaves only 2.6% of the momentum premium unexplained, which is the same as for the Carhart model. The HXZ4 model (4.1%) and the FF5 model (9.1%) result in larger pricing errors for the momentum portfolios. The MAPE is lowest for the Carhart model (74 bps), followed by our model (76 bps), the HXZ4 (83 bps), and the FF5 (1.1%) models. The differences for the first three models are small, but, as before, the other models result in larger pricing errors for the bond portfolios than our model.

Table A.V in Appendix B repeats the analysis using investment, size, and return-to-equity deciles instead of the B/M, size, and momentum portfolios. The investment and return-to-equity deciles are from Hou, Xue, and Zhang (2015) from January 1967 to December 2015. All models are able to explain the spread in size-sorted portfolios. While our model is able to price the bond portfolios better for any cross-section of test assets that we consider, our model results in larger pricing errors for investment and return-to-equity sorted portfolios compared to the other pricing models. The HXZ4 and FF5 models perform substantially better for the largest cross-section that we consider in Appendix B.

Table 4: Pricing Errors and MPR - Book-to-Market, Size, and Momentum Deciles

This table reports pricing errors on 10 book-to-market sorted stock portfolios, the value-weighted market portfolio, and five bond portfolios of maturities 1, 2, 5, 7, and 10 years in columns (1)-(5). In columns (6)-(10), we add 10 size-sorted portfolios. In columns (11)-(15), we add 10 momentum portfolios. Pricing errors are expressed in percent per year (monthly numbers multiplied by 1200). We also report the mean absolute pricing error across all securities (MAPE) and estimates of the prices of risk. We compare our model extended with the UMD factor, the Carhart model, the model of Hou, Xue, and Zhang (2015) (HXZ4), and the 5-factor model of Fama and French (2015) (FF5). Panel C reports asset pricing tests that either all risk prices are zero (top row) or that all pricing errors are zero (bottom row). The sample is from January 1967 to December 2015.

	(1) RN	(2) CP SDF +UMD	(3) HXZ4	(4) Carhart	(5) FF5	(6) RN	(7) CP SDF +UMD	(8) HXZ4	(9) Carhart	(10) FF5	(11) RN	(12) CP SDF +UMD	(13) HXZ4	(14) Carhart	(15) FF5
Panel A: Pricing Errors (% per year)															
10-yr	2.35	-0.14	1.89	0.68	1.01	2.35	0.20	1.99	0.78	1.84	2.35	0.08	1.62	1.46	2.20
7-yr	2.55	0.32	2.23	1.26	1.54	2.55	0.59	2.33	1.33	2.25	2.55	0.42	2.03	1.90	2.44
5-yr	2.08	-0.04	1.78	1.09	1.38	2.08	0.07	1.86	1.14	1.86	2.08	-0.22	1.65	1.55	1.87
2-yr	1.32	-0.29	1.08	0.98	0.96	1.32	-0.56	1.12	0.97	1.16	1.32	-1.18	1.06	1.05	1.11
1-yr	1.03	0.01	0.87	0.87	0.85	1.03	-0.26	0.90	0.86	0.93	1.03	-0.76	0.90	0.88	0.87
Market	5.94	-0.71	-0.56	-0.78	-0.04	5.94	-0.93	-0.51	-0.73	-0.44	5.94	-0.95	-0.17	-0.78	-0.35
BM1	5.10	-0.30	0.33	0.56	-0.60	5.10	-0.32	0.28	0.51	-0.37	5.10	-0.28	0.26	0.14	1.86
BM2	6.51	0.26	0.13	0.40	1.19	6.51	0.22	-0.05	0.43	0.39	6.51	0.36	-0.63	0.34	-0.92
BM3	7.20	0.39	0.08	0.11	0.74	7.20	0.05	-0.10	0.15	0.22	7.20	-0.18	-0.81	0.22	-0.75
BM4	6.88	0.15	-0.65	-0.49	-0.45	6.88	0.49	-0.76	-0.51	-0.61	6.88	1.50	-1.31	-0.53	-1.05
BM5	6.39	-0.92	-1.23	-0.32	-1.32	6.39	-0.75	-1.10	-0.51	-0.91	6.39	0.10	-1.32	-0.79	-1.10
BM6	7.63	1.18	-0.13	-1.08	0.11	7.63	1.02	-0.21	-0.94	0.17	7.63	1.11	-0.96	-0.28	-0.69
BM7	7.23	-1.36	-0.99	-0.08	-1.07	7.23	-1.53	-0.67	-0.30	-0.63	7.23	-1.06	-0.29	-0.66	-1.07
BM8	8.07	-0.13	-0.19	-1.01	-0.75	8.07	-0.26	-0.33	-0.93	-0.55	8.07	0.24	-0.40	-0.74	0.74
BM9	10.34	1.36	2.04	1.60	1.07	10.34	1.28	2.00	1.58	1.27	10.34	2.02	2.53	1.36	3.17
BM10	10.59	0.08	0.62	0.61	0.62	10.59	0.31	0.47	0.63	-0.20	10.59	1.99	1.58	-0.01	1.62
S1						8.83	0.90	0.14	0.08	0.53	8.83	0.22	0.47	0.37	-0.71
S2						8.52	-0.02	-0.64	-0.98	-0.53	8.52	-0.82	-0.69	-0.58	-0.80
S3						9.35	0.79	0.31	0.55	0.25	9.35	0.54	-0.03	0.60	0.49
S4						8.52	0.31	-0.20	-0.34	-0.40	8.52	-0.27	-0.89	-0.05	0.48
S5						8.92	0.18	0.41	0.54	0.35	8.92	-0.46	0.12	0.57	0.86
S6						8.19	-0.17	0.20	0.28	0.02	8.19	-0.64	-0.18	0.29	0.94
S7						8.32	-0.22	0.26	0.19	0.53	8.32	-1.07	-0.02	0.35	-0.06
S8						7.70	-0.66	0.11	0.24	0.43	7.70	-0.83	0.42	0.07	-0.44
S9						6.95	-0.37	-0.08	-0.11	0.12	6.95	-0.17	0.10	-0.18	-0.07
S10						5.47	-0.67	-0.12	-0.35	-0.10	5.47	-0.47	0.32	-0.49	0.09
MOM1											0.82	-1.72	-1.49	-2.08	-4.84
MOM2											3.88	-0.54	0.14	0.61	-1.73
MOM3											6.03	1.63	1.27	2.04	0.12
MOM4											6.62	1.54	1.27	1.88	0.52
MOM5											5.92	0.50	-0.06	0.43	-0.06
MOM6											5.78	-0.46	-0.74	-0.45	-0.74
MOM7											6.55	-0.98	-0.48	-0.37	-0.39
MOM8											6.82	-0.69	-0.97	-1.11	-0.49
MOM9											8.25	-0.30	-0.22	-0.75	1.01
MOM10											11.23	0.91	2.59	0.51	4.30
MAPE	5.70	0.48	0.93	0.74	0.86	6.62	0.51	0.66	0.61	0.66	6.50	0.76	0.83	0.74	1.14
Panel B: Prices of Risk Estimates λ_0															
MKT		1.24	4.37	5.69	2.73		1.49	4.16	5.48	2.92		1.98	4.95	4.23	7.01
LVL/SMB/ME		-17.17	-0.12	1.87	-1.58		-16.98	2.43	0.94	2.80		-19.09	5.78	0.66	0.57
CP/HML/IA		23.82	10.50	9.54	15.23		48.70	10.39	8.85	6.78		88.65	13.72	6.13	-14.82
UMD/ROE/RMW		-6.16	0.53	13.21	6.73		-2.43	2.94	11.19	5.16		5.60	11.50	4.55	1.68
CMA					-19.72					-5.51					39.35
Panel C: Test on Risk Prices and Pricing Errors (p-values in %)															
$H_0: \lambda_0 = 0$		0.24	1.50	3.79	2.16	0.00	1.19	1.74	3.18	2.96	0.00	1.63	0.21	0.07	1.08
H_0 : Pr. error $= 0$		2.42	0.03	0.09	0.33	0.00	6.81	0.05	0.21	0.06	0.00	0.09	0.00	0.00	0.00

4.2 GMM versus Fama-MacBeth

Much of the cross-sectional asset pricing literature uses the method of Fama and MacBeth (1973). Appendix B.2 finds that are results are nearly unchanged under the Fama-MacBeth methodology. However, it is important to estimate the CP-betas with enough data. Using 60-month rolling-windows leads to imprecise estimates and results in a deterioration of the pricing performance of our model.

5 Model with Business Cycle Risk

In the last part of the paper, we propose a simple asset pricing model that connects our empirical findings. It formalizes the relationships between the returns on value and growth stocks, the bond risk premium (CP), and the state of the macroeconomy. It does so in a pricing framework that quantitatively accounts for the observed risk premia on stock and bond portfolios, the dynamics of dividend growth rates, inflation, and basic properties of the term structure of interest rates. Its role is to clarify the minimal structure necessary to account for the observed moments. In the interest of space, we relegate a full model description, solution, and calibration to Online Appendix F.

The model has one key state variable, s, which is as a leading business cycle indicator. It follows an autoregressive process and its innovations ε^s_{t+1} are the first priced source of risk. Real dividend growth for value (V), growth (G), and market (M) equity portfolios are given by:

$$\Delta d^i_{t+1} = \gamma_{0i} + \gamma_{1i} s_t + \sigma_{di}\varepsilon^d_{t+1} + \sigma_i \varepsilon^i_{t+1},\ \forall i = \{V, G, M\}.$$

The shock ε^d_{t+1} is an aggregate dividend shock, the second priced source of risk, while ε^i_{t+1} is a non-priced idiosyncratic shock. The market portfolio has no idiosyncratic risk; $\sigma_M = 0$. The key parameter configuration is $\gamma_{1V} > \gamma_{1G}$ so that value stocks are more exposed to shocks in macroeconomic activity than growth stocks. As in the data (Section 2.2.1), a low value for s is associated with lower future dividend growth on V minus G. Our calibration matches the frequency and duration of recessions and chooses γ_{1V} and γ_{1G} to match the decline in dividend growth value minus growth over the course of recessions. Inflation is the sum of a autoregressive process which captures expected inflation and an unexpected inflation shock. Expected inflation, x_t, is the second state variable in the model; its innovation ε^x_{t+1} the third and last priced source of risk. Inflation and dividend growth parameters are chosen to match the unconditional mean and volatility of dividend growth and inflation, as well as the volatility and persistence of one- through five-year nominal bond yields.

Investor preferences are summarized by a real stochastic discount factor (SDF), whose log m evolves according to the process: $-m_{t+1} = y + \frac{1}{2}\Lambda_t'\Lambda_t + \Lambda_t'\varepsilon_{t+1}$, where the vector $\varepsilon_{t+1} \equiv \left(\varepsilon^d_{t+1}, \varepsilon^x_{t+1}, \varepsilon^s_{t+1}\right)'$ and y is the constant real rate of interest. The market prices of risk are chosen to match the equity risk premium (the one associated with ε^d), slope of the yield curve (ε^x), and value risk premium (ε^s).

The model generates an affine nominal term structure of log nominal interest rates: $y^\$_t(n) = -\frac{A^\$_n}{n} - \frac{B^\$_n}{n}s_t - \frac{C^\$_n}{n}x_t$. It also generates a one-factor model for the nominal bond risk premium: All variation in bond risk premia comes from cyclical variation in the economy, s_t.

$$E_t\left[rx^\$_{t+1}(n)\right] = \underbrace{\lambda_0(2)C^\$_{n-1}\sigma_x + \lambda_0(3)B^\$_{n-1}\sigma_s}_{\text{Constant component bond risk premium}} + \underbrace{\lambda_1(2)C^\$_{n-1}\sigma_x s_t}_{\text{Time-varying component bond risk premium}}.$$

Thus, the CP factor which measures the bond risk premium is perfectly positively correlated with s_t, the leading indicator of macroeconomic activity. Innovations to the CP factor are innovations to s (ε^s), lending a structural interpretation to CP shocks which is consistent with our empirical evidence. The constant component of the bond risk premium reflects compensation for exposure to expected inflation risk (first term) and cyclical risk (second term). Exposure to the cyclical shock contributes negatively to excess bond returns: A positive ε^s shock lowers bond prices and returns, and more so for long than for short bonds. Exposure to expected inflation shocks contributes positively to excess bond returns: A positive ε^x shock lowers bond prices and returns but the price of expected inflation risk is negative. Since common variation in bond yields is predominantly driven by the inflation shock in the model, the latter acts like (and provides a structural interpretation for) a shock to the level of the term structure (LVL). Long bonds are more sensitive to level shocks, the traditional duration effect.

Turning to stock pricing, the log price-dividend ratio on stocks is affine in the state s_t: $pd^i_t = A_i + B_i s_t$. The equity risk premium provides compensation for exposure to aggregate dividend growth risk (ε^d) and for cyclical risk (ε^s):

$$E_t\left[rx_{t+1}^i\right] = \underbrace{\lambda_0(1)\sigma_{di} + \lambda_0(3)\kappa_{1i}B_i\sigma_s}_{\text{Constant component equity risk premium}} + \underbrace{\lambda_1(1)\sigma_{di}s_t}_{\text{Time-varying component equity risk premium}}.$$

Shocks to the market return (MKT) are a linear combination of ε^d and ε^s shocks. Like bond risk premia, equity risk premia vary over time with the state of the economy s_t. The model generates both an equity risk premium and a value premium. The reason for the value premium can be traced back to the fact that value stocks' dividends are more sensitive to cyclical shocks than those of growth stocks: B_i increases in γ_{1i}. Because the price of cyclical risk is naturally positive, the second term delivers the value premium. Put differently, in the model, as in the data, returns on value stocks are more exposed to bond risk premium shocks than returns on growth stocks.

For each asset, we can then compute covariances of unexpected returns with the MKT, LVL, and CP shocks, as defined inside the model. Interestingly, we are able to replicate the three-factor risk premium decomposition we uncovered in Section 3. Figure 6 is the model's counterpart to Figure 5 in the data. It shows a good quantitative match for the relative contribution of each of the three sources of risk to the risk premia for growth, value, and market equity portfolios, as well as for maturity-sorted government bond portfolios. This fit is not a forgone conclusion, but results from the richness of the model and the choice of parameters.[14] The model also generates interesting asset pricing dynamics over the business cycle as detailed in the Online Appendix.

The model delivers a structural interpretation for the MKT, LVL, and CP shocks. CP shocks reflect (transitory) cyclical shocks to the real economy, which naturally carry a positive price of risk. LVL shocks capture changes in expected inflation. MKT shocks mostly capture (permanent) changes in dividend growth. The model quantitatively replicates the unconditional risk premium on growth, value, and market equity portfolios, and bond portfolios of various maturities, as well as the decomposition of these risk premia in terms of their MKT, LVL, and CP shock exposures. Furthermore, it matches some simple features of nominal term structure of interest rates and bond risk premia. It does so for plausibly calibrated dividend growth and inflation processes.

[14]For example, differential exposure to the market factor could have well been the source of the value risk premium in the model given that the market shocks are linear combinations of permanent dividend growth and transitory cyclical shocks. Or, bonds of different maturity could have differential exposure to the market factor shocks. The data show no heterogeneity in both types of exposures. The model has just enough richness to replicate these patterns.

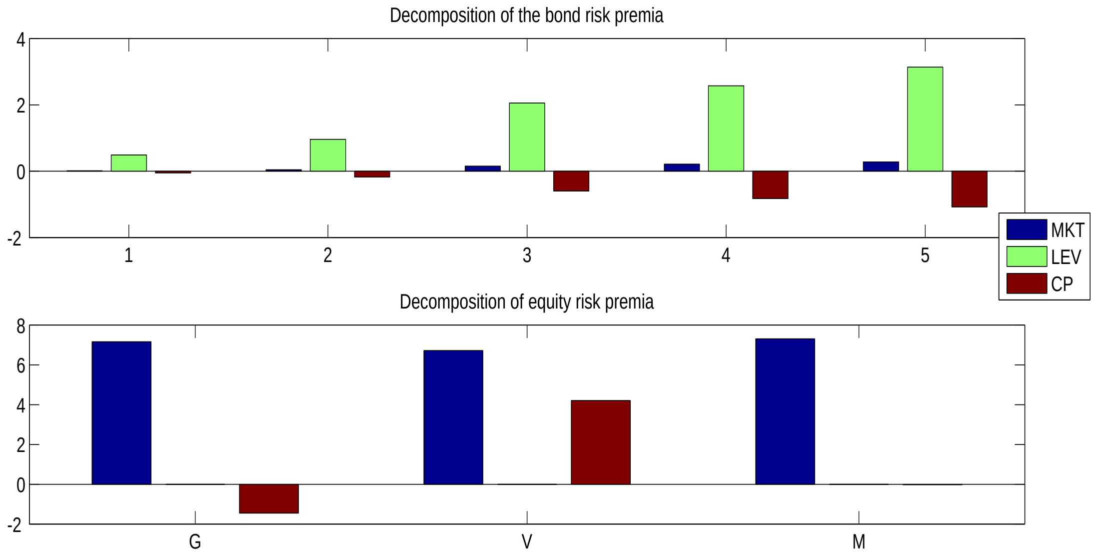

Figure 6: Decomposition of annualized excess returns in model.

The figure plots the risk premium (expected excess return) decomposition into risk compensation for exposure to the CP factor, the LVL factor, and the MKT factor. Risk premia, plotted against the left axis, are expressed in percent per year. The top panel is for the five bond portfolios (1-yr, 2-yr, 5-yr, 7-yr, and 10-yr) whereas the bottom panel is for growth (G), value (V), and market (M) stock portfolios. The results are computed from a 10,000 month model simulation under the calibration described in detail in Online Appendix F.3.

6 Conclusion

We provide new evidence that the value premium reflects compensation for macroeconomic risk. Periods of low returns on value stocks versus growth stocks are times when future economic activity is low and future cash-flows on value stocks are low relative to those on growth stocks. We find that several bond market variables such as the Cochrane-Piazzesi (CP) factor and the slope of the yield curve are leading indicators of these business cycle turning points. Innovations to these factors are contemporaneously highly positively correlated with returns on value stocks, but uncorrelated with returns on growth stocks.

Based on this connection, we estimate a parsimonious three-factor pricing model that can be used to explain return differences between average excess returns on book-to-market sorted stock portfolios, the aggregate stock market portfolio, government bond portfolios sorted by maturity, and corporate bond portfolios. The first factor in our three-factor model is the traditional market return factor, the second one is the level of the term structure, and the third factor is the CP factor or the yield spread. We estimate a positive market price of risk for the latter risk factor, consistent with the notion that positive innovations represent good news about future economic activity.

Our results suggest that transitory shocks to the real economy operating at business cycle frequencies play a key role in accounting for the cross-section of stock returns. Future work on structural Dynamic Asset Pricing Models should bring the business cycle explicitly inside the model as a key state variable. The model we solved in Section 5 is a starting point in this research agenda. More work is needed to help us fully understand why the market compensates exposure to innovations in this state variable so generously.

References

Baker, Malcolm, and Jeffrey Wurgler, 2012, Comovement and predictability relationships between bonds and the cross-section of stocks, *Review of Asset Pricing Studies* June. 5

Bakshi, Gurdip, and Zhiwu Chen, 2005, Stock valuation in dynamic economies, *Journal of Financial Markets* 8, 111–151. 3

Bansal, Ravi, and Bruce N. Lehman, 1997, Growth optimal portfolio restrictions on asset pricing models, *Macroeconomic Dynamics* 1, 333–354. 4

Bansal, Ravi, and Ivan Shaliastovich, 2013, A long-run risks explanation of predictability puzzles in bond and currency markets, *Review of Financial Studies* 26, 1–33. 4, 18

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset prizing puzzles, *The Journal of Finance* 59, 1481–1509. 4

Basu, Susanto, 1977, The investment performance of common stocks in relation to their price to earnings ratios: A test of the efficient market hypothesis, *Journal of Finance* 32, 663–82. 1

Bekaert, Geert, Eric Engstrom, and Steven Grenadier, 2010, Stock and bond returns with moody investors, *Journal of Empirical Finance* 17 (5), 867–894. 18

Bekaert, Geert, Eric Engstrom, and Yuhang Xing, 2009, Risk, uncertainty and asset prices, *Journal of Financial Economics* 91, 59–82. 3, 18

Berk, Jonathan B., Richard C. Green, and Vasant Naik, 1999, Optimal investment, growth options and security returns, *Journal of Finance* 54, 1153–1607. 4

Binsbergen, Jules van, and Ralph Koijen, 2010, Predictive regressions: A present-value approach, *Journal of Finance* 65 (4). 6

Brennan, Michael J., Ashley W. Wang, and Yihong Xia, 2004, Estimation and test of a simple model of intertemporal capital asset pricing, *The Journal of Finance* 59 (4), 1743–1775. 5

Brooks, Jordan, 2011, Unspanned risk premia in the term structure of interest rates, Ph.D. thesis New York University. 8

Campbell, John, Adi Sunderam, and Luis Viceira, 2013, Inflation bets or deflation hedges? the changing risk of nominal bonds, Working Paper Harvard University. 3, 18, 25

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251. 4

Campbell, John Y., Christopher Polk, and Tuomo Vuolteenaho, 2010, Growth or glamour? fundamentals and systematic risk in stock returns., *The Review of Financial Studies* 23(1), 305–344. 5

Campbell, John Y., and Tuomo Vuolteenaho, 2004, Good beta, bad beta, *American Economic Review* 94(5), 1249–1275. 5

Carhart, Mark M., 1997, On the persistence of mutual fund performance, *Journal of Finance* 52, 57–82. 29

Chen, Long, 2009, On the reversal of return and dividend predictability: A tale of two periods, *Journal of Financial Economics* 92, 128–151. 7

Chen, Nai-Fu, Richard Roll, and Stephen A. Ross, 1986, Economic forces and the stock market, *Journal of Business* 59, 383–403. 4

Cochrane, John H., and Lars Peter Hansen, 1992, Asset pricing explorations for macroeconomics, in Olivier J. Blanchard, and Stanley Fischer, ed.: *NBER Macroeconomics Annual: 1992* . pp. 115–165 (MIT Press: Cambridge, MA). 4

Cochrane, John H., and Monika Piazzesi, 2005, Bond risk premia, *American Economic Review* 95, 138–160. 1, 8, 19, 21, 2

———, 2008, Decomposing the yield curve, Working Paper, University of Chicago. 2, 17, 23

Cooper, Ilan, and Richard Priestley, 2008, Time-varying risk premiums and the output gap, *Review of Financial Studies* 22, 2801–2833. 8

David, Alexander, and Pietro Veronesi, 2013, What ties return volatilities to price valuations and fundamentals?, *Journal of Policitical Economy* 121, 682–746. 18

De Bondt, Werner F.-M., and Richard Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 40, 793–805. 1

Duffie, Darrell, and Raymond Kan, 1996, A yield factor model of interest rates, *Mathematical Finance* 6, 379–406. 15

Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465. 1, 17, 25

———, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56. 25

———, 2015, A five-factor asset pricing model, *Journal of financial economics* 116, 1–22. 30, 31, 10

Fama, Eugene F., and James MacBeth, 1973, Risk, return and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636. 32

Ferson, Wayne E., and Campbell R. Harvey, 1991, The variation of economic risk premiums, *Journal of Political Economy* 99, 385–415. 4

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro finance, *Quarterly Journal of Economics* 127(2), 645–700. 3

Gilchrist, Simon, and Egon Zakrajsek, 2012, Credit spreads and business cycle fluctuations, *American Economic Review* Forthcoming. 8

Gomes, Joao, Leonid Kogan, and Lu Zhang, 2003, Equilibrium cross section of returns, *Journal of Policitical Economy* 111, 693–732. 4

Hansen, Lars Peter, John C. Heaton, and Nan Li, 2008, Consumption strikes back? measuring long-run risk, *Journal of Political Economy* 116 (2), 260–302. 4, 6

Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting anomalies: An investment approach, *Review of Financial Studies* 28, 650–705. 4, 29, 30, 31, 10

Joslin, Scott, Michael Priebsch, and Kenneth Singleton, 2014, Risk premiums in dynamic term structure models with unspanned macro risks, *Journal of Finance* 69, 1197–1233. 8

Kazemi, Hossein, 1992, An intertemporal model of asset prices in a markov economy with a limiting stationary distribution, *Review of Financial Studies* 5, 85–104. 4

Kogan, Leonid, and Dimitris Papanikolaou, 2014, Growth opportunities, technology shocks and asset prices, *Journal of Finance* 69, 675–718. 4

Kung, Howard, 2015, Macroeconomic linkages between monetary policy and the term structure of interest rates, *Journal of Financial Economcis* 115, 42–57. 4

Lakonishok, Josef, Andrei Schleifer, and Robert W. Vishny, 1994, Contrarian investment, extrapolation, and risk, *The Journal of Finance* 49 (5), 1541–1578. 1

Larrain, Borja, and Motohiro Yogo, 2007, Does firm value move too much to be justified by subsequent changes in cash flow?, *Journal of Financial Economics* 87(1), 200–226. 6

Lettau, Martin, and Jessica Wachter, 2007, Why is long-horizon equity less risky? a duration-based explanationof the value premium, *Journal of Finance* 65, 55–92. 1, 3, 18, 25

———, 2011, The term structures of equity and interest rates, *Journal of Financial Economics* 101, 90–113. 3, 18

Liu, Laura, Toni Whited, and Lu Zhang, 2009, Investment-based expected stock returns, *Journal of Policitical Economy* 117, 1105–1139. 4

Ludvigson, Sydney, and Serena Ng, 2009, Macro factors in bond risk premia, *Forthcoming Review of Financial Studies*. 8

Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan, 2013, The wealth-consumption ratio, *Review of Asset Pricing Studies* 3(1), 38–94. Review of Asset Pricing Studies. 3

Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding predictability, *Journal of Political Economy* 112, 1–47. 4

Papanikolaou, Dimitris, 2007, Investment-specific technology shocks and asset prices, Working Paper Kellogg. 4

Petkova, Ralitsa, 2006, Do the fama-french factors proxy for innovations in predictive variables?, *Journal of Finance* 61(2), 581–612. 5

Piazzesi, Monika, and Martin Schneider, 2006, Equilibrium yield curves, *National Bureau of Economic Analysis Macroeconomics Annual*. 4

Santos, Jesus, and Pietro Veronesi, 2010, Habit formation, the cross-section of stock returns and the cash flow risk puzzle, *Journal of Financial Economics* 98, 385–413. 4

Wachter, Jessica, 2006, A consumption-based model of the term structure of interest rates, *Journal of Financial Economics* 79, 365–399. 4, 18

Yoon, Pyung Sig, and Laura T. Starks, 1995, Signaling, investment opportunities, and dividend annoucements, *Review of Financial Studies* 8 (4), 995–1018. 7

Zhang, Lu, 2005, The value premium, *Journal of Finance* 60 (1), 67–103. 1, 4

Online Appendix

A. Additional Results for Section 2

A.1. Summary Statistics

	CP	YSP	Log real dividend growth	Log dividend growth V-G	CFNAI
Mean	0.9%	0.7%	2.3%	4.6%	0.00
St.dev.	1.5%	0.8%	7.0%	26.1%	1.01
Skewness	0.38	-0.20	0.60	0.38	-1.14
Kurtosis	3.50	3.19	4.82	8.65	6.55
AC (12m)	0.49	0.51	0.32	-0.14	0.07
AC (24m)	0.18	0.10	0.07	-0.06	-0.18

Table A.I: Summary statistics.

The table reports the mean, standard deviation, skewness, kurtosis, and both the 12-month and 24-month autocorrelation for CP, the yield spread YSP, log real dividend growth on the aggregate stock market, the difference in log dividend growth between value and growth firms, and CFNAI. The sample is from June 1952 until December 2015. The data for CFNAI start in March 1967.

A.2. Dividends Around NBER Recessions pre-1952

The main text shows the behavior of log annual real dividends on value (fifth book-to-market), growth (first book-to-market), and market portfolios for the sample 1952-2015 in the left panel of Figure 1 as well as the difference in dividend growth between value and growth portfolios in the right panel of Figure 1. Figure A.1 shows the corresponding evidence for the period 1926 until 1952. The message of these figures is very much consistent with the discussion in the main text. The massive decline in dividends of value stocks relative to that of dividends of growth stocks in the Great depression is noteworthy.

A.3. Predicting economic activity and dividend growth

Table A.II reports the slopes of time-series predictability regressions of the form

$$y_{t+k} = c_k + \beta_k Z_t + \varepsilon_{t+k}, \tag{A.1}$$

where Z_t is one of three different bond factors and y is one of three measures of economic growth. The left three columns (Panel A) use the economic activity measure $CFNAI$ as the outcome variable. The bond factor is the CP factor in the first column, the slope of the yield curve YSP in the second column, and the linear combination of bond yields that best forecasts CFNAI 12 months ahead, YGR, in the third columns. For ease of comparability, YSP and YGR have been rescaled to have the same standard deviation as CP. The next three columns (Panel B) use dividend growth on the market portfolio as outcome variable y. The last three columns (Panel C) use the difference between dividend growth on the value portfolio (fifth book-to-market quintile portfolio) and the growth portfolio (first book-to-market quintile) as the y variable. Figure 2 in the main text contains a visual representation of the results with CP as predictor.

YGR, in the third column, has the highest possible slope coefficient (28.1), t-statistic (4.9), and R^2 at the 12-month forecast horizon (17.9%) by construction. The predictive ability of YGR deteriorates with the horizon. At 24 months, the slope is 10.4 and the point estimate is no longer significantly different from zero. The yield spread, in the second column, is a slightly stronger predictor of economic activity 12 months out (slope of 19.8) than CP (slope of 18.0), but the R^2 values are about half of those for the best linear combination of yields (8.9% and 7.4%, respectively). The predictability of YSP peaks at 18 months, with a slope of 21.0, a t-stat of 3.6, and an R^2 of 9.9%. The predictability

Table A.II: Predicting economic activity and dividend growth

This table reports slope coefficients from predictive regressions. The predictors Z are listed in the first row. They are the CP factor, the yield spread YSP, and the best linear forecaster of $CFNAI_{t+12}$, YGR. The forecast horizon is listed in the first column. All predictors have the same standard deviation over the sample so that the slope coefficients within each panel are directly comparable. In Panel A, the bond market variables forecast $CFNAI$. In Panel B, they forecast real dividend growth on the market portfolio. In Panel C, they forecast dividend growth on the value minus the growth portfolio. The data are monthly from March 1967 through December 2015. Numbers in bold have Newey-West t-statistics (absolute values) in excess of 1.96.

	CP	YSP	YGR	CP	YSP	YGR	CP	YSP	YGR
k (months)	Panel A: CFNAI			Panel B: Div. Growth M			Panel C: Div. Growth $V-G$		
12	**17.98**	**19.84**	**28.06**	0.31	**1.13**	0.31	**-3.60**	-3.34	**-3.28**
15	**21.67**	**20.15**	**26.05**	0.64	**1.49**	0.49	-1.85	-0.96	-1.51
18	**23.19**	**20.96**	**24.39**	0.81	**1.78**	0.67	-0.21	1.21	0.43
21	**24.03**	**19.02**	**16.94**	**0.99**	**2.02**	**0.88**	1.08	3.05	2.02
24	**20.78**	**14.48**	10.40	**1.08**	**2.18**	**1.03**	1.92	4.27	3.44
27	**17.49**	12.02	7.73	**1.10**	**2.22**	**1.06**	2.38	4.77	**4.56**
30	**14.13**	8.91	4.62	**1.26**	**2.32**	**1.17**	2.23	4.74	4.18
33	11.88	5.93	1.45	**1.37**	**2.33**	**1.14**	2.53	4.49	3.03
36	7.40	2.35	-1.89	**1.41**	**2.23**	1.01	3.01	4.01	1.72

is statistically significant for horizons from 2 months to 25 months. After that same 24-month horizon, YSP also loses its predictive ability. The CP factor, in contrast, is a much stronger predictor than YSP or YGR 24 months out. In fact, CP is close to the best linear predictor at that 24-month horizon.

Panel B shows that both CP and YSP predict future dividend growth on the market significantly at (nearly) all horizons. YGR predicts future dividend growth at horizons beyond 18 months. In terms of size of the coefficient and R-squared, YSP has the strongest predictive ability and YGR the weakest, with CP in between. Panel C shows that there is some predictability of value-minus growth dividends, but it is weaker than aggregate dividend growth.

A.4. Predicting GDP growth with CP

In the main text we show that the bond factors Z forecast future economic activity, as measured by the $CFNAI$ index. As an alternative to $CFNAI$, we consider real gross domestic product (GDP) growth (seasonally adjusted annual rates) from the National Income and Product Accounts. The GDP data are available only at quarterly frequency, but go back to 1952 when the CP series starts. This gives us a longer sample than for CFNAI, which starts in 1967. When $Z = CP$, our results update a regression that appears in the working paper version of Cochrane and Piazzesi (2005). The yield factor Z in a given quarter is set equal to the value in the last month of the quarter. We estimate

$$\Delta GDP_{t+k} = c_k + \beta_k Z_t + \varepsilon_{t+k}, \tag{A.2}$$

where k is the forecast horizon expressed in quarters. Table A.III shows the coefficient estimates β_k in Panel B. For comparison, Panel A predicts $CFNAI$ with the same variables using the same quarterly frequency. $CFNAI$ then refers to the last month of the quarter. The predictors have been scaled to have the same standard deviation within each sample so that the point estimates are directly comparable for different predictors Z. In addition to CP and YSP, we also consider the best linear forecaster of GDP growth 5 quarters out, the horizon over which we get the highest overall predictability. We call this yield curve predictor $YGDP$.

We find that all three predictors strongly forecast annual GDP growth 4 to 8 quarters ahead. That is, they predict GDP growth over the following year and over the year thereafter. The yield spread YSP is again a stronger predictor at shorter horizons while CP is a stronger predictor at longer horizons. All variables lose statistical significance for

Table A.III: Predicting quarterly $CFNAI$ and GDP Growth

This table reports slope coefficients from predictive regressions. The predictors Z are listed in the first row. They are the CP factor, the yield spread YSP, and the best linear forecaster of real GDP growth 5 quarters ahead, $YGDP$. The forecast horizon is listed in the first column. All predictors have the same standard deviation over the sample so that the slope coefficients within each panel are directly comparable. In Panel A, the bond market variables forecast $CFNAI$ (last month of the quarter). In Panel B, they forecast real four-quarter GDP growth, measured quarterly. The data in Panel A are quarterly for 1967.I through 2015.IV while the data in Panel B are quarterly for 1952.III until 2015.IV.

	CP	YSP	$YGDP$	CP	YSP	$YGDP$
k (quarters)	Panel A: $CFNAI$			Panel B: GDP Growth		
4	**15.80**	**19.21**	**30.80**	**0.35**	**0.47**	**0.72**
5	**19.33**	**20.23**	**32.57**	**0.40**	**0.50**	**0.77**
6	**25.03**	**20.70**	**27.20**	**0.39**	**0.44**	**0.70**
7	**28.22**	**20.44**	**20.84**	**0.42**	**0.38**	**0.59**
8	**24.78**	**16.46**	11.45	**0.43**	**0.32**	**0.49**
9	**21.92**	13.77	8.67	**0.38**	0.21	**0.27**
10	**19.56**	11.84	5.51	**0.35**	0.15	0.16
11	**17.68**	9.12	1.46	0.30	0.07	0.05
12	11.11	4.44	-3.39	0.22	0.01	0.02

horizons beyond 10 quarters. The results in Panel A confirm what we learned in the main text: CP predicts economic activity strongly, and more strongly so at longer horizons. The YSP predicts $CFNAI$ about as well as $YGDP$ at intermediate horizons.

A.5. CP, YSP, and NBER Recessions

Figure A.2 plots the CP and YSP factors over time (right axis) while drawing in NBER recessions (shaded areas). Consistent with the economic forecasting regressions, the CP and YSP factors are low before the start of most recessions in the post-1952 sample. They subsequently increases over the course of a recession, especially towards the end of the recession when better times are around the corner. In nearly every recession, the CP and YSP factors are substantially higher at the end than at the beginning of the recession. In the three deepest post-war recessions, the 1974, 1982, and 2008 recessions, CP dips in the middle of the recession -suggesting that bond markets fear a deterioration of future economic prospects- before recovering.

A.6. Real GDP in CP-event Time

We also study the behavior of real annual GDP growth in CP-event time. GDP growth rates are available over the entire post-war sample, whereas $CFNAI$ only starts in 1967. Figure A.3 is the same as Figure 3 in the main text, except that real GDP growth is plotted in the bottom right-hand side panel instead of $CFNAI$. Like $CFNAI$, GDP growth also shows a clean cycle around low-CP events. GDP grows at a rate that is 1.2% point above average two quarters before the event, the growth rate slows down to the average in the event quarter, and growth further falls to a rate of 1.4% points below average five quarters after the event. The amplitude of this cycle (2.6% points) is economically large, representing 1.1 standard deviations of GDP growth.

A.7. One-factor Model

One may wonder whether the facts our paper documents are consistent with a one-factor model that differentially affects cash flow growth rates and therefore returns on value versus growth stocks. The data suggest that they are not.

An adequate description of dividend dynamics contains at least two shocks: one shock that equally affects dividend growth rates on all portfolios and a second shock (to the Z factor) that affects value dividends relative to growth dividends.

To see this, we orthogonalize value-minus-growth dividend growth to the dividend growth rate on the market portfolio. Figure A.4 compares the dynamics of dividend growth on value minus growth around low-CP events (left panel, which repeats the bottom left panel of Figure 3) to those of dividend growth on the market portfolio (middle panel), and of the orthogonal component of value-minus-growth dividend growth (right panel). All three dividend growth series are demeaned over the full sample. The figure shows that the dividend growth on the market portfolio falls in the aftermath of a low-CP event, consistent with the facts on aggregate economic activity or GDP growth. The dividend growth rate on the market portfolio falls by much less than the value-minus-growth dividend growth. Furthermore, the part of value-minus-growth dividend growth that is orthogonal to the market dividend growth, in the right panel, has qualitatively and quantitatively similar dynamics around CP events as the raw value-minus-growth dividend growth in the left panel. Hence, a lot of the dynamics in dividend growth on value-minus-growth is unaccounted for by dividend growth on the market portfolio.

The theoretical model of Appendix 5 articulates this two-shock structure of cash flow growth. It features a common and permanent cash-flow shock that affects all portfolios alike, and a business-cycle frequency shock that differentially affects dividend growth rates of value and growth stocks.

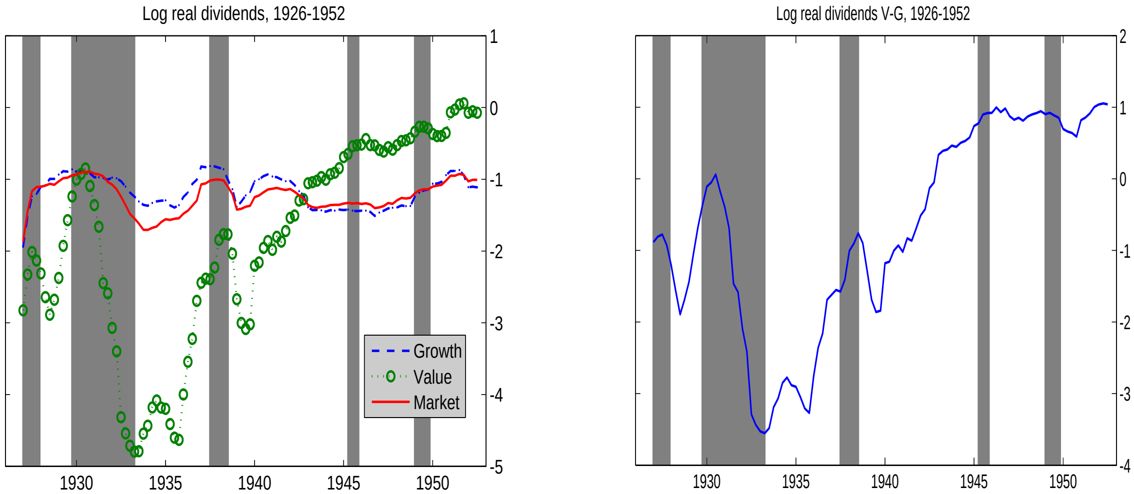

Figure A.1: Dividends on value, growth, and market portfolios pre-1952.

The left panel plots the log real dividend on book-to-market quintile portfolios 1 (growth, dashed line with squares) and 5 (value, dotted line with circles) and on the CRSP value-weighted market portfolio. The right panel plots the log real dividend on book-to-market quintile portfolios 5 (value) minus the log real dividend on the boot-to-market portfolio 1 (growth). Dividends are constructed form cum- and ex-dividend returns on these portfolios. Monthly dividends are annualized by summing dividends received during the year. The data are monthly from December 1926 until June 1952 and are sampled every three months in the figure.

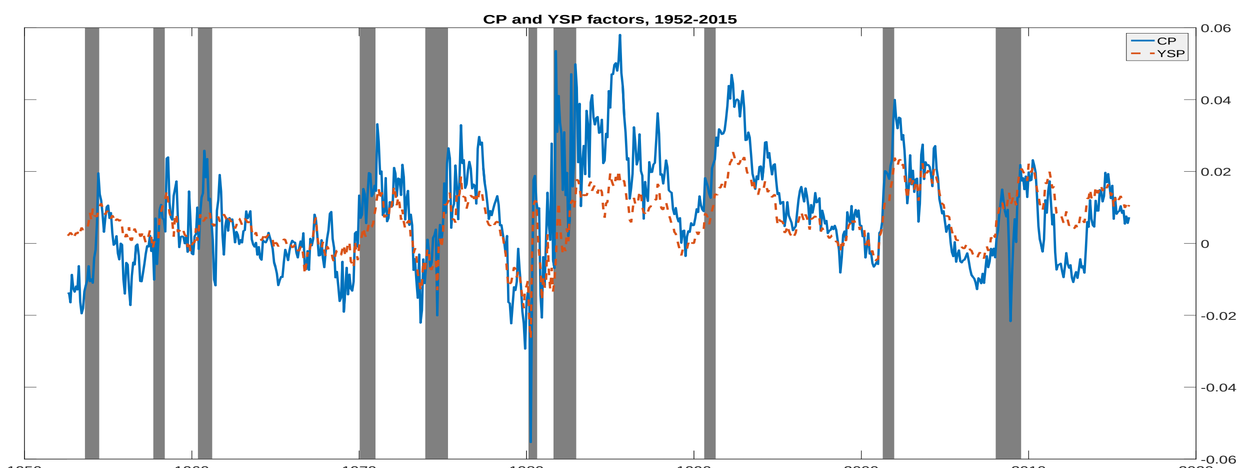

Figure A.2: CP factor and NBER recessions.

The figure plots the CP factor (solid line, against the right axis) and the NBER recessions (shaded areas). The sample is July 1952 until December 2015.

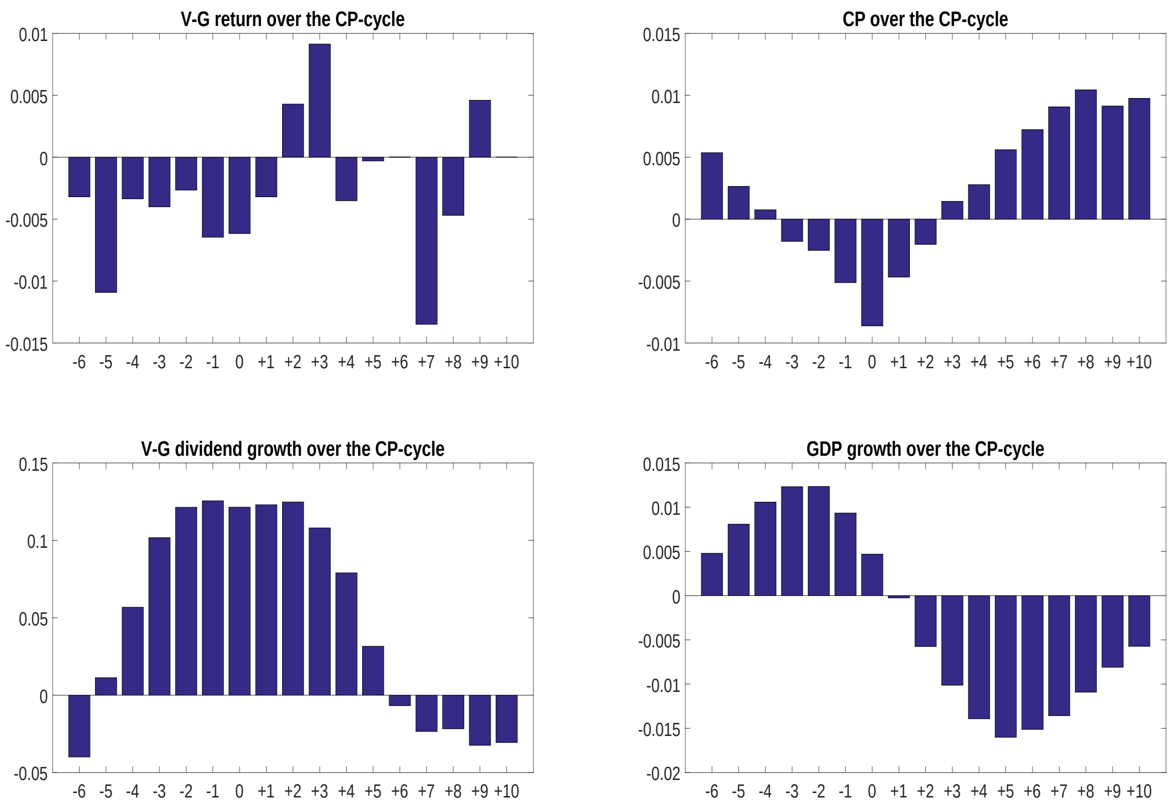

Figure A.3: Low-CP events with GDP growth

The figure plots four quarterly series in event time. The event is defined as a quarter in which the quarterly CP factor in its respective lowest 25% of observations. This selection leads to 63 events out of 254 quarters. The sample runs from 1953.III until 2015.IV. In each panel, the period labeled '0' is the quarter in which the event takes place. The labels -1, -2, -3, etc refer to one, two, three, etc quarters before the event whereas the labels +1, +2, +3, etc. refer to one, two, three, etc quarters after the event. The top left panel plots the realization of the quarterly log return on value minus growth. The bottom left panel reports annual log dividend growth on value minus growth. The top right panel plots the CP factor. The bottom right panel plots real GDP growth. Real GDP growth is demeaned over the full sample.

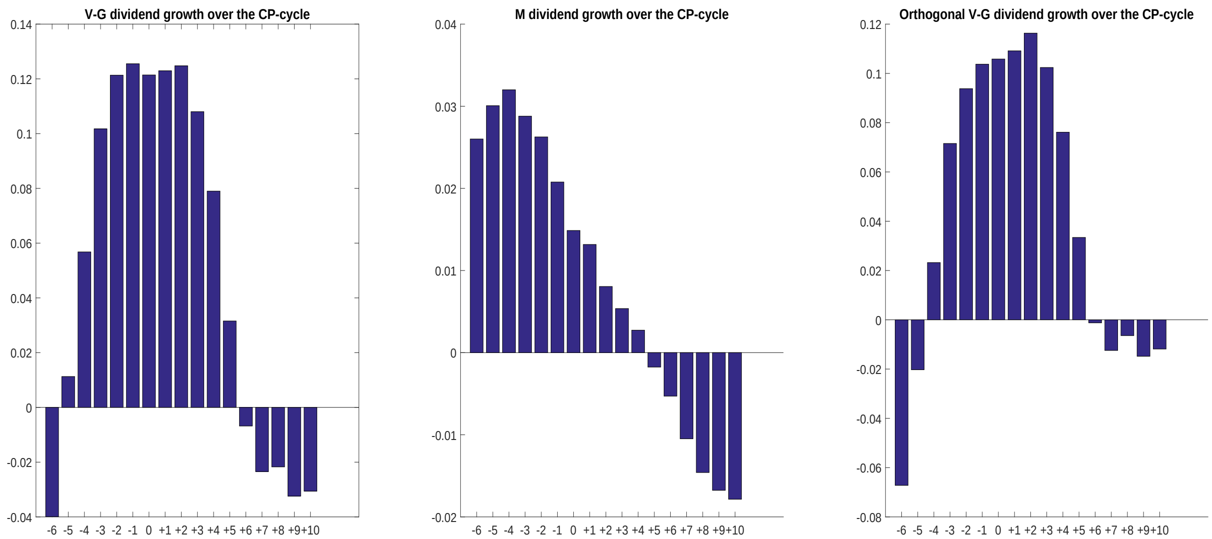

Figure A.4: Dividend growth around Low-CP events

The figure plots three quarterly series in event time. The event is defined as a quarter in which the quarterly CP factor in its respective lowest 25% of observations. This selection leads to 63 events out of 254 quarters. The sample runs from 1953.Q3 until 2015.Q4. In each panel, the period labeled '0' is the quarter in which the event takes place. The labels -1, -2, -3, etc refer to one, two, three, etc quarters before the event whereas the labels +1, +2, +3, etc. refer to one, two, three, etc quarters after the event. The left panel plots annual log dividend growth on value minus growth, the middle panel plots annual log dividend growth on the market portfolio, and the right panel plots annual log dividend growth on value minus growth, orthogonalized to annual log dividend growth on the market portfolio. All three series have mean-zero over the full sample.

B. Additional Results for Factor Models

This section considers several exercises investigating the robustness of our empirical results in Section 3 and 4. First, we study additional sets of test assets. Second, we re-estimate our pricing model using Fama-MacBeth as opposed to GMM.

B.1. Other Test Assets

Table A.IV shows three exercises where we replace the book-to-market sorted equity portfolios by other equity portfolios. In the first three columns we use ten market capitalization-sorted portfolios alongside the bond portfolios and the market. The first column shows the risk premia to be explained (RN). Small firms (S1) have 3.2% higher risk premia than large stocks (S10). Our CP SDF model in the second column manages to bring the overall mean absolute pricing error down from 6.41% per year to 30 bps per year. The market prices of risk are similar when we use size- or book-to-market sorted portfolios. The MAPEs are lower than the 57 bps in the Fama-French model in the third column. The Fama-French model does better eliminating the spread between small and large stocks, whereas our model does better pricing the bond portfolios.

The next three columns use earnings-price-sorted decile stock portfolios. The highest earnings-price portfolio has an average risk premium that is 5.7% higher per year than the lowest earnings-price portfolio. Our CP SDF model reduces this spread in risk premia to 1.7% per year, while continuing to price the bonds reasonably well. The MAPE is 1.0% per year compared to 76 bps in the Fama-French model.

The last three columns use the five-by-five market capitalization and book-to-market double sorted portfolios. Our CP SDF model manages to bring the overall mean absolute pricing error down from 7.9% per year to 1.3% per year. This is again comparable to the three-factor Fama-French model's MAPE of 1.2%.

The market price of risk estimates λ_0 in Panel B of Table A.IV are comparable to those we found for the book-to-market portfolios in Table 1. Panel C shows that we reject the null hypothesis that all market prices of risk are zero for all three sets of test assets. We fail to reject the null hypothesis that all pricing errors are zero on the earnings-to-price portfolios. We conclude that these results are in line with our benchmark results and that they further strengthen the usefulness of our empirical three-factor model.

In Table A.V, we replace the 10 book-to-market deciles by 10 investment portfolios in column (1)-(5). We add 10 size-sorted portfolios in column (6)-(10), and 10 return-to-equity portfolios in column (11)-(15). We consider the same factor models as in Table 4. The decile spread in average returns for investment portfolios is 5%. Our model struggles to price the extreme portfolios. All other models reduce the spread by 2-3%, but leave larger pricing errors for the bond portfolios. The MAPE of all models is similar and ranges from 84 bps (FF5) to 88 bps (Carhart). All models are able to explain a large part of the spread in size-sorted portfolios and the MAPE are again very similar. For ROE-sorted portfolios, our model (augmented with UMD), can only explain 3.5% of the 8% spread. The other models are more successful on these portfolios. The MAPE of our model is 99 bps, compared to 78 bps for the HXZ4 model and 77 bps for both the Carhart and FF5 models. However, as before, the pricing errors of the bond portfolios are substantially smaller in our model compared to the other models.

B.2. GMM vs Fama-MacBeth

This appendix provides a comparison of the GMM and Fama-MacBeth (FMB) estimation methodologies using book-to-market and size decile portfolios over the 1967.1-2015.12 sample. We find that the two estimation procedures give similar results. That conclusion carries over to adding momentum portfolios as well as to studying size, investment, and return-to-equity portfolios instead (unreported).

We study three versions of the FMB approach that differ in the fist-stage estimation of the factor exposures. In the first one, factor betas are estimated on the full time series. This is the closest comparison to our benchmark GMM results, where we only change the estimation methodology. In the second one, factor betas are estimated using expanding windows. In the third one, factor betas are estimate using 60-month rolling-windows. For each of these specifications,

Table A.IV: Other Stock Portfolios - Pricing Errors

This table reports robustness with respect to different stock market portfolios, listed in the first row. Panel A of this table reports pricing errors (in % per year) on various stock portfolios, the value-weighted market portfolio, and five bond portfolios of maturities 1, 2, 5, 7, and 10 years. Each column corresponds to a different stochastic discount factor (SDF) model. The first column contains the risk-neutral SDF. The second column presents our CP SDF model with three priced factors. The third column refers to the three-factor model of Fama and French (FF). The last row of Panel A reports the mean absolute pricing error across all securities (MAPE). Panel B reports the estimates of the prices of risk. The data are monthly from June 1952 through December 2015.

Panel A: Pricing Errors (in % per year)											
	Size Portfolios				EP Portfolios				Size and Value Portfolios		
	RN	CP SDF	FF3		RN	CP SDF	FF3		RN	CP SDF	FF3
10-yr	1.74	0.15	1.38	10-yr	1.74	0.70	1.22	10-yr	1.74	0.24	1.56
7-yr	2.03	0.43	2.01	7-yr	2.03	0.56	1.71	7-yr	2.03	0.18	2.03
5-yr	1.69	-0.07	1.46	5-yr	1.69	-0.62	1.38	5-yr	1.69	-0.88	1.61
2-yr	1.19	-0.44	0.84	2-yr	1.19	-1.79	0.93	2-yr	1.19	-1.80	1.03
1-yr	0.94	-0.25	0.66	1-yr	0.94	-1.42	0.75	1-yr	0.94	-1.37	0.81
Market	6.83	-0.61	-0.34	Market	6.83	-1.33	-0.57	Market	6.83	0.22	0.08
ME1	9.58	0.89	-0.08	EP1	6.43	-0.07	1.28	S1B1	3.88	-4.16	-5.17
ME2	9.33	-0.18	-0.18	EP2	5.67	-1.52	-0.88	S1B2	9.67	-0.67	0.29
ME3	10.02	0.70	0.34	EP3	6.84	-0.71	-0.12	S1B3	9.36	-1.31	-0.10
ME4	9.20	0.09	0.20	EP4	6.93	-1.01	-0.66	S1B4	12.28	3.16	2.48
ME5	9.47	-0.01	0.28	EP5	7.38	-1.02	-0.87	S1B5	13.28	2.99	1.90
ME6	9.00	-0.01	0.22	EP6	8.80	0.97	-0.21	S2B1	5.99	-1.99	-1.79
ME7	8.86	-0.25	0.23	EP7	9.53	0.23	0.47	S2B2	9.43	-0.19	0.51
ME8	8.46	-0.49	-0.27	EP8	10.05	2.02	0.11	S2B3	10.81	0.92	1.44
ME9	7.83	-0.03	-0.61	EP9	11.12	0.95	0.28	S2B4	11.28	1.50	1.16
ME10	6.36	-0.22	-0.07	EP10	12.08	1.64	0.70	S2B5	12.39	2.48	0.05
								S3B1	7.18	-1.70	0.37
								S3B2	9.90	0.77	1.37
								S3B3	9.42	-0.51	0.22
								S3B4	10.91	0.46	0.95
								S3B5	12.01	1.37	0.29
								S4B1	7.69	0.02	1.54
								S4B2	7.90	-0.67	-0.30
								S4B3	9.34	0.23	0.16
								S4B4	10.39	-0.74	0.76
								S4B5	9.88	-2.05	-1.92
								S5B1	6.57	2.49	2.05
								S5B2	6.85	1.44	0.19
								S5B3	7.77	1.87	0.49
								S5B4	6.63	-2.08	-3.00
								S5B5	8.70	-0.84	-2.36
MAPE	6.41	0.30	0.57		6.20	1.03	0.76		7.87	1.33	1.23
Panel B: Prices of Risk Estimates λ_0											
MKT		2.68	6.14			2.20	4.89			1.53	3.81
LVL/SMB		-16.54	0.03			-26.58	-1.93			-28.44	1.48
CP/HML		65.87	20.08			163.85	8.38			154.25	7.18
Panel C: Test on Risk Prices and Pricing Errors (p-values in %)											
$H_0 : \lambda_0 = 0$		0.41	0.55			0.55	0.00			0.27	0.00
H_0 : Pr. error = 0		0.75	1.75			7.91	0.01			0.05	0.00

Table A.V: Pricing Errors and MPR - Investment, Size, and Return-to-Equity Deciles

This table reports pricing errors on 10 investment-sorted stock portfolios, the value-weighted market portfolio, and five bond portfolios of maturities 1, 2, 5, 7, and 10 years in columns (1)-(5). In columns (6)-(10), we add 10 size-sorted portfolios. In columns (11)-(15), we add 10 return-to-equity portfolios. Pricing errors are expressed in percent per year (monthly numbers multiplied by 1200). We also report the mean absolute pricing error across all securities (MAPE) and estimates of the prices of risk. We compare our model extended with the UMD factor, the Carhart model, the model of Hou, Xue, and Zhang (2015) (HXZ4), and the 5-factor model of Fama and French (2015) (FF5). Panel C reports asset pricing tests that either all risk prices are zero (top row) or that all pricing errors are zero (bottom row). The sample is from January 1967 to December 2015.

	(1) RN	(2) CP SDF +UMD	(3) HXZ4	(4) Car-hart	(5) FF5	(6) RN	(7) CP SDF +UMD	(8) HXZ4	(9) Car-hart	(10) FF5	(11) RN	(12) CP SDF +UMD	(13) HXZ4	(14) Car-hart	(15) FF5
Panel A: Pricing Errors (% per year)															
10-yr	2.35	0.04	1.19	0.83	1.16	2.35	0.15	1.58	0.97	1.69	2.35	-0.21	1.56	0.58	1.51
7-yr	2.55	0.40	1.57	1.36	1.49	2.55	0.50	1.99	1.50	2.14	2.55	0.16	1.97	1.18	1.99
5-yr	2.08	-0.18	1.31	1.13	1.23	2.08	-0.11	1.63	1.23	1.69	2.08	-0.42	1.62	0.99	1.74
2-yr	1.32	-0.99	0.90	0.89	0.89	1.32	-0.98	1.07	0.95	1.01	1.32	-1.29	1.06	0.89	1.14
1-yr	1.03	-0.59	0.80	0.80	0.80	1.03	-0.60	0.90	0.83	0.82	1.03	-0.84	0.89	0.81	0.94
Market	5.94	-0.75	-0.93	-0.92	-0.83	5.94	-0.72	-0.57	-0.83	-0.93	5.94	-0.39	-0.60	-0.86	-0.30
INV1	8.89	1.64	1.64	0.61	1.53	8.89	1.65	-0.08	-0.39	0.32	8.89	2.06	-0.17	-0.58	0.23
INV2	8.48	-0.17	0.25	1.57	0.25	8.48	-0.13	0.54	1.76	0.42	8.48	1.48	0.42	2.21	0.21
INV3	7.78	1.28	-0.46	-0.57	-0.29	7.78	1.29	-0.58	-0.55	-0.07	7.78	1.26	-0.72	-1.00	-1.02
INV4	6.57	0.62	-0.58	-0.13	-1.27	6.57	0.64	0.27	0.10	0.33	6.57	1.12	0.18	0.01	1.14
INV5	6.71	-0.45	-0.38	-0.49	-0.11	6.71	-0.45	0.19	-0.32	-0.27	6.71	-0.56	0.11	-0.56	-0.23
INV6	6.64	0.55	-0.50	-0.94	-0.44	6.64	0.57	-0.75	-0.80	-0.56	6.64	0.59	-0.82	-1.13	-1.20
INV7	7.25	1.51	-0.20	-0.31	0.04	7.25	1.52	0.42	0.06	0.47	7.25	1.21	0.36	-0.34	-0.04
INV8	6.30	-0.25	0.12	0.08	0.05	6.30	-0.25	0.02	0.30	0.57	6.30	-0.74	0.03	0.16	-0.08
INV9	7.09	0.83	1.91	1.94	1.82	7.09	0.84	2.31	1.95	2.25	7.09	0.35	2.40	2.02	2.52
INV10	3.93	-3.60	-1.31	-1.17	-1.23	3.93	-3.58	-2.23	-1.37	-1.74	3.93	-3.55	-2.12	-0.98	-2.48
S1						8.83	0.68	0.95	0.25	0.67	8.83	0.75	1.07	0.48	1.07
S2						8.52	-0.29	-0.30	-0.74	-0.32	8.52	-0.40	-0.18	-0.60	-0.15
S3						9.35	0.87	0.32	0.57	0.34	9.35	1.31	0.42	0.81	0.24
S4						8.52	0.18	-0.56	-0.20	-0.21	8.52	0.21	-0.46	-0.11	-0.58
S5						8.92	0.00	0.30	0.51	0.25	8.92	0.16	0.37	0.68	0.38
S6						8.19	-0.25	-0.10	0.22	-0.14	8.19	-0.02	-0.06	0.33	-0.10
S7						8.32	-0.55	0.00	0.14	0.15	8.32	-0.64	0.02	0.13	0.14
S8						7.70	-0.55	0.25	0.01	-0.26	7.70	-0.02	0.25	0.13	0.66
S9						6.95	-0.02	-0.12	-0.38	-0.71	6.95	0.62	-0.16	-0.39	0.09
S10						5.47	-0.33	-0.29	-0.51	-0.72	5.47	0.09	-0.36	-0.58	-0.13
ROE1											1.01	-3.14	-1.23	-1.99	-0.33
ROE2											4.01	-1.59	0.02	0.32	-1.28
ROE3											4.80	2.20	1.70	1.34	0.13
ROE4											5.47	0.13	0.29	0.07	0.33
ROE5											6.56	1.80	0.70	0.70	0.87
ROE6											5.36	-1.67	-0.93	-1.34	-0.25
ROE7											6.90	0.95	0.13	0.63	-0.10
ROE8											6.63	-1.33	-1.12	-0.73	-0.53
ROE9											6.93	-1.13	0.19	-0.54	0.51
ROE10											9.00	1.40	0.47	1.72	0.52
MAPE	5.31	0.87	0.88	0.86	0.84	6.37	0.68	0.70	0.67	0.73	6.18	0.99	0.70	0.78	0.70
Panel B: Prices of Risk Estimates λ_0															
MKT		1.68	5.20	5.55	5.38		1.70	4.70	5.47	4.08		2.43	4.80	6.20	4.97
LVL/SMB/ME		-18.92	-7.64	-2.34	-8.58		-18.47	4.34	0.62	1.24		-18.88	4.19	0.41	5.37
CP/HML/IA		73.57	8.32	8.53	-5.44		75.72	10.74	10.06	11.16		98.47	11.18	11.52	1.52
UMD/ROE/RMW		2.38	-2.35	6.31	-3.53		2.77	9.77	8.62	2.86		11.85	9.73	11.91	15.82
CMA					13.90					-4.39					6.67
Panel C: Test on Risk Prices and Pricing Errors (p-values in %)															
$H_0 : \lambda_0 = 0$		1.90	0.07	0.18	0.09	0.00	6.44	0.30	0.31	0.74	0.00	2.81	0.00	0.04	0.00
H_0 : Pr. error $= 0$		2.31	0.00	0.00	0.02	0.00	1.13	0.00	0.01	0.01	0.00	7.27	0.00	0.00	0.02

we report the corresponding pricing errors and market prices of risk. In all three cases, we estimate CP and LVL factors on the full sample because it facilitates comparison across results (we only vary estimation methodology and beta estimation), and because the CP factor is the outcome of a bond return predictability exercise. A 60-month window is too short for bond return predictability purposes.

The results for the first-stage betas are in Table A.VI. The first three columns estimate betas on the full sample. The second three columns use expanding windows, and average the resulting betas in the time series. The first window is calculated with 60 months of data. The last three columns use 60-month rolling windows, and average the betas in the time series. The first message from this table is that the results are generally consistent across beta estimation methodologies. Specifically, high book-to-market portfolios have a large and positive CP-beta while low book-to-market portfolios have a low, negative CP-beta. The book-to-market portfolios have little systematic relationship with the other two pricing factors. The size portfolios have a declining LVL-beta moving from the small (S1) to the large (S10) decile. The second observation is that the expanding-window and full-sample betas are particularly close, but the rolling-window beta estimates differ non-trivially from the other two. The CP-beta spread of the book-to-market decile portfolios is smaller and the CP-beta spread of the size portfolios is larger. As we shall see below, this has non-trivial implications for the estimated market price of CP risk. Estimating CP-betas based on sixty months of data introduces measurement error, which is natural given that CP is not a return-based factor. The difference between full-sample and rolling-window based betas is larger than it is when using return-based factors.

Table A.VI: Fama-McBeth: First-stage Beta Estimates

	Full-sample betas			Expanding-windows betas			Rolling-windows betas		
	β_{MKT}	β_{LVL}	β_{CP}	β_{MKT}	β_{LVL}	β_{CP}	β_{MKT}	β_{LVL}	β_{CP}
10-yr	0.01	-2.41	-0.25	0.01	-2.23	-0.22	0.02	-2.32	-0.26
7-yr	0.00	-2.15	-0.16	0.01	-1.93	-0.17	0.01	-2.03	-0.19
5-yr	-0.01	-1.82	-0.02	0.00	-1.68	-0.03	0.00	-1.71	-0.05
2-yr	0.00	-0.90	0.18	0.00	-0.84	0.10	0.00	-0.81	0.15
1-yr	0.00	-0.46	0.17	0.00	-0.43	0.11	0.00	-0.40	0.14
Market	1.00	0.00	0.00	0.93	-0.01	-0.03	0.92	-0.06	0.10
BM1	1.07	0.09	-0.31	1.03	0.26	-0.22	1.01	0.15	-0.03
BM2	0.99	-0.14	-0.12	0.93	-0.10	-0.08	0.93	-0.07	-0.07
BM3	0.98	-0.08	0.05	0.92	-0.13	-0.03	0.91	-0.08	0.05
BM4	0.97	-0.05	-0.11	0.89	-0.15	-0.18	0.90	-0.12	0.01
BM5	0.91	-0.10	0.10	0.83	-0.16	0.12	0.83	-0.20	0.24
BM6	0.88	-0.21	0.06	0.83	-0.31	0.03	0.80	-0.14	-0.03
BM7	0.92	-0.13	0.39	0.87	0.02	0.28	0.85	-0.11	0.39
BM8	0.93	-0.08	0.29	0.92	0.05	0.23	0.86	-0.13	0.33
BM9	0.97	0.03	0.43	0.90	0.10	0.26	0.89	-0.07	0.45
BM10	1.13	0.42	0.52	1.03	0.42	0.30	1.03	0.12	0.69
S1	1.07	0.51	0.30	0.99	0.25	0.12	1.01	0.27	0.78
S2	1.17	0.45	0.32	1.03	0.31	0.23	1.09	0.40	0.65
S3	1.17	0.33	0.24	1.05	0.10	0.07	1.10	0.26	0.57
S4	1.14	0.30	0.23	1.03	0.16	0.11	1.07	0.29	0.55
S5	1.14	0.28	0.33	1.02	0.12	0.19	1.06	0.23	0.61
S6	1.09	0.07	0.25	0.99	0.00	0.12	1.02	0.02	0.52
S7	1.09	0.05	0.27	0.99	0.00	0.14	1.01	0.02	0.43
S8	1.08	-0.07	0.23	0.97	-0.18	0.15	1.00	-0.11	0.33
S9	1.00	-0.17	0.06	0.91	-0.24	0.01	0.92	-0.21	0.17
S10	0.94	-0.05	-0.11	0.89	0.02	-0.10	0.86	-0.08	-0.10

The results for the market prices of risk are in Table A.VII. Panel A reports the pricing errors. The first column restates the pricing errors that are to be explained. The MAPE is 7.13% per year. The second column repeats our benchmark estimation results for this sample using GMM. The MAPE is 53 bps. The last three columns are for the Fama-MacBeth method, where betas are estimated on the full sample, from expanding windows, and from rolling-windows, respectively. Panel B reports the market prices of risk. We denote the market price of risk estimates in the Fama-MacBeth exercise that are associated with the first-stage betas by ℓ, and report their point estimates standard errors (s.e.). The last three rows transform these estimates into the same units as our market prices of risk associated with covariances, λ_0. The latter are comparable across GMM and Fama-MacBeth approaches. We use the delta method to compute the corresponding standard errors.

Three main findings stand out. First, and most importantly, the GMM and FMB full-sample approaches deliver nearly identical results. The MAPE are 53 bps and 54 bps, respectively. The implied risk prices λ_0 are nearly identical as well. In particular, the market price for CP risk is 78.5 in the GMM and 77.4 in the FMB approach, statistically indistinguishable from one another. The prices of MKT and LVL risk are also very close. We conclude that our main results are invariant to whether we use GMM or FMB.

Second, the expanding-window FMB results remain close to the GMM results. The MAPE increases slightly to 58 bps and the risk price estimate for CP increases slightly to 88.1. Third, the results are weaker with 60-month rolling-window FMB. The MAPE rises further to 65 bps, while the price of CP risk drops to 31.5, which is insignificant. The deteriorated performance suggests that the estimation error in betas is substantial when using 60-month rolling windows. We have repeated our estimation for the other factor models and we generally find very similar results when using GMM or FMB.

Table A.VII: Fama-McBeth: Pricing Errors and Market Prices of Risk

	RN	GMM	Full sample	Exp.-window	Rol.-window
			Panel A: Pricing Errors		
10-yr	1.74	0.15	0.13	0.11	0.35
7-yr	2.03	0.39	0.38	0.65	1.08
5-yr	1.69	-0.19	-0.20	-0.12	0.46
2-yr	1.19	-0.64	-0.63	-0.33	0.37
1-yr	0.94	-0.41	-0.40	-0.27	0.25
Market	6.83	-0.72	-0.69	-0.86	-1.07
BM1	6.31	-0.21	-0.25	-1.11	-1.30
BM2	7.00	-0.09	-0.06	-0.52	-0.32
BM3	7.35	-0.39	-0.38	-0.39	-0.34
BM4	7.33	0.45	0.48	0.58	0.09
BM5	7.76	0.30	0.33	-0.48	-0.32
BM6	8.72	1.66	1.68	1.01	1.32
BM7	7.57	-1.32	-1.32	-1.11	-1.48
BM8	9.41	0.95	0.94	0.06	-0.10
BM9	10.63	1.37	1.41	1.87	1.44
BM10	10.84	0.48	0.42	1.20	0.84
S1	9.58	0.60	0.65	0.90	1.57
S2	9.33	-0.50	-0.56	-0.46	0.24
S3	10.02	0.43	0.40	1.09	1.19
S4	9.20	-0.18	-0.21	0.27	0.28
S5	9.47	-0.36	-0.37	0.03	0.30
S6	9.00	-0.31	-0.32	0.07	0.29
S7	8.86	-0.56	-0.56	-0.08	0.51
S8	8.46	-0.78	-0.80	-0.62	-0.13
S9	7.83	-0.21	-0.18	-0.17	-0.25
S10	6.36	-0.24	-0.22	-0.63	-1.09
MAPE	7.13	0.53	0.54	0.58	0.65
			Panel B: Risk Prices		
ℓ_{MKT}			7.52	7.07	7.48
ℓ_{LVL}			-1.13	-1.25	-0.92
ℓ_{CP}			4.66	5.06	3.08
s.e. ℓ_{MKT}			0.55	0.58	0.58
s.e. ℓ_{LVL}			0.11	0.12	0.12
s.e. ℓ_{CP}			0.48	0.52	0.27
$\lambda_{0,MKT}$		2.62	2.58	3.04	3.60
$\lambda_{0,LVL}$		-18.26	-18.25	-9.75	5.50
$\lambda_{0,CP}$		78.51	77.43	88.07	31.46
s.e. $\lambda_{0,MKT}$		1.15	0.88	1.14	1.34
s.e. $\lambda_{0,LVL}$		7.91	5.56	8.94	23.80
s.e. $\lambda_{0,CP}$		33.95	27.80	34.52	37.03

C. How Pricing Stocks and Bonds Jointly Can Go Wrong

Consider two factors F_t^i, $i = 1, 2$, with innovations η_{t+1}^i. We normalize $\sigma\left(\eta_{t+1}^i\right) = 1$. Let $cov\left(\eta_{t+1}^1, \eta_{t+1}^2\right) = \rho = corr\left(\eta_{t+1}^1, \eta_{t+1}^2\right)$. We also have two cross-sections of test assets with excess, geometric returns r_{t+1}^{ki}, $i = 1, 2$ and $k = 1, ..., K_i$, with innovations ε_{t+1}^{ki}. We assume that these returns include the Jensen's correction term. Suppose that both cross-sections exhibit a one-factor pricing structure:

$$E\left(r_{t+1}^{ki}\right) = cov\left(\varepsilon_{t+1}^{ki}, \eta_{t+1}^i\right)\lambda_i,\ i = 1, 2.$$

The first factor perfectly prices the first set of test assets, whereas the second factor prices the second set of test assets. We show below that this does *not* imply that there exists a single SDF that prices both sets of assets.

Consider the following model of unexpected returns for both sets of test assets:

$$\begin{aligned}
\varepsilon_{t+1}^{k1} &= E\left(r_{t+1}^{k1}\right)\eta_{t+1}^1, \\
\varepsilon_{t+1}^{k2} &= E\left(r_{t+1}^{k2}\right)\eta_{t+1}^2 + \alpha_{2k}\eta_{t+1}^3,
\end{aligned}$$

with $cov\left(\eta_{t+1}^2, \eta_{t+1}^3\right) = 0$. Unexpected returns on the first set of test assets are completely governed by innovations to the first factor, whereas unexpected returns on the second set of test assets contain a component $\alpha_{2k}\eta_{t+1}^3$ that is orthogonal to the component governed by innovations to the second factor. These η^3 shocks are not priced (by assumption). We assume that they are correlated with the η^1 shocks: $cov\left(\eta_{t+1}^1, \eta_{t+1}^3\right) \neq 0$.

This structure implies:

$$cov\left(\varepsilon_{t+1}^{ki}, \eta_{t+1}^i\right) = E\left(r_{t+1}^{ki}\right) var\left(\eta_{t+1}^i\right) = E\left(r_{t+1}^{ki}\right),$$

and hence $\lambda_i = 1$, $i = 1, 2$. Then we have:

$$\begin{aligned}
&cov\left(\varepsilon_{t+1}^{k1}, \eta_{t+1}^1\right) = E\left(r_{t+1}^{k1}\right), && cov\left(\varepsilon_{t+1}^{k1}, \eta_{t+1}^2\right) = E\left(r_{t+1}^{k1}\right)\rho, \\
&cov\left(\varepsilon_{t+1}^{k2}, \eta_{t+1}^1\right) = E\left(r_{t+1}^{k2}\right)\rho + \alpha_{2k} cov\left(\eta_{t+1}^1, \eta_{t+1}^3\right), && cov\left(\varepsilon_{t+1}^{k2}, \eta_{t+1}^2\right) = E\left(r_{t+1}^{k2}\right).
\end{aligned}$$

The main point is that, if α_{2k} is not proportional to $E\left(r_{t+1}^{k2}\right)$, then there exist no λ_1 and l_2 such that:

$$E\left(r_{t+1}^{ki}\right) = cov\left(\varepsilon_{t+1}^{ki}, \eta_{t+1}^1\right)\lambda_1 + cov\left(\varepsilon_{t+1}^{ki}, \eta_{t+1}^2\right)\lambda_2.$$

On the other hand, if there is proportionality and $\alpha_{2k} = \alpha E\left(r_{t+1}^{k2}\right)$, then we have:

$$cov\left(\varepsilon_{t+1}^{k2}, \eta_{t+1}^1\right) = E\left(r_{t+1}^{k2}\right)\left(\rho + \alpha cov\left(\eta_{t+1}^1, \eta_{t+1}^3\right)\right) = E\left(r_{t+1}^{k2}\right)\xi,$$

and λ_1 and λ_2 are given by:

$$\lambda_1 = \frac{1-\rho}{1-\xi\rho}, \text{and } \lambda_2 = \frac{1-\xi}{1-\xi\rho}.$$

This setup is satisfied approximately in our model, where the first set of test assets are the book-to-market portfolios, η^1 are CP innovations, the second set of test assets are the bond portfolios, and η^2 are LVL innovations. Unexpected bond returns contain a component η^3 that is uncorrelated with LVL innovations, but that is correlated with CP innovations. Unexpected book-to-market portfolio returns, in contrast, are largely uncorrelated with LVL innovations. The result above illustrates that consistent risk pricing is possible because unexpected bond returns' exposure to CP shocks has a proportionality structure. This can also be seen in the top panel of Figure 5.

D. Order of the VAR Model

In our benchmark model, we use a first-order VAR for the bond factors, consistent with most of the term structure literature. In this appendix, we explore the robustness to the order of the VAR model. We use three model selection criteria, namely AIC, HQ, and SC. The statistics are reported in Table A.VIII. The preferred model is the one that minimizes a given statistic. The different selection criteria point to VAR of an order between one and three, suggesting that our current model provides a reasonable description of the various time series.

Table A.VIII: Lag Selection Criteria for VAR Model

Order	AIC	HQ	SC
1	-25.95	-25.93	-25.82
2	-26.02	-25.98	-25.75
3	-26.05	-25.99	-25.66
4	-26.05	-25.97	-25.52
5	-26.04	-25.93	-25.37
6	-26.04	-25.91	-25.24
7	-26.04	-25.89	-25.10
8	-26.04	-25.87	-24.97
9	-26.02	-25.83	-24.82
10	-26.02	-25.81	-24.69
11	-26.00	-25.77	-24.54
12	-26.03	-25.77	-24.43
Selected order	3	3	1

We are ultimately interested in the model's ability to price the cross-section of expected returns. In Table A.IX, we report the pricing errors, the MAPE, and the risk prices for the different VAR models of order one to three. We find that along all dimensions, the models look very similar. The MAPE increases slightly when we increase the number of lags from one to either two or three. The difference arises from slightly larger pricing errors for the BM portfolios, while the pricing error of the aggregate stock market decreases. The pricing of the bond portfolios is virtually unaffected. Our main pricing results are robust to changing the order of the VAR model.

Table A.IX: Asset Pricing With Different VAR Models

We estimate the pricing model for VAR models of order 1 to 3. The first column reports the risk premium, which corresponds to a VAR model of order 0. we report the pricing errors, the MAPE, and the risk prices. The sample is from June 1952 to December 2015.

	Panel A: Pricing Errors (% per year)			
		Order of the VAR model		
	RN	1	2	3
10-yr	1.74	0.20	0.22	0.22
7-yr	2.03	0.41	0.39	0.37
5-yr	1.69	-0.22	-0.26	-0.26
2-yr	1.19	-0.73	-0.74	-0.73
1-yr	0.94	-0.49	-0.51	-0.48
Market	6.83	-0.79	-0.85	-0.87
BM1	6.49	-0.20	-0.56	-0.60
BM2	7.20	-0.19	-0.03	-0.05
BM3	8.24	0.90	1.21	1.29
BM4	8.43	-0.45	-0.29	-0.29
BM5	10.64	0.79	0.65	0.64
MAPE	5.04	0.49	0.52	0.53
	Panel B: Prices of Risk Estimates λ_0			
MKT		2.60	2.80	2.82
LVL/SMB		-18.90	-21.74	-20.59
CP/HML		85.63	89.47	88.40

E. Robustness to Time-varying Second Moments

Throughout the paper, we make the simplifying assumption that second moments of returns are constant. We now explore the impact of time-varying second moments. By relaxing the assumption of homoscedasticity, conditional expected returns are given by:

$$E_t\left[rx_{t+1}^j\right] = \Sigma_{Xj,t}\Lambda_t, \tag{A.3}$$

where

$$\Sigma_{Xj,t} = Cov_t\left[\eta_{t+1}^j, \epsilon_{t+1}'\right] = E_t\left[\eta_{t+1}^j \epsilon_{t+1}'\right]. \tag{A.4}$$

We can rewrite (A.3) as

$$E_t\left[rx_{t+1}^j\right] = \Sigma_{Xj,t}\lambda_0 + \Sigma_{Xj,t}\lambda_1 CP_t^\star. \tag{A.5}$$

Unconditional expected returns are given by

$$E\left[rx_{t+1}^j\right] = \Sigma_{Xj}\lambda_0 + E\left[\eta_{t+1}^j \epsilon_{t+1}' CP_t^\star\right]\lambda_1. \tag{A.6}$$

This derivation shows that it does not matter how we model the conditional variance of returns for pricing the cross-section of expected returns. It does matter, however, how we model the conditional *covariance* of returns with the pricing factors, $E\left[\eta_{t+1}^j \epsilon_{t+1}' CP_t^\star\right]$. The bias caused by assuming constant second moments depends on the covariance between CP_t and $\eta_{t+1}^j \epsilon_{t+1}$. To make this point precise, we define the bias for portfolio j as the difference in expected returns in case of the model with time-varying second moments and our model with constant expected returns, which equals

$$bias_j = \left(E\left[\eta_{t+1}^j \epsilon_{t+1}' CP_t^\star\right] - E\left[\eta_{t+1}^j \epsilon_{t+1}'\right] E\left[CP_t^\star\right]\right)\lambda_1 = Cov\left[\eta_{t+1}^j \epsilon_{t+1}', CP_t^\star\right]\lambda_1, \tag{A.7}$$

assuming λ_1 is the same across both models. We can estimate the bias *without* assuming a particular model for second moments as we can estimate the unconditional covariance between CP_t and $\eta_{t+1}^j \epsilon_{t+1}'$ directly in the data. Our estimates imply that the bias is smaller than 0.05bp for all of the test assets in Table 1. We assume that our results are unaffected by time-varying second moments.

F. Structural Model with Business Cycle Risk

This appendix sets up and calibrates a structural asset pricing model that connects our empirical findings in a transparent way. The model formalizes the relationships between the returns on value and growth stocks, the CP factor, and the state of the macro-economy. It does so in a pricing framework that can quantitatively account for the observed risk premia on stock and bond portfolios, while being consistent with the observed dynamics of dividend growth rates, inflation, and basic properties of the term structure of interest rates. Its role is largely pedagogical: to clarify the minimal structure necessary to account for the observed moments. We start by describing the setup and provide the derivations of the asset pricing expressions. We also discuss the parameters used in the numerical example, and how they were chosen.

F.1. Setup

The model has one key state variable, s, which measures macroeconomic activity. One interpretation of s is as a leading business cycle indicator. This state variable follows an autoregressive process, with modest persistence, and its innovations ε^s_{t+1} are the first priced source of risk.

$$s_{t+1} = \rho_s s_t + \sigma_s \varepsilon^s_{t+1}.$$

Higher values of s denote higher economic activity. The model permits an interpretation of s as a signal about *future* economic activity. Since this variable moves at business cycle frequency, the persistence ρ_s is moderate.

Real dividend growth on asset $i = \{G, V, M\}$ (Value, Growth, and the Market) is given by:

$$\Delta d^i_{t+1} = \gamma_{0i} + \gamma_{1i} s_t + \sigma_{di} \varepsilon^d_{t+1} + \sigma_i \varepsilon^i_{t+1}. \tag{A.8}$$

If $\gamma_{1i} > 0$, dividend growth is pro-cyclical. The shock ε^d_{t+1} is an aggregate dividend shock, while ε^i_{t+1} is an (non-priced) idiosyncratic shock; the market portfolio has no idiosyncratic risk; $\sigma_M = 0$. The key parameter configuration is $\gamma_{1V} > \gamma_{1G}$ so that value stocks are more exposed to cyclical risk than growth stocks. As is the data (Section 2.2.1), a low value for s is associated with lower future dividend growth on V minus G. Below, we will calibrate γ_{1V} and γ_{1G} to capture the decline in dividend growth value minus growth over the course of recessions.

Inflation is the sum of a constant, a mean-zero autoregressive process which captures expected inflation, and an unexpected inflation term:

$$\begin{aligned}\pi_{t+1} &= \bar{\pi} + x_t + \sigma_\pi \varepsilon^\pi_{t+1}, \\ x_{t+1} &= \rho_x x_t + \sigma_x \varepsilon^x_{t+1}.\end{aligned}$$

All shocks are cross-sectionally and serially independent and standard normally distributed. It would be straightforward to add a correlation between inflation shocks and shocks to the business cycle variable. This inflation process is common in the literature (e.g., Wachter, 2006; Bansal and Shaliastovich, 2013).

To simplify our analysis, we assume that market participants' preferences are summarized by a real stochastic discount factor (SDF), whose log evolves according to the process:

$$-m_{t+1} = y + \frac{1}{2}\Lambda_t'\Lambda_t + \Lambda_t'\varepsilon_{t+1}.$$

where the vector $\varepsilon_{t+1} \equiv \left(\varepsilon^d_{t+1}, \varepsilon^x_{t+1}, \varepsilon^s_{t+1}\right)'$ and y is the real interest rate. The risk price dynamics are affine in the state of the economy s_t:

$$\Lambda_t = \lambda_0 + \lambda_1 s_t$$

For similar approaches to the SDF, see Bekaert, Engstrom, and Xing (2009), Bekaert, Engstrom, and Grenadier (2010), Lettau and Wachter (2007, 2011), Campbell, Sunderam, and Viceira (2013), and David and Veronesi (2013).

As in the reduced form model in the main text, the structural model features three priced sources of risk: aggregate dividend growth risk, which carries a positive price of risk ($\lambda_0(1) > 0$), inflation risk ($\lambda_0(2) < 0$), and cyclical risk ($\lambda_0(3) > 0$). Choosing $\lambda_1(2) < 0$ makes the price of inflation risk counter-cyclical. As we show below, this makes bond risk premia pro-cyclical. We also set $\lambda_1(1) > 0$ resulting in a pro-cyclical price of aggregate dividend risk. The log nominal SDF is given by $m^\$_{t+1} = m_{t+1} - \pi_{t+1}$.

F.2. Asset Prices

We now study the equilibrium bond and stock prices in this model. The model generates an affine nominal term structure of interest rates. It also generates a one-factor model for the nominal bond risk premium: All variation in bond risk premia comes from cyclical variation in the economy, s_t. Thus, the CP factor which measures the bond risk premium in the model is perfectly positively correlated with s_t, the (leading) indicator of macroeconomic activity.

F.2.1. Bond Prices and Risk Premia

It follows immediately from the specification of the real SDF that the real term structure of interest rates is flat at y. Nominal bond prices are exponentially affine in the state of the economy and in expected inflation:

$$P_t^{\$}(n) = \exp\left(A_n^{\$} + B_n^{\$} s_t + C_n^{\$} x_t\right),$$

with coefficients that follow recursions described in the proof below. As usual, nominal bond yields are $y_t^{\$}(n) = -\log(P_t^{\$}(n))/n$.

Proof. The nominal SDF is given by:

$$\begin{aligned}
m_{t+1}^{\$} &= m_{t+1} - \pi_{t+1} \\
&= -y - \bar{\pi} - x_t - \frac{1}{2}\Lambda_t'\Lambda_t - \Lambda_t'\varepsilon_{t+1} - \sigma_\pi \varepsilon_{t+1}^{\pi}
\end{aligned}$$

The recursion of nominal bond prices is given by:

$$\begin{aligned}
P_t^{\$}(n) &= E_t\left(P_{t+1}^{\$}(n-1) M_{t+1}^{\$}\right) \\
&= E_t\left(\exp\left(A_{n-1}^{\$} + B_{n-1}^{\$} s_{t+1} + C_{n-1}^{\$} x_{t+1} - y - \bar{\pi} - x_t - \frac{1}{2}\Lambda_t'\Lambda_t - \Lambda_t'\varepsilon_{t+1} - \sigma_\pi \varepsilon_{t+1}^{\pi}\right)\right) \\
&= \exp\left(A_{n-1}^{\$} - y - \bar{\pi} - x_t - \frac{1}{2}\Lambda_t'\Lambda_t + B_{n-1}^{\$}\rho_s s_t + C_{n-1}^{\$}\rho_x x_t\right) \times \\
&\quad E_t\left(\exp\left(B_{n-1}^{\$}\sigma_s \varepsilon_{t+1}^{s} + C_{n-1}^{\$}\sigma_x \varepsilon_{t+1}^{x} - \Lambda_t'\varepsilon_{t+1} - \sigma_\pi \varepsilon^{\pi} t + 1\right)\right) \\
&= \exp\left(A_{n-1}^{\$} - y - \bar{\pi} - x_t + B_{n-1}^{\$}\rho_s s_t + C_{n-1}^{\$}\rho_x x_t\right) \times \\
&\quad \exp\left(\frac{1}{2}[B_{n-1}^{\$}]^2\sigma_s^2 + \frac{1}{2}[C_{n-1}^{\$}]^2\sigma_x^2 - B_{n-1}^{\$}\sigma_s\Lambda_t(3) - C_{n-1}^{\$}\sigma_x\Lambda_t(2) + \frac{1}{2}\sigma_\pi^2\right),
\end{aligned}$$

which implies:

$$\begin{aligned}
A_n^{\$} &= A_{n-1}^{\$} - y - \bar{\pi} + \frac{1}{2}[B_{n-1}^{\$}\sigma_s]^2 + \frac{1}{2}[C_{n-1}^{\$}\sigma_x]^2 + \frac{1}{2}\sigma_\pi^2 - B_{n-1}^{\$}\sigma_s\lambda_0(3) - C_{n-1}^{\$}\sigma_x\lambda_0(2), \\
B_n^{\$} &= B_{n-1}^{\$}\rho_s - C_{n-1}^{\$}\sigma_x\lambda_1(2), \\
C_n^{\$} &= -1 + C_{n-1}^{\$}\rho_x.
\end{aligned}$$

The starting values for the recursion are $A_0^{\$} = 0$, $B_0^{\$} = 0$, and $C_0^{\$} = 0$. □

The expression for $C_n^{\$}$ can be written more compactly as:

$$C_n^{\$} = -\frac{1-\rho_x^n}{1-\rho_x}\rho_x < 0,$$

implying that bond prices drop -and nominal interest rates increase- when expected inflation increases ($C_n^{\$} < 0$). Consistent with the data, we assume that $\lambda_1(2) < 0$. It follows that $B_n^{\$} < 0$, implying that nominal bond prices fall -and nominal interest rates rise- with the state of the economy (s_t). Both signs seem consistent with intuition.

The nominal bond risk premium, the expected excess log return on buying an n-period nominal bond and selling it one period later (as a $n-1$-period bond), is given by:

$$\begin{aligned}
E_t\left[rx^{\$}_{t+1}(n)\right] &= -cov_t\left(m^{\$}_{t+1}, B^{\$}_{n-1}s_{t+1} + C^{\$}_{n-1}\pi_{t+1}\right) \\
&= cov_t\left(\Lambda_t'\varepsilon_{t+1}, B^{\$}_{n-1}s_{t+1} + C^{\$}_{n-1}x_{t+1}\right) \\
&= \Lambda_t(2)C^{\$}_{n-1}\sigma_x + \Lambda_t(3)B^{\$}_{n-1}\sigma_s \\
&= \underbrace{\lambda_0(2)C^{\$}_{n-1}\sigma_x + \lambda_0(3)B^{\$}_{n-1}\sigma_s}_{\text{Constant component bond risk premium}} + \underbrace{\lambda_1(2)C^{\$}_{n-1}\sigma_x s_t}_{\text{Time-varying component bond risk premium}}.
\end{aligned}$$

In this model, all of the variation in bond risk premia comes from cyclical variation in the economy, s_t. This lends the interpretation of CP factor to s_t which is consistent with our empirical evidence. Innovations to the CP factor are innovations to s (ε^s). Because $C^{\$}_{n-1} < 0$, $\lambda_1(2) < 0$ generates lower bond risk premia when economic activity is low ($s_t < 0$). The effect is stronger for longer-maturity bonds.

The constant component of the bond risk premium partly reflects compensation for cyclical risk and partly exposure to expected inflation risk. Exposure to the cyclical shock contributes negatively to excess bond returns: A positive ε^s shock lowers bond prices and returns, and more so for long than for short bonds. Exposure to expected inflation shocks contributes positively to excess bond returns: A positive ε^x shock lowers bond prices and returns but the price of expected inflation risk is negative. Since common variation in bond yields is predominantly driven by the inflation shock in the model, the latter acts like (and provides a structural interpretation for) a shock to the level of the term structure (LVL). Long bonds are more sensitive to level shocks, the traditional duration effect.

F.2.2. Stock Prices, Equity Risk Premium, Value Premium

The log price-dividend (pd) ratio on stock (portfolio) i is affine in s_t:

$$pd^i_t = A_i + B_i s_t,$$

where

$$B_i = \frac{\gamma_{1i} - \lambda_1(1)\sigma_{di}}{1 - \kappa_{1i}\rho_s},$$

and the expression for A_i is given in the proof below.

Proof. The return definition implies:

$$\begin{aligned}
r_{t+1} &= \ln\left(\exp\left(pd_{t+1}\right) + 1\right) + \Delta d_{t+1} - pd_t \\
&\simeq \ln\left(\exp\left(\overline{pd}\right) + 1\right) + \frac{\exp\left(\overline{pd}\right)}{\exp\left(\overline{pd}\right) + 1}\left(pd_{t+1} - \overline{pd}\right) + \Delta d_{t+1} - pd_t \\
&= \kappa_0 + \kappa_1 pd_{t+1} + \Delta d_{t+1} - pd_t,
\end{aligned}$$

where:

$$\begin{aligned}
\kappa_0 &= \ln\left(\exp\left(\overline{pd}\right) + 1\right) - \kappa_1\overline{pd}, \\
\kappa_1 &= \frac{\exp\left(\overline{pd}\right)}{\exp\left(\overline{pd}\right) + 1}.
\end{aligned}$$

We conjecture that the log price-dividend ratio is of the form:

$$pd_t = A + B s_t,$$

The price-dividend ratio coefficients are obtained by solving the Euler equation:

$$E_t\left(M_{t+1}^{\$}R_{t+1}^{\$}\right)=1.$$

We suppress the dependence on i in the following derivation:

$$\begin{aligned}
1 &= E_t\left(\exp\left(m_{t+1}-\pi_{t+1}+\kappa_0+\kappa_1 pd_{t+1}+\Delta d_{t+1}-pd_t+\pi_{t+1}\right)\right)\\
0 &= E_t\left(m_{t+1}\right)+\frac{1}{2}V_t\left(m_{t+1}\right)+E_t\left(\kappa_0+\Delta d_{t+1}+\kappa_1 pd_{t+1}-pd_t\right)\\
&\quad +\frac{1}{2}V_t\left(\Delta d_{t+1}+\kappa_1 pd_{t+1}\right)+Cov_t\left(-\Lambda_t'\varepsilon_{t+1},\Delta d_{t+1}+\kappa_1 pd_{t+1}\right)\\
&= -y+\kappa_0+\gamma_0+\gamma_1 s_t+\left(\kappa_1-1\right)A+\left(\kappa_1\rho_s-1\right)Bs_t\\
&\quad +\frac{1}{2}\sigma_d^2+\frac{1}{2}\sigma^2+\frac{1}{2}\kappa_1^2B^2\sigma_s^2-\Lambda_t(1)\sigma_d-\Lambda_t(3)\kappa_1 B\sigma_s.
\end{aligned}$$

This results in the system:

$$\begin{aligned}
0 &= -y+\kappa_0+\gamma_0+\left(\kappa_1-1\right)A+\frac{1}{2}\sigma_d^2+\frac{1}{2}\sigma^2+\frac{1}{2}\kappa_1^2B^2\sigma_s^2-\lambda_0(1)\sigma_d-\lambda_0(3)\kappa_1 B\sigma_s,\\
0 &= \left(\kappa_1\rho_s-1\right)B-\lambda_1(1)\sigma_d+\gamma_1,
\end{aligned}$$

Rearranging terms, we get the following expressions for the pd ratio coefficients, where we make the dependence on i explicit:

$$\begin{aligned}
A_i &= \frac{\frac{1}{2}\sigma_{di}^2+\frac{1}{2}\sigma_i^2+\frac{1}{2}\kappa_{1i}^2B_i^2\sigma_s^2-\lambda_0(1)\sigma_{di}-\lambda_0(3)\kappa_{1i}B_i\sigma_s-y+\kappa_{0i}+\gamma_{0i}}{1-\kappa_{1i}},\\
B_i &= \frac{\gamma_{1i}-\lambda_1(1)\sigma_{di}}{1-\kappa_{1i}\rho_s}.
\end{aligned}$$

□

We note that B_i can be positive or negative depending on the importance of dividend growth predictability (γ_{1i}) and fluctuations in risk premia ($\lambda_1(1)\sigma_{di}$). Stock i's price-dividend ratios is pro-cyclical ($B_i>0$) when dividend growth is more pro-cyclical than the risk premium for the aggregate dividend risk of asset i: $\gamma_{1i}>\sigma_{di}\Lambda_1(1)$.

The equity risk premium on portfolio i can be computed to be:

$$\begin{aligned}
E_t\left[rx_{t+1}^i\right] &= cov_t\left(-m_{t+1}^{\$},r_{t+1}^i+\pi_{t+1}\right)\\
&= cov\left(\Lambda_t'\varepsilon_{t+1},\kappa_{1i}B_i\sigma_s\varepsilon_{t+1}^s+\sigma_{di}\varepsilon_{t+1}^d\right)\\
&= \underbrace{\lambda_0(1)\sigma_{di}+\lambda_0(3)\kappa_{1i}B_i\sigma_s}_{\text{Constant component equity risk premium}}+\underbrace{\lambda_1(1)\sigma_{di}s_t}_{\text{Time-varying component equity risk premium}}.
\end{aligned}$$

The equity risk premium provides compensation for aggregate dividend growth risk (first term, ε^d) and for cyclical risk (second term, ε^s). Like bond risk premia, equity risk premia vary over time with the state of the economy s_t (third term). The model generates both an equity risk premium and a value premium. The reason for the value premium can be traced back to the fact that value stocks' dividends are more sensitive to cyclical shocks than those of growth stocks. As we showed above, the data suggest that value stocks' dividends fall more in recessions than those of growth stocks ($\gamma_{1V}>\gamma_{1G}$). With $\sigma_{dV}\approx\sigma_{dG}$, this implies that $B_V>B_G$. Because the price of cyclical risk $\Lambda_0(3)$ is naturally positive, the second term delivers the value premium. Put differently, in the model, as in the data, returns on value stocks are more exposed to bond risk premium shocks than returns on growth stocks.

F.2.3. Link with Reduced-form Model

To make the link with the reduced-form model of Section 3 clear, we study the link between the structural shocks and the reduced form shocks. In the model, shocks to the market return (MKT) are given a linear combination of ε^d and ε^s shocks:

$$\varepsilon^{MKT}_{t+1} \equiv r^M_{t+1} - E_t[r^M_{t+1}] = \sigma_{dM}\varepsilon^d_{t+1} + \kappa_{1M} B_M \sigma_s \varepsilon^s_{t+1}$$

We construct the CP factor in the same way as in the data, from yields on 1- through 5-year yields and average excess bond returns. Since the model has a two-factor structure for bond yields and forward rates, we use only the two- and the five-year forward rates as independent variables in the CP regression of average excess returns on forward rates. The model's CP factor is perfectly correlated with the process s, and has a innovations that differs by a factor σ^{CP}: $\varepsilon^{CP}_{t+1} = \varepsilon^s_{t+1}\sigma^{CP}$. Finally, since expected inflation drives most of the variation in bond yields in the model, LVL shocks in the model are proportional to expected inflation shocks: $\varepsilon^{LVL}_{t+1} = \varepsilon^x_{t+1}\sigma^L$. Denote $\tilde{\varepsilon} = [\varepsilon^{MKT}, \varepsilon^{LVL}, \varepsilon^{CP}]'$. Associated with $\tilde{\varepsilon}$, we can define market prices of risk $\tilde{\Lambda}$, such that SDF innovations remain unaltered: $\Lambda'_t\varepsilon_{t+1} = \tilde{\Lambda}'_t\tilde{\varepsilon}_{t+1}$. It is easy to verify that $\tilde{\lambda}_0(1) = \lambda_0(1)/\sigma_{dM}$, $\tilde{\lambda}_0(2) = \lambda_0(2)/\sigma^L$, and $\tilde{\lambda}_0(3) = \lambda_0(3)/\sigma^{CP} - \kappa_{1M}B_M\sigma_s\lambda_0(1)/(\sigma_{dM}\sigma^{CP})$.

For each asset, we can compute covariances of unexpected returns with the MKT, LVL, and CP shocks inside the model. In the model that first covariance is given by:

$$cov(r^i_{t+1} - E_t[r^i_{t+1}], \varepsilon^{MKT}_{t+1}) = \sigma_{dM}\sigma_{di} + \kappa_{1M}B_M\kappa_{1i}B_i\sigma^2_s.$$

A calibration where $B_M \approx 0$ and $\sigma_{dV} \approx \sigma_{dG}$ will replicate the observed pattern (the linearization constant κ_{1i} will be close to 1 for all portfolios). Second, the covariance of stock portfolio returns with CP shocks is given by:

$$cov(r^i_{t+1} - E_t[r^i_{t+1}], \varepsilon^{CP}_{t+1}) = \kappa_{1i}B_i\sigma_s\sigma^{CP}.$$

The model generates a value premium because of differential exposure to CP shocks when $B_V > B_G$. When $\sigma_{dV} \approx \sigma_{dG}$, the stronger loading of expected dividend growth of value stocks to s_t ($\gamma_{1V} > \gamma_{1G}$) makes $B_V > B_G$. Put differently, in the model -as in the data- returns on value stocks are more exposed to bond risk premium shocks than returns on growth stocks. Third, stock return innovations have a zero covariance with LVL shocks in the model by construction, similar to the small exposures in the data.

Likewise, we can compute covariances of bond return innovations with the MKT, LVL, and CP shocks. In that order, they are:

$$B^{\$}_n\kappa_{1M}B_M\sigma_s, \quad C^{\$}_n\sigma_x\sigma^L, \quad B^{\$}_n\sigma_s\sigma^{CP}.$$

When $B_M \approx 0$, exposure of bond returns to the market factor shocks is close to zero. Exposure to level shocks is negative: an increase in the level of interest rates reduces bond prices and returns. Exposure to CP shocks is also negative: an increase in the bond risk premium reduces bond prices and returns. Both exposures become more negative with the horizon because $B^{\$}_n$ and $C^{\$}_n$ increase in absolute value with maturity n.

F.3. Calibration

This section describes our calibration. We start by describing how we define recessions in the model. We construct recessions in the model in a procedure that mimics the NBER dating algorithm and that matches the frequency and duration of recessions. Second, we describe the calibration of dividends and inflation processes. Third, we describe the choice of market price of risk parameters.

Recessions in the Model In order to measure how dividends change over the recession, we have to define recessions in the model. Our algorithm mimics several of the features of the NBER dating procedure: (i) The recession is determined by looking back in time at past real economic activity (s_t in the model) and its start is not known in real time, (ii) there is a minimum recession length, and (iii) it captures the notion that the economy went through a sequence of negative shocks and that economic activity is at a low level. We split each recession into three equal periods and refer to the last month of each period as the first, second, and third stage of the recession. The s process is negative at the start of the recession, falls considerably in the first stage of a recession, continues to fall in the second stage, and partially recovers in the last stage. Our recession dating procedure is novel, matches the empirical distribution of recession duration, and generates interesting asset pricing dynamics during recessions, to which we return to below. We now describe the recession dating procedure in detail.

Recessions in the model are determined by the dynamics of the state process s_t. Define the cumulative shock process $\chi_t \equiv \sum_{k=0}^{K} \varepsilon^s_{t-k}$, where the parameter K governs the length of the backward-looking window. Let $\underline{\chi}$ and $\overline{\chi}$ be the p_1^{th} and p_2^{th} percentiles of the distribution of χ_t, respectively, and let $\underline{s}$ be the be the p_3^{th} percentile of the distribution of the s process. Whenever $\chi_t < \underline{\chi}$, we find the first negative shock between $t-K$ and t; say it occurs in month $t-j$. If, in addition, $s_{t-j} < \underline{s}$, we say that the recession started in month $t-j$. We say that the recession ends the fist month that $\chi_{t+i} > \overline{\chi}$, for $i \geq 1$. We assume that a new recession cannot start before the previous one has ended.

We find the recession parameters (K, p_1, p_2, p_3) by matching features of the fifteen recessions in the 1926-2009 data. In particular, we consider the fraction of recession months (19.86% in the data), the average length of a recession (13.3 months), the minimum length of a recession (6 months), the 25^{th} percentile (8 months), the median (11 months), the 75^{th} percentile (14.5 months), and the maximum length (43 months). We simulate the process for s_t for 10,000 months, determine recession months as described above, and calculate the weighted distance between the seven moments in the simulation and in the data. We iterate on the procedure to find the four parameters that minimize the distance between model and data. The weighting matrix is diagonal and takes on the following values: .9, .9, .7, .5, .7, .5, and .5, where the weights are described in the same order as the moments in the text. We use an extensive grid search and limit ourselves to integer values for the parameters. The best fit has 19.70% of months in recession, an average length of 12.0 months, a minimum of 6, 25^{th} percentile of 8, median of 11, 75^{th} percentile of 14, and maximum of 43 months. The corresponding parameters are $K = 7$ months, $p_1 = 17$, $p_2 = 37$, and $p_3 = 29$.

To describe how the variables of interest behave over the course of a recession, it is convenient to divide each recession into three equal stages, and to keep track of the value in the last month of each stage. More precisely, we express the variable in percentage difference from the peak, which is the month before the recession starts. For example, if a recession lasts 9 (10) months, we calculate how much lower dividends are in months 3, 6, and 9 (10) of the recession, in percentage terms relative to peak. Averaging these numbers over recessions indicates the typical change of the variable of interest in three stages of a recession. The third-stage number summarizes the behavior of the variable over the entire course of the recession. We apply this procedure equally to the data and the model simulation.

We set $\rho_s = .9355$ to exactly match the 12-month autocorrelation of the CP factor of .435. This low annual autocorrelation is consistent with the interpretation of s as a business-cycle frequency variable. We set $\sigma_s = 1$; this is an innocuous normalization. The s process is negative at the start of the recession (1.6 standard deviations below the mean), falls considerably in the first stage of a recession (to 3.2 standard deviations below the mean), continues to fall in the second stage (to -3.9 standard deviations), and partially recovers in the last stage (to -2.9 standard deviations).

Dividend and Inflation Parameters We calibrate parameters to match moments of real dividend growth on the market portfolio, value portfolio (fifth book-to-market quintile), and growth portfolio (first quintile) for 1927-2009 (997 months). Since nominal bond yields are unavailable before 1952, we compare our model's output for nominal bond yields and associated returns to the average for 1952-2009. In our model simulation, we reinvest monthly dividends at the risk-free rate to compute an annual real dividend series, replicating the procedure in the data. We calculate annual inflation as the twelve-month sum of log monthly inflation, as in the data.

The most important parameter is γ_{1i}, which measures how sensitive dividend growth is to changes in real economic activity. In light of the empirical evidence presented in Section II.A of the main paper, we choose γ_{1i} to match the log change in annual real dividends between the peak of the cycle and the last month of the recession. In the data, the corresponding change is -21.0% for value stocks (the fifth BM portfolio), + 2.2% for growth stocks (first BM portfolio), and -5.2% for the market portfolio (CRSP value-weighted portfolio). Given the parameters governing the s dynamics and the recession determination described above, the model matches these changes exactly for $\gamma_{1G} = -.4e-4$, $\gamma_{1V} = 97.6e-4$, and $\gamma_{1M} = 24.8e-4$. Note that $\gamma_{1V} > \gamma_{1G}$ delivers the differential fall of dividends on value and growth stocks. This is the central mechanism behind the value premium.

The rest of the dividend growth parameters are chosen to match the observed mean and volatility. We choose $\gamma_{0G} = .0010$, $\gamma_{0V} = .0044$, and $\gamma_{0M} = .0010$ to exactly match the unconditional mean annual log real dividend growth of 1.23% on growth, 5.26% on value, and 1.23% on the market portfolio. We choose $\sigma_{dM} = 2.09\%$ to exactly match the unconditional volatility of annual log real dividend growth of 10.48%. We set $\sigma_{dG} = 1.94\%$ and $\sigma_{dV} = 2.23\%$ in order to match the fact that the covariance of growth stocks with market return innovations is slightly higher than that of value stocks. However, the difference needs to be small to prevent the value premium from being due to differential exposure to market return shocks. To be precise, this difference makes the contribution of the market factor to the value premium equal to 0.44% per year, the same as in the data. We set the idiosyncratic volatility parameter for growth $\sigma_G = 3.48\%$ to match exactly the 13.75% volatility of dividend growth on growth stocks, given the other parameters. We set $\sigma_V = 10.94\%$ to match the volatility of dividend growth on value stocks of 48.93%. The 12-month autocorrelation of annual log real dividend growth resulting from these parameter choices is -.01 for G, .21 for V, and .29 for M in the model, close to the observed values of .11, .16, and .29, respectively.

Inflation parameters are chosen to match mean inflation, and the volatility and persistence of nominal bond yields. We choose $\bar{\pi} = .0026$ to exactly match average annual inflation of 3.06%. We choose $\rho_x = .989$ and $\sigma_x = .03894\%$ to match the unconditional volatility and 12-month autocorrelation of nominal bond yields of maturities 1- through 5-years (1952-2009 Fama-Bliss data). In the model, volatilities decline from 3.13% for 1-year to 2.58% for 5-year bonds. In the data, volatilities decline from 2.93% to 2.72%. The 12-month autocorrelations of nominal yields range from .88 to .84 in the model, and from .84 to .90 in the data. Our parameters match the averages of the autocorrelations and volatilities across these maturities. We choose the volatility of unexpected inflation $\sigma_\pi = .7044\%$ to match the volatility of inflation of 4.08% in the data. The 12-month autocorrelation of annual inflation is implied by these parameter choices and is .59 in the model, close to the .61 in the data. We set the real short rate $y = .0018$, or 2.1% per year, to match the mean 1-year nominal bond yield of 5.37% exactly, given all other parameters.

Market Prices of Risk We set $\lambda_0(1) = .2913$ to match the unconditional equity risk premium on the market portfolio of 7.28% per year (in the 1927-2009 data). The market price of expected inflation risk $\lambda_0(1) = -.0986$ is set to match the 5-1-year slope of the nominal yield curve of 0.60%. The term structure behaves nicely at longer horizons with 10-year yields equal to 6.27% per year, and 30-year yields equal to 6.49% per year. The average of the annual bond risk premium on 2-year, 3-year, 4-year, and 5-year bond returns, which is the left-hand side variable of the CP regression, is 0.75% in the model compared to 0.87% in the data. The mean CP factor is .0075 in model and .0075 in the data. We set the market price of cyclical risk $\lambda_0(3) = .0249$ in order to match the 5.22% annual value premium (in the 1927-2009 data).

We set $\lambda_1(1) = .1208$ in order to generate a slightly negative $B_M = -0.000624$. As argued above, the near-zero B_M prevents the value premium from arising from exposure to market return shocks, and it prevents bond returns from being heavily exposed to market risk. The slight negative sign delivers a slightly positive contribution of exposure to market return shocks to bond excess returns, as in the data. In particular, it generates a 15 basis point spread between ten-year and 1-year bond risk premia coming from market exposure, close to the 30 basis points in the post-1952 data. Finally, we set $\lambda_1(2) = -0.0702$ in order to exactly match the volatility of the CP factor of 1.55%. The volatility of the average annual bond risk premium on 2-year, 3-year, 4-year, and 5-year bonds is 3.93% in the model and 3.72% in the data. As mentioned above, ρ_s is chosen to match the persistence of CP. Thus the model replicates the mean, volatility, and persistence of the CP factor and the nominal bond risk premium. The maximum annualized log Sharpe ratio implied by the model, $E[\sqrt{\Lambda_t'\Lambda_t}]\sqrt{(12)}$ is 1.44.

Risk Premium Decomposition The main result from the calibration exercise is that we are able to replicate the three-factor risk premium decomposition we uncovered in Section 3. Figure 6 is the model's counterpart to Figure 5 in the data. It shows a good quantitative match for the relative contribution of each of the three sources of risk to the risk premia for growth, value, and market equity portfolios, as well as for maturity-sorted government bond portfolios. This fit is not a forgone conclusion, but results from the richness of the model and the choice of parameters. For example, differential exposure to the market factor could have well been the source of the value risk premium in the model given that the market shocks are linear combinations of permanent dividend growth and transitory cyclical shocks. Or, bonds of different maturity could have differential exposure to the market factor shocks. The data show no heterogeneity in both types of exposures. The model has just enough richness to replicate these patterns.

We conclude that the model delivers a structural interpretation for the MKT, LVL, and CP shocks. CP shocks reflect (transitory) cyclical shocks to the real economy, which naturally carry a positive price of risk. The LVL shock captures an expected inflation shock, and the MKT shock mostly captures a (permanent) dividend growth shock. The model quantitatively replicates the unconditional risk premium on growth, value, and market equity portfolios, and bond portfolios of various maturities, as well as the decomposition of these risk premia in terms of their MKT, LVL, and CP shock exposures. Furthermore, it matches some simple features of nominal term structure of interest rates and bond risk premia. It does so for plausibly calibrated dividend growth and inflation processes.

F.4. Asset Pricing Dynamics over the Cycle

Finally, our model implies interesting asset pricing dynamics over the cycle. The CP factor, or nominal bond risk premium, starts out negative at the start of the recession, falls substantially in the first stage of the recession, falls slightly more in the second stage, before increasing substantially in the third stage of the recession. This pattern for bond risk premia is reflected in realized bond returns. In particular, the negative risk premium shocks at the start of a recession increase bond prices and returns, and more so on long-term than short-term bonds. An investment of \$100 made at the peak in a portfolio that goes long the 30-year and short the 3-month nominal bond gains \$8.0 in the first stage of the recession. The gain further increases to \$11.7 in the second stage, before falling back to a \$7.4 gain by the last month of the recession. The latter increase occurs as consequence of the rising bond risk premium. Taken over the entire recession, long bonds gain in value so that they are recession hedges Campbell, Sunderam, and Viceira (2013). The same is true in the data, where the gain on long-short bond position is \$6.1 by the last month of the recession. Value stocks are risky in the model. Their price-dividend ratio falls by 21% in the first stage compared to peak, continues to fall to -34%, before recovering to -29% by the end of the recession. In the data, the pd ratio on value stocks similarly falls by 16% in the first stage, falls further to -26%, before recovering to +4%. Value stocks perform poorly, losing more during the recession than growth stocks, both in the model and in the data.

One important feature the model (deliberately) abstracts from are discount rate shocks to the stock market. As a result, the price-dividend ratio and stock return are insufficiently volatile and reflect mostly cash-flow risk. While obviously counter-factual, this assumption is made to keep the exposition focussed on the main, new channel: time variation in the bond risk premium, the exposure to cyclical risk, and its relationship to the value risk premium. One could write down a richer model to address this issues, but only at the cost of making the model more complicated. Such a model would feature a market price of aggregate dividend risk which varies with some state variable z. The latter would follow an AR(1) process with high persistence, as in Lettau and Wachter (2007). All price-dividend ratios and expected stock returns would become more volatile and more persistent, generating a difference between the business-cycle frequency behavior of the bond risk premium and the generational-frequency behavior of the pd ratio. This state variable could differentially affect value and growth stocks, potentially lead to a stronger increase in the pd ratio of value than that of growth in the last stage of a recession. This would shrink the cumulative return gap between value and growth stocks during recessions, which the model now overstates.