Essays in Expectation Formation and Asset Pricing

by

Zhengyang Xu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2020

Doctoral Committee:

Professor Stefan Nagel, Co–Chair
Professor Tyler Shumway, Co–Chair
Professor Robert Dittmar
Assistant Professor Lindsey Gallo
Assistant Professor Pablo Ottonello

Zhengyang Xu

zhengyxu@umich.edu

ORCID iD: 0000-0003-2826-2881

In memory of my father, Xinliang Xu.

ACKNOWLEDGEMENTS

First and foremost, I would like to express my deepest gratitude to my committee members. My co-chair, Stefan Nagel, has inspired my research interests in asset pricing and investors' expectation, which lead to this dissertation. I am extremely grateful for his continuous guidance and encouragement throughout my doctoral studies even after he moved to Chicago. My other co-chair, Tyler Shumway, has dedicated much time and effort to sharpening my dissertation and provided numerous important insights. Robert Dittmar has advised me on many of my research ideas and helped me tremendously in the process of trial and error. My sincere thanks also goes to Lindsey Gallo and Pablo Ottonello, who have provided many constructive and invaluable suggestions.

I would also like to take this opportunity to thank my friends, fellow PhD students, and faculty members both within and outside the University of Michigan, who have provided many insightful suggestions and comments on my dissertation. I have greatly benefited from discussions with Vadim Elenev, Samuel Hartzmark, Steve Hou, Serhiy Kozak, George Malikov, Indrajit Mitra, Taeuk Seo, Stijn Van Nieuwerburgh, Yifei Wang, and Zach Yan.

Last but not least, I would like to thank my mother, Hua Fang, my girlfriend, Yiqing Zhu, and Yiqing's parents, Peiying Ji and Ye Zhu. They have been there for me during the most difficult times of this long journey. Their love, understanding, and belief in me have been the major emotional support of my academic pursuit so far. I dedicate my PhD to them.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

Figure

LIST OF APPENDICES

Appendix

ABSTRACT

This dissertation examines the role of investors' belief formation in asset valuation.

In the first chapter, I document that subjective bond risk premia implied by survey forecasts of future Treasury yields are acyclical at the one-year horizon. This is in stark contrast to large countercyclical variation in objective risk premia fitted from in-sample predictive regressions of future bond excess returns. This difference in risk premia implies a wedge between subjective and objective expectations of future short rates, which I show is predictable by trend and cycle components of macroeconomic forecasts. I show that these empirical findings can be explained with a learning model in which the agent filters latent trend and cycle components of fundamentals in real time, while an econometrician analyzing the data ex-post has full knowledge of the data-generating processes. The model also yields predictions, consistent with the data, on the joint behavior of the unconditional yield curve slope, the cyclicality of short-rate and macroeconomic expectation wedges, and the cyclicality of objective risk premia. My results suggest that equilibrium models of bond risk premia should target acyclical subjective risk premia and expectation formation, rather than ex-post in-sample fitted risk premia from predictive regressions.

The second chapter (co-authored with Stefan Nagel) builds on recent evidence that lifetime experience shapes individuals' macroeconomic expectations and it explores the asset-pricing implications of this evidence for the aggregate US stock market. We study an economy in which a representative agent learns—but with fading memory—about

the constant underlying endowment growth rate. The agent downweighs observations in the distant past but is otherwise Bayesian in evaluating uncertainty. The model explains both standard asset pricing facts and investor expectations within a simple and tractable framework, in which subjective belief dynamics are constrained by survey data. In the model, fading memory implies perpetual learning and permanently high subjective uncertainty about long-run growth, but the subjective equity premium is virtually constant. In contrast, an econometrician who knows the true long-run growth will find a high and strongly countercyclical objective equity premium which is predictable. Consistent with this theory, we show empirically that experienced payout growth (an exponentially weighted average of past growth rates) is negatively related to future stock market excess returns, predicts survey expectation errors, and is positively related to aggregate analyst forecasts of long-run earnings growth.

In the third chapter, I argue that econometricians find high returns from trading on the profitability anomaly because investors in real time failed to spot profitable firms that are difficult to analyze. I document that the Fama-French three-factor alphas of the profitability anomaly only exist among firms with high information frictions, proxied by young age, high forecast dispersion, high past return volatility, and/or high option-implied volatility. The results are robust to excluding micro-firms and using different measures of profitability. Short-sale constraints, liquidity, and financial distress do not fully account for the alphas. I show that the empirical pattern is consistent with a noisy rational expectations equilibrium model in which investors use profitability as a noisy signal to learn about future firm payoffs.

CHAPTER I

Expectation Formation in the Treasury Bond Market

1.1 Introduction

The nature of predictable variation in excess returns on US Treasury bonds remains to be understood.[1] Previous studies exploring predictive regressions of future bond excess returns generally conclude that bond risk premia (fitted future excess returns) are time-varying and countercyclical (Cochrane and Piazzesi 2005; Cieslak and Povala 2015). Leading rational expectations models interpret this evidence as investors requiring a time-varying risk compensation and provide explanations of these cyclical variations by introducing either time-varying prices of risk (Wachter 2006) or quantities of risk (Bansal and Shaliastovich 2013). A key assumption of this interpretation is that the in-sample regressions fitted by an econometrician with hindsight accurately capture the expectations investors have when they price bonds in real time. A number of studies have challenged this assumption by documenting large persistent yield survey forecast errors made by professional forecasters (Froot 1989; Piazzesi et al. 2015; Cieslak 2018), which indicate very different pricing dynamics perceived in real time.

[1] For example, the one-year excess return on a 10-year Treasury bond is earned from borrowing at the one-year rate, buying a 10-year Treasury bond, and selling it after one year.

In this paper, I show that allowing for a belief wedge in macroeconomic expectations between real-time investors and econometricians has the potential to reconcile survey and predictive regression evidence. I distinguish *subjective* pricing dynamics, which are determined by the subjective macroeconomic expectations of investors, from *objective* pricing dynamics, which are measured by econometricians with access to future realizations of data. Empirically, I show that the macroeconomic expectation wedge manifests itself in predictable pricing errors from the econometrician's point of view. Theoretically, I study a learning model that allows for this expectation wedge and show that its asset-pricing implications are consistent with the empirical findings.

I start the analysis by studying the behavior of subjective risk premia implied by real-time yield forecasts. An immediate observation is that subjective risk premia are less volatile and cyclical than their objective counterparts (Piazzesi et al. 2015), even more so after adjusting for potential measurement errors. To study the possible determinants of movements in subjective risk premia, I regress them on a host of macro and yield variables related to trend, cycle, volatility, and uncertainty. Notably, the long-term trend in inflation levels captures most of the variation, while the evidence is mixed for other categories of variables. This is in stark contrast to objective risk premia, for which trend variables capture substantially less variation than cycle variables.

To see how much cyclical variation we should explain in subjective risk premia, I test whether yield forecasts reject the Expectation Hypothesis (EH) under the subjective measure, à la Froot (1989). The subjective EH states that under the null of constant subjective risk premia, forecasts of future long yields should (approximately) move one-for-one with current long yields. For the period 1987 to 2018, the subjective EH cannot be rejected for long yields with maturities of five years or above at any forecast horizon and the regression coefficients are very close to the subjective null. In addition, at the one-year horizon—the typical holding period studied in the literature—the subjective

EH cannot be rejected at the 10%-significance level for all maturities above one year. Thus, subjective risk premia at the one-year horizon are acyclical. To the extent that survey forecasts do not systematically deviate from market expectations, this suggests that models of bond risk premia should not target cyclical variation.

The question now arises as to why risk premia behave differently under subjective and objective expectations. I use the insight that long-maturity yields can be decomposed into expectations about future short rates and future excess returns (Campbell and Ammer 1993). Since real-time investors and the econometrician observe the same current long-maturity yields, but have very different expectations of future excess returns, they must have very different expectations of future short rates. Using long-term survey forecasts, I show that subjective expectations of long-term short rates are tightly linked to subjective expectations of long-term real GDP growth and inflation, which are mainly characterized by low-frequency trend movements. In contrast, realized future short rates over long horizons are much less affected by current expectations of long-term real GDP growth, but are driven by a multi-decade-long decline in long-run inflation levels that is not anticipated by real-time investors. The trend components in macroeconomic expectations also show up in short-term short-rate survey forecasts and have predictive power for short-rate forecast errors.

These empirical findings emphasize the importance of recognizing the wedge between subjective and objective expectations, which is assumed away in rational expectations models. This leads me to construct a learning model that highlights this expectation wedge and to study its asset-pricing implications. In the economy, exogenous real growth consists of a constant long-run trend and an ARMA(1, 1) cycle component that is negatively auto-correlated at all lags (Clark 1987). I assume that the econometrician has full knowledge of the data-generating processes. A representative agent, however, cannot directly observe the trend and cycle components in aggregate growth and has to rely on standard Kalman filtering to form her beliefs about future dynamics. In addition, to capture the time-varying long-term forecasts observed in the data, I assume that the agent subjectively perceives a stochastic trend that follows an AR(1) process. This difference in perceived trend process between the agent and the econometrician leads to a perpetual belief wedge. In the steady state, subjective uncertainties of future real growth do not change over time, which implies constant subjective risk premia. Thus, variations in bond prices are all driven by time-varying subjective growth expectations. With full knowledge, the econometrician anticipates future revisions in subjective expectations and finds bond price variation to be predictable.[2] Thus, the setting of the model is consistent with the evidence from survey data and predictive regressions.

The auto-correlation structures of subjective trend and cycle components play important roles in model dynamics. For a concrete example, consider a bad current realization of aggregate real growth. In this unobserved components model, the agent only perceives a single bad aggregate belief shock and revises both her current trend and cycle expectations downwards. Since trend growth is positively auto-correlated, this makes the agent pessimistic about all future trend growth rates. The intertemporal smoothing motive thus makes long-maturity bonds more attractive. In addition, since the agent perceives future bad news from trend expectation shocks to always happen in a high marginal utility state, it implies that risk accumulates over time. This leads to a stronger precautionary saving motive at longer horizons, which makes long-maturity bonds even more attractive. This trend effect is stronger if trend expectations are more sensitive to the initial bad aggregate news or if the perceived trend process is more persistent. The same logic holds true for negatively auto-correlated cycle expectations in the opposite direction, as the initial bad news makes the agent optimistic about future

[2] A similar setting is applied to an exogenous inflation process whose shocks are uncorrelated with real growth shocks. For parsimony, I do not discuss it here.

cycle growth while perceived risk dissipates over time. This cycle effect makes long-maturity bonds less attractive and is stronger if cycle expectations are more sensitive to the initial bad aggregate news or if the perceived cycle process is more persistent. The relative strength of the trend and cycle effects determines both the agent's level of optimism about future aggregate growth and the amount of risk perceived in the long-term.

This economic mechanism yields novel predictions on the joint behavior of the unconditional yield curve slope, the cyclicality of short-rate and macroeconomic expectation wedges, and the cyclicality of objective risk premia under the econometrician's measure. Specifically, if the cycle effect outweighs the trend effect, the agent perceives less risk and has a lower hedging demand at longer horizons. This depresses prices of long-maturity bonds and raises their yields relative to short-maturity bonds. Therefore, the unconditional term structure of yields is upward-sloping.[3] In addition, when the agent is overly pessimistic relative to the objective forecast, she overestimates the amount of mean-reversion and becomes overly optimistic about future growth. The intertemporal smoothing motive pushes up her expectations for future short rates and depresses current prices of long-maturity bonds. The econometrician anticipates that the agent will be disappointed next period, indicating a lower short rate and higher prices of long-maturity bonds going forward. Thus, the short-rate expectation wedge (objective minus subjective) is procyclical, the objective risk premium is countercyclical, and both are predictable by macroeconomic expectation wedges.

The relative strength of the trend and cycle effects depends on model parameters of subjective dynamics, which I identify using short- and long-term macroeconomic survey forecasts. Identified parameter values suggest that the agent's cycle expectations are

[3]Under the econometrician's measure, the agent *on average* makes the correct forecasts of future aggregate growth. Thus the intertemporal smoothing motive plays no role in determining the *unconditional* yield curve slope.

much more sensitive than trend expectations to aggregate belief shocks , while having a lower persistence. Overall, the parameters support a stronger cycle effect. Thus, the model generates an upward-sloping yield curve, procyclical short-rate expectation wedges, and countercyclical objective risk premia, all of which are consistent with empirical evidence from survey and yield data. The model has several additional predictions—on the determinants of yield movements, the correlation between short-rate and macroeconomic expectation wedges, and the signs of macro trend expectations in bond return predictive regressions—for which I find support in the data.

This paper is mainly connected to three strands of the literature. On the empirical side, it builds on previous work that studies investors' real-time expectations using survey data. Piazzesi et al. (2015) document that subjective risk premia are less volatile and cyclical than their objective counterparts. I complement their study by showing that subjective risk premia are acyclical and that their movements are not captured by the various volatility and uncertainty measures examined. In addition, I study how short-rate expectations and expectation errors are linked to macroeconomic expectations and how these findings can be reconciled in an equilibrium model. Cieslak (2018) documents persistent forecast errors in federal fund rates and shows that these forecast errors induce ex-post return predictability, especially for short-maturity bonds. By studying long-term survey forecast data, I further document the role of macro trend expectations in short-rate expectations under the subjective measure. I also explain why short-rate forecast errors (objective minus subjective) are procyclical. Froot (1989) shows that the subjective EH cannot be rejected for long-maturity yield survey forecasts from 1969 to 1986. I provide *true* out-of-sample evidence that the subjective EH still holds for the non-overlapping period from 1987 to 2018, indicating that absence of cyclicality is a robust feature of subjective risk premia. Wang (2019) documents that short rate forecasts underreact to new information while long yield forecasts overreact and provides an explanation based on forecasters' misestimation of auto-correlations.

Several papers also study equilibrium models of expectation formation and their implications for yield movements. Piazzesi and Schneider (2006) consider a learning model in an affine state-space setting, but do not investigate trend expectations. Compared to their framework, my model yields tractable closed-form solutions that can be used to jointly study risk premia and expectation wedges. Hasler et al. (2019) study the unconditional term structure of interest rates by examining a model of real endowment growth in which the agent learns about the latent stochastic trend and cycle components in a continuous-time Bayesian setting. An important distinction here is that trend growth rates under the data-generating process are actually constant. Thus, there is a belief wedge between the econometrician and the representative agent, even though the learning is optimal under the subjective measure. This belief wedge is crucial to reconciling both survey evidence and predictive regression results. Zhao (2019) uses a constant-gain learning model to link trend and cycle movements in Treasury yields to learning from stable and transitory components in macroeconomic variables, but he studies neither bond risk premia nor forecast errors. In addition, the unconditional term structure of interest rates in his model is generated from ambiguity aversion, which is a different mechanism. More broadly, my paper fits into the growing literature of general equilibrium asset-pricing models that deviate from the rational expectations framework (Timmermann 1993; Lewellen and Shanken 2002; Collin-Dufresne et al. 2016b; Nagel and Xu 2019).

My model also provides a potential explanation for the role of macro trends in bond return predictive regressions. A number of studies (Cieslak and Povala 2015; Jørgensen 2018; Bauer and Rudebusch 2019) have documented that including long-term trends in inflation and real GDP growth strongly increases the predictive power of future bond excess returns. My model generates this predictability by linking macro trends to expectation wedges. It also makes testable predictions for the signs of macro trends in these regressions.

The rest of the paper proceeds as follows. Section 1.2 describes the data and its treatment. Section 1.3 documents the distinct behavior of subjective and objective risk premia. Section 1.4 links this difference to short-rate forecast errors, which are predictable by macroeconomic expectations. Section 1.5 builds a model that highlights the role of macro expectation wedges in reconciling the empirical findings. Section 1.6 identifies model parameters and tests additional predictions. Section 1.7 concludes.

1.2 Data

1.2.1 Survey forecasts

For short-term interest rate expectations, I use yield forecasts from the Blue Chip Financial Forecasts (BCFF) survey.[4] This survey publishes monthly forecasts of about 45 professional forecasters from leading financial institutions. I use forecasts for 3-month, 6-month, and 1-year Treasury bills (tbill), as well as 2-year, 5-year, and 10-year Treasury notes. These forecasts have been consistently published since January 1988, with forecast horizons up to four quarters ahead.[5] I use the consensus forecasts from BCFF, defined as the average of individual forecasts.

The publisher of BCFF also conducts another survey, Blue Chip Economic Indicators (BCEI), which targets a similar group of professional forecasters. BCEI primarily focuses on macroeconomic variables, including real GDP growth and inflation, and has published forecasts of 3-month Treasury bill rates since October 1980.[6] Forecasts

[4]Other studies that use yield forecast data from BCFF include Kim and Orphanides (2012), Piazzesi et al. (2015), Cieslak (2018), Giacoletti et al. (2018), and Buraschi et al. (2019).

[5]Forecasts for 3-month tbill rates have been published since November 1982.

[6]Before January 1982, BCEI used the 6-month commercial paper rate as a proxy for the short-

up to four quarters ahead have been consistently available at monthly frequency since October 1980.

For long-term forecasts, I use the semi-annual survey results from both BCFF and BCEI starting with October 1983, which ask respondents to forecast the annual averages of the following five calendar years, as well as the 6- to 10-years-ahead (or 7- to 11-years-ahead) five-year averages (long-term average) on selected yields and macroeconomic variables. Before December 1996, both surveys were conducted in March and October. BCFF then switched the survey months to June and December. Thus, by combining data from BCFF and BCEI, I effectively have quarterly long-term forecast observations after 1996. For the period before 1996, I interpolate the bi-annual data to get quarterly observations.

In some of the following exercises, monthly zero-coupon yield forecasts with constant forecast horizons are needed. However, two complications arise due to the features of the BCFF survey. First, forecasts are made for average realizations over a calendar quarter instead of a specific month. This introduces a time-varying forecast horizon inside the same survey quarter. For example, the one-quarter-ahead forecasts published in January and February are both made for average realizations over April, May, and June. Second, BCFF only provides par yield forecasts for Treasury securities with maturities above one year. To deal with these issues, I interpolate available yield forecasts along both maturities and forecast horizons and bootstrap the zero-coupon forecasts. More details are given in Appendix A.4. Table 1.2, Panel A, shows that bootstrapped yield forecasts have first and second moments similar to those of observed yields during the same period.

Short-term forecasts of real GDP growth and inflation are supplemented with mean forecasts from the Survey of Professional Forecasters (SPF) to extend the sample period

term interest rate. I show in the online supplementary appendix that 3-month Treasury bill rates and 6-month commercial paper rates during that period are very similar.

back to 1968. SPF also provides quarterly 3-month tbill forecasts beginning in 1981.

Finally, I supplement short-term forecasts on 3-month tbill rates using data from the Goldsmith-Nagan (GN) survey.[7] This is a quarterly survey of a selected panel of approximately 50 market professionals who have subscribed to the *Goldsmith-Nagan Bond and Money Market Letter*. The survey was conducted, from 1969 to 1986, late in the last month of each quarter for forecasts of the last business day of the following two quarters. The online supplementary appendix shows that at the one-quarter horizon, forecasts are very similar across GN, BCFF, and SPF during the overlapping period.

Table 1.1 summarizes the frequency, survey timing, and forecast dates across surveys. Appendix A.1 provides more details on each survey.

1.2.2 Interest rate data

For monthly zero-coupon nominal Treasury yields, I use data from Liu and Wu (2019) with maturities ranging from one month to 30 years at one-month intervals.[8] If daily zero-coupon nominal Treasury yields are needed, I use data from Gürkaynak et al. (2007), available from the Federal Reserve Board. The Liu and Wu (2019) data has the advantage that it simultaneously captures the very short and very long ends of the yield curve. In addition, the maturity availability at one-month intervals is crucial for a bootstrap method that will be introduced later. Other data on observed yields and bond portfolio returns are from standard sources, described in Appendix A.1.

1.2.3 Macroeconomic data

To avoid look-ahead bias introduced by revisions, I use vintage series from the real-time dataset for macroeconomics available from the Philadelphia Fed to construct year-on-year real GDP growth, employment growth, and inflation. I assume that forecasters

[7]I thank Kenneth Froot for kindly sharing the data with me.
[8]I thank the authors for making the data available on their website.

in surveys published in the middle month of a calendar quarter have access to the first estimate of real GDP for the previous quarter, which is released by the Bureau of Economic Analysis.[9] For more details, see Appendix A.1.

1.3 Behavior of Subjective Risk Premia

1.3.1 A comparison of objective and subjective risk premia

I denote the continuously compounded nominal yield on an n-maturity zero-coupon bond as y_t^n. The one-period short rate is denoted with $i_t \equiv y_t^1$. The realized m-period excess return from holding an n-maturity zero-coupon bond from t to $t+m$ is

$$rx_{t,t+m}^n = ny_t^n - (n-m)y_{t+m}^{n-m} - my_t^m. \tag{1.1}$$

The m-period subjective risk premia on an n-maturity zero-coupon bond can be calculated by taking the subjective expectations of both sides of Equation (1.1):

$$\tilde{\mathbb{E}}_t rx_{t,t+m}^n = ny_t^n - (n-m)\tilde{\mathbb{E}}_t y_{t+m}^{n-m} - my_t^m, \tag{1.2}$$

where $\tilde{\mathbb{E}}_t[\cdot]$ denotes the expectation of a real-time agent. I use bootstrapped m-period-ahead forecasts on n-maturity zero-coupon yields as the measure of $\tilde{\mathbb{E}}_t y_{t+m}^n$. The sample period is from December 1987 to December 2018. In terms of time t yield information, I assume that forecasters observe realizations of end-of-the-third-week yields in survey months, based on the survey timing described in Table 1.1.[10] However, becasue the zero-coupon yields from Liu and Wu (2019) are only available at monthly frequency, I match the survey forecast *published* in month t to end-of-month yields in month $t-1$.

[9] This assumption is consistent with the documentation of SPF.

[10] BCFF also publishes this information along with yield forecasts. For example, in the February 2015 publication, the history of yield average is reported until the week ending January 23.

Using daily data from Gürkaynak et al. (2007) and matching forecasts to the third-week values of yields in month $t-1$ do not materially change the conclusions of the following analysis.

The objective risk premia are fitted from predictive regressions with future excess returns as dependent variables using data from November 1971 to December 2018, for which zero-coupon yields for maturities up to 15 years are available. Using $\mathbb{E}_t$ to denote the expectations of an econometrician, objective risk premia are

$$\mathbb{E}_t rx^n_{t,t+m} = ny^n_t - (n-m)\mathbb{E}_t y^{n-m}_{t+m} - my^m_t. \tag{1.3}$$

For predictors, I use the cycle factor from Cieslak and Povala (2015). To improve the fitting performance for short-maturity bonds, I also include the macro factor from Ludvigson and Ng (2009), which has been shown to have the strongest predictive power for two-year bonds. Appendix A.5 documents the construction and performance of these factors.

The first two columns of Table 1.2, Panel B, provide summary statistics of one-year-holding-period risk premia. Across maturities, subjective risk premia are much smaller in magnitude. For example, subjective risk premia on 10-year zero-coupon bonds are 1.58% (annualized) on average , while objective risk premia during the same period are 5.31% (annualized). Subjective risk premia are also less volatile, with standard deviations that are roughly one-half those of their objective counterparts. These results are consistent with Piazzesi et al. (2015). Figure 1.1 compares the maturity-weighted average of subjective and objective risk premia for maturities from 2 to 10 years, defined as $\tilde{\mathbb{E}}_t \overline{rx}_{t+1} \equiv \frac{1}{9}\sum_{n=2}^{10} \tilde{\mathbb{E}}_t rx^n_{t+1}/n$ and $\mathbb{E}_t \overline{rx}_{t+1} \equiv \frac{1}{9}\sum_{n=2}^{10} \mathbb{E}_t rx^n_{t+1}/n$, respectively. It is visible that average subjective risk premia are less cyclical, especially during recessions.

1.3.1.1 Measurement errors in subjective risk premia

Since subjective risk premia are calculated from small spreads between current long yields and long-yield survey forecasts, which are then scaled by maturity, small measurement errors in survey forecasts can induce large variation in subjective risk premia. The measurement errors can come from multiple sources. First, treating consensus survey forecasts as the market expectations of a single agent can introduce aggregation errors.[11] Second, due to the limited availability of maturities and lack of knowledge of true forecast horizons, the interpolated zero-coupon yield survey forecasts are contaminated by interpolation errors. Third, forecasters in the same survey month are unlikely to make forecasts on the same day, making the conditioning yield information vary across individual forecasts. From matching yield forecasts to a single current yield, some errors are introduced in implied risk premia.

To alleviate this concern, I first follow Froot (1989) to compute a "cleaned" measure of subjective risk premia by regressing them on the current 1- and 10-year yields (including a constant) and using the fitted values. Column (3) in Table 1.2, Panel B, shows that this approach captures around 60% of the variation in subjective risk premia across maturities. For a fair comparison, I run a similar regression for the cycle factor and use the fitted values. The top panel in Figure 1.2 shows that the business-cycle variation in the cycle factor is still preserved, while the cyclical variation in subjective risk premia is largely eliminated and the "cleaned" measure exhibits a strong downward trend.

A related concern is raised in Cieslak and Povala (2015) who conclude, based on a strong positive correlation between average subjective risk premia and the cycle factor, that subjective risk premium is "a noisier measure of the bond risk premium". In my

[11] See, for example, Rubinstein (1974) and Buraschi et al. (2019). The aggregation errors are distinct from rounding errors introduced by averaging individual forecasts, which, as I show in simulations, are very small.

sample, subjective risk premia and the cycle factor do have a positive correlation of 0.57. However, Column (4) in Table 1.2, Panel B, shows that the cycle factor only explains around 30% of variation in subjective risk premia and that fitted components are even less volatile, with standard deviations further reduced by one-half. The bottom panel in Figure 1.2 shows that the cycle factor does not capture the trend component in subjective risk premia, which is what drives most variation.

Finally, I directly compare yield forecasts under the subjective and objective expectations. The objective yield forecasts are calculated from predicted future excess returns, $\mathbb{E}_t rx^n_{t,t+m}$, using Equation (1.3)

$$\mathbb{E}_t y^n_{t+m} = \frac{(n+m)y^{n+m}_t - my^m_t - \mathbb{E}_t rx^{n+m}_{t,t+m}}{n}. \tag{1.4}$$

If subjective risk premia are simply objective risk premia plus measurement errors, we should expect $\Delta^n_{t,m} \equiv \tilde{\mathbb{E}}_t y^n_{t+m} - \mathbb{E}_t y^n_{t+m}$ not to be systematically different from zero. The solid lines in Figure 1.3 examine $\Delta^n_{t,m}$ for 2- and 10-year yields at one-year horizon. We see that there is a persistent wedge between subjective and objective yield forecasts, especially during recessions. To partially address the concern that consensus forecasts may be a biased measure of market expectations, I also compare the top- and bottom-10 average forecasts from BCFF to objective forecasts.[12] Their differences provide the boundaries of the shaded areas in Figure 1.3. We see that even for these relatively extreme forecasts, there is a sizable wedge during recessions.

[12] Note that BCFF provides only top- and bottom-10 average forecasts for par yields, which are different from zero-coupon yields. I make the assumption that the distance between top/bottom average forecasts and consensus par yield forecasts is equal to the distance between hypothetical top/bottom average zero-coupon yield forecasts and consensus zero-coupon yield forecasts.

1.3.1.2 What explains movements in subjective risk premia?

I next study the possible economic determinants of movements in subjective risk premia. I run regressions of maturity-weighted average subjective risk premia, $\tilde{\mathbb{E}}_t\overline{rx}_{t+1}$, on a host of variables from previous literature. Motivated by Piazzesi et al. (2015), who use a band pass filter to decompose risk premia into statistical "trend" and "cycle" components, I first study economic measures of trends and cycles, then the variables that relate to volatilities (Duffee 2002) and uncertainty (Bansal and Shaliastovich 2013). All explanatory variables are standardized to have a unit standard deviation in order to compare coefficients across regressions. I perform the same regressions for average objective risk premia, $\mathbb{E}_t\overline{rx}_{t+1}$. To examine a longer sample period, I combine six-month-ahead forecasts from BCFF and GN on 1-year yields. This enables me to construct the six-month-holding-period subjective risk premia on 1.5-year zero-coupon bonds since 1969 at quarterly frequency. Objective risk premia for the same maturity and holding period are constructed from quarterly predictive regressions of six-month-ahead future excess returns.

Table 1.3, Panel A, reports the regression results of trend-related variables, including yield level and trend inflation from Cieslak and Povala (2015).[13] Average subjective risk premia since 1987 are strongly positively related to the yield level and trend inflation, which both explain more than 50% of the variation. These trend variables also capture around 20% of the variation in 1.5-year bond risk premia for the longer sample period since 1969. In comparison, objective risk premia do not have a significant trend component.

The second set of variables relates to cyclical movements. I consider yield slope and year-on-year employment growth.[14] To isolate the cyclical movements, I include trend

[13]Yield level is defined as the first principle component of 1-, 2-, 5-, 7-, 10-, and 15-year zero-coupon yields. Trend inflation is constructed as the exponentially weighted average of past inflation.

[14]Yield slope is the second principle component of 1-, 2-, 5-, 7-, 10-, and 15-year zero-coupon

inflation in all regressions (Cochrane 2015). Table 1.3, Panel B, shows that subjective risk premia do display some countercyclical movement based on the statistical significance of coefficients at conventional levels. However, the explained R^2 does not increase much (only around 3%) and the magnitudes of the coefficients are much smaller than those from trend regressions. In stark contrast, the inclusion of cyclical measures in regressions for objective risk premia drastically increase the explained R^2 by 40% to 50% and the coefficients are both statistically and economically significant.

I next consider volatility-related variables. Since the mid-1980s, we have observed a reduction in macroeconomic volatility (Bernanke 2004), which may contribute to the decline in subjective risk premia. Table 1.3, Panel C, fails to find strong evidence that subjective risk premia are related to measures of interest rate volatility, real volatility, or inflation volatility.[15] The coefficients are all small in magnitude and the explained R^2 is at most 12 percent.[16]

Finally, I examine measures that relate to uncertainty or forecast dispersion, including the Economic Policy Uncertainty Index from Baker et al. (2016), macro forecast uncertainty from Bansal and Shaliastovich (2013), and macro forecast dispersion from SPF.[17] Table 1.3, Panel D, shows mixed evidence on whether subjective risk premia are correlated with forecast uncertainty or forecast dispersion. For example, forecast uncertainty on inflation captures 16% of the variation, while uncertainty on real GDP

---

yields. I use the negative of employment growth in the regression to make the independent variable countercyclical.

[15]I use two measures of interest rate volatility. The first is physical volatility, defined as the first principle component of intra-month sums of daily squared zero-coupon yield changes for maturities from 1 to 10 years. The second measure is risk-neutral volatility, defined as the 10-year Treasury VIX obtained from Choi et al. (2017) and the Chicago Board Options Exchange (CBOE). Real and inflation volatility are rolling sums of monthly squared values of the "real" factor ("F1," as in Ludvigson and Ng 2009) and core inflation over a window of 24 months, respectively.

[16]In untabulated results, coefficients on physical interest rate volatility and real volatility do become statistically significant at the 95%-confidence level once trend inflation is controlled, but the magnitudes of the coefficients are still small, with no significant improvements in R^2.

[17]For macroeconomic variables, I consider real GDP growth and inflation. Forecast dispersion is defined as the difference between the 75th and 25th percentiles of individual forecasts in the cross-section.

growth only accounts for 5 percent. When it comes to forecast dispersion, however, dispersion on real GDP growth now explains a much larger fraction of the variation than dispersion on inflation does. The explanatory power of forecast uncertainty for objective risk premia is consistent with Bansal and Shaliastovich (2013).

To summarize, the regression results suggest that subjective risk premia are mainly driven by a trend component, which closely tracks the trend in long-run inflation *levels* and is not captured by the measures of volatility or uncertainty examined here.

1.3.2 Amount of cyclical variation in subjective risk premia: The Expectation Hypothesis test

I further quantify the cyclical variations in subjective risk premia through the lens of the Expectation Hypothesis(EH), which has been widely tested in the literature (e.g., Campbell and Shiller 1991). The EH states that if expected excess returns on a long-maturity bond are constant, the expected changes in future long yields should move one-for-one with changes in current yield spread. To see this, I rearrange Equation (1.2) to obtain

$$\tilde{\mathbb{E}}_t y_{t+m}^{n-m} - y_t^m = \underbrace{\frac{n}{n-m}(y_t^n - y_t^m)}_{\text{Scaled yield spread}} - \frac{\tilde{\mathbb{E}}_t rx_{t,t+m}^n}{n-m} \equiv S_t^{n,m} + c_{n,m}. \tag{1.5}$$

I test the above form of the EH by running the following regression

$$\tilde{\mathbb{E}}_t y_{t+m}^{n-m} - y_t^m = \beta_0 + \beta_1 S_t^{n,m} + \varepsilon_t. \tag{1.6}$$

Under the null that m-period subjective risk premia on an n-maturity zero-coupon bond are constant over time, the econometrician should find $\beta_1 = 1$.[18] Note that becasue this is a contemporaneous regression, the finite-sample Stambaugh (1999) bias studied in Bekaert et al. (1997) does not apply.

Given that I have only a small sample, inference based on the asymptotic distribution of a test statistic may have an incorrect size. I address this issue using a nonparametric bootstrap approach from Crump and Gospodinov (2019) and make inference based on bootstrapped p-values. Appendix C.2 provides more details of this bootstrap approach.

Based on survey data availability, I consider $m = 3$, 6, and 12 months and $n = 6$, 12, 24, 60, and 120 months.[19] Table 1.4, Panel A, suggests that the EH is only rejected for short-maturity yields at short horizons of less than one year. For long-maturity yields with maturities of five years or above, the EH is not rejected at any horizon based on bootstrapped p-values, with coefficients very close to 1. In addition, at one-year horizon, which is the typical holding period studied in the literature, the EH is not rejected at the 10%-significance level for all maturities above one year. However, we cannot directly claim that subjective risk premia are constant, since this is not a necessary condition for $\beta_1 = 1$ in Equation (1.6). By not rejecting $\beta_1 = 1$, we can only conclude that the covariance of subjective risk premia with the yield spread is zero; that is, subjective risk premia are acyclical. To see how reasonable it is to describe subjective

[18]An equivalent form of the EH is

$$\tilde{\mathbb{E}}_t y_{t+m}^{n-m} - y_t^n = \frac{m}{n-m}(y_t^n - y_t^m) + c_{n,m},$$

which is the form typically tested in the literature (Wachter 2006; Gabaix 2012; Bansal and Shaliastovich 2013). The online supplementary appendix shows that, given a small sample, the slope coefficient from the above regression is more sensitive to measurement errors in the dependent variable than Euqation (1.6) (but size and power are the same for these two forms).

[19]For $m = 3$, I need interpolated zero-coupon yield forecasts with maturities at three-month intervals. I do not use the approximation as in Campbell and Shiller (1991), which Bekaert et al. (1997) show can lead to an upward bias even asymptotically. Rather, I again estimate a Diebold and Li (2006) model for maturities at three-month intervals.

risk premia as constants, I construct the counterfactual one-year-ahead survey forecasts implied by constant subjective risk premia

$$\tilde{\mathbb{E}}_t^* y_{t+1}^n \equiv \frac{(n+1)y_t^{n+1} - y_t^1 - \tilde{\mathbb{E}}^* rx_{t+1}^{n+1}}{n}, \tag{1.7}$$

where $\tilde{\mathbb{E}}^* rx_{t+1}^n$ is the sample average of one-year-holding-period subjective risk premia on an n-maturity bond. Figure B.1 suggests that $\tilde{\mathbb{E}}_t^* y_{t+1}^n$ provides a reasonable fit of consensus forecasts. Froot (1989) provides similar evidence that long-maturity yield survey forecasts from GN do not reject the EH during a non-overlapping sample period from 1969 to 1986. Thus, the analysis here provides *true* out-of-sample evidence that absence of cyclicality is a robust feature of subjective risk premia.

For comparison, I also test the EH under the objective measure using future realized yields as dependent variables. To correct for the finite-sample bias, I use the first-order approximation method from Bekaert et al. (1997).[20] Table 1.4, Panel B, shows that the bias-adjusted coeffcients are uniformly negative across horizons and yield maturities, indicating a strong rejection of the EH under the objective measure and large cyclical variation in objective risk premia.

1.4 Difference in Risk Premia: The Role of Expectation Errors

The natural question then is: what gives rise to the distinct behaviors of subjective and objective risk premia? I use the insight that bond yields can be decomposed into short-rate expectations and average risk premia over the life of the bond (Campbell and Ammer 1993; Duffee 2018) and link the distinction in risk premia to the difference in

[20]This method assumes an AR(1) process for the short rates (one-month zero-coupon yields). Table B.2 reports the estimation results of the short-rate process. The formulas for the first-order approximation of biases are referred to Bekaert et al. (1997).

short-rate expectations. For concreteness, consider the following yield decomposition:

$$y_t^n = \underbrace{\frac{1}{n}\widehat{\mathbb{E}}_t\left(\sum_{k=0}^{n-1} i_{t+k}\right)}_{\text{Average future short rates}} + \underbrace{\frac{1}{n}\widehat{\mathbb{E}}_t\left(\sum_{k=0}^{n-1} rx_{t+k+1}^{n-k}\right)}_{\text{Average risk premia}}, \tag{1.8}$$

where $\widehat{\mathbb{E}}_t$ denotes a generic expectation (Appendix A.2 provides detailed derivations). Performing this decomposition under both subjective and objective expectations and taking difference, we have

$$\underbrace{\tilde{\mathbb{E}}_t\left(\sum_{k=0}^{n-1} i_{t+k}\right) - \mathbb{E}_t\left(\sum_{k=0}^{n-1} i_{t+k}\right)}_{\text{Difference in short rate expectations}} = \underbrace{\mathbb{E}_t\left(\sum_{k=0}^{n-1} rx_{t+k+1}^{n-k}\right) - \tilde{\mathbb{E}}_t\left(\sum_{k=0}^{n-1} rx_{t+k+1}^{n-k}\right)}_{\text{Difference in expected future excess returns}}. \tag{1.9}$$

Thus, difference in expected future excess returns should be perfectly negatively correlated with difference in short-rate expectations. Cieslak (2018) provides such evidence using survey forecasts of the federal funds rate from BCFF. She documents that short-term forecast errors in short rates are predictable by cyclical real-activity variables and that the fitted short-rate wedge has strong predictive power for future excess returns on short-maturity bonds. Thus, Equation (1.9) provides a good description for a small n.

For a large n, however, short-rate expectation errors in the long-term also matter. In the following sections, I study long-term survey forecasts of short rates from BCFF and BCEI to shed light on how long-term short-rate expectation behaves and its possible economic determinants.

1.4.1 Trend and cycle in subjective short-rate expectations

To measure subjective short-rate expectations, I use survey forecasts of 3-month tbill rates, for which both short- and long-term forecasts are available from BCFF and

BCEI. Figure 1.4 plots the one-year-ahead and long-term average survey forecasts, with a comparison to prevailing 3-month tbill rates in the survey month. Short-term survey forecasts at the one-year horizon have strong cyclical variations that track prevailing rates. In comparison, long-term average survey forecasts mainly exhibit a strong time-varying trend that is less affected by current economic conditions. For example, in December 2008, the 3-month tbill rate effectively reached zero (0.03%), while the survey still predicted an average 4.1% rate for 7 to 11 years ahead. This trend expectation becomes increasingly important for survey forecasts beyond a one-year horizon. Table B.3, Panel A, reports regression results of survey forecasts at different horizons on long-term average survey forecasts and current-quarter survey forecasts (nowcasts), with the latter used as a proxy for cycle expectations. The magnitude of coefficients and the statistical significance of long-term average survey forecasts increase monotonically as forecast horizon increases. Thus, it is useful to understand how this trend expectation of future short rates is formed for a real-time agent.

1.4.1.1 Links to subjective macroeconomic expectations

Most consumption-based asset-pricing models link short-rate expectations to expected conditional means of real endowment growth and inflation (Wachter 2006; Bansal and Shaliastovich 2013). I follow Bansal and Shaliastovich (2013) in measuring subjective real endowment growth expectations with survey forecasts of real GDP growth.[21] Long-term average survey forecasts of real GDP growth and inflation, beginning with 1983, are available for the same set of forecasters from BCFF and BCEI. The following analyses focus on subjective expectations measured by survey forecasts.

Figure 1.5, top panel, suggests that the downward trend in long-term short-rate expectations is accompanied by a similar downward trend in long-term inflation ex-

[21] To my knowledge, consumption forecasts are available from SPF and BCEI only for personal consumption expenditures, which include durable goods.

pectations. This observation echoes the findings of Cieslak and Povala (2015), who document that long-term survey inflation expectations correlate with low-frequency movements in yield levels. Figure 1.5, bottom panel, plots the residuals of regressing long-term short-rate expectations on long-term inflation expectations.[22] The residuals still exhibit some low-frequency movements that seem to co-move with long-term real GDP growth expectations (with a correlation of 0.56).[23] Table B.3, Panel B, shows that, even for short-rate forecasts at a one-year horizon, long-term macroeconomic expectations capture a large fraction of the variation, with nowcasts capturing the remaining higher-frequency movements.

1.4.2 Predicting short-rate forecast errors using subjective macroeconomic expectations

The previous section shows that trend and cycle components in subjective macroeconomic expectations are reflected in subjective short-rate expectations. Here, I examine short-rate forecast errors to study how those components enter into short-rate dynamics under the objective expectations. I first study one-year-ahead 3-month tbill rate forecast errors. With a slight abuse of notation, the quarterly forecast errors are defined as

$$FE_t(i_{t+4}) \equiv \bar{i}_{t+4} - \tilde{\mathbb{E}}_t \bar{i}_{t+4}, \tag{1.10}$$

where $\bar{i}_{t+4}$ is the within-quarter average tbill rate realized four calendar quarters ahead of quarter t and $\tilde{\mathbb{E}}_t \bar{i}_{t+4}$ is the four-quarter-ahead quarter-average forecast taken directly

[22]The coefficient on long-term inflation expectations is 1.47. The adjusted R^2 of the regression is 0.81.

[23]Table B.3, Panel B, provides regression-based evidence that long-term survey forecasts of short rates are linked to long-term survey forecasts of real GDP growth and inflation. Information in forecasts of near-term real GDP growth and inflation does not add significant explanatory power once long-term macroeconomic expectations are controlled. The caveat here is that long-term expectations may be non-stationary, which gives rise to the spurious regression problem. I therefore rely on visualizations to show the correlated movements.

from BCFF. Survey forecasts are sampled in the middle month, when advance estimates of last-quarter real GDP growth are released from the Bureau of Economic Analysis. The sample period is from 1983Q3 to 2018Q4, for which long-term macroeconomic forecasts are available.

Instead of directly using long-term macroeconomic forecasts, I follow Orphanides and Wei (2012) and Cieslak and Povala (2015) in using proxies for long-term subjective expectations. This is based on two considerations. First, survey forecasts—especially long-term forecasts—are inevitably contaminated with measurement errors. It is possible that measurement errors in short-rate and macroeconomic forecasts from the same survey are contemporaneously correlated. Thus, in regressions with survey forecasts at both sides, the coefficients are biased in an unknown way. In this sense, the proxies are used as instrumental variables. Second, the use of proxies extends the available sample period for later analyses, to the extent that they are also informative in early periods. The proxies for long-term subjective expectations of real GDP growth, g_t^{LT}, and inflation, π_t^{LT}, are constructed as the exponentially weighted average of past data realizations:

$$g_t^{LT} = \nu_g g_t + (1-\nu_g) g_{t-1}^{LT} = \nu_g \sum_{k=0}^{\infty} (1-\nu_g)^k g_{t-k}, \tag{1.11}$$

$$\pi_t^{LT} = \nu_\pi \pi_t + (1-\nu_\pi) \pi_{t-1}^{LT} = \nu_\pi \sum_{k=0}^{\infty} (1-\nu_\pi)^k \pi_{t-k}, \tag{1.12}$$

where ν_g and ν_π are constant weights placed on the most recent observations for real GDP growth and for inflation, respectively. The above specifications map to a constant-gain learning rule, which is widely used in the literature (Orphanides and Williams 2005; Branch and Evans 2006). Appendix A.7 provides more details on estimating weights and constructing proxies. As a measure of cycle expectations, I use current-quarter forecasts (nowcasts) from SPF.

I run regressions of the form

$$FE_t(i_{t+4}) = \gamma_0 + \gamma_1 \text{LT}_t + \gamma_2 \text{Nowcast}_t + \varepsilon_t, \tag{1.13}$$

for real GDP growth and inflation. To improve the small-sample inference for predictive regressions, I also provide bootstrapped p-values, using the bootstrap method from Greenwood and Vayanos (2014). More details about that method are provided in Appendix A.8.

Table 1.5, Panel A, shows that trend and cycle components of subjective real GDP growth expectations do have predictability for short-term short-rate forecast errors. When both components are included in the regressions, they can jointly explain about 12% of the variation and their coefficients are both statistically significant. Univariate regressions using trend and cycle components individually produce similar R^2. In addition, the slope coefficient of the trend component in the univariate regression is statistically significant at the 10% level based on bootstrapped p-values. To give an idea of the economic significance, Figure 1.6 plots the predicted forecast errors from univariate regressions, with a comparison to the measure from Cieslak (2018).[24] The figure shows that the subjective real GDP growth trend captures the low-frequency movements in short-term short-rate forecast errors. However, subjective inflation expectations do not seem to have predictability for short-term short-rate forecast errors. One possible explanation is that survey inflation forecasts in the short-term are accurate compared to model-based forecasts using historical inflation and other nonsurvey information (Ang et al. 2007; Chernov and Mueller 2012).

To investigate long-term forecast errors, I use five-year-ahead annual average forecasts of 3-month tbill rates from BCFF and BCEI. The choice of a five-year horizon is to minimize the effect of time-varying forecast horizons (see Section 1.2.1). I use

[24]Cieslak (2018) uses lagged short rates and year-on-year employment growth as predictors.

20-quarter-ahead within-quarter average tbill rates for future realizations.[25] Columns (1) and (2) in Table 1.5, Panel B, show that the subjective long-term inflation expectation, π^{LT}, has strong predictability for short-rate realizations in the distant future. This predictability comes from the multi-decade-long decline in long-run inflation levels observed by the econometrician. In comparison, subjective real GDP growth expectation in the short and long terms do not have any predictability for distant-future short rates. Turning to forecast errors, Columns (4) and (5) in Table 1.5, Panel B, show that π^{LT} continues to be a strong predictor, indicating that the real-time forecaster does not anticipate the same future decline in long-run inflation levels that the econometrician does; that is, the trend inflation wedge strongly predicts long-term short-rate forecast errors. Column (3) in Table 1.5, Panel B, shows that the subjective trend expectation of real GDP growth is a marginally significant predictor of long-term short-rate forecast errors.

1.5 A Model of Macroeconomic Expectation Formation

The above empirical evidence emphasizes the importance of recognizing the wedge between subjective and objective expectations. Specifically, subjective risk premia at the one-year horizon are acyclical, with little co-movement with the volatility and uncertainty measures examined so far. At the same time, long-term subjective expectations of macroeconomic variables and short rates strongly co-move, with subjective macroeconomic expectations having predictive power for short-rate forecast errors. Since rational expectations models assume away such an expectation wedge, I construct a learning model in this section to highlight the role of the expectation wedge in reconciling the evidence from survey data and predictive regressions.

[25] I do not use the calendar year average which introduces a severe overlapping problem in predictive regressions. But results are similar when calendar year averages are used.

The model is set up in an endowment economy with a representative agent and is built on the widely adopted, unobserved components (UC0) model from Clark (1987) in order to introduce trend and cycle components as observed from macroeconomic forecasts.

1.5.1 Real endowment growth and inflation

Following Clark (1987), the data-generating process (DGP) of real endowment *level*, G_t, is described as

$$G_t = \psi_t + C_t, \tag{1.14}$$

where

$$\psi_t = \mu + \psi_{t-1} + \xi_t, \quad \xi_t \sim \mathcal{N}(0, \sigma_\xi^2), \tag{1.15}$$

$$C_t = \rho_c C_{t-1} + \zeta_t, \qquad \zeta_t \sim \mathcal{N}(0, \sigma_\zeta^2), \tag{1.16}$$

are the levels of local trend and AR(1) cycle, respectively.[26] The implied real endowment growth, g_t, under the DGP is:

$$g_t \equiv G_t - G_{t-1} = \mu + c_t + \xi_t, \tag{1.17}$$

where $c_t \equiv C_t - C_{t-1}$ is the cycle growth. I assume that the econometrician has full knowledge of the DGP and can directly observe C_t.[27]

The representative agent, however, cannot directly observe the trend and cycle components from aggregate realizations. In addition, she perceives a stochastic drift in the trend level. This is motivated by the observed time-varying long-term macroeco-

[26]The terminology comes from Clements et al. (2012).

[27]This is mainly a simplifying assumption to focus on the learning problem of the representative agent.

nomic forecasts in Figure 1.5 and is similar to the "shifting endpoints" assumption in Kozicki and Tinsley (2001). The dynamics under the representative agent's *subjective* expectations are:

$$G_t = \tilde{\psi}_t + \tilde{C}_t, \tag{1.18}$$

where

$$\tilde{\psi}_t = \tilde{\mu}_t + \tilde{\psi}_{t-1} + \tilde{\xi}_t, \qquad \tilde{\xi}_t \sim \mathcal{N}(0, \tilde{\sigma}_\xi^2), \tag{1.19}$$

$$\tilde{\mu}_t = \phi_\mu \tilde{\mu}_{t-1} + (1 - \phi_\mu)\mu + \tilde{\omega}_t, \quad \tilde{\omega}_t \sim \mathcal{N}(0, \tilde{\sigma}_\omega^2), \tag{1.20}$$

$$\tilde{C}_t = \tilde{\rho}_c \tilde{C}_{t-1} + \tilde{\zeta}_t, \qquad \tilde{\zeta}_t \sim \mathcal{N}(0, \tilde{\sigma}_\zeta^2). \tag{1.21}$$

Note that $\tilde{\rho}_c$ is not necessarily equal to ρ_c.[28] The implied real endowment growth under subjective expectations is:

$$g_t = \tilde{\mu}_t + \tilde{c}_t + \tilde{\xi}_t, \tag{1.22}$$

where $\tilde{c}_t \equiv \tilde{C}_t - \tilde{C}_{t-1}$ is the perceived cycle growth.

Two key assumptions of subjective dynamics are addressed here. First, I assume that the perceived cycle level, $\tilde{C}_t$, follows an AR(1) process. This implies that the perceived cycle growth follows an ARMA(1, 1) process with negative auto-correlations at any given lag.[29] In previous literature, it is more standard to assume an AR(2) process for $\tilde{C}_t$ (Harvey 1985), which aims to also capture the positive auto-correlations at short lags (usually one lag) of $\tilde{c}_t$. For the purpose of my model, because negative auto-correlations in perceived cycle growth are more crucial, I choose the AR(1) specification for parsimony.[30] The second assumption is that the perceived stochastic

[28] Piazzesi et al. (2015) show that in an affine state-space model, $\rho_c = \tilde{\rho}_c$ if and only if subjective forecast errors are state-independent.

[29] The cycle growth is stationary but not invertible. The auto-correlation function is $acf(k) = \tilde{\rho}_c^{k-1}\frac{\tilde{\rho}_c - 1}{2} < 0$ for any lag k.

[30] In the online supplementary appendix, I have solved a version of the model with cycle levels following an AR(2) process. The model yields qualitatively similar predictions.

drift, $\tilde{\mu}_t$, follows a stationary AR(1) process which mean-reverts around the true trend growth rate, μ. In contrast to a random walk setting in which future trend growth rates follow a martingale and can drift unboundedly, the AR(1) specification anchors the very long-run expectation of trend growth rate. Real-world examples that could justify this assumption include Fed forward guidance and *a priori* judgments of plausible ranges (for example, that post-war real GDP growth never exceeds 15%). In a Bayesian learning setting, μ can be interpreted as the mean of the informative prior, with ϕ_μ measuring its precision.[31] With a lower ϕ_μ, the prior is more informative and posterior expectations put less weight on $\tilde{\mu}_t$.

Since trend and cycle components are not directly observable, the agent applies the standard Kalman filter (Hamilton 1994) to form her beliefs about each component, based on past observations of aggregate growth. For simplicity, I assume that the agent is endowed with an infinite history of observations on g_t, $\mathcal{H}_t \equiv \{g_{t-j}\}_{j=0}^{\infty}$, and that the Kalman filter is in the steady state.[32] I show in Appendix A.3.1 that the subjective optimal forecasts of the latent variables evolve as

$$\widehat{\mu}_t = \phi_\mu \widehat{\mu}_{t-1} + (1-\phi_\mu)\mu + \nu_\mu \left(g_t - \widehat{\mu}_{t-1} - \widehat{c}_{t-1}\right), \tag{1.23}$$

$$\widehat{c}_t = \tilde{\rho}_c \widehat{c}_{t-1} - \nu_c \left(g_t - \widehat{\mu}_{t-1} - \widehat{c}_{t-1}\right), \tag{1.24}$$

where $\widehat{\mu}_t \equiv \tilde{\mathbb{E}}[\tilde{\mu}_{t+1}|\mathcal{H}_t]$ and $\widehat{c}_t \equiv \tilde{\mathbb{E}}[\tilde{c}_{t+1}|\mathcal{H}_t]$ are the perceived conditional means of future trend and cycle growth, respectively. The Kalman gain parameters, ν_μ and ν_c, are both positive.[33] This is because a lower-than-expected realization of g_t indicates a lower-than-expected current trend growth ($\tilde{\mu}_t$) or a lower-than-expected current cycle growth ($\tilde{c}_t$) or both. Given that these two components are not directly observable, the

[31] Nagel and Xu (2019) provide such an interpretation in a constant-gain learning setting.

[32] This assumption also avoids the need to specify an initial starting point from a given sample period.

[33] Note that there is a negative sign before ν_c.

agent perceives only a single aggregate bad shock and will revise both expectations downwards. Since trend growth is positively auto-correlated, this will drive down all future trend growth expectations, implying $\nu_\mu > 0$. For cycle growth, however, the negative auto-correlations will drive up future expectations of cycle growth and $\nu_c > 0$ after adjusting for its negative sign in Equation (1.24). The resulting predictive distribution of next-period real endowment growth is

$$g_{t+1}|\mathcal{H}_t \sim \mathcal{N}\left(\widehat{\mu}_t + \widehat{c}_t\ ,\ \tilde{\sigma}_g^2\right), \tag{1.25}$$

where $\tilde{\sigma}_g^2 > \tilde{\sigma}_\xi^2$.

To price nominal assets, I introduce an exogeneous inflation process with the DGP:

$$\pi_t = \tau + \kappa_t + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \tag{1.26}$$

$$K_t = \rho_\kappa K_{t-1} + \iota_t, \quad \iota_t \sim \mathcal{N}(0, \sigma_\iota^2), \tag{1.27}$$

where $\kappa_t \equiv K_t - K_{t-1}$. This specification nests the IMA(1, 1) model studied in Stock and Watson (2007) with $\rho_\kappa = 1$. The agent similarly perceives a stochastic drift in the trend growth rate,

$$\pi_t = \tilde{\tau}_t + (\tilde{K}_t - \tilde{K}_{t-1}) + \tilde{\eta}_t, \qquad \tilde{\eta}_t \sim \mathcal{N}(0, \tilde{\sigma}_\eta^2) \tag{1.28}$$

$$\tilde{\tau}_t = \phi_\tau \tilde{\tau}_{t-1} + (1 - \phi_\tau)\tau + \tilde{\varepsilon}_t, \quad \tilde{\varepsilon}_t \sim \mathcal{N}(0, \tilde{\sigma}_\varepsilon^2), \tag{1.29}$$

$$\tilde{K}_t = \tilde{\rho}_\kappa \tilde{K}_{t-1} + \tilde{\iota}_t, \qquad \tilde{\iota}_t \sim \mathcal{N}(0, \tilde{\sigma}_\iota^2). \tag{1.30}$$

With the steady-state Kalman filter, the subjective optimal forecasts, given an infinite history of observations $\mathcal{I}_t \equiv \{\pi_{t-k}\}_{k=0}^{\infty}$, evolve as

$$\widehat{\tau}_t = \phi_\tau \widehat{\tau}_{t-1} + (1-\phi_\tau)\tau + \nu_\tau \left(\pi_t - \widehat{\tau}_{t-1} - \widehat{\kappa}_{t-1}\right), \tag{1.31}$$

$$\widehat{\kappa}_t = \tilde{\rho}_\kappa \widehat{\kappa}_{t-1} - \nu_\kappa \left(\pi_t - \widehat{\tau}_{t-1} - \widehat{\kappa}_{t-1}\right), \tag{1.32}$$

where $\widehat{\tau}_t \equiv \tilde{\mathbb{E}}[\tau_{t+1}|\mathcal{I}_t]$ and $\widehat{\kappa}_t \equiv \tilde{\mathbb{E}}[\kappa_{t+1}|\mathcal{I}_t]$. Both Kalman gain parameters, ν_τ and τ_κ, are positive. The predictive distribution of π_{t+1} is

$$\pi_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\widehat{\tau}_t + \widehat{\kappa}_t \ , \ \tilde{\sigma}_\pi^2\right), \tag{1.33}$$

where $\tilde{\sigma}_\pi^2 > \tilde{\sigma}_\eta^2$. More details are given in Appendix A.3.1.

1.5.2 Stochastic discount factor and short rates

I assume that the representative agent evaluates real payoffs with Epstein–Zin recursive preferences (Epstein and Zin 1989). The value function is

$$V_t = \left[(1-\delta)G_t^{1-\frac{1}{\psi}} + \delta\tilde{\mathbb{E}}_t[V_{t+1}^{1-\gamma}]^{\frac{1-\frac{1}{\psi}}{1-\gamma}}\right]^{\frac{1}{1-\frac{1}{\psi}}}, \tag{1.34}$$

where δ denotes the time discount factor, γ the risk aversion parameter, and ψ the elasticity of intertemporal substitution. Note that the continuation value is evaluated under the agent's subjective expectation $\tilde{\mathbb{E}}_t$.

I follow Bansal and Yaron (2004) to solve the model with log-linearization as in Campbell and Shiller (1988). Appendix A.3.2 shows that the real log stochastic discount factor (SDF) can be written as

$$m_{t+1} = \tilde{\mu}_m - \frac{1}{\psi}\left(\widehat{\mu}_t + \widehat{c}_t\right) - \xi\tilde{\sigma}_g\tilde{\varepsilon}_{t+1}, \tag{1.35}$$

where $\tilde{\mu}_m$ and $\xi > 0$ are constants and $\{\tilde{\varepsilon}_t\}$ are standard normal shocks under the agent's subjective expectations. The nominal log SDF is

$$m_{t+1}^{\$} = m_{t+1} - \pi_{t+1}. \tag{1.36}$$

The real and nominal short rates are obtained from the conditional means of real and nominal SDF, respectively, as

$$i_t = -\tilde{\mu}_m + \frac{1}{\psi}\left(\widehat{\mu}_t + \widehat{c}_t\right) - \frac{1}{2}\xi^2\tilde{\sigma}_g^2, \tag{1.37}$$

$$i_t^{\$} = -\tilde{\mu}_m + \frac{1}{\psi}\left(\widehat{\mu}_t + \widehat{c}_t\right) + \left(\widehat{\tau}_t + \widehat{\kappa}_t\right) - \frac{1}{2}\xi^2\tilde{\sigma}_g^2 - \frac{1}{2}\tilde{\sigma}_\pi^2. \tag{1.38}$$

Thus, real short rates are high when the agent perceives a high aggregate real growth in the next period (intertemporal smoothing). The nominal short rates are additionally affected by perceived future inflation.

1.5.3 Bond pricing

I conjecture that log prices of the real and nominal bonds with maturity n, p_t^n, and $p_t^{\$,n}$ are affine functions of the state variables $(\widehat{\mu}_t, \widehat{c}_t, \widehat{\tau}_t, \widehat{\kappa}_t)$:

$$p_t^n = a_n + b_n\widehat{\mu}_t + c_n\widehat{c}_t, \tag{1.39}$$

$$p_t^{\$,n} = a_n^{\$} + b_n^{\$}\widehat{\mu}_t + c_n^{\$}\widehat{c}_t + d_n^{\$}\widehat{\tau}_t + e_n^{\$}\widehat{\kappa}_t. \tag{1.40}$$

The coefficients are solved recursively by applying the subjective Euler equations as

$$b_n = b_n^{\$} = -\frac{1-\phi_\mu^n}{1-\phi_\mu}\frac{1}{\psi} \quad , \quad c_n = c_n^{\$} = -\frac{1-\tilde{\rho}_c^n}{1-\tilde{\rho}_c}\frac{1}{\psi} \ , \tag{1.41}$$

$$d_n^{\$} = -\frac{1-\phi_\tau^n}{1-\phi_\tau} \quad , \quad e_n^{\$} = -\frac{1-\tilde{\rho}_\kappa^n}{1-\tilde{\rho}_\kappa} \ , \tag{1.42}$$

and the expressions for a_n and $a_n^\$$ are given in Appendix A.3.3. Note that all coefficients are negative due to the intertemporal smoothing motive.

1.5.4 Term structure of interest rates

The log real and nominal yields are obtained from log bond prices as

$$y_t^n = -\frac{1}{n}p_t^n \ , \ y_t^{\$,n} = -\frac{1}{n}p_t^{\$,n}, \tag{1.43}$$

and we can calculate the real and nominal term spread, $s_t^n \equiv y_t^n - y_t^1$ and $s_t^{\$,n} \equiv y_t^{\$,n} - y_t^{\$,1}$, as

$$s_t^n = q_n + \left[1 - \frac{1-\phi_\mu^n}{n(1-\phi_\mu)}\right]\frac{1}{\psi}(\mu - \widehat{\mu}_t) - \left[1 - \frac{1-\tilde{\rho}_c^n}{n(1-\tilde{\rho}_c)}\right]\frac{1}{\psi}\widehat{c}_t, \tag{1.44}$$

$$s_t^{\$,n} = s_t^n + q_n^\$ + \left[1 - \frac{1-\phi_\tau^n}{n(1-\phi_\tau)}\right](\tau - \widehat{\tau}_t) - \left[1 - \frac{1-\tilde{\rho}_\kappa^n}{n(1-\tilde{\rho}_\kappa)}\right]\widehat{\kappa}_t, \tag{1.45}$$

where q_n and $q_n^\$$ are maturity-specific constants. For large n, the real term spread (without the constant term) is approximately $-(\widehat{\mu}_t + \widehat{c}_t)/\psi$, or $-\tilde{\mathbb{E}}_t g_{t+1}/\psi$. Thus the model predicts a countercyclical term spread, which is high when the real endowment growth expectation is low. The same cyclicality applies to the nominal term spread, with additional variation driven by the inflation expectation.

As we will see later, the model has two key parameters

$$\lambda_{g,n} \equiv b_n\nu_\mu - c_n\nu_c = \frac{1}{\psi}\left(\frac{1-\tilde{\rho}_c^n}{1-\tilde{\rho}_c}\nu_c - \frac{1-\phi_\mu^n}{1-\phi_\mu}\nu_\mu\right), \tag{1.46}$$

$$\lambda_{\pi,n} \equiv d_n^\$\nu_\tau - e_n^\$\nu_\kappa = \frac{1-\tilde{\rho}_\kappa^n}{1-\tilde{\rho}_\kappa}\nu_\kappa - \frac{1-\phi_\tau^n}{1-\phi_\tau}\nu_\tau. \tag{1.47}$$

For conciseness, I focus on discussion of Equation (1.46), as the same intuition applies to Equation (1.47). Recall that trend growth expectations are positively auto-correlated while cycle growth expectations are negatively auto-correlated. Given a worse-than-expected realization of aggregate real growth, the agent revises downwards both trend and cycle expectations for the current period, as discussed in Section 1.5.1. Since trend growth is positively auto-correlated, this makes the agent pessimistic about all future trend growth rates. As a result, the prices of long-maturity bonds increase due to the intertemporal smoothing motive. The coefficient $b_n\nu_\mu$ measures the sensitivity of n-maturity bond prices to the trend expectation shock, which is a fraction, ν_μ, of the initial bad aggregate shock. In addition, the agent perceives future trend expectation shocks to be positively auto-correlated, implying that risk accumulates over time. The coefficient $b_n\nu_\mu$ also measures the risk accumulation from $n-1$ to n periods ahead. In this sense, $b_n\nu_\mu$ represents an additional precautionary saving motive from period to period, which also drives up long-maturity bond prices. Both motives are stronger if trend expectations are more sensitive to the initial bad aggregate news (ν_μ is high) or if the perceived trend process is more persistent (ϕ_μ is high). In contrast, the negative auto-correlations in perceived cycle growth make both motives to work in the opposite direction, with a higher $c_n\nu_c$ indicating lower long-term saving and hedging needs. This depresses long-maturity bond prices. The coefficient $\lambda_{g,n}$ measures the net effect of trend and cycle expectation shocks on intertemporal smoothing and precautionary saving motives. In the case of $\lambda_{g,n} > 0$, long-maturity bonds are less attractive than short-maturity bonds.

Thus, the model is able to generate upward-sloping unconditional real and nominal term structures under the *objective* expectation, as shown by the following theorems (proofs are given in Appendix A.3.6).

Theorem I.1. *If*

$$0 < \lambda_{g,n} < 2\xi \ , \quad \forall n \geq 1, \tag{1.48}$$

we have

$$\mathbb{E}[s_t^n] > 0, \quad \forall n > 1; \tag{1.49}$$

that is, the unconditional real term structure is upward-sloping. In addition,

$$0 < \lambda_{g,\infty} < 2\xi \tag{1.50}$$

is a sufficient and necessary condition for the limiting unconditional real term spread, $\lim_{n\to\infty} \mathbb{E}[s_t^n]$*, to be positive.*[34]

Under the objective expectation, the agent *on average* makes the correct forecasts. Thus, from the econometrician's point of view, the intertemporal smoothing motive does not play a role in determining the unconditional yield curve slope. Similarly, for the nominal term structure, we have

Corollary I.2. *If*

$$0 < \lambda_{\pi,n} < 2, \quad \forall n \geq 1, \tag{1.51}$$

we have

$$\mathbb{E}[s_t^{\$,n}] > \mathbb{E}[s_t^n], \quad \forall n > 1; \tag{1.52}$$

that is, the unconditional nominal term structure is more upward-sloping than the real one. In addition,

$$0 < \lambda_{\pi,\infty} < 2 \tag{1.53}$$

is a sufficient and necessary condition for the limiting unconditional nominal term spread, $\lim_{n\to\infty} \mathbb{E}[s_t^{\$,n}]$*, to be larger than the real counterpart.*

Corollary I.2 shows that parameters of subjective inflation processes determine the additional slope in the nominal term structure relative to the real one. Thus, this

[34] Note that for reasonable parameter values, $\lambda_{g,n}$ is approximately monotone in n, and Equation (1.50) can be used as a simple measure to judge the unconditional slope of the real term structure.

model is flexible in generating different patterns of nominal and real term structure; for example, a flat real term structure and an upward-sloping nominal term structure at the same time.

1.5.5 Cyclicality of macroeconomic and short-rate expectation wedges

The difference in perceived pricing dynamics between the representative agent and the econometrician is purely driven by expectation wedges of real growth and inflation. In this section, I study the cyclicality of macroeconomic expectation wedges, then of short-rate expectation wedges.

Given the assumed symmetry between real endowment growth and inflation dynamics, I focus, for conciseness, on real endowment growth, g_t. The expectation wedge of g_t is defined as

$$\mathbb{E}_t g_{t+1} - \tilde{\mathbb{E}}_t g_{t+1} = \mu - \widehat{\mu}_t + c_t^e - \widehat{c}_t \equiv \Delta_{\mu,t} + \Delta_{c,t}, \tag{1.54}$$

where $c_t^e \equiv \mathbb{E}[c_{t+1}|\mathcal{H}_t] = (\rho_c - 1)c_t$ and where $\Delta_{\mu,t} = \mu - \widehat{\mu}_t$ and $\Delta_{c,t} = c_t^e - \widehat{c}_t$ are the trend and cycle wedge, respectively. If there were no cycle components in the model, objective expectations of future growth would be determined only by trend growth expectations which are the constant μ. Subjective expectations, however, would be low during bad times and the expectation wedge $\mathbb{E}_t g_{t+1} - \tilde{\mathbb{E}}_t g_{t+1} = \Delta_{\mu,t}$ would be countercyclical. With cycle components, both objective and subjective growth expectations are time-varying, and the expectation wedge has an additional component that relates to cycle wedge $\Delta_{c,t}$.

The cyclicality of $\Delta_{c,t}$ depends on the subjective and objective persistence of cycle processes, $\tilde{\rho}_c$ and ρ_c, and on the Kalman gain parameter ν_c. For a concrete example, assume that $\Delta_{c,t-1} = 0$ holds. With a low realization of g_t or, equivalently, c_t

both the agent and the econometrician will expect a future mean-reversion in cycle growth, the magnitude of which depends (negatively) on cycle persistence. While the econometrician perceives

$$c_t^e = (\rho_c - 1)c_t = \underbrace{\rho_c c_{t-1}^e}_{\text{reversion from old belief}} + \underbrace{(\rho_c - 1)\zeta_t}_{\text{reversion from new shock}}, \tag{1.55}$$

the agent perceives

$$\widehat{c}_t = \underbrace{\tilde{\rho}_c \widehat{c}_{t-1}}_{\text{reversion from old belief}} + \underbrace{(-\nu_c)\zeta_t^*}_{\text{reversion from new shock}}. \tag{1.56}$$

Thus, the mean-reversion comes from two sources: previous belief and newly observed bad shock. In the special case of $\rho_c = \tilde{\rho}_c$, the contribution of mean-reversion from previous belief is the same for the econometrician and the agent, and cycle wedge is determined by how much news is incorporated in updating beliefs. If the agent does not update her cycle belief strongly—that is, if ν_c is small—she will underestimate the magnitude of mean-reversion relative to the econometrician and the cycle wedge will appear to be countercyclical. If $\tilde{\rho}_c$ is further increased, the agent will also underestimate mean-reversion from previous belief and make the cycle wedge more countercyclical, provided $\widehat{c}_{t-1}$ is low.

To summarize, the trend wedge is always countercyclical, while the cycle wedge can be either procyclical or countercyclical depending on model parameters. Thus, the cyclicality of the aggregate growth expectation wedge is not uniquely determined. However, based on parameter estimations and empirical evidence that will be introduced later, both the model and the data support a countercyclical expectation wedge at least for the real endowment growth.

Next, I study short-rate expectation wedges. In the model, forecasting short rates is equivalent to forecasting state variables $(\widehat{\mu}, \widehat{c}, \widehat{\tau}, \widehat{\kappa})$. Expressions of subjective forecasts can be easily derived from the Kalman updating rules, which evolve as AR(1) processes. The expressions for objective forecasts, however, are complicated and need to be derived from recursions (see Appendix A.3.4.2). This is because the econometrician has to keep track of how the representative agent incorporates aggregate news into her trend and cycle expectations over time. I therefore examine here the one-period-ahead short-rate forecast error, for which a simple expression can be derived:

$$\mathbb{E}_t i^{\$}_{t+1} - \tilde{\mathbb{E}}_t i^{\$}_{t+1} = -\lambda_{g,1}\left(\mathbb{E}_t g_{t+1} - \tilde{\mathbb{E}}_t g_{t+1}\right) - \lambda_{\pi,1}\left(\mathbb{E}_t \pi_{t+1} - \tilde{\mathbb{E}}_t \pi_{t+1}\right). \tag{1.57}$$

If $\lambda_{g,1} > 0$ or $\nu_c > \nu_\mu$, the one-period short-rate expectation wedge is negatively correlated with the real growth expectation wedge. This happens through the intertemporal smoothing motive, whose direction is determined by the reversal effect from cycle expectation shocks. When the agent is overly pessimistic about next-period real growth—that is, when $\mathbb{E}_t g_{t+1} - \tilde{\mathbb{E}}_t g_{t+1}$ is high—she will overestimate the amount of mean-reversion and be overly optimistic about g_{t+2}, implying too high a short-rate expectation, $\tilde{\mathbb{E}}_t i^{\$}_{t+1}$. When time $t+1$ arrives, however, the agent recognizes her growth expectation error and prices a lower short rate instead. This revision in subjective expectations is anticipated by the econometrician, who has full knowledge.

For long-term short-rate expectation wedges, I rely on calibration, since simple closed-form expressions are not available. Appendix A.3.4.2 shows that, generally, the expectation wedge is a function of $(c^e_t, \kappa^e_t, \widehat{c}_t, \widehat{\mu}_t, \widehat{\kappa}_t, \widehat{\tau}_t)$. Figure B.3 plots the coefficients of $\widehat{\mu}_t$, $\widehat{c}_t$, and the negative of c^e_t as an illustration.[35] The coefficients of $\widehat{c}_t$ and $-c^e_t$ are very similar, indicating a cancellation effect if subjective and objective cycle expectations are similar. The coefficient of $\widehat{\mu}_t$ is positive at shorter horizons and turns

[35] The parameter values are close to estimations that will be introduced later. I choose $\phi_\mu = 0.96$, $\tilde{\rho}_c = 0.82$, $\rho_c = 0.75$, $\nu_\mu = 0.02$, and $\nu_c = 0.2$.

negative at longer horizons, due to model assumptions that trend growth rates are more persistent under the subjective expectation. In addition, cycle coefficients decay more quickly than trend coefficients, implying a more important role for trend expectation wedges at longer horizons.

1.5.6 Subjective and objective risk premia

Finally, I study model-implied risk premia on nominal bonds, which are studied empirically in this paper. Because shocks in the model are homoskedastic, the subjective uncertainties perceived by the agent remain constant over time in the assumed steady state. As a result, the agent requires constant subjective risk premia on zero-coupon bonds:

$$\tilde{\mathbb{E}}_t r_{t+1}^n - i_t^{\$} + \frac{1}{2}\tilde{\text{var}}(r_{t+1}^n) = \xi \lambda_{g,n-1}\tilde{\sigma}_g^2 + \lambda_{\pi,n-1}\tilde{\sigma}_\pi^2. \tag{1.58}$$

In comparison, the objective risk premia are time-varying, and co-move with macroeconomic expectation wedges between the econometrician and the representative agent:

$$\mathbb{E}_t r_{t+1}^n - i_t^{\$} + \frac{1}{2}\text{var}(r_{t+1}^n) = C_n + \lambda_{g,n-1}\left(\mathbb{E}_t g_{t+1} - \tilde{\mathbb{E}}_t g_{t+1}\right) + \lambda_{\pi,n-1}\left(\mathbb{E}_t \pi_{t+1} - \tilde{\mathbb{E}}_t \pi_{t+1}\right), \tag{1.59}$$

where C_n is a maturity-specific constant (more details are in Appendix A.3.5). Thus, if $\lambda_{g,n} > 0$, objective risk premia are countercyclical.[36] The same intuition invoked in previous sections applies here. When $\lambda_{g,n} > 0$, the net effect of trend and cycle expectation shocks decreases the agent's long-term saving needs, given that she is overly pessimistic relative to the objective forecast. This is because she overestimates the amount of mean-reversion and becomes overly optimistic about future growth. The intertemporal smoothing motive depresses current prices of long-maturity bonds. With

[36] Provided $\mathbb{E}_t g_{t+1} - \tilde{\mathbb{E}}_t g_{t+1}$ is countercyclical. See discussion in Section 1.5.5.

full knowledge, the econometrician anticipates that the agent will be disappointed next period, indicating higher prices of long-maturity bonds going forward and higher future excess returns.

1.5.7 Summary of model predictions

To summarize, the model yields novel predictions on the joint behavior of the unconditional term structure of interest rates, the cyclicality of macroeconomic and short-rate expectation wedges, and the cyclicality of objective risk premia, all of which are testable in the data.

Specifically, if $\lambda_{n,g} > 0$, the unconditional term structure of real interest rates is upward-sloping, the short-term short-rate expectation wedge is negatively correlated with the real growth expectation wedge, and objective risk premia are countercyclical. A similar conclusion applies to inflation dynamics.

1.6 Identifying Parameters and Testing Additional Predictions

In this section, I identify model parameters, using only macroeconomic variables and survey forecasts, to examine whether my model's predictions are consistent with bond prices.

For parameters under the objective measure, I consider the estimates in Grant and Chan (2017) of a UC0-AR(2) model, which is the closest to my setting.[37] The model assumes that real GDP levels can be described, as in Equations (1.14) and (1.15), but with cycle levels following an AR(2) process. To map AR(2) estimates to my AR(1) setting, I consider an approximation described in Appendix A.9, which maps

[37] I thank the authors for making the code available on their website.

the persistence coefficients by ignoring the auto-correlation structure in the shocks.[38] I also apply the above approximation to the estimates of log CPI. Table 1.6, Panel A, reports the original AR(2) estimates and their mapping to an AR(1) setting.

Parameters under the subjective measure are identified from both short- and long-term survey forecasts of real GDP growth and inflation at quarterly frequency. I use a step-by-step GMM estimation approach to reduce the burden of numerical optimization. I first identify trend-related parameters from long-term annual average forecasts from BCFF and BCEI. Using real GDP growth as an example, Appendix A.3.4 shows that if $\tilde{\rho}_c$ is much smaller than ϕ_μ, we approximately have

$$\tilde{\mathbb{E}}_t g_{t+j} = \phi_\mu^{j-1}\widehat{\mu}_t + (1-\phi_\mu^{j-1})\mu \tag{1.60}$$

for large j. This implies the following moment conditions for long-term forecasts:

$$\tilde{\mathbb{E}}_t \bar{g}_{t+n+m}^{annual} - \phi_\mu^{4m}\tilde{\mathbb{E}}_t \bar{g}_{t+n}^{annual} - (1-\phi_\mu^{4m})\mu = 0, \tag{1.61}$$

where $\tilde{\mathbb{E}}_t \bar{g}_{t+n}^{annual}$ is the n-year-ahead annual average forecast in quarter t and the parameters are in quarterly values. I estimate the model with 2-, 3-, 4-, and 5-year-ahead annual averages and the long-term average. These averages yield four moment conditions to be used to identify two parameters. Table 1.6, Panel B, shows that the model provides a reasonable fit to long-term annual average forecasts. The model is not rejected for either the real GDP growth or inflation forecasts, based on p-values. The true means of real GDP growth and inflation are estimated to be annualized 2.61% and 2.76%, respectively. For inflation, the estimated mean is higher than the Fed's 2% target rate set in 2012, which is reasonable given that my sample encompasses the high-inflation period in the early 1980s.

[38] The accuracy of this approximation does not affect the later analyses.

I next identify cycle-related parameters, using short-term survey forecasts from BCEI and SPF. Still from Appendix A.3.4, I identify cycle expectations using

$$\tilde{\mathbb{E}}_t\tilde{c}_{t+h} \equiv \rho_c^{h-1}\widehat{c}_t = \tilde{\mathbb{E}}_t g_{t+h} - \phi_\mu^{h-1}\widehat{\mu}_t - (1-\phi_\mu^{h-1})\mu, \quad h \geq 1, \tag{1.62}$$

with estimated trend expectations. The moment conditions for cycle expectations satisfy

$$\tilde{\mathbb{E}}_t\tilde{c}_{t+h+j} - \tilde{\rho}_c^j\tilde{\mathbb{E}}_t\tilde{c}_{t+h} = 0. \tag{1.63}$$

However, in the finite sample, it is unlikely that $\mathbb{E}_T[\tilde{\mathbb{E}}_t\tilde{c}_{t+h}] = 0$. Thus I allow a constant term in the moment conditions and estimate instead

$$\tilde{\mathbb{E}}_t\tilde{c}_{t+h+j} - \tilde{\rho}_c^j\tilde{\mathbb{E}}_t\tilde{c}_{t+h} - (1-\tilde{\rho}_c^j)x_c = 0, \tag{1.64}$$

which is implied by an AR(1) model with a drift. Table 1.6, Panel B, shows that the AR(1) model of cycle expectations is not rejected. In addition, the cycle processes under subjective expectations appear to be more persistent than their objective counterparts.

Finally, I identify Kalman gain parameters, using belief shocks $g_t - \tilde{\mathbb{E}}_{t-1}g_t$, which are directly observable. From Equations (1.23) and (1.24), I run the following regressions:

$$\widehat{\mu}_t - \phi_\mu\widehat{\mu}_{t-1} = \kappa_{\mu,0} + \kappa_{\mu,1}\left(g_t - \tilde{\mathbb{E}}_{t-1}g_t\right) + \eta_t, \tag{1.65}$$

$$\widehat{c}_t - \tilde{\rho}_c\widehat{c}_{t-1} = \kappa_{c,0} + \kappa_{c,1}\left(\tilde{\mathbb{E}}_{t-1}g_t - g_t\right) + \iota_t. \tag{1.66}$$

Note that I use $\tilde{\mathbb{E}}_{t-1}g_t - g_t$ in Equation (1.66) to be consistent with the negative sign of ν_c in Equation (1.24). Since $\widehat{\mu}_t$, ϕ_μ, $\widehat{c}_t$, and $\tilde{\rho}_c$ are all estimated, the dependent variables of the above regressions are inevitably contaminated with measurement errors. The estimation noise is especially large for Equation (1.65), given the high persistence of $\widehat{\mu}_t$.

However, Table 1.6, Panel B, shows that the estimated value, $\hat{\nu}_\mu = 0.022$, is reasonable compared to the model-free estimate in Appendix A.7. In addition, ν_c is fairly precisely estimated to be 0.182. Since Kalman gain parameters are positively correlated with volatilities of subjective shocks, this suggests that subjectively the cycle growth is much more volatile than the trend growth. Parameters of the subjective and objective inflation dynamics are identified similarly.

Based on the identified parameters from Table 1.6, a simple calculation shows that

$$\psi\lambda_{g,\infty} = 0.41 \quad , \quad \lambda_{\pi,\infty} = 2.77. \tag{1.67}$$

Thus, Theorem I.2 and Corollary I.2 predict an upward-sloping real term structure and a more upward-sloping nominal term structure. This is consistent with the observation that from January 1999 to December 2018, the average real spread between 20- and 2-year yields is 1.43%, using the TIPS data from Gürkaynak et al. (2010), and that the average nominal spread for the same maturities is 2.06%.[39]

1.6.1 Testing additional predictions

1.6.1.1 Yield movements

Appendix A.3.3 shows that nominal bond yields are functions of trend and cycle expectations of real growth and inflation:

$$y_t^{\$,n} \equiv -\frac{1}{n}p_t^{\$,n} = -\frac{a_n^{\$}}{n} + \frac{1-\phi_\mu^n}{n(1-\phi_\mu)}\frac{1}{\psi}\widehat{\mu}_t + \frac{1-\tilde{\rho}_c^n}{n(1-\tilde{\rho}_c)}\frac{1}{\psi}\widehat{c}_t + \frac{1-\phi_\tau^n}{n(1-\phi_\tau)}\widehat{\tau}_t + \frac{1-\tilde{\rho}_\kappa^n}{n(1-\tilde{\rho}_\kappa)}\widehat{\kappa}_t. \tag{1.68}$$

[39] Ang et al. (2008) estimate the unconditional real (nominal) spread between 5-year and 3-month yields to be around 1.3% (6.5%) using data from 1952Q2 to 2004Q4.

Thus, the coefficients of regressing nominal yields on trend expectations should decline with maturity and should decline more quickly if trend expectation persistence ϕ is low. Trend parameter estimations from Table 1.6 suggest that the coefficients of $\widehat{\mu}_t$ should decline more quickly with maturity than those of $\widehat{\tau}_t$. Using proxies for long-term expectations constructed earlier, as in Appendix A.7, I perform the following regression at quarterly frequency:

$$y_t^{\$,n} = \beta_0 + \beta_{1,n} g_t^{LT} + \beta_{2,n} \pi_t^{LT} + \varepsilon_t \tag{1.69}$$

for nominal yields with maturities of 1 to 15 years from 1971Q4 to 2018Q4 (yields are sampled in the middle month of a calendar quarter). Figure B.5 shows that both estimates, $\hat{\beta}_{1,n}$ and $\hat{\beta}_{2,n}$, are statistically significantly different from zero for all maturities considered. In addition, $\hat{\beta}_{1,n}$ declines more quickly than $\hat{\beta}_{2,n}$, consistent with the observation that $\phi_\mu < \phi_\tau$.[40]

Equation (1.68) also shows that regression residuals from Equation (1.69) should contain important information about cycle expectations. To separate real growth and inflation cycle expectations, I do not use nominal yields, but rather perform a regression of 2-year TIPS yields on g^{LT} and use the regression residuals as real-yield-implied cycle expectations of real GDP growth. Figure B.6 shows that this yield-implied measure co-moves with survey-implied cycle expectations, except during the 2007–2008 financial crisis.[41]

[40] Jørgensen (2018) documents a similar pattern and interprets inflation trend as a level factor and real growth trend as a slope factor. But he does not explain why this is the case.

[41] The distinct behavior during the 2007–2008 financial crisis could be due to a large liquidity premium in TIPS yield (D'Amico et al. 2018).

1.6.1.2 Cyclicality of short-rate and macroeconomic expectation wedges

Given the identified parameters, we should anticipate a negative relation between short-rate and macroeconomic expectation wedges in the short term, as suggested by Equation (1.57).

I consider the short-term expectation wedges of real GDP growth but not inflation given the ample evidence that short-term survey forecasts of inflation are accurate from the econometrician's point of view (for a review, see Duffee 2018); that is, I assume that the short-term inflation expectation wedge is zero. To construct objective expectations of real GDP growth, $\mathbb{E}_t g_{t+1}$, I follow a large literature in using yield spread as a predictor (for a review, see Stock and Watson 2003) while controlling for lagged real GDP growth realizations:

$$g_{t+1} = \gamma_0 + \gamma_1 g_t + \gamma_2 s_t^{\$,n} + \varepsilon_{t+1}. \tag{1.70}$$

Note that this is also the approach used by the Cleveland Fed.[42] Under the null of the model, Equation (1.70) also works since yield spread reflects the agent's macroeconomic expectations, which incorporate information on real growth and inflation. To avoid seasonality effects, I use year-on-year growth and run one-year-ahead predictive regressions at quarterly frequency. The yield spread is defined as the difference between 7-year and 3-month zero-coupon yields, for which observations back to 1964 are available, and is sampled in the middle month of a calendar quarter.[43] Table B.4 shows that yield spread is a significant predictor of future real GDP growth. To be consistent with the year-on-year growth specification, I use the average of one- to four-quarter-ahead survey forecasts as the measure of subjective expectations, $\tilde{\mathbb{E}}_t g_{t+1}$.

[42] `https://www.clevelandfed.org/en/our-research/indicators-and-data/yield-curve-and-gdp-growth`

[43] Fama and Bliss (1987) suggest that long-maturity yields are not reliable before 1964.

Following Cieslak (2018), I construct the objective expectations of short rates, $\mathbb{E}_t i^{\$}_{t+1}$, from one-year-ahead predictive regressions of within-quarter average 3-month tbill rates with lagged average rates and year-on-year employment growth as predictors. The subjective short-rate expectations, $\tilde{\mathbb{E}}_t i^{\$}_{t+1}$, are simply the one-year-ahead survey forecasts of 3-month tbill rates taken from BCEI.

Figure 1.7 shows that short-rate and real GDP growth expectation wedges are strongly negatively correlated at a one-year horizon. The regression

$$\mathbb{E}_t i^{\$}_{t+1} - \tilde{\mathbb{E}}_t i^{\$}_{t+1} = \alpha_0 + \alpha_1 \left(\mathbb{E}_t g_{t+1} - \tilde{\mathbb{E}}_t g_{t+1} \right) + \varepsilon_t \tag{1.71}$$

yields an estimate of $\hat{\alpha}_1 = -0.25$ with a Newey-West t-statistic (6-quarter lag) of -2.73 and an adjusted R^2 of 0.14. Thus, the evidence is consistent with Equation (1.57), which also suggests that real GDP growth expectation wedges are countercyclical.

1.6.1.3 Objective risk premia predictability

A number of studies (Cieslak and Povala 2015; Jørgensen 2018; Bauer and Rudebusch 2019) have documented that including long-term trends in inflation and real GDP growth strongly increases the predictive power of future bond excess returns in addition to yield curve principle components (PCs); that is, the spanning hypothesis is rejected. My model generates this predictability by linking the macroeconomic trends to expectation wedges as shown in Equation (1.59) (ignoring the constant term):

$$\mathbb{E}_t r^n_{t+1} - i^{\$}_t + \frac{1}{2}\mathrm{var}(r^n_{t+1}) = \lambda_{g,n-1} \left(\mu - \widehat{\mu}_t + c^e_t - \widehat{c}_t \right) + \lambda_{\pi,n-1} \left(\tau - \widehat{\tau}_t + \kappa^e_t - \widehat{\kappa}_t \right). \tag{1.72}$$

Thus, if $\lambda_{g,n-1} > 0$, the subjective real growth trend expectation enters into the return predictability regression with a negative sign. Since the spanning hypothesis holds in my model, I test whether subjective trend expectations have predictability for future bond excess returns without controlling for yield curve PCs. Given the low-frequency nature of trend expectations, I use a long history of bond excess returns from Fama bond portfolios and use g^{LT} and π^{LT} as proxies (see Appendix A.7). The predictive regression is performed at a one-year horizon, using data from 1964 to 2018.

Table 1.7 shows that the real GDP growth trend expectation is a significant predictor of future excess returns, especially for bonds with longer maturities. Also, the predictive sign is negative, consistent with the model's prediction. In comparison, inflation trend expectation does not have predictability. However, given the earlier observation that subjective risk premium co-moves positively with the inflation trend, this implies that the *difference* between objective and subjective risk premium is still negatively correlated with the inflation trend expectation. Thus, this pattern can still be consistent with an extended version of the model where this trend in subjective risk premium is accounted for.

1.7 Conclusion

This paper shows that allowing for a wedge between subjective and objective macroeconomic expectations can reconcile bond prices and survey data. Subjective bond risk premia implied by real-time survey forecasts of future yields are acyclical, while short-rate forecast errors are predictable by trend and cycle components of macroeconomic forecasts. These findings challenge the assumption of rational expectations models that the ex-post risk premia fitted from in-sample predictive regressions accurately reflect the risk compensation required by investors in real time. Departing from the canonical rational expectations framework, I focus on modeling the expectation wedge between the real-time agents and the econometrician in a highly tractable equilibrium. In my model, the risk compensation required by the agent is constant by construction and variation in yields is driven purely by time-varying subjective macroeconomic expectations. Given that the agent has to learn about macroeconomic dynamics in real time while the econometrician has the advantage of hindsight, future revisions in the agent's expectations are anticipated by the econometrician. Thus, future yield changes appear to be predictable from the econometrician's point of view, as they are in the data. Also, the model can generate predictions that are consistent with several stylized facts in the Treasury bond market.

Overall, my results suggest that with the help of survey data, equilibrium models of bond pricing could explore differences between subjective and objective dynamics and provide a unified explanation of many empirical patterns we have observed in the Treasury bond market.

1.8 Table and Figures

Table 1.1: Description of Surveys

This table summarizes the basic information of the surveys. For concreteness, I use the survey published in 1984Q2 (and 1984M6 for monthly surveys) as an illustrative example.

	Frequency	Survey time	Forecast date
BCFF	Monthly	1984M5 last week	Quarterly average
BCEI	Monthly	1984M6 first week	Quarterly average
SPF	Quarterly	1984M5 first half	Quarterly average
GN	Quarterly	1984M6	Last business day

Table 1.2: Moments of Yield Forecasts and Implied Risk Premia

Panel A reports the moments of bootstrapped constant-horizon zero-coupon yield forecasts, with sample means in the first row and sample standard deviations in parentheses. All values are in annualized percentage points. 3-month tbill rate forecasts are from October 1982 to December 2018. Yield forecasts for other maturities are from December 1987 to December 2018. "Data" reports moments calculated from observed yields during each sample period that is matched to the survey sample.

Columns (1) and (2) in Panel B report means and standard deviations of monthly subjective and objective risk premia for bond maturities with 2, 5, and 10 years. All values are in annualized percentage points. The sample period is from December 1987 to December 2018. Subjective risk premia are calculated from consensus yield forecasts using Equation (1.2), while objective risk premia are fitted from predictive regressions using the cycle factor from Cieslak and Povala (2015) and the macro factor from Ludvigson and Ng (2009) as predictors. Columns (3) and (4) report moments of "cleaned" subjective risk premia, which are fitted values from regressing subjective risk premia on current long/short yields and the cycle factor, respectively.

A. Yield Forecast Moments

	Survey					Data
	1Q	2Q	3Q	4Q	5Q	
$n = 3M^*$	4.08	4.15	4.23	4.33	4.43	4.26
	(3.17)	(3.07)	(2.98)	(2.89)	(2.79)	(3.59)
$n = 1Y$	3.39	3.49	3.61	3.74	3.87	3.41
	(2.44)	(2.38)	(2.30)	(2.20)	(2.10)	(2.59)
$n = 5Y$	4.27	4.36	4.46	4.56	4.66	4.33
	(2.14)	(2.08)	(2.00)	(1.93)	(1.86)	(2.33)
$n = 10Y$	4.83	4.91	4.99	5.07	5.15	4.98
	(1.90)	(1.84)	(1.78)	(1.72)	(1.67)	(2.08)

B. Risk Premia Moments

	(1)		(2)		(3)			(4)		
	Subjective		Objective		Cleaned sub.			Cycle fitted		
	Mean	S.D.	Mean	S.D.	Mean	S.D.	R^2	Mean	S.D.	R^2
$n = 2Y$	0.21	0.54	0.90	0.85	0.21	0.42	0.60	0.21	0.27	0.25
$n = 5Y$	0.38	1.86	2.96	3.00	0.38	1.47	0.63	0.38	1.06	0.32
$n = 10Y$	1.58	3.86	5.31	6.04	1.58	3.00	0.60	1.58	2.20	0.32

Table 1.3: Subjective and Objective Risk Premia Movements

This table reports the regression results of subjective and objective risk premia on yield and macroeconomic variables, $\hat{E}_t rx^n_{t+1} = \beta_0 + \beta_1' X_t + \varepsilon_t$. For "average", I use the maturity-weighted average of monthly 1-year-holding-period risk premia on 2- to 10-year zero-coupon bonds. The sample period is from December 1987 to December 2018. For "$n = 1.5Y$," I use quarterly 6-month-holding-period risk premia on 1.5-year zero-coupon bonds. The sample period is from 1969Q4 to 2018Q4. All explanatory variables are standardized to have a unit standard deviation. Detailed definitions of explanatory variables are given in Section 1.3.1.2. Newey-West t-statistics with 18-month (6-quarter for "$n = 1.5Y$") lags are reported in parentheses.

	Subjective $\tilde{\mathbb{E}}_t rx^n_{t+1}$				Objective $\mathbb{E}_t rx^n_{t+1}$			
	Average		$n = 1.5Y$		Average		$n = 1.5Y$	
	Coeff.	R^2	Coeff.	R^2	Coeff.	R^2	Coeff.	R^2
Panel A: Trend								
Yield level	0.26	0.53	0.58	0.19	0.18	0.09	-0.04	-0.00
	(8.48)		(5.94)		(1.76)		(-0.34)	
Trend inflation	0.26	0.55	0.58	0.19	0.16	0.07	-0.11	0.01
	(7.22)		(5.49)		(1.51)		(-0.96)	
Panel B: Cycle (controlling for trend inflation)								
Yield slope	0.06	0.58	0.17	0.20	0.44	0.62	0.51	0.39
	(2.93)		(1.48)		(9.02)		(8.73)	
(-)Employ. growth	0.06	0.58	0.02	0.18	0.38	0.47	0.34	0.18
	(3.06)		(0.19)		(6.67)		(3.20)	
Panel C: Volatility								
Yield volatility	0.08	0.04	0.20	0.02	0.19	0.10	0.06	-0.00
	(1.75)		(2.47)		(2.47)		(0.54)	
Treasury VIX	0.13	0.12			0.27	0.22		
	(3.24)				(3.54)			
Real volatility	0.07	0.03	0.31	0.05	0.30	0.27	0.06	0.00
	(1.11)		(2.12)		(3.51)		(0.71)	
Inflation volatility	0.10	0.08	0.30	0.05	0.18	0.10	0.00	-0.01
	(1.49)		(1.77)		(1.90)		(0.01)	
Panel D: Uncertainty/Forecast Dispersion								
Policy uncertainty	-0.04	0.01			0.01	-0.00		
	(-1.04)				(0.18)			
rGDP uncertainty	0.08	0.05	0.10	0.00	-0.28	0.23	-0.43	0.27
	(1.63)		(0.57)		(-3.89)		(-3.74)	
CPI uncertainty	0.14	0.16	0.25	0.03	-0.12	0.04	-0.34	0.17
	(2.78)		(1.51)		(-1.37)		(-2.85)	
rGDP dispersion	0.15	0.18	0.20	0.02	0.14	0.06	-0.36	0.19
	(4.78)		(1.36)		(1.60)		(-4.76)	
CPI dispersion	0.08	0.05			0.12	0.04		
	(1.99)				(1.72)			

Table 1.4: Expectation Hypothesis Tests

This table reports the regression results of the EH test for survey yield forecasts and future realized yields, as in Equation (1.6). Panel A reports the results for survey yield forecasts. Bootstrapped p-values following Crump and Gospodinov (2019) are in brackets. The sample period is from December 1987 to December 2018. Panel B reports the results for future realized yields. The bias-adjusted slope coefficients following the first-order approximation method in Bekaert et al. (1997) are reported in brackets. Newey-West standard errors with 18-month lags are reported in parentheses for both panels.

A. Survey Expectation $\tilde{\mathbb{E}}_t y^{n-m}_{t+m} - y^m_t = \beta_0 + \beta_1 \frac{n}{n-m}(y^n_t - y^m_t) + \eta_t$.

	$m = 3M$	$m = 6M$	$m = 12M$
$n = 6M$	0.34		
(s.e.)	(0.14)		
[p-value]	[0.04]		
$n = 12M$	0.56	0.24	
(s.e.)	(0.07)	(0.10)	
[p-value]	[0.01]	[0.00]	
$n = 24M$	0.75	0.64	0.53
(s.e.)	(0.04)	(0.07)	(0.12)
[p-value]	[0.02]	[0.03]	[0.11]
$n = 60M$	0.92	0.90	0.92
(s.e.)	(0.04)	(0.05)	(0.07)
[p-value]	[0.53]	[0.50]	[0.66]
$n = 120M$	0.93	0.94	0.98
(s.e.)	(0.03)	(0.04)	(0.05)
[p-value]	[0.42]	[0.56]	[0.86]

B. Objective Expectation $y^{n-m}_{t+m} - y^m_t = \beta_0 + \beta_1 \frac{n}{n-m}(y^n_t - y^m_t) + \varepsilon_{t+m}$.

	1972M1 - 2018M12			1987M12 - 2018M12		
	$m = 3M$	$m = 6M$	$m = 12M$	$m = 3M$	$m = 6M$	$m = 12M$
$n = 6M$	0.32			0.81		
[adj. β]	[-1.70]			[-2.23]		
(s.e.)	(0.23)			(0.21)		
$n = 12M$	0.53	0.26		0.89	0.89	
[adj. β]	[-1.47]	[-1.74]		[-2.13]	[-2.14]	
(s.e.)	(0.13)	(0.19)		(0.14)	(0.31)	
$n = 24M$	0.73	0.52	0.23	0.87	0.74	0.71
[adj. β]	[-1.25]	[-1.46]	[-1.76]	[-2.12]	[-2.25]	[-2.30]
(s.e.)	(0.08)	(0.12)	(0.27)	(0.12)	(0.24)	(0.42)
$n = 60M$	0.86	0.73	0.52	0.89	0.78	0.69
[adj. β]	[-1.04]	[-1.17]	[-1.38]	[-1.98]	[-2.09]	[-2.19]
(s.e.)	(0.03)	(0.06)	(0.14)	(0.05)	(0.10)	(0.17)
$n = 120M$	0.90	0.82	0.67	0.93	0.86	0.77
[adj. β]	[-0.88]	[-0.96]	[-1.12]	[-1.76]	[-1.83]	[-1.93]
(s.e.)	(0.02)	(0.04)	(0.08)	(0.03)	(0.05)	(0.08)

Table 1.5: Predicting Short-rate Expectation Errors with Subjective Macroeconomic Expectations

This table reports regression results of 3-month Treasury bill (tbill) rate forecast errors on macroeconomic expectations. Long-term expectations (LT) are proxied by exponentially weighted averages of past data realizations, as in Equations (1.11) and (1.12). Nowcasts are the current-quarter forecasts from SPF. Expectations related to real GDP growth and inflation are indicated by g and π, respectively. In Panel A, I run the 1-year-ahead short-rate forecast error (realization minus forecast) on LT and on nowcasts of real GDP growth and inflation. In Panel B, the realization, i_{t+5}, is defined as the 20-quarter-ahead within-quarter average tbill rate. The forecast, $\tilde{\mathbb{E}}_t i_{t+5}$, is the 5-year-ahead tbill rate forecast from BCFF and BCEI. The sample period is from 1983Q3 to 2018Q4. Newey-West t-statistics with 6-quarter lags are reported in brackets. Bootstrapped p-values following Greenwood and Vayanos (2014) are reported in parentheses.

Panel A: 1-year-ahead Forecast Error $i_{t+1} - \tilde{\mathbb{E}}_t i_{t+1}$

	g			π		
	(1)	(2)	(3)	(4)	(5)	(6)
LT	-1.25		-1.11	0.05		0.06
[t-stat]	[-1.49]		[-2.15]	[0.30]		[0.30]
(b.s. p)	(0.12)		(0.08)	(0.69)		(0.74)
Nowcast		0.23	0.30		0.02	-0.01
[t-stat]		[2.37]	[2.95]		[0.19]	[-0.09]
(b.s. p)		(0.05)	(0.03)		(0.82)	(0.93)
Obs.	138	138	138	138	138	138
Adj. R^2	0.05	0.05	0.12	0.00	0.00	0.00

Panel B: 5-year-ahead Forecast Error $i_{t+5} - \tilde{\mathbb{E}}_t i_{t+5}$

	Realizations i_{t+5}		Forecast error		
	(1)	(2)	(3)	(4)	(5)
LT g	-0.82		-2.52		-1.90
[t-stat]	[-0.77]		[-2.33]		[-2.13]
(b.s. p)	(0.59)		(0.12)		(0.13)
LT π	1.46			0.74	0.70
[t-stat]	[9.04]			[4.38]	[4.46]
(b.s. p)	(0.00)			(0.02)	(0.01)
Nowcast g		0.00			
[t-stat]		[0.01]			
(b.s. p)		(0.99)			
Nowcast π		0.87			
[t-stat]		[3.84]			
(b.s. p)		(0.02)			
Obs.	122	122	122	122	122
Adj. R^2	0.72	0.17	0.16	0.35	0.43

Table 1.6: Identifying Parameters

This table reports the identified parameters under objective and subjective expectations. Panel A reports estimated trend and cycle parameters using the UC0-AR(2) model from Grant and Chan (2017) for real GDP and CPI from 1947 to 2018. ρ_c and ρ_κ are implied persistence coefficients using an AR(1) approximation described in Appendix A.9. Panel B reports the iterative GMM estimations of trend and cycle parameters using quarterly short- and long-term forecasts of real GDP growth and inflation from 1983 to 2018. The moment conditions are outlined in Section 1.6. Gain parameters are identified with the following regressions:

$$\widehat{\mu}_t - \phi_\mu \widehat{\mu}_{t-1} = \alpha_0 + \alpha_1 \left(g_t - \tilde{\mathbb{E}}_{t-1} g_t \right) + \varepsilon_t$$

and

$$\widehat{c}_t - \tilde{\rho}_c \widehat{c}_{t-1} = \beta_0 + \beta_1 \left(\tilde{\mathbb{E}}_{t-1} g_t - g_t \right) + \eta_t$$

and similarly for inflation. Estimated $\hat{\alpha}_1$ and $\hat{\beta}_1$ correspond to ν_μ and ν_c, respectively. Newey-West standard errors with 6-quarter lags are reported in parentheses.

Panel A: Objective Measure

	Trend	Cycle			Volatility	
g	μ	$\rho_{c,1}$	$\rho_{c,2}$	ρ_c	σ_ξ	σ_ζ
	0.811	1.525	-0.555	0.763	0.574	0.655
	(0.039)	(0.088)	(0.093)		(0.107)	(0.108)
π	τ	$\rho_{\kappa,1}$	$\rho_{\kappa,2}$	ρ_κ	σ_η	σ_ι
	0.903	1.739	-0.750	0.870	0.442	0.532
	(0.033)	(0.072)	(0.0720)		(0.066)	(0.075)

Panel B: Subjective Measure

	Trend		Cycle	Gain	
g	μ	ϕ_μ	$\tilde{\rho}_c$	ν_μ	ν_c
	0.653	0.963	0.818	0.022	0.182
	(0.085)	(0.116)	(0.069)	(0.017)	(0.065)
p-value	[0.27]		[0.24]		
Adj. R^2				0.00	0.06
π	τ	ϕ_τ	$\tilde{\rho}_\kappa$	ν_τ	ν_κ
	0.690	0.982	0.939	0.029	0.267
	(0.002)	(0.042)	(0.044)	(0.012)	(0.054)
p-value	[0.25]		[0.71]		
Adj. R^2				0.01	0.27

Table 1.7: Predicting Returns with Subjective Trend Expectations

Dependent variable is the 1-year-holding-period excess returns of Fama bond portfolios, which are calculated from the monthly Fama bond portfolio returns from CRSP. "Short Maturity," "Medium Maturity," and "Long Maturity" correspond to portfolios with maturities of 1 to 2 years, 5 to 10 years, and above 10 years, respectively. Portfolio returns are scaled by average bond maturities for comparison. g^{LT} and π^{LT} are proxies for subjective trend expectations of real GDP growth and inflation, respectively. The details of their construction are given in Appendix A.7. All explanatory variables are standardlized to have a unit standard deviation. Coefficients are multiplied with 100. The sample period is from 1964Q1 to 2018Q4. Newey-West t-statistics with 6-quarter lags are reported in brackets. Bootstrapped p-values following Greenwood and Vayanos (2014) are reported in parentheses.

	Short Maturity			Medium Maturity			Long Maturity		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
g^{LT}	-0.30		-0.27	-1.61		-1.62	-2.49		-2.39
[t-stat]	[-1.79]		[-1.59]	[-2.93]		[-2.73]	[-2.33]		[-2.43]
(b.s. p)	(0.08)		(0.13)	(0.00)		(0.01)	(0.02)		(0.01)
π^{LT}		0.22	0.18		0.08	-0.13		-0.99	-0.57
[t-stat]		[0.91]	[0.74]		[0.08]	[-0.21]		[-0.66]	[-0.52]
(b.s. p)		(0.37)	(0.47)		(0.94)	(0.83)		(0.49)	(0.56)
Observations	216	216	216	216	216	216	185	185	185
Adj. R^2	0.03	0.01	0.03	0.06	0.00	0.06	0.05	0.00	0.05

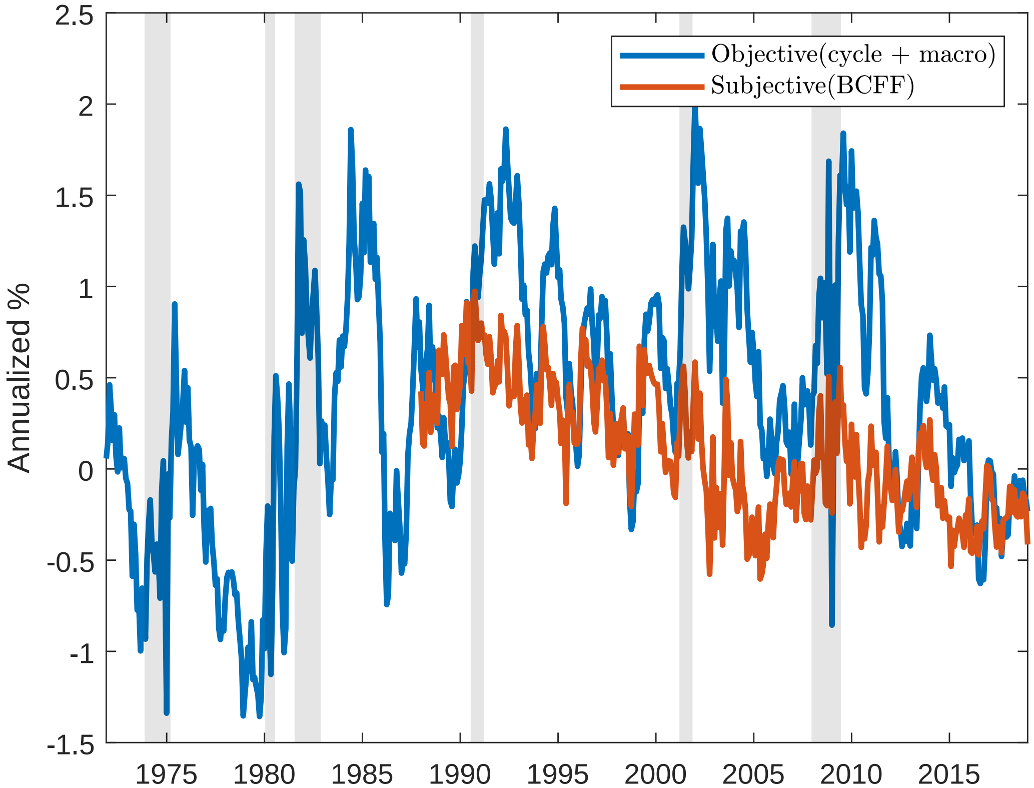

Figure 1.1: Comparison of Subjective and Objective Risk Premia

This figure plots the maturity-weighted average of 1-year-holding-period subjective and objective risk premia, defined as

$$\tilde{\mathbb{E}}_t\overline{rx}_{t+1} \equiv \frac{1}{9}\sum_{n=2}^{10}\tilde{\mathbb{E}}_t rx^n_{t+1}/n \ , \ \mathbb{E}_t\overline{rx}_{t+1} \equiv \frac{1}{9}\sum_{n=2}^{10}\mathbb{E}_t rx^n_{t+1}/n.$$

The maturities of bonds included in the average are from 2 to 10 years. The red line plots $\tilde{\mathbb{E}}_t\overline{rx}_{t+1}$. Individual subjective risk premia are constructed using yield forecasts from BCFF:

$$\tilde{\mathbb{E}}_t rx^n_{t+1} \equiv ny^n_t - (n-1)\tilde{\mathbb{E}}_t y^{n-1}_{t+1} - y^1_t.$$

The sample period is from December 1987 to December 2018. The blue line plots $\mathbb{E}_t\overline{rx}_{t+1}$. Individual objective risk premia, $\mathbb{E}_t rx^n_{t+1}$, are constructed as the fitted values from regressing 1-year-ahead individual bond excess returns on the cycle factor from Cieslak and Povala (2015) and the macro factor from Ludvigson and Ng (2009), using data from November 1971 to December 2018.

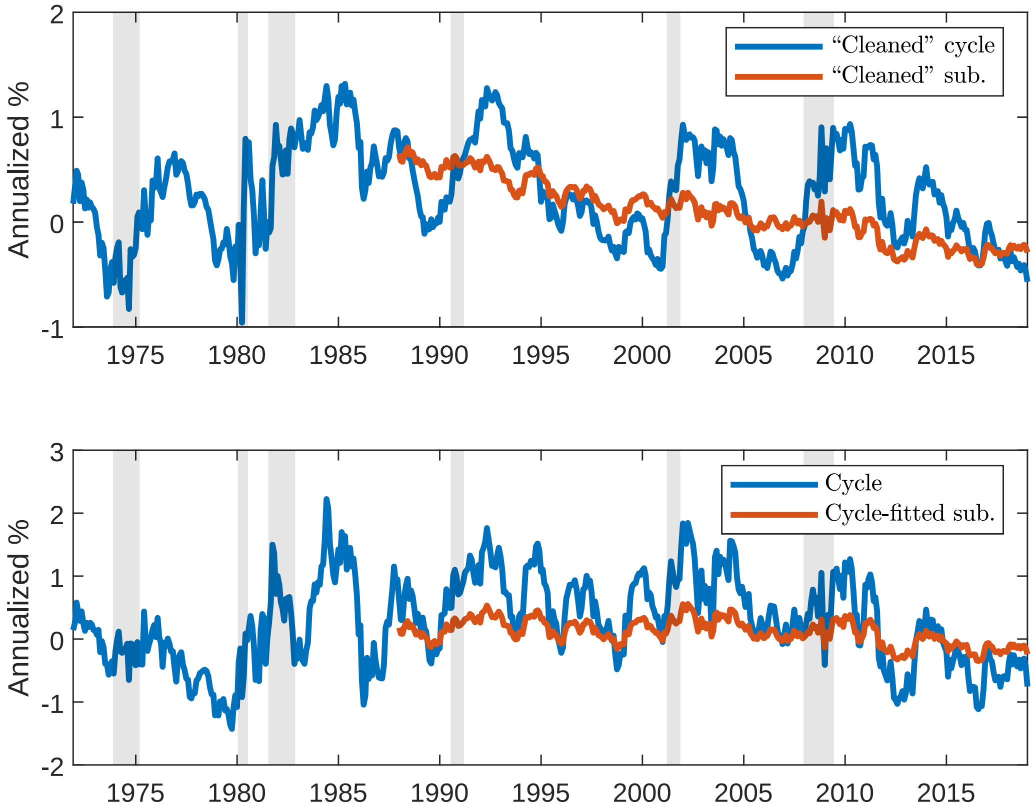

Figure 1.2: Comparison of "Cleaned" Subjective and Objective Risk Premia

In the top panel, the red line plots the fitted values from regressing the maturity-weighted average subjective risk premia, $\tilde{\mathbb{E}}_t \overline{rx}_{t+1}$, on contemporaneous 1- and 10-year yields (including a constant). The blue line plots the fitted values from a similar regression but with the cycle factor as the depedent variable. In the bottom panel, the red line plots the fitted values from regressing the maturity-weighted average subjective risk premia on the cycle factor. The blue line plots the cycle factor for comparison. The sample period is from November 1971 to December 2018 for blue lines and from December 1987 to December 2018 for red lines.

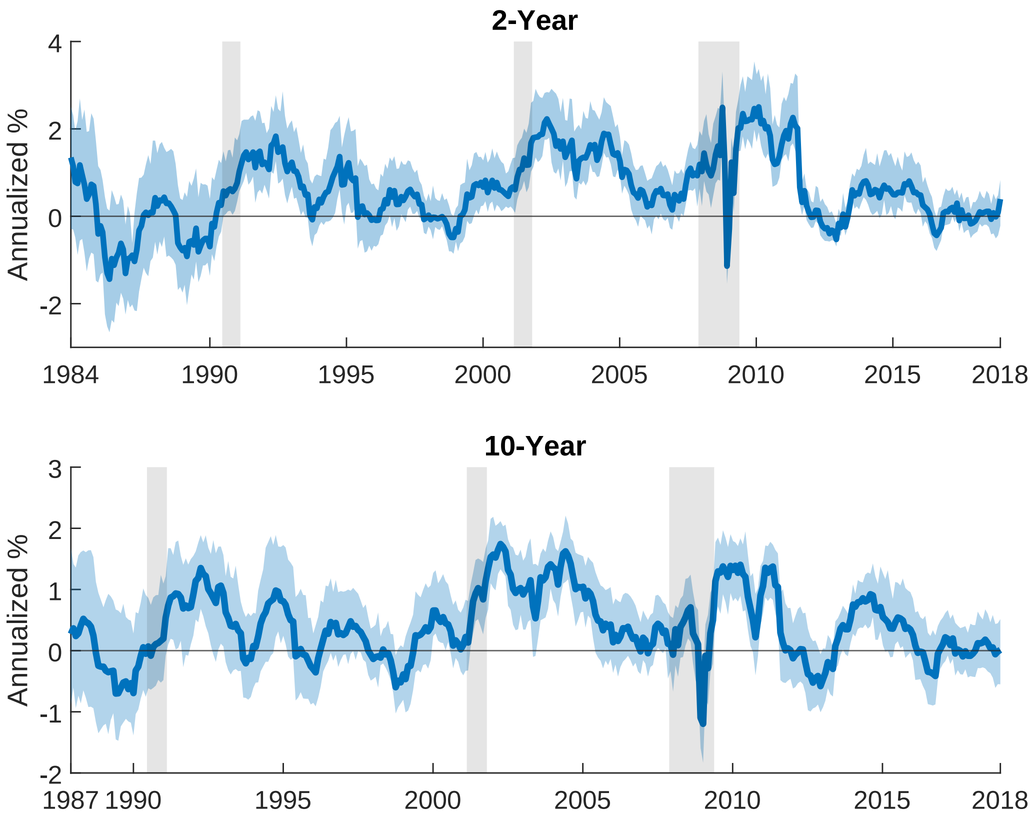

Figure 1.3: Forecasts Implied by Objective Risk Premia

The solid blue lines plot the difference between survey forecasts of 1-year-ahead n-maturity yields from BCFF, $\tilde{\mathbb{E}}_t y_{t+1}^n$, and counterfactual 1-year-ahead forecasts implied by objective risk premia:

$$\mathbb{E}_t y_{t+1}^n \equiv \frac{(n+1)y_t^{n+1} - y_t^1 - \mathbb{E}_t rx_{t+1}^{n+1}}{n}.$$

Objective risk premia on an n-maturity zero-coupon bond, $\mathbb{E}_t rx_{t+1}^n$, are fitted using the cycle factor from Cieslak and Povala (2015) and the macro factor from Ludvigson and Ng (2009). The blue shaded areas are bounded by the difference between top-/bottom-10 average forecasts (also from BCFF) and $\mathbb{E}_t y_{t+1}^n$. The upper panel uses 2-year yield (3-year yield before December 1987) and the bottom uses 10-year yield. The sample period is from December 1984 to December 2018 and December 1987 to December 2018 for the upper and bottom panels, respectively.

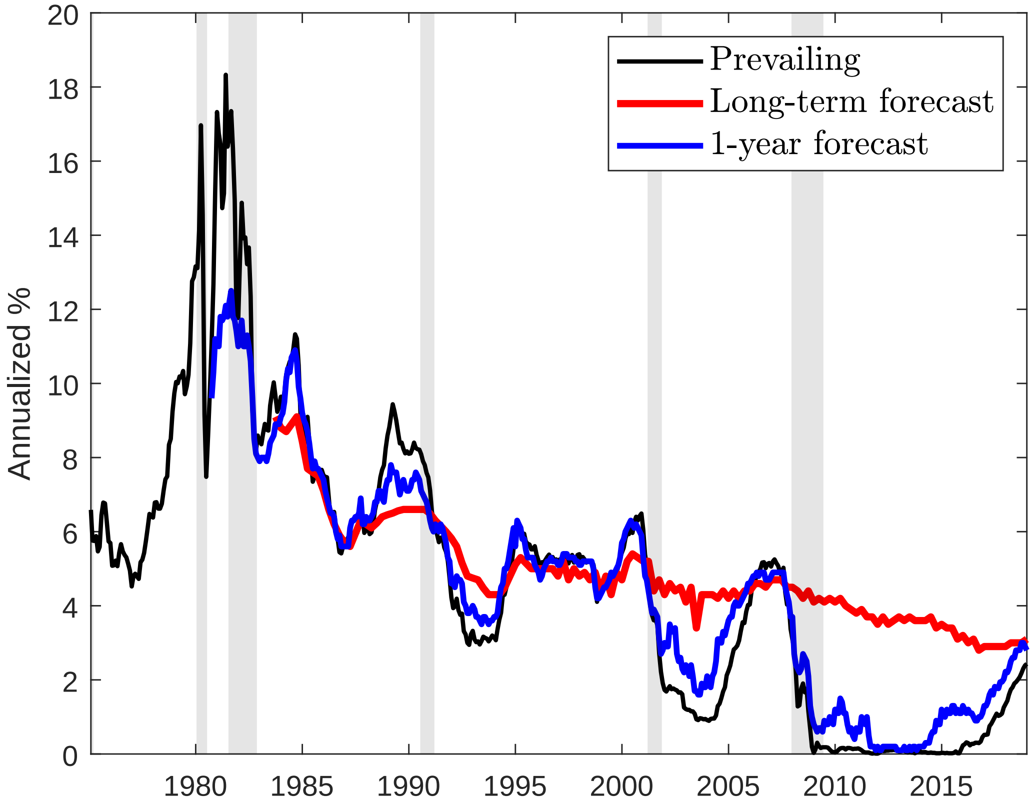

Figure 1.4: Short- and Long-term Forecasts of Short Rates

The figure plots the short- and long-term forecasts of 3-month Treasury bill (tbill) rates from BCFF and BCEI. The black line plots the prevailing 3-month tbill rates in the *survey* month (month of making forecasts). The red line plots forecasts of the 5-year averages for 7 to 11 years ahead. The blue line plots the 1-year-ahead forecasts. The sample period for long-term average forecasts is from 1983Q3 to 2018Q4. The sample period for 1-year-ahead forecasts is from September 1980 to December 2018.

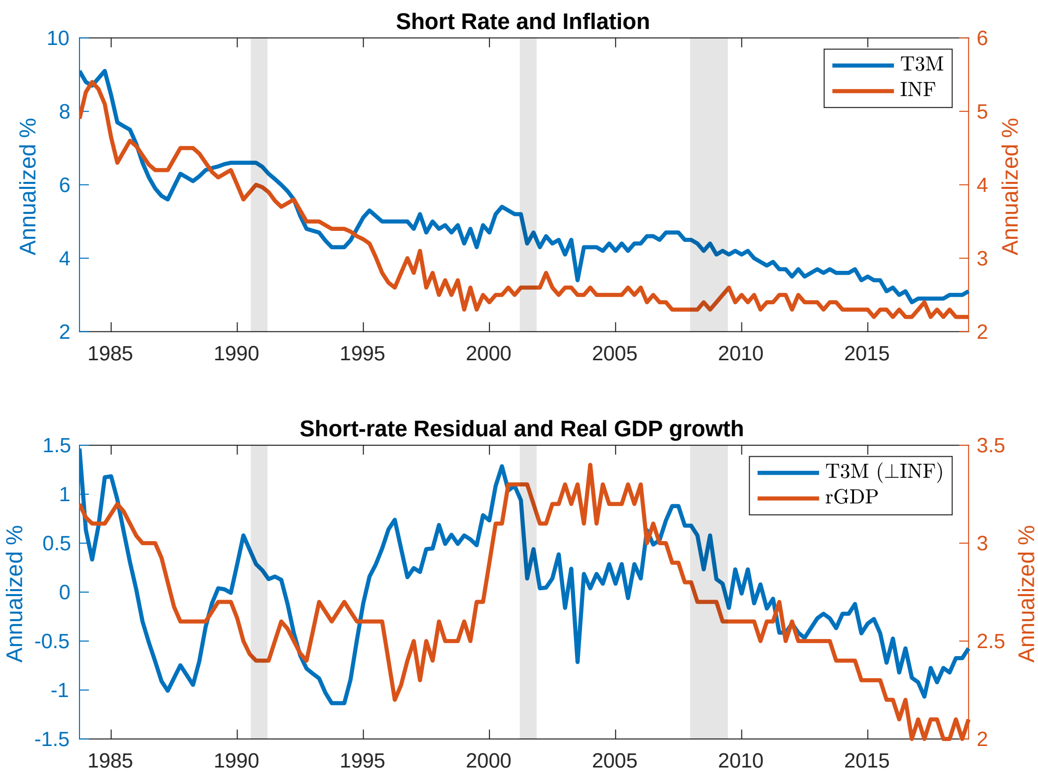

Figure 1.5: Long-term Forecasts of Short Rates and Macroeconomic Variables

In the top panel, the blue and red lines plot the long-term average (LT) forecasts of 3-month Treasury bill (tbill) rates and inflation, respectively. In the bottom panel, the red line plots the LT forecasts of real GDP growth. The blue line plots the residuals of regressing LT forecasts of tbill rates on LT forecasts of inflation. The sample period is from 1983Q3 to 2018Q4.

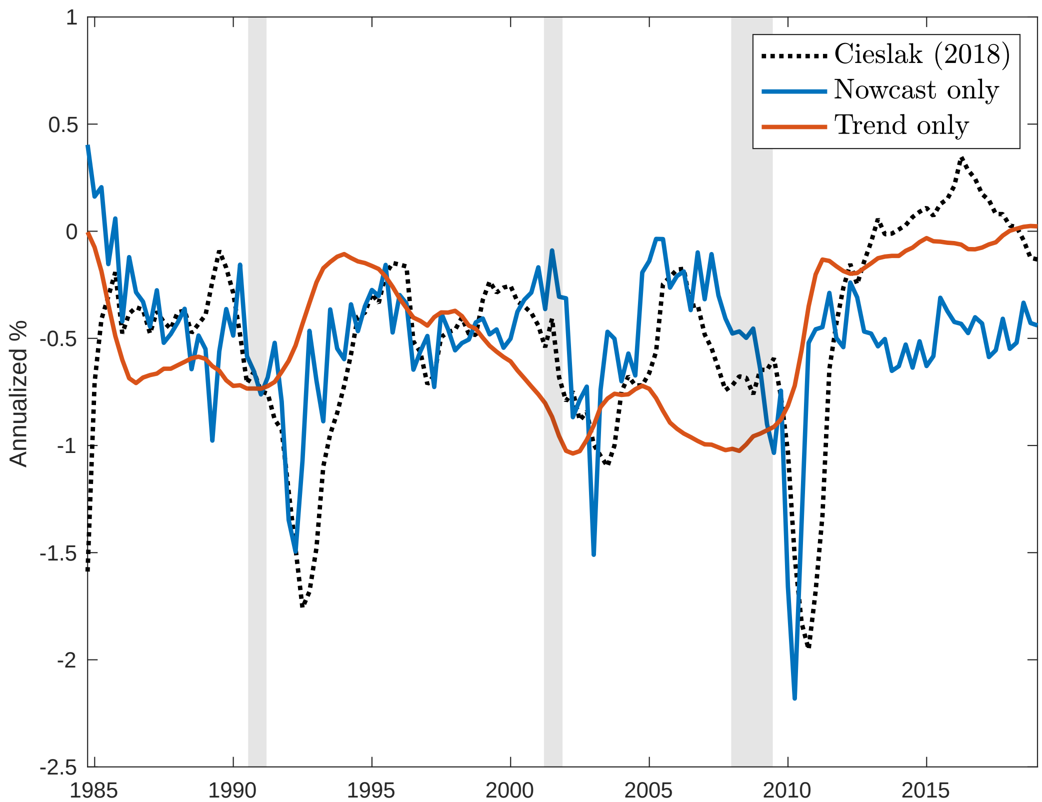

Figure 1.6: Fitted Short-rate Forecast Errors Using Macroeconomic Expectations

The figure plots the fitted values of regressing the forecast error of the 1-year-ahead 3-month Treasury bill (tbill) rate, $FE_t(i_{t+1})$, on subjective expectations of real GDP growth. The forecast error $FE_t(i_{t+1})$ is defined as the within-quarter average tbill rate realized 4 calendar quarters ahead of quarter t minus the 1-year-ahead quarter-average forecast directly taken from BCFF. The blue line uses nowcasts from SPF as the indepedent variable. The red line uses g^{LT} (the exponentially weighted average of past real GDP growth rates) as the independent variable. The black dotted line uses the lagged tbill rate and the year-on-year employment growth rate as independent variables following Cieslak (2018). The sample period is from 1983Q3 to 2018Q4.

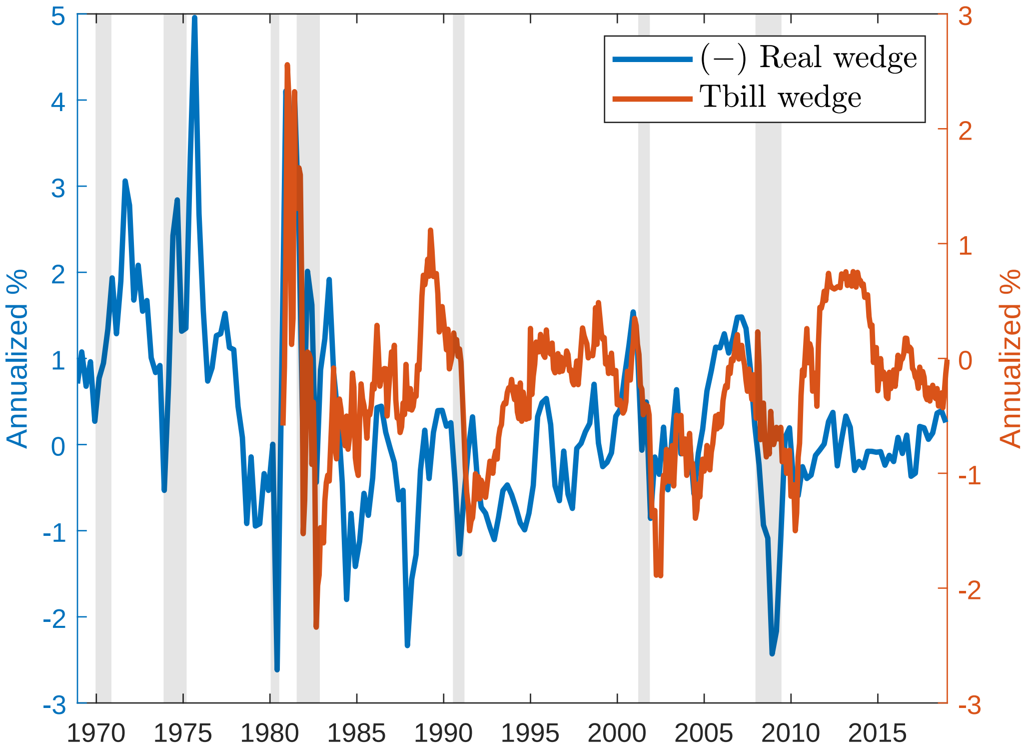

Figure 1.7: Short-rate and Macroeconomic Expectation Wedges

The figure plots the expectation wedges of 3-month Treasury bill (tbill) rate, $\mathbb{E}_t i^{\$}_{t+1} - \tilde{\mathbb{E}}_t i^{\$}_{t+1}$, and of real GDP growth, $\mathbb{E}_t g_{t+1} - \tilde{\mathbb{E}}_t g_{t+1}$ (taking the negative). $\mathbb{E}_t g_{t+1}$ is the fitted value from the following regression:

$$g_{t+1} = \gamma_0 + \gamma_1 g_t + \gamma_2 s_t^{\$,n} + \varepsilon_{t+1}.$$

$\tilde{\mathbb{E}}_t g_{t+1}$ is the average of 1- to 4-quarter-ahead survey forecasts of real GDP growth. The sample period is from 1968Q4 to 2018Q4. $\mathbb{E}_t i^{\$}_{t+1}$ is constructed from 1-year-ahead predictive regressions of within-quarter average 3-month tbill rates with lagged average rates and year-on-year employment growth as predictors, as in Cieslak (2018). $\tilde{\mathbb{E}}_t i^{\$}_{t+1}$ is the 1-year-ahead survey forecasts of 3-month tbill rates taken from BCEI. The sample period is from 1980Q3 to 2018Q4.

CHAPTER II

Asset Pricing with Fading Memory

2.1 Introduction

The predictably counter-cyclical nature of the equity risk premium continues to be a major challenge in asset pricing. Researchers have proposed rational expectations models that generate time-variation in the equity premium by introducing modifications into the representative agent's utility (Campbell and Cochrane 1999; Barberis et al. 2001) or by introducing persistence and stochastic volatility into the endowment growth process (Bansal and Yaron 2004). A key feature of these rational expectations models is that the representative agent knows the objective probability distribution she faces in equilibrium: subjective and objective expectations are the same. Therefore, the agent is fully aware of the counter-cyclical nature of the equity premium and knows the values of the parameters driving this process. This is a troubling feature of these models on two levels—conceptual and empirical. Conceptually, it is not clear how an agent could come to possess so much knowledge about parameters when econometricians struggle to estimate such parameters with much precision even from very long time-series samples (Hansen 2007). Empirically, surveys of investor return expectations from a number of sources fail to find evidence that investors' return expectations are counter-cyclical. If anything, the survey data indicate pro-cyclicality

(Vissing-Jorgensen 2004; Amromin and Sharpe 2014; Greenwood and Shleifer 2014).

Learning about the parameters of the asset payoff process offers a potential solution to both parts of this conundrum. Learning allows for subjective parameter uncertainty and a wedge between the subjective expectations of agents within the model and the objective expectations that an econometrician outside of the model could estimate from a large sample of data generated by the economy ex post. In such a learning model, the dynamics of subjective beliefs and asset prices depend crucially on the specification of agents' memory. The modeling of memory determines how experienced historical data feeds into beliefs, decisions, and prices. A standard assumption would be that agents retain full memory of all observations experienced in the past. But full memory is, in an application, neither well defined—does it start at the start of today's electronic records of stock prices and dividends? when the first stock market opened? some other time?—nor necessarily empirically plausible.

In this paper, we study asset prices when the representative agent learns with gradually fading memory. Our approach is grounded in micro-evidence from household portfolio choice and survey expectations data showing that individuals learn from experience—that is, their expectations are shaped by data realized during their lifetimes, and most strongly by recently experienced data (Malmendier and Nagel 2011; Malmendier and Nagel 2015).[1] Malmendier and Nagel (2015) further show that when individuals learn from life-time experiences, the dynamics of their *average* expectation can be approximated very closely by a constant-gain learning scheme in which a data point's influence on beliefs gradually fades over time as it recedes into the past. Abstracting from generational heterogeneity, we use this insight and consider a representative agent who learns with constant gain. This allows us to obtain a highly

[1] A growing body of evidence also suggests that extrapolation from experience helps understand the expectations and behavior of professionals. See, e.g., Greenwood and Nagel (2009) for mutual fund managers, Andonov and Rauh (2019) for pension fund return expectations, and Malmendier et al. (2019a) for inflation expectations of members of the Federal Reserve's Open Market Committee.

tractable model that nevertheless captures what may be, for the purposes of asset pricing, the essential aspect of learning from experience: the gradual loss of memory.

The decaying memory of observations in the past is the only modification to an otherwise standard Bayesian parameter learning model. As a consequence of the memory decay, learning is perpetual and there is a persistent time-varying wedge between the agent's subjective beliefs and the objective beliefs implied by the true parameters of the process generating the asset payoffs. Importantly, however, by retaining everything else from the standard Bayesian set up, we are able to analyze the asset pricing effects of posterior subjective uncertainty under constant-gain learning. Based on this approach, the dynamics of asset prices and survey expectations can be reconciled within a simple setting with IID endowment growth, recursive utility with constant risk aversion, and a representative agent who learns with fading memory about the mean endowment growth rate.

We start the analysis by documenting several new facts about equity market returns and subjective stock return expectations. We look at the data through the lens of a simple reduced-form framework that combines a present-value identity with constant-gain learning about the growth rate of dividends. Constant-gain learning implies that investors' subjective expectation of long-run dividend growth is equal to an exponentially weighted average of past dividend growth rates, which we label *experienced payout growth*. For the purpose of this initial analysis, we exogenously fix the subjective risk premium demanded by investors and the risk-free rate to be constant. Under objective expectations, then, the resulting equity premium is counter-cyclical. For example, after a string of positive growth innovations, investors are subjectively optimistic about the mean growth rate, the equity price is high, and subsequent returns are low because the investors' optimistic expectations are likely disappointed ex post.

How quickly memory of past realized growth rate observations decays is determined by a gain parameter in the belief-updating rule. This parameter plays a key role in our analysis as it determines the volatility and persistence of the price-dividend ratio and the strength of return predictability. We do not tweak this parameter to fit asset prices. Instead, we rely on the estimates in Malmendier and Nagel (2011, 2015) from survey data to pin down the value of this parameter at 0.018 for quarterly data. This means that the agent's posterior mean growth rate in the current quarter puts a weight of 0.018 on the most recent quarterly growth rate surprise and $1 - 0.018$ on the posterior mean from the prior quarter. Experienced payout growth is therefore a slow-moving variable. We construct it empirically with dividend data going back to the 19th century.

We use the experienced payout growth series to uncover three novel empirical facts that are consistent with the predictions of this constant-gain learning model. First, experienced payout growth is strongly negatively related to subsequent stock returns in excess of the risk-free rate. Remarkably, unlike most existing equity return predictors in the literature, this one does not use information on price levels, just information from the past history of asset payouts. Second, using data on individuals' subjective expectations of stock market excess returns, we find that they are basically unrelated to experienced payout growth. As a consequence, there is a wedge between subjective and objective expectations of excess returns that generates subjective expectations errors that are predictable by experienced payout growth. Third, subjective expectations of growth in fundamentals, proxied by long-term earnings forecasts of stock market analysts, are strongly positively related to experienced payout growth.

We then construct a structural asset-pricing model that matches these empirical regularities as well as other standard stylized asset pricing facts. A representative agent with recursive utility and constant risk aversion learns with fading memory about the mean of IID endowment growth. We keep the gain parameter fixed at the same value as in our reduced-form analysis. Equity is a levered claim to the endowment. Because of fading memory, learning is perpetual, the economy is stationary, and the agent's subjective uncertainty about long-run growth is high, which generates a high equity premium. The real risk-free rate varies slowly over time, but its volatility is low, consistent with empirical data.

The dynamics of risk premia in the model are in line with our empirical findings. Objectively, excess returns are predictable by experienced payout growth or the price-dividend ratio—without the subtle persistent components in the endowment process (and the agent's knowledge of these) in Bansal and Yaron (2004) (BY) or time-varying risk aversion built into the agent's preferences as in Campbell and Cochrane (1999) (CC). From the agent's subjective viewpoint, the world looks different. Subjective expected excess returns are virtually unrelated to experienced payout growth, consistent with the survey data. This wedge between subjective and objective expectations generates a strong negative relationship between subjective expectations errors and experienced payout growth. Thus, our model reconciles the evidence from returns data and surveys of investor return expectations.

Our model also addresses the Sharpe Ratio variability puzzle highlighted in Lettau and Ludvigson (2010). The leading rational expectations models by CC and BY imply that the conditional equity premium and conditional market return variance are almost perfectly positively correlated. In CC, this happens because at times when risk aversion is high, it is also very volatile. In BY, the reason is that stochastic volatility in endowment growth is the driver of the time-varying equity premium. As a consequence, variation in conditional market return variance dampens the variability of the Sharpe Ratio relative to the variability of the conditional equity premium. Empirically, however, the first and second moments of market excess returns conditioned on the price-dividend ratio or experienced payout growth are basically uncorrelated. Our model is consistent with this empirical observation: subjective belief dynamics generate predictable variation in the objective equity premium, but without simultaneous variation in the volatility of equity returns, which generates a volatile Sharpe Ratio.

The model also matches the lack of out-of-sample predictability found empirically in Welch and Goyal (2007). In our model simulations, standard out-of-sample tests show no out-of-sample stock return predictability for empirically realistic sample sizes, even though returns are truly predictable under the objective distribution. While an econometrician can find predictable returns in sample by studying data ex post, it would be difficult, even with full memory, for the econometrician to construct a viable trading strategy in real time that takes advantage of the fading-memory agent.

To understand the economic mechanism generating investor belief and asset price dynamics in the model, it is useful to contrast it with a model in which constant-gain learning is an optimal Bayesian approach for tracking a random-walk component of growth rates. For appropriately chosen volatility of the random-walk increments, this alternative model would produce the same dynamics of investor beliefs about long-run growth. Thus, one can think of investors in our model as forming beliefs under the *perception* that endowment growth has a random-walk component. However, for the model's predictions about predictability of stock returns and subjective forecast errors it is crucial that the *actual* endowment growth rate is IID. The resulting wedge between subjective and objective expectations is key and makes the model predictions very different from models in which constant-gain learning is optimal based on the actual law of motion.

This highlights an important difference relative to the large literature on constant-gain learning in macroeconomics (see Evans and Honkapohja (2001) for an overview). The typical motivation for constant-gain learning in this literature is structural change that renders long-distant historical data irrelevant, but it has been difficult to relate the gains that best explain survey forecasts of macroeconomic data to the gains that would be optimal given the degree of structural change in the forecasted time series (Branch and Evans 2006; Berardi and Galimberti 2017). In our case, fading memory reflects the fact that individuals rely on their life-time experiences when forming expectations, which is why we tie our gain parameter to earlier estimates from survey expectations microdata, not to the degree of structural change in the underlying time series. Our paper shares this focus on memory formation with a recent emerging literature on the topic in economics (e.g., Azeredo da Silveira and Woodford 2019; Bordalo et al. 2019b; and Kahana and Wachter 2019).

Our model builds on earlier work that has used different approaches to investigate the asset pricing implications of learning from experience. Collin-Dufresne et al. (2016a), Ehling et al. (2017), Malmendier et al. (2019b), Schraeder (2015), and Nakov and Nuño (2015) use an overlapping generations (OLG) approach to study learning-from-experience effects in asset pricing. The advantage of the OLG approach is that it maps very closely to the empirical work in Malmendier and Nagel (2011, 2015) that studied between-cohort heterogeneity in experiences, expectations, and choices. Moreover, these models produce interesting implications for cross-sectional heterogeneity in portfolio choices and wealth. The downside is that model solution requires simplifications that make a quantitative mapping to empirical data difficult.[2] By abstracting from cross-cohort heterogeneity, we also employ a simplified approach, but one that delivers quantitatively realistic asset-pricing predictions. Our model is highly tractable and should therefore be well suited for further extensions such as, for example, to study production and real investment.

Our model shares many elements with full-memory Bayesian parameter learning

[2]For example, Collin-Dufresne et al. (2016a) use two overlapping cohorts and objective and subjective risk premia jump every 20 years when there is a generational shift; Ehling et al. (2017) assume log utility, Schraeder (2015) and Malmendier et al. (2019b) work with CARA preferences in partial equilibrium with an exogenous risk-free rate, and the model in Nakov and Nuño (2015) has risk-neutral agents.

models, especially the IID-normal model in Collin-Dufresne et al. (2016b), but the fading memory feature avoids the arguably unrealistic implication of these models that learning effects disappear and risk premia decrease deterministically over time as the agent acquires more data (see, also, Timmermann (1993) and Lewellen and Shanken (2002) for partial equilibrium models with full memory). The fading memory model also avoids the problem of having to take a stand on what "year zero" is in an empirical implementation of a full-memory Bayesian learning model. And the gain parameter that determines memory can be pinned down based on microdata estimates.

Our model is also related to, but also in important ways different from recent models with extrapolative expectations. In Barberis et al. (2015) some investors extrapolate from stock price changes in recent years, which helps match the evidence in Greenwood and Shleifer (2014) that lagged stock market returns from the past few years are positively related to subjective expected returns. Jin and Sui (2019) build a representative agent model with return extrapolation. While these models can match the strong correlation of survey measures of subjective expected returns with lagged one-year stock market returns, they produce the counterfactual prediction that stock market excess returns are predictable by lagged one-year returns and that the price-dividend ratio quickly mean-reverts. The experienced growth measures in our setting put much greater weight on more distant observations in the past. As a consequence, the price-dividend ratio and objective expected returns in our model vary at a lower frequency, much close to the high persistence of the price-dividend ratio observed in the data. But our model cannot produce the correlation of one-year past returns with subjective expected returns. In Hirshleifer et al. (2015) and Choi and Mertens (2019), extrapolation occurs at low frequency, like in our model. A key difference is that in our model the agent perceives and prices subjective long-run growth rate uncertainty which allows us to generate a high equity premium in an IID economy. In Adam et al. (2017), agents know the expected growth rate of dividends, but they don't know the pricing function that maps expected fundamentals into prices, and they use an exponentially-weighted average of past price growth to forecast future prices. Matching the empirical equity risk premium in their model requires that subjective volatility of one-period ahead consumption growth far exceeds the actual volatility in the data. In our model, perceived short-run consumption volatility is very close to the objective volatility. The riskiness of equity in our model instead arises from subjective *long-run* growth rate uncertainty.

2.2 Facts about Subjective and Objective Expectations of Returns and Payoffs

Before looking at asset pricing with learning from experience within a structural asset-pricing framework, we start by laying out some empirical facts about stock market returns and investor return expectations from survey data that we want our asset-pricing model to match.

We consider a setting in which investors are learning about the mean growth rate μ_d of log real stock market payouts, d,

$$\Delta d_t = \mu_d + \epsilon_t, \tag{2.1}$$

where ϵ is an IID shock. The microdata evidence in Malmendier and Nagel (2011, 2015) suggests that individuals form expectations from data they observe throughout their lifetimes and, within their life-time data set, with more weight on relatively recent data. In our analysis, we focus on the dynamics of the average individual's expectation in such a learning-from-experience setting. Malmendier and Nagel (2015) show that if individuals in different birth cohorts learn from their life-time experience, their average belief is described very well by a constant-gain learning rule. Applied in our setting here, this means that the perceived growth rate $\tilde{\mu}_d$ evolves as

$$\tilde{\mu}_{d,t+1} = \tilde{\mu}_{d,t} + \nu(\Delta d_{t+1} - \tilde{\mu}_{d,t}), \tag{2.2}$$

and where ν is the (constant) gain parameter (see, e.g., Evans and Honkapohja (2001)).

As this expression shows, $\tilde{\mu}_d$ is updated every period based on the observed surprise $\Delta d_{t+1} - \tilde{\mu}_{d,t}$. How much this surprise shifts the growth rate expectation depends on ν. Malmendier and Nagel (2015) show that $\nu = 0.018$ for quarterly data fits the dynamics of the average belief in inflation expectations microdata (and this value is also within the range of estimates obtained from microdata on household investment decisions in Malmendier and Nagel (2011)). Iterating on (2.2) one can see that $\tilde{\mu}_{d,t}$ is an exponentially-weighted average of past Δd observations. In this way, the constant-gain updating scheme (2.2) captures the memory-loss implied by learning from experience and generational turnover. The more observations recede into the past, the lower the weight on these observations. In contrast, with full-memory Bayesian learning, the posterior mean would be formed by taking an equal-weighted average of all observed growth-rate realizations.

As a preliminary step, we explore some basic asset pricing implications when investors form expectations as in (2.2). At time t, they price stocks based on their growth rate expectation $\tilde{\mu}_{d,t}$. For now, we further assume that they price in a constant risk premium θ and a constant real risk-free rate r_f under their subjective beliefs. As we will show later, these assumptions are very close to the subjective belief dynamics that we obtain for a representative agent in a fully specified asset-pricing model with constant-gain learning and priced subjective uncertainty about long-run growth.

Now apply a Campbell and Shiller (1988) approximate present-value identity, used as in Campbell (1991) to decompose return innovations into changes in expectations about future growth rates and changes in return expectations. Under the investors' subjective expectations, denoted $\tilde{\mathbb{E}}[.]$, the innovation in stock returns is

$$r_{t+1} - \tilde{\mathbb{E}}_t r_{t+1} = (\tilde{\mathbb{E}}_{t+1} - \tilde{\mathbb{E}}_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} \tag{2.3}$$

$$= \frac{\rho}{1-\rho}(\tilde{\mu}_{d,t+1} - \tilde{\mu}_{d,t}) + \Delta d_{t+1} - \tilde{\mu}_{d,t} \tag{2.4}$$

$$= \left(1 + \frac{\rho\nu}{1-\rho}\right)(\Delta d_{t+1} - \tilde{\mu}_{d,t}). \tag{2.5}$$

Under the investors' subjective beliefs there is no term for the revision of return expectations because subjective return expectations stay fixed at $\theta + r_f$. Under these subjective beliefs, all variance of unexpected returns is due to revisions in forecasts of future cash flows. Adding investors' subjectively expected return we obtain total realized returns

$$r_{t+1} = \left(1 + \frac{\rho\nu}{1-\rho}\right)(\Delta d_{t+1} - \tilde{\mu}_{d,t}) + \theta + r_f. \tag{2.6}$$

Now consider an econometrician who knows (from a large sample of data) the true growth rate μ_d. Taking expectations of (2.6) under these objective beliefs yields

$$\mathbb{E}_t r_{t+1} - r_f = \theta + \left(1 + \frac{\rho\nu}{1-\rho}\right)(\mu_d - \tilde{\mu}_{d,t}), \tag{2.7}$$

where the term in parentheses times $\mu_d - \tilde{\mu}_{d,t}$ represents the subjective growth-rate expectations revision that the econometrician anticipates, on average, in the next period, given her knowledge of μ_d. This expression shows that the econometrician should find returns to be predictable. Specifically, $\tilde{\mu}_{d,t}$ should predict future excess returns negatively.

Moreover, while subjective excess return expectations are constant, the expectations error $\mathbb{E}_t r_{t+1} - \tilde{\mathbb{E}}_t r_{t+1}$ should be predictable by $\tilde{\mu}_{d,t}$. We can see this by subtracting the subjective equity premium $\tilde{\mathbb{E}}_t r_{t+1} - r_f = \theta$ from (2.7). We obtain

$$\mathbb{E}_t r_{t+1} - \tilde{\mathbb{E}}_t r_{t+1} = \left(1 + \frac{\rho\nu}{1-\rho}\right)(\mu_d - \tilde{\mu}_{d,t}). \tag{2.8}$$

In summary, this reduced-form analysis suggests three empirical relationships that we now investigate before moving on to explaining these within a structural asset-pricing framework: (i) excess returns should be predictable by $\tilde{\mu}_{d,t}$; (ii) subjective expectations errors should be predictable by $\tilde{\mu}_{d,t}$; and, most basically, (iii) subjective cash-flow growth expectations should be positively related to $\tilde{\mu}_{d,t}$.

2.2.1 Measurement of experienced growth

To estimate the relationship between a slow-moving predictor like $\tilde{\mu}_{d,t}$ and future returns in (2.7), we want to use the full history of returns back to the start of the CRSP database in 1926. And to implement the constant-gain learning scheme in (2.2), we need a long history of past observations on stock market fundamentals. For example, to compute $\tilde{\mu}_{d,t}$ in 1926, we then need data on stock market payout growth, Δd, stretching back at least an additional 50 years, up to the point where the weights become close to negligible.

From 1926 onwards, we obtain quarterly observations of aggregate Δd on the CRSP value-weighted index. We use a payout series, constructed as in Bansal et al. (2005), that includes repurchases in addition to dividends. Shifts from dividends to stock repurchases (e.g., motivated by tax changes) in the last few decades of the sample could otherwise distort the link between our payout measure and the stock market fundamentals that we want to proxy for. In the early decades in the sample, the role of repurchases is negligible. Prior to 1926, we use data on household dividend receipts from tax data in Piketty et al. (2017) for the period 1913 to 1926, and data on aggregate corporate non-farm non-financial dividends from Wright (2004) for the period 1900 to 1913. We deflate payout growth with CPI inflation rates and calculate per-capita real growth rates. For the period from 1871 to 1900 we use per-capita real GDP growth rates from Barro and Ursúa (2008) as proxy for Δd. Appendix C.1.1 provides more details on the construction of the Δd times series. We then use this series to calculate experienced payout growth based on the recursion in (2.2) and we label it $\tilde{\mu}_{d,t}$.

The experienced growth measure based on corporate payouts is likely imperfect. While the inclusion of repurchases may help alleviate distortions in the time-series properties caused by shifts in payout policies, some distortions likely remain. Moreover, the pre-1926 payout data is of lower quality than the CRSP data. For these reasons, we also construct an alternative measure, experienced real returns, $\tilde{\mu}_{r,t}$, that represents a weighted average of past log real stock market index returns. From the point it becomes available in 1926, we use quarterly returns on the CRSP value-weighted stock market index. Before that, we use data from Shiller (2005) back to 1871 to construct quarterly returns on the S&P Composite index up to 1926. We deflate returns with the CPI inflation series from Shiller (2005). For averages taken over long periods, real payout growth and real returns should be highly correlated and hence $\tilde{\mu}_{r,t}$ should capture similar information as $\tilde{\mu}_{d,t}$.[3] While experienced returns have some advantages as a measure of experienced fundamentals growth, there are also some potential shortcomings because asset price movements unrelated to fundamentals could contaminate $\tilde{\mu}_r$. For example, sentiment shocks that are orthogonal to fundamentals could cause

[3] To check this, we simulated dividend growth and returns from eqs. (2.1) and (2.5) in 1,000 samples of 360 quarters plus 400 quarters as a burn-in period to compute $\tilde{\mu}_d$ and $\tilde{\mu}_r$ at the start of the estimation sample. How well $\tilde{\mu}_{r,t}$ tracks $\tilde{\mu}_t$ depends on the gain parameter ν. For the value $\nu = 0.018$ that we work with here, the correlation is a very high 0.82. Other parameters like μ_d, the variance of ϵ, θ, r_f, or ρ do not influence this correlation. Thus, the approach of using $\tilde{\mu}_r$ to capture the time-series dynamics of $\tilde{\mu}_d$ should work well. We confirm this again below when we study $\tilde{\mu}_r$ in data simulated from our asset-pricing model.

movements in asset prices (that we accumulate in $\tilde{\mu}_r$) and future objective expected returns (a dependent variable in our regressions below). For this reason, we use both measures, $\tilde{\mu}_d$ and $\tilde{\mu}_r$, in our tests.

2.2.2 Return predictability

Table 2.1 presents predictive regressions along the lines suggested by (2.7). In Panel A, we use $\tilde{\mu}_d$ as a predictor, in Panel B we use $\tilde{\mu}_r$. Both predictors are constructed with data up to the end of quarter t. The dependent variable is the quarterly return on the CRSP value-weighted index in quarter $t+1$ in excess of the three-month T-bill yield at the end of quarter t. Based on eq. (2.7), the present-value model in (2.7) would predict an OLS slope coefficient of -2.78 (with $\nu = 0.018$ and $\rho = 0.99$, which is the quarterly value implied by the value of $\rho = 0.964$ for annual data reported in Campbell (2000)).

As Panel A shows, our estimates in the full 1927-2016 period are roughly in line with this predicted value. To account for small-sample biases in predictive regressions, we run bootstrap simulations as in Kothari and Shanken (1997) to compute a bias adjustment and a bootstrap p-value. Appendix C.2 provides details on these bootstrap simulations. With $\tilde{\mu}_d$ as the only predictor, we get an OLS coefficient estimate of -5.79. Bias-adjustment shrinks the coefficient only slightly to -5.71, which is bigger in magnitude than the predicted value of -2.78. Based on the bootstrapped p-value of < 0.01, we can reject the no-predictability null at high levels of confidence.

One potential issue with these regressions is that the experienced real growth variables could be distorted by recent unexpected inflation. If companies are sluggish to adjust nominal payout growth one-for-one with inflation, a burst of recently high inflation would temporarily depress the real experienced payout growth that we measure with our simple exponentially-weighted average, but not necessarily the real fundamentals growth that investors truly experience. For this reason, column (2) therefore adds the average log CPI inflation rate during quarters $t-3$ to t to the regression. The coefficient for experienced real payout growth gets somewhat more negative, but not by much.

Column (3) adds the log price-dividend ratio to the regression. As a test of the economic story that we propose here, adding the price-dividend ratio, or other fundamentals-price ratios, to the regression is not really meaningful. The price-dividend ratio should—following the usual present-value identity logic—pick up essentially any variation in objective expected returns, and so it should also absorb predictability associated with $\tilde{\mu}_d$. However, as a purely descriptive empirical matter, it is useful to know whether experienced dividend growth adds any forecasting power over and above the log price-dividend ratio, $p-d$. Column (3) suggests that it does. In fact, in the presence of experienced payout growth and inflation in the regression, $p-d$ is not a significant predictor and does not raise the R^2 compared with column (2).

Column (4) re-runs the regressions of column (2) for the post-World War II sample to address a potential concern that the results could be driven by the Great Depression period. The estimated slope coefficient of -2.99 is now very close to the predicted value of -2.78. Adding $p-d$ in column (5) takes away a substantial part of the predictability in this shorter sample. Of course, as we have noted above, $\tilde{\mu}_d$ or $p-d$ should capture the same information and the learning-from-experience theory does not imply that $\tilde{\mu}_d$ should necessarily have incremental predictive power over and above $p-d$.

Panel B uses experienced real returns, $\tilde{\mu}_r$, as a proxy for experienced stock market fundamentals growth. For the full sample in column (1), we obtain a bias-adjusted point estimate of -1.74 that is highly statistically significant. The remaining columns show that adding inflation and $p-d$ and focusing on the post-1945 sample slightly strengthens the predictive relationship between experienced real stock returns and future excess stock returns.

Figure 2.1 shows that experienced real payout growth and future excess returns are also strongly correlated at much longer prediction horizons. In this figure, we plot the predicted 5-year excess log return based on the bias-adjusted fitted values from column (1) in Panel A of Table 2.1, and iterating on it using the AR(1) dynamics of the experienced growth updating rule (2.2) with AR coefficient $1 - \nu = 0.982$. For comparison, we then plot the actual future 5-year cumulative excess log returns in quarters $t + 1$ to $t + 20$. As the figure shows, there is a strong positive correlation. Time periods in which predicted returns were low also tend to be periods when subsequent five-year excess returns were poor. For example, low experienced payout growth correctly forecasted high excess returns following the last three recessions in the early 90s, early 2000s, and in the financial crisis. That the cycles in expected excess returns line up so well with cycles in experienced growth is remarkable because we did not pick the ν parameter value—and hence the degree of smoothing implied by the weights—to match asset prices, but we fixed it at a value obtained from earlier microdata evidence. Moreover, unlike most return predictors in the literature, $\tilde{\mu}_d$ does not include price level information and hence it does not automatically pick up expected return shocks that affect prices.

Overall, the evidence indicates that experienced growth of stock market fundamentals—whether proxied for by experienced real payout growth or experienced real returns—is strongly predictive of stock market excess returns, consistent with a model in which investors use experienced growth to forecast future growth in fundamentals.

2.2.3 Subjective expectation error predictability

Subjective belief dynamics are a key feature of the economic effects we explore in this paper. For this reason, we want to confront our model with data on individual investor return expectations from surveys. We need a relatively long time series of survey expectations because the experienced growth variables that we focus on to explain dynamics change only slowly over time. For this reason, we put together survey data from several sources that spans the period 1972 to 1977 and 1987 to 2016. We focus on surveys that target a representative sample of the U.S. population, supplemented with two surveys of brokerage and investment firm customers.

Several of the surveys in our data elicit respondents' expected stock market returns, in percent, over a one-year horizon (UBS/Gallup survey, 1998-2007, monthly; Vanguard Research Initiative survey of Vanguard customers in Ameriks et al. (2016), one survey in 2014; surveys of Lease et al. (1974) and Lewellen et al. (1977), annual, 1972 and 1973). To extend these series, we bring in data from three additional surveys that don't elicit the percentage expected return but ask respondents to provide the probability of a rise in the stock market over a one-year horizon (Michigan Survey of Consumers, monthly 2002-2016) or the categorial opinion whether they expect stock prices to rise, or stay about where they are, or decline over the next year (Conference Board Survey, monthly 1987-2016;[4] Roper Center Surveys, annual, 1974-1977). We impute a time-series of implied percentage expected return from these alternative series. Roughly, the approach involves projecting the average expected returns each period from the first set of surveys on the coarser expectations measures in the second set of surveys, using periods of overlapping coverage to estimate the projection. Appendix C.1.2 provides more detail.[5] We average the expectations within calendar quarters to obtain

[4]The data was kindly provided by The Conference Board.

[5]Greenwood and Shleifer (2014) use two different data sources to cover time periods prior to the 1990s. From the mid-1980s onwards, they use the American Association of Individual Investors Investor (AAII) Sentiment Survey. The AAII survey is conducted among members of the AAII and it records responses of members that self-select into participation. Respondents state whether they are "bullish" or "bearish" about the stock market. We prefer the Conference Board survey for this time period as it is based on a representative sample of the U.S. population. For the early part of their sample starting in the 1960s, Greenwood and Shleifer use the Investors' Intelligence newsletter sentiment. For consistency over time, we prefer to stick to individual investor surveys in all time periods. The Roper and Lewellen et al. surveys give us at least partial coverage of the 1970s.

a quarterly series.

We start the analysis by looking at the time-series relationship between experienced growth and subjective expected excess returns. In the present-value model we have sketched above, the level of asset prices is affected by the experience-driven optimism or pessimism of investors. But the subjective expected excess return on the stock market, $\tilde{\mathbb{E}}_t r_{t+1} - r_f$, is constant. At each point in time, assets are priced such that subjective expected returns equal the (constant) equity premium required by investors. This will also be approximately true in our full model below, though not exactly.

As Panel A of Table 2.2 shows, subjective expected excess returns also seem to be approximately constant empirically at the frequencies that are relevant for our theory. In this table, we show the results of regressions of one-year excess return expectations in quarter t on experienced real payout growth or experienced real returns up to the end of quarter $t-1$. We calculate subjective expected excess returns by subtracting the average one-year Treasury yields measured at the beginning of the survey months within each quarter. As Panel A shows there is only a weak, and statistically not significant, positive relationship between $\tilde{\mu}_d$ and subjectively expected excess returns. Column (4) repeats this analysis with experienced real returns as the key explanatory variable. The results are very similar to those in column (1).

Looking at past returns over a much shorter time window, Greenwood and Shleifer (2014) find that survey return expectations are positively related to returns. As column (2) shows, we also find this in our data (which partly overlaps with Greenwood and Shleifer's) when we introduce the past 12-month return on the CRSP value-weighted index as an explanatory variable. The estimated coefficient on this lagged return is about three standard errors bigger than zero and the R^2 is substantially higher than in column (1). Column (3) and (5) show that when experienced growth variables and lagged returns are used jointly, the experienced growth effect remains very weak.

The important take-away is that in terms of the lower frequency movements that are captured by the experienced growth variables and that we focus on in our analysis, the subjective equity premium in the survey data is close to acyclical. That there are short-run fluctuations in subjective return expectations with one-year lagged returns is also interesting, but this is not a fact that we try to explain in this paper.

We now turn to the prediction, based on equation (2.7), that $\tilde{\mu}_d$ and $\tilde{\mu}_r$ should predict expectation errors. We calculate the expectation error $r_{t+1} - \tilde{\mathbb{E}}_t r_{t+1}$ on the left-hand side of (2.7) by subtracting the one-year survey expected return from the realized one-year return from the beginning of quarter $t+1$ to the end of quarter $t+4$. The fact that survey expectations in Panel A are unrelated to experienced growth combined with the fact in Table 2.1 that future returns are negatively related to experienced growth implies that the expectations error should be negatively related to experienced growth. However, since the survey data is restricted to the 1970s and 1987-2016, the samples in Table 2.1 and 2.2 cover very different periods. For this reason, it is still useful to check whether there is actually a negative relationship in the part of the sample in which survey data is available.

Panel B of Table 2.2 shows that this is the case. There is a strong negative relationship between $\tilde{\mu}_d$ or $\tilde{\mu}_r$ and the expectations error. When experienced growth is high, return expectations are predictably too optimistic. Since the prediction horizon is one year rather than the one-quarter horizon in the return prediction regressions in Table 2.1, the coefficient that we would expect, if these relations are stable across samples, is about four times the coefficient in Table 2.1. The results in Panel B show that this is approximately true.

2.2.4 Subjective long-run growth expectations

In the last part of our empirical analysis, we look at proxies for subjective payout growth expectations. The surveys that provide subjective return expectations unfortunately do not elicit respondents' views about future cash-flow growth. For this reason, we follow Sharpe (2002) and use financial analysts' long-term earnings growth forecasts aggregated at the market level as a proxy for investors' subjective payout growth expectations.[6] We start with analysts' monthly median forecasts of long-run earnings per-share (EPS) growth. According to I/B/E/S, they represent forecasts over a horizon of between three to five years.

We aggregate across stocks by forming a value-weighted average each month, using current-year forecasted total earnings for each stock as weights. We then average the monthly aggregated growth rate forecasts within each quarter, which yields a time series from quarter 1984:3 to 2016:4. We construct expected long-term real EPS growth, expressed in terms of per-quarter growth rates, by subtracting CPI inflation expectations from the Survey of Professional Forecasters issued in the quarter prior to the issuance of the analyst forecast.[7] Appendix C.1.3 provides more detail on the procedure to construct these measures.

According to the constant-gain updating scheme in (2.2), subjective long-term growth expectations should be positively related to experienced payout growth. This is what we find. Figure 2.2 provides a visual impression of the positive co-movement of the two series. Column (1) in Table 2.3 shows the results of regressions of analysts' long-term expectations in quarter t on experienced payout growth constructed from

[6] La Porta (1996) and Bordalo et al. (2019a) also use analyst long-term earnings growth expectations to measure investor expectations of future stock fundamentals, but they focus on a cross-sectional analysis while here we analyze time-series relationships at the aggregate level.

[7] We use the SPF forecast from the prior quarter to make sure analysts have access to this information. Stock analysts are presumably not experts in inflation forecasting and rely on other forecasters for macro forecast inputs.

realized growth rates up to quarter $t-1$. We find that there is a positive relationship with a coefficient of 0.30 that is statistically significant at conventional levels (s.e. 0.15).

While the results in column (1) are qualitatively consistent with constant-gain learning from payout growth, the coefficient estimate is lower than predicted. According to the constant-gain updating scheme in (2.2), subjective long-term growth rate expectations should be equal to experienced payout growth and hence the coefficient should be unity. As we will show below, in our full asset-pricing model, where dividends have some tendency to mean-revert, we actually expect a coefficient slightly lower than unity (about 0.89 with $\tilde{\mu}_d$ as predictor and 0.68 with $\tilde{\mu}_r$), but still substantially higher than the point estimate in column (1). There is one complication, however, in mapping the constant-gain learning scheme to the data. The setup in (2.1) and (2.2) is based on growth rates that are perceived as IID. In contrast, empirically, aggregate earnings changes have a sizable transitory component at short horizons of one to four quarters that is predictable by past earnings (Kothari et al. 2006) and by stock returns in recent quarters (Sadka and Sadka 2009). Analysts' aggregate short-term earnings forecasts vary over time to capture such short-horizon earnings growth predictability (Choi et al. 2016). If analysts' long-term forecasts, g_t^{LT}, are also influenced to some extent by such short-run predictability that is absent from our constant-gain learning framework, this could distort our estimates.[8] To be more precise, suppose

$$g_t^{LT} = (1-\beta)g_t^{\infty} + \beta g_t^{ST}, \tag{2.9}$$

where g_t^{∞} is the unobserved long-run forecast undistorted by short-run expectations,

[8] Whether long-term forecasts are influenced by these short-term components is an open question. According to I/B/E/S, analysts should report long-term forecasts that represent expected growth "over the company's next full business cycle" (Reuters 2009), but it is not clear to what extent analysts are actually implementing this guidance.

which receives weight $1-\beta$, and g_t^{ST} is the short-run forecast, which receives weight β. Our constant-gain learning framework suggests $g_t^{\infty} = \tilde{\mu}_{d,t}$, but if g_t^{ST} is correlated with $\tilde{\mu}_{d,t}$, a regression of g_t^{LT} on $\tilde{\mu}_{d,t}$ will yield biased estimates.

For this reason, column (2) in Table 2.3 adds the aggregate one-year forecast as a control variable to the regression. The one-year forecast represents the expected real growth rate from current fiscal year to next fiscal year EPS, aggregated in the same way as we described for long-term growth expectations, and similarly expressed in terms of per-quarter growth rates. By controlling for the one-year forecast, we ask whether experienced payout growth explains long-term growth expectations holding fixed the potentially distorting short-term forecast components. In this specification, we obtain a much higher coefficient of 0.61 (s.e. 0.18) for the experienced payout growth variable and the one-year forecast obtains a coefficient of 0.28 (s.e. 0.09) which suggests that the long-term forecast puts a plausible weight of around one quarter on the short-term forecast. The implied coefficient in a regression of g_t^{∞} on $\tilde{\mu}_{d,t}$ is $0.61/(1-0.28) = 0.85$, which suggests that the results are very close to the predictions of the constant-gain learning model once one controls for the distorting effect of short-horizon predictability.

Column (3) adds the most recent one-year lagged return to the regression, which has little effect. Columns (4) to (6) repeat the analysis with experienced real returns in place of experienced payout growth and the results are very similar. In column (5), we obtain an implied coefficient in a regression of g_t^{∞} on $\tilde{\mu}_{r,t}$ of $0.52/(1-0.05) = 0.64$. Overall, the evidence on subjective expectations of growth in stock market fundamentals is broadly in line with the predictions of the constant-gain learning scheme.

2.3 Asset Pricing Model

We now develop these ideas more fully in a representative-agent endowment economy.

2.3.1 Learning with fading memory

Endowment growth follows an IID law of motion

$$\Delta c_{t+1} = \mu + \sigma \varepsilon_{t+1}, \tag{2.10}$$

where $\{\varepsilon_t\}$ is a series of IID standard normal shocks. The agent knows that Δc_{t+1} is IID, and she also knows σ, but not μ. The agent relies on the history of past endowment growth realizations, $H_t \equiv \{\Delta c_t, \Delta c_{t-1}, \dots\}$, to form an estimate of μ.

We assume that the agent learns with constant gain and hence fading memory. Unlike in standard constant-gain learning models, however, we retain the modeling of the full posterior distribution—and hence the agent's subjective uncertainty—of the Bayesian approach. To do so, we use a weighted likelihood that has been used in the theoretical biology literature to model memory decay in organisms (Mangel 1990). An agent who has seen an infinite history of observations on Δc, but with fading memory, forms a posterior

$$p(\mu|H_t) \propto p(\mu) \prod_{j=0}^{\infty} \left[\exp\left(-\frac{(\Delta c_{t-j} - \mu)^2}{2\sigma^2} \right) \right]^{(1-\nu)^j}, \tag{2.11}$$

where $1 - \nu$ is a positive number close to one ($\nu = 0$ is the standard full memory Bayesian case where observations are equally weighted). Thus, $(1 - \nu)^j$ represents a (geometric) weight on each observation. This weighting scheme assigns smaller weights the more the observation recedes into the past. With a prior

$$\mu \sim \mathcal{N}(\mu_0, \sigma_0^2) \tag{2.12}$$

held before seeing any data we then obtain the posterior

$$\mu|H_t \sim \mathcal{N}\left(\frac{\sigma_0^2}{\nu\sigma^2+\sigma_0^2}\tilde{\mu}_t + \frac{\nu\sigma^2}{\nu\sigma^2+\sigma_0^2}\mu_0, \left(\frac{1}{\sigma_0^2}+\frac{1}{\nu\sigma^2}\right)^{-1}\right), \tag{2.13}$$

where

$$\tilde{\mu}_t = \nu\sum_{j=0}^{\infty}(1-\nu)^j\Delta c_{t-j}. \tag{2.14}$$

The variance of the posterior is the same as if the agent had observed, and retained fully in memory with equal weights, $S \equiv 1/\nu$ realized growth rate observations. In our case, the *actual* number of observed realizations is infinite, but the loss of memory implies that the *effective* sample size is finite and equal to S.

Due to the limited effective sample size, the prior beliefs retain influence on the posterior. For now, however, we work with an uninformative prior ($\sigma_0 \to \infty$) and hence the posterior

$$\mu|H_t \sim \mathcal{N}(\tilde{\mu}_t, \nu\sigma^2). \tag{2.15}$$

We will return to the informative prior case when we consider versions of the model that generalize our baseline assumption about the elasticity of intertemporal substitution. With this uninformative prior, the posterior mean is updated recursively as

$$\tilde{\mu}_t = \tilde{\mu}_{t-1} + \nu(\Delta c_t - \tilde{\mu}_{t-1}). \tag{2.16}$$

Thus, the $\tilde{\mu}_t$ resulting from this weighted-likelihood approach with an uninformative prior is identical to the perceived μ that one obtains from the constant-gain updating scheme (2.2) with gain ν. However, in contrast to standard constant-gain learning specifications in macroeconomics that focus purely on the first moment, we obtain a full posterior distribution. For the purpose of asset pricing, the subjective uncertainty implied by the posterior distribution can be crucial.

We further get the predictive distribution

$$\Delta c_{t+j}|H_t \sim \mathcal{N}\left(\tilde{\mu}_t, (1+\nu)\sigma^2\right), \qquad j = 1, 2, ..., \tag{2.17}$$

where the variance of the predictive distribution reflects not only the uncertainty due to future ϵ_{t+j} shocks, but also the uncertainty about μ. We denote expectations under the predictive distribution with $\tilde{\mathbb{E}}_t[\cdot]$. To understand better how the stochastic nature of the endowment process looks like from the agent's subjective viewpoint, we can rewrite (2.16) as

$$\tilde{\mu}_{t+1} = \tilde{\mu}_t + \nu\sigma\sqrt{1+\nu}\tilde{\varepsilon}_{t+1}, \qquad \text{where} \quad \tilde{\varepsilon}_{t+1} = \frac{\Delta c_{t+1} - \tilde{\mu}_t}{\sigma\sqrt{1+\nu}}, \tag{2.18}$$

and $\tilde{\varepsilon}_{t+1}$ is $\mathcal{N}(0,1)$ distributed and hence unpredictable under the agent's time-t predictive distribution.

Under Bayesian learning with full memory, the agent's information would be represented by a filtration and posterior beliefs would follow a martingale under this filtration. With fading memory, however, the posterior in periods $t+j > t$ will be formed based on information that is different, but not more informative about μ than the information available to the agent at time t. Hence, the information structure is not a filtration. For this reason, the time-t agent anticipates that $\tilde{\mu}_{t+j}$ in the future may vary from period to period, but she knows that this variation will be stationary and there is no convergence to μ in the long run. Consistent with stationarity, the time-t agent perceives future increments $\tilde{\varepsilon}_{t+j}$, $j = 1, 2, ...$ in (2.18) as negatively serially correlated (see Appendix C.3 for more details) and not as martingale differences.[9]

[9] At time t, the agent however cannot make use of this serial correlation by using $\tilde{\varepsilon}_t$ to forecast $\tilde{\varepsilon}_{t+1}$, because $\tilde{\varepsilon}_t$ is not observable to the agent. To figure it out, the agent would need full memory to compare $\tilde{\mu}_t$ with $\tilde{\mu}_{t-1}$, but under constant-gain learning this is not possible. As a consequence, $\tilde{\mathbb{E}}_t[\tilde{\mu}_{t+1}] = \tilde{\mu}_t$ remains the agent's posterior mean.

When we calculate asset prices in this model, we therefore cannot directly rely on certain laws like the law of iterated expectations (LIE) that require a filtration or results that presume martingale posteriors. This requires special care in evaluating valuation equations.

2.3.2 Valuation

The valuation approach we use throughout the paper is a "resale" valuation approach. To illustrate, consider the valuation at date t of a claim to consumption at date $t+2$. Under resale valuation, the agent at t prices the asset under the time-t predictive distribution of the stochastically discounted $t+1$ asset value,

$$P_{R,t} = \tilde{\mathbb{E}}_t \left[M_{t+1|t} \tilde{\mathbb{E}}_{t+1} \left[M_{t+2|t+1} C_{t+2} \right] \right], \tag{2.19}$$

where we use $M_{t+j|t}$ to denote the one-period stochastic discount factor (SDF) from $t+j-1$ to $t+j$ that applies given the agent's predictive distribution at t.

An alternative way of valuing this claim would be a "buy-and-hold" valuation, where the agent values the asset based on the stochastically discounted payoff under the time-t predictive distribution

$$P_{H,t} = \tilde{\mathbb{E}}_t [M_{t+1|t} M_{t+2|t} C_{t+2}], \tag{2.20}$$

In a full-memory Bayesian setting, the LIE would apply in the valuation equation of $P_{R,t}$ with the result that $P_{R,t} = P_{H,t}$, but with fading memory the information structure is not a filtration and the LIE typically fails to hold.[10] As a consequence $P_{R,t} \neq P_{H,t}$.

The valuation discrepancy between the two valuation approaches arises because the

[10]For subjective expectations of linear functions of Δc, we still get an LIE, e.g., $\tilde{\mathbb{E}}_t \tilde{\mathbb{E}}_{t+1} \Delta c_{t+2} = \tilde{\mathbb{E}}_t \Delta c_{t+2}$, but the LIE does not hold for nonlinear functions of Δc, e.g., $\tilde{\mathbb{E}}_t \tilde{\mathbb{E}}_{t+1} [\exp(a + b\Delta c_{t+2})] \neq \tilde{\mathbb{E}}_t [\exp(a + b\Delta c_{t+2})]$.

agent at t and at $t+1$ sees the statistical properties of the shock $\tilde{\varepsilon}_{t+2}$ differently. The buy-and-hold valuation incorporates the negative serial correlation of $\tilde{\varepsilon}_{t+1}$ and $\tilde{\varepsilon}_{t+2}$. In contrast, the resale valuation at t is based on the anticipation that the value of the asset at date $t+1$ will be determined by an agent—or a future self of the agent—who perceives $\tilde{\varepsilon}_{t+2}$ as unpredictable. Thus, the resale valuation is based on a chain of valuations that each views the one-period ahead $\tilde{\varepsilon}$ shock as unpredictable.

We work with the resale valuation approach below, for two reasons. First, the resale valuation is time-consistent. In contrast, if the asset was priced at time t at the buy-and-hold valuation and the anticipation of a predictable $\tilde{\varepsilon}_{t+2}$, and time moves on to $t+1$, the agent would, after memory loss, suddenly find $\tilde{\varepsilon}_{t+2}$ unpredictable. Thus, the agent would then agree with a valuation based on an unpredictable $\tilde{\varepsilon}_{t+2}$, but this is not consistent with the buy-and-hold valuation at t. Second, the resale valuation also fits with the underlying motivation of our model as an approximation for experience-based learning in an overlapping generations model in which actual resale would occur when generations turn over.

2.3.3 Kalman filtering interpretation

The updating scheme in (2.16) is reminiscent of optimal filtering in the case of a latent stochastic trend. Indeed, if the agent *perceived* μ to follow a random walk—i.e., as $\mu_t = \mu_{t-1} + \zeta_t$, where ζ is an IID-normal shock—rather than a constant as in (2.10), application of the steady-state Kalman filter yields exactly the same posterior distribution as in (2.15). With appropriate choice of the volatility of the ζ shocks, and as long as the *actual* law of motion is still (2.10) with constant μ, the dynamics of the posterior beliefs would be the same as in our fading memory model (see Appendix C.4 for more details). The predictive distribution of one-period ahead endowment growth, and, as a consequence, asset prices under resale valuation would be the same as well.[11] This is useful for technical purposes because it allows us to map our framework into one in which the information structure is a filtration and it is Markovian. Through this mapping, we can use results from Hansen and Scheinkman (2012) to determine parameter restrictions that are sufficient for existence of equilibrium. While this reinterpretation is convenient for technical reasons, optimal Kalman filtering does not explain the micro-evidence on learning from experience that motivates our fading memory approach. For this reason, we stick to the fading memory interpretation in the discussion of our model.

2.3.4 Stochastic discount factor

We assume that the representative agent evaluates payoffs under recursive utility as in Epstein and Zin (1989), with value function

$$V_t = \left[(1-\delta)C_t^{1-\frac{1}{\psi}} + \delta\tilde{\mathbb{E}}_t[V_{t+1}^{1-\gamma}]^{\frac{1-\frac{1}{\psi}}{1-\gamma}}\right]^{\frac{1}{1-\frac{1}{\psi}}}, \tag{2.21}$$

where δ denotes the time discount factor, γ relative risk aversion for static gambles and ψ the elasticity of intertemporal substitution (EIS). Note that the agent evaluates the continuation value under her subjective expectations $\tilde{\mathbb{E}}_t[.]$. We apply the same resale valuation approach that we use for asset pricing to this continuation value as well.

In our baseline model, we set $\psi = 1$. Iterating on the value function as in Hansen et al. (2008), but here under the agent's predictive distribution, we then obtain the log SDF that prices assets under the agent's subjective beliefs,

$$m_{t+1|t} = \tilde{\mu}_m - \tilde{\mu}_t - \xi\sigma\tilde{\varepsilon}_{t+1}, \tag{2.22}$$

[11] However, for the asset-pricing predictions to remain the same, it is crucial that (2.10), with constant μ, remains the *actual* law of motion. As we will show below, the time-varying wedge $\tilde{\mu}_t - \mu$ between subjective and objective beliefs plays an important role in generating volatile asset prices and predictable excess returns. Without this wedge, there would be no return predictability.

with

$$\tilde{\mu}_m = \log\delta - \frac{1}{2}(1-\gamma)^2(\nu U_v + 1)^2(1+\nu)\sigma^2, \tag{2.23}$$

$$\xi = [1-(1-\gamma)(\nu U_v + 1)]\sqrt{1+\nu}, \tag{2.24}$$

$$U_v = \frac{\delta}{1-\delta}. \tag{2.25}$$

Details are in Appendix C.5.1. This SDF implies the risk-free rate

$$r_{f,t} = -\tilde{\mu}_m + \tilde{\mu}_t - \frac{1}{2}\xi^2\sigma^2. \tag{2.26}$$

2.3.5 Pricing the consumption claim

We can now solve for the consumption-wealth ratio, $\zeta \equiv W_t/C_t$, and the subjective risk premium for the consumption claim. Using (2.18), we can express the log of the return on the consumption claim as

$$r_{w,t+1} = \tilde{\mu}_t + \sqrt{1+\nu}\sigma\tilde{\varepsilon}_{t+1} + \log\left(\frac{\zeta}{\zeta-1}\right). \tag{2.27}$$

Applying the subjective pricing equation $\tilde{\mathbb{E}}_t[M_{t+1|t}R_{W,t+1}] = 1$, we can solve for the wealth-consumption ratio

$$\zeta = \frac{1}{1-\delta}. \tag{2.28}$$

Thus, as in the rational expectations case, $\psi = 1$ implies a constant and finite consumption-wealth ratio. In the posterior distribution in (2.15), extremely large values of μ have greater than zero probability. The agent therefore also assigns some probability mass to extremely large future $\tilde{\mu}_{t+j}$. However, since $\psi = 1$ implies that $r_{f,t+j}$ moves one-for-one with $\tilde{\mu}_{t+j}$, the effect of high subjectively expected growth rates on the value of the consumption claim is exactly offset by a high future risk-free rate. As a consequence, the wealth-consumption ratio is constant and finite.

Evaluating the subjective pricing equation for $R_{W,t+1}$ again, now using the fact that ζ is constant, we can solve for the subjective risk premium of the consumption claim

$$\log \tilde{\mathbb{E}}_t[R_{w,t+1}] - r_{f,t} = \xi\sqrt{1+\nu}\sigma^2, \tag{2.29}$$

which is constant over time. In contrast, the objective risk premium under the econometrician's measure, generated by data sampled from this economy, is time-varying: taking the objective and subjective expectations and variance of (2.27), we can calculate the wedge between subjective and objective expectations, and combining with (2.29), we obtain the objective risk premium

$$\log \mathbb{E}_t[R_{w,t+1}] - r_{f,t} = \xi\sqrt{1+\nu}\sigma^2 - \frac{1}{2}\nu\sigma^2 + \mu - \tilde{\mu}_t, \tag{2.30}$$

where the time-varying wedge $\mu - \tilde{\mu}_t$ reflects the disagreement between the econometrician and the agent about the conditional expectation of $r_{w,t+1}$. The wedge is observable to the econometrician who knows μ, but since $\tilde{\mathbb{E}}_t[\mu] = \tilde{\mu}_t$ the wedge is zero from the viewpoint of the agent at time t.

2.3.6 Pricing the dividend claim

We now turn to pricing the dividend claim, which is the main focus of our analysis. Dividends in our model are a levered claim to the endowment. We assume that dividends and endowment are cointegrated. Specifically, we assume

$$\Delta d_{t+1} = \lambda \Delta c_{t+1} - \alpha(d_t - c_t - \mu_{dc}) + \sigma_d \eta_{t+1}, \quad \alpha > 0, \tag{2.31}$$

similar to Bansal et al. (2007). We assume that μ_{dc}, λ, and α are known to the agent. The agent's learning problem is focused on the unknown μ.

Cointegration is economically realistic, and it is of particular importance in a model like ours with subjective growth rate uncertainty. Since the price of a dividend claim is convex in dividend growth rates, the subjective growth rate uncertainty in this model could cause the price to be infinite. For the consumption claim this issue was resolved by setting $\psi = 1$. However, leverage magnifies the convexity effect and without sufficiently strong cointegration, the price of the equity claim explodes even if the consumption claim has a finite price. In our quantitative implementation, we will assume that α is very small and so dividends and consumption can drift away from each other quite far, but we keep α sufficiently big to yield a finite price-dividend ratio with empirically reasonable moments. Appendix C.5.3 provides more detail.

The price of the n-period dividend strip is

$$P_t^n \equiv \tilde{\mathbb{E}}_t[M_{t+1}|_t\tilde{\mathbb{E}}_{t+1}[\cdots\tilde{\mathbb{E}}_{t+n-1}[M_{t+n}|_{t+n-1}D_{t+n}]]]. \qquad (2.32)$$

As we discussed earlier, when we evaluate these expectations, we do so by iterating backwards from the payoff at $t+n$, evaluating one conditional expectation at a time without relying on the LIE. Taking logs and evaluating (2.32), we obtain

$$p_t^n - d_t = \left[1-(1-\alpha)^n\right]\left(c_t - d_t + \mu_{dc} + \frac{\lambda-1}{\alpha}\tilde{\mu}_t\right) + n\tilde{\mu}_m + \frac{1}{2}(A_n\sigma^2 + B_n\sigma_d^2), \qquad (2.33)$$

where, for very large n, approximately

$$A_n \approx n\left[\sqrt{1+\nu}\left(1+\nu\frac{\lambda-1}{\alpha}\right) - \xi\right]^2, \qquad (2.34)$$

and B_n, which does not grow with n, becomes very small relative to A_n (see Appendix

C.5.3).

The analytical solution for dividend strip prices is useful for understanding the behavior of subjective and objective risk premia in this model. Consider the one-period return on the "infinite-horizon" dividend strip

$$R_{t+1}^{\infty} \equiv \lim_{n\to\infty} P_{t+1}^{n-1}/P_t^n. \tag{2.35}$$

As we show in Appendix C.5.3, we can use equation (2.33) to find the one-period subjective risk premium for this claim

$$\log \tilde{\mathbb{E}}_t[R_{t+1}^{\infty}] - r_{f,t} = \left[1 + \nu\frac{\lambda - 1}{\alpha}\right] \xi\sqrt{1+\nu}\sigma^2. \tag{2.36}$$

For $\gamma \geq 1$, ξ is a positive constant. We observe from the above that lowering α raises the subjective risk premium because it enhances the persistence of the leverage effect by weakening the forces of cointegration. The subjective risk premium is positively related to ν because higher ν implies a smaller effective sample size used to estimate μ and hence higher subjective uncertainty about μ.

While the subjective risk premium is constant, the objective risk premium is

$$\begin{aligned}\log \mathbb{E}_t[R_{t+1}^{\infty}] - r_{f,t} = &\left[1 + \nu\frac{\lambda - 1}{\alpha}\right] \xi\sqrt{1+\nu}\sigma^2 - \frac{1}{2}\nu\left(1 + \nu\frac{\lambda-1}{\alpha}\right)^2\sigma^2 \\ &+ \left(1 + \nu\frac{\lambda-1}{\alpha}\right)(\mu - \tilde{\mu}_t),\end{aligned} \tag{2.37}$$

and hence time-varying with the wedge $\tilde{\mu}_t - \mu$: the more optimistic the agent relative to the econometrician, the lower the objective expected excess return. Thus, learning induces return predictability. And unlike Bayesian learning with full memory as in Collin-Dufresne et al. (2016b) where return predictability dies out in the long-run, learning with constant gain means that the learning effects (and hence return predictability) are perpetual.

As equation (2.37) shows, leverage $\lambda > 1$ magnifies the time-variation in the objective risk premium. With $\lambda = 1$, or for a consumption strip, all variation would come purely from excess variation in the risk-free rate: objective expected returns on the consumption strip are constant because the ratio of its price to consumption is constant and the objective expected growth rate of Δc is constant, while the risk-free rate rises with $\tilde{\mu}_t$. With leverage, however, the price of a dividend claim rises with $\tilde{\mu}_t$, which produces additional variation in the objective risk premium.

Like the long-horizon claim, shorter-horizon claims also have a constant subjective risk premium. For example, a one-period claim with return $R^1_{t+1} \equiv D_{t+1}/P^1_t$, has the constant subjective risk premium

$$\log \tilde{\mathbb{E}}_t[R^1_{t+1}] - r_{f,t} = \lambda\xi\sqrt{1+\nu}\sigma^2. \tag{2.38}$$

The ex-dividend price of the equity claim, i.e., the claim to the entire stream of dividends, is simply the sum of prices of dividend strips

$$P_t = \sum_{n=1}^{\infty} P^n_t. \tag{2.39}$$

We compute the sum in (2.39) numerically using the analytical solution for dividend strips. Details are in Appendix C.5.4.

Our result that the subjective risk premium is constant for short and long-maturity dividend strips does not imply that the subjective risk premium for the whole stream of dividends is constant. We solve for the subjective risk premium of the equity claim numerically using methods from Pohl et al. (2018) (see Appendix C.5.4). As we report below, we find a slightly positive relationship between $\tilde{\mu}_t$ and the subjective equity risk premium. This arises from the fact that the contribution of long-horizon claims to the overall value of the portfolio of strips gets bigger when $\tilde{\mu}_t$ is higher: due to exponentiating, the effect of a rise of $\tilde{\mu}_t$ on long-horizon equity is bigger than on short-horizon equity. As a consequence, the claim on the whole stream behaves more like long-horizon equity when $\tilde{\mu}_t$ is high and is subjectively priced more like long-horizon equity, i.e., with a higher risk premium if $\alpha < \nu$.

2.3.7 Solving the model with $\psi > 1$ and an informative prior

When $\psi \neq 1$, and the prior is diffuse, the consumption-wealth ratio is no longer finite. For example, if $\psi > 1$, the effect of high subjectively expected growth rates on the value of the consumption claim is no longer fully offset by a high future risk-free rate, which causes the consumption-wealth to explode. Earlier work has resolved this through truncation of the state space (Collin-Dufresne et al. 2016a) or limiting the time horizon over which growth is uncertain (Pastor and Veronesi 2003a; Pastor and Veronesi 2006). We take a different approach by endowing the agent with an informative prior. In our fading memory model, this approach is effective in preventing the explosion of the consumption-wealth ratio because future agents never gain more precise information about μ than the current agent has. As a consequence, the weight on the prior does not decay and the current agent anticipates that the posterior means of the agents pricing the asset at times in the future will always have a similarly strong tilt towards the prior mean.

Our approach is similar in spirit to the state-space truncation approach—both methods effectively pull the perceived distribution of future posterior means towards economically plausible growth rates—but it is far more tractable in our setting. We center the prior distribution at the true mean μ, but this is not essential. Since we work with high prior variance, the prior will remain almost uninformative. Therefore, setting the prior mean to any value in an economically plausible neighborhood around

μ would deliver similar results.

Given the prior in (2.12) with $\mu_0 = \mu$, we obtain the posterior

$$\mu | H_t \sim \mathcal{N}\left(\phi \tilde{\mu}_t + (1-\phi)\mu, \left(\frac{1}{\sigma_0^2} + \frac{1}{\nu \sigma^2} \right)^{-1} \right), \qquad \text{where} \quad \phi \equiv \frac{\sigma_0^2}{\nu\sigma^2 + \sigma_0^2}. \tag{2.40}$$

The perceived consumption growth can be represented as

$$\Delta c_{t+1} = \phi \tilde{\mu}_t + (1-\phi)\mu + \sqrt{1+\phi\nu}\sigma \tilde{\varepsilon}_{t+1}, \tag{2.41}$$

where $\tilde{\varepsilon}_{t+1}$ is $\mathcal{N}(0,1)$ distributed under the agent's time-t predictive distribution. With an informative prior we have $\phi < 1$ and so the volatility of the subjectively unexpected endowment growth is lower than in the diffuse prior case where $\phi = 1$. We solve this version of the model with log-linearization. Details, including parameter restrictions sufficient for existence of equilibrium, are provided in Appendix C.6.

2.4 Calibration and Evaluation

Table 2.4 summarizes the parameter values we use in our baseline quantitative analysis. We fix the gain parameter ν at the value that Malmendier and Nagel (2015) estimated from survey data on inflation expectations. For the endowment process and preferences, we set most parameters to the same values as in Bansal et al. (2012) and Collin-Dufresne et al. (2016a).

We set σ_d to a relatively low value of 1% quarterly. At this value, dividend volatility in the model will be smaller than in the data. However, it will allow us to roughly match the volatility of $\tilde{\mu}_d$ in the data, which is more important for our purposes. We cannot match both at the same time because in our model Δd is IID whereas in the data Δd has substantial negative autocorrelation, which implies that a lot of this dividend growth volatility cancels out when we form the long-run weighted average $\tilde{\mu}_d$. For this reason, it makes more sense to calibrate our IID dividend process to the volatility of $\tilde{\mu}_d$, which reflects permanent shocks, rather than the volatility of Δd in the data, which is influenced by a substantial transitory component.

In our baseline specification, we set $\psi = 1$, but we also report some results for the $\psi = 1.5$ case. We choose the remaining parameters γ and α to get a realistic equity premium and equity volatility. We work with a relatively low risk aversion of $\gamma = 4$. The value $\alpha = 0.001$ satisfies the condition required for a finite price of the dividend claim in Appendix C.5.3 and it implies dividends can wander quite far away from consumption (but not as far as in models without cointegration such as Bansal et al. (2012)).

2.4.1 Unconditional moments

We simulate the model at a quarterly frequency. Table 2.5 reports the annualized population moments estimated from an extremely long sample simulated from the model. We also show empirical moments for the 1927 to 2016 period for comparison.

As we anticipated, the volatility of Δd in the model is lower than in the data, but the volatility of $\tilde{\mu}_d$ is close to the data, and actually even a bit higher. This reinforces our earlier point that much of the volatility of Δd in the data is due to transitory components.

In terms of unconditional asset pricing moments, the model produces a high equity premium (7.16%) and Sharpe Ratio (0.54) that are quite close to the empirical estimates in the first column. Return volatility and the volatility of the log price-dividend ratio are lower than in the data. The version of the model with $\psi = 1.5$ in the third column gets somewhat closer to the empirical values.

The model also produces a low subjective real risk-free rate with low volatility.[12] The volatility of r_f in the data (2.47%) is higher than in the model (0.51%), but one should keep in mind that the inflation expectations are estimated with error and this measurement error contributes at least some of the empirically observed volatility in r_f. The low volatility of r_f is a virtue of the model (which is why Campbell and Cochrane (1999), for example, specifically reverse-engineer their model to produce a constant risk-free rate).

Overall, the model provides a reasonably good fit to standard unconditional asset pricing moments. However, the most interesting predictions of the model concern time-variation in objective and subjective conditional moments, which we turn to next.

2.4.2 Predictability of excess returns

We now evaluate time-variation in the objective equity premium. In our model, this time-variation is induced by subjective belief dynamics rather than time-varying risk aversion or time-varying objective risk that generate time-varying risk premia in rational expectations models and this mechanism leads us to construct new return predictors $\tilde{\mu}_d$ and $\tilde{\mu}_r$.

Table 2.6 presents the results from predictive regressions of log excess returns on the equity claim in data simulated from the model. The estimates in this table are the model-implied counterpart to the empirical predictive regression results in Table 2.1. The first block of rows presents mean coefficients and adj. R^2 from regressions with $\tilde{\mu}_d$ as predictor variable. In column (1), the prediction horizon is one quarter, as in Table 2.1. We find coefficients on $\tilde{\mu}_d$ that are about half as big as those we found in

[12] For the purpose of this moments comparison, we calculate the subjective real risk-free rate in the data using expected inflation expectation an AR(1) constant-gain learning inflation forecast with gain $\nu = 0.018$. Malmendier and Nagel (2015) show that this AR(1) constant-gain learning forecast fits household inflation expectations from the Michigan Survey of Consumers well (and we can construct it in the decades before survey data becomes available).

the empirical data including the Great Depression, but quite close to the estimate from the post-WWII sample. Since the volatility of $\tilde{\mu}_d$ is somewhat higher in the calibrated model than in the full empirical sample, this means that the variation in the objective risk premium in our model is at least roughly similar to the variation in the data. Similar comments apply to the estimates with $\tilde{\mu}_r$ as predictor presented in the second block of rows.

Columns (2) and (3) show the regression coefficients when returns are measured over longer horizons of one and five years, respectively. These estimate the persistence in the expected returns in the model. For example, the coefficients in column (3) where the prediction horizon is 20 times longer than in column (1) are only slightly smaller than 20 times the coefficients in column (1). On this dimension, the model also fits well with the empirical data. We didn't explicitly report the empirical predictive regressions for horizons beyond one quarter, but Figure 2.1 earlier showed how predicted 5-year returns line up well with realized 5-year returns.

The bottom block of rows shows regressions with $p - d$ as predictor. The regression coefficients of around -0.06 are somewhat larger than in the empirical data, but certainly of the right order of magnitude.

Columns (4) to (6) repeat this analysis with $\psi = 1.5$ and an informative prior with $\phi = 0.99$. Return predictability with $\tilde{\mu}_d$ gets somewhat stronger and closer to the empirical magnitudes that we found in the full sample in Table 2.1. But overall, the effects of changing the EIS are quite small.

2.4.3 Predictability of subjective expectations errors

To replicate the subjective expectations error regressions that we ran on the empirical data, we use the model-implied subjective equity premium as dependent variable, calculated as explained in Section 2.3.6. Table 2.7 is the simulated counterpart to the empirical results in Table 2.2. Panel A shows that the model-generated data yields a relationship between subjective expected excess returns and experienced returns that is weakly positive and quantitatively similar to the empirical estimates. The empirical point estimate in column (1) of Table 2.2, with $\tilde{\mu}_d$ as predictor, is 0.31, while the regression on model-generated data in column (1) of Table 2.7 yields a mean coefficient of 0.83. With $\tilde{\mu}_r$ as predictor in column (4), the simulated data yields a mean coefficient of 0.60, close to the empirical estimate of 0.86 in Table 2.2, column (4). But it is also useful to keep in mind that the volatility of subjective expected returns in the simulated data is tiny relative to the volatility of objective expected returns. Therefore, the take-away from this analysis is that subjective expected excess returns both in the data and the model are nearly constant.[13]

In Panel B, with subjective expectations errors as dependent variable, the mean coefficient on $\tilde{\mu}_d$ in column (1) is very similar to its empirical counterpart of -12.34 in Table 2.2. With $\tilde{\mu}_r$ as a predictor in column (4), the model-implied coefficients are about half as big as the empirical ones, so the empirical experienced real returns variable may capture some additional expectations errors beyond those implied by our model.

Summing up, the dynamics of subjective and objective expected returns in the model are broadly consistent with the empirical data. Objectively, excess returns are strongly predictable by $\tilde{\mu}_d$ and $\tilde{\mu}_r$, while subjective expected excess returns are largely acyclical.

[13] A comparison of the predictive regression coefficients in Tables 2.6 and 2.7 shows that objective expected returns at a one-year horizon have a volatility that is more than 10 times as high (absolute regression coefficient of 9.24 times volatility of quarterly $\tilde{\mu}_d$ of $0.39\% \approx 3.60\%$) as the volatility of subjective expected returns (regression coefficient of 0.83 times volatility of quarterly $\tilde{\mu}_d \approx 0.32\%$).

2.4.4 Subjective long-run growth expectations

As the final piece of our comparison of the model with our earlier empirical estimates from the reduced-form framework, we look at the dynamics of subjective dividend growth expectations.

By iterating on eqs. (2.16) and (2.31) we can construct, at every point in time in our simulations, the representative agent's subjective expectation of average log dividend growth over the next 20 quarters (1,000 simulations of the model for 50,000 periods). The 20-quarter forecast horizon corresponds roughly to the analyst forecast horizon in our analysis in Table 2.3. In each simulation run, we regress this subjective dividend growth expectation at time t on $\tilde{\mu}_{d,t}$ or $\tilde{\mu}_{r,t}$. We obtain coefficients of 0.899 and 0.683. These values are extremely close to the implied coefficients of 0.85 and 0.64 that we obtained in Table 2.3 after adjusting for the distortionary effects of transitory predictable components in earnings on the analyst forecasts. Thus, our model quantitatively matches the dynamics of long-term forecasts of growth in stock market fundamentals by stock market analysts.

2.5 Relationship between First and Second Moments of Equity Returns

Another interesting implication of our model that we have not discussed so far is that its predictions about the relationship between conditional first and second moments of equity market returns are sharply different from those of the leading rational expectations asset pricing models by CC and BY. In CC and BY, the conditional equity premium is approximately linearly and positively related to conditional equity return variance. In CC, this happens because at times when risk aversion is high, it is also very volatile. In BY, the reason is that stochastic volatility in endowment growth is the driver of the time-varying equity premium.[14] Therefore, variables like $p - d$ that predict excess returns should also predict equity market return variance with the same sign—but empirical data does not support this prediction. This gives rise to what Lettau and Ludvigson (2010) have termed the *Sharpe Ratio variability puzzle*: These models make the empirically counterfactual prediction that positive co-movement of conditional first and second moments should dampen the variability of the Sharpe Ratio compared with the variability in expected excess returns.

Table 2.8 shows a similar empirical result when we use experienced payout growth, $\tilde{\mu}_d$, as a predictor variable. The dependent variable in these regressions is the sum of squared daily log returns of the CRSP value-weighted index in quarter $t + 1$, i.e., an estimate of quarterly realized variance. Recall that $\tilde{\mu}_d$ predicts returns with a negative sign. However, as column (1) and (4) show, there is a weakly positive, but not statistically significant relationship between $\tilde{\mu}_d$ and next-quarter variance both in the full sample and the post-WW II sample. The magnitude of the point estimate is small: a one standard deviation increase in $\tilde{\mu}_d$ is associated with an increase of 0.15 percentage points in conditional quarterly return variance, which is about 10% of the sample standard deviation of the return variance. Using $p-d$ as a predictor in columns (2) and (5), the picture is mixed. In the full sample, $p-d$ predicts volatility negatively, but the coefficient estimate is not statistically significant. In the post-WW II period, the estimate is statistically significant, but it is positive, which is inconsistent with the CC and BY models. Columns (3) and (6) control for the risk-free rate because $p - d$ and the real risk-free rate together should span the state variables in the BY model, but this has little effect on the coefficient of the $p-d$ ratio. Overall, there is no support for the prediction of the CC and BY models that the conditional equity premium has strong positive correlation with conditional market return variance. At low frequencies

[14]For the BY model, compare equations (A13) and (A14) in the appendix of BY. For the CC model one can infer the positive and close to linear relationship by comparing Figures 4 and 5 in CC.

captured by the slow-moving predictors $\tilde{\mu}_d$ and $p - d$, there is not any co-movement between equity premium and volatility.

In contrast to CC and BY, our model is consistent with this empirical result. Learning with fading memory generates predictable variation in the objective equity premium, but without simultaneous variation in equity market return variance.[15] The decoupling objective and subjective beliefs allows the model to also decouple conditional equity return volatility and the objective conditional equity premium.

2.6 Lack of Out-of-Sample Return Predictability

Welch and Goyal (2007) show that the simple trailing sample mean of past returns often beats an out-of-sample predictive regression forecast as a predictor of future returns. Since the representative agent in our model discards historical information at a relatively high rate (the half-life in terms of the observation's weight in the log likelihood is about 10 years), one might suspect that a predictive regression run in real time, but with full memory of past data, should be able to identify the agent's errors and hence predict returns out-of-sample better than the sample mean. However, as we show now, this is not the case.

We apply the Goyal and Welch analysis to simulated data from our model. We run 10,000 simulations of a 360-quarter sample period with a 400-quarter burn-in period to compute $\tilde{\mu}_d$ at the start of each sample. Within each 360-quarter sample, we then examine the in-sample explanatory power of the predictive regression by plotting the cumulative squared demeaned excess returns minus the cumulative squared full-sample regression residual from the beginning to the end of the sample. The predictive regres-

[15] Our model actually produces a positive relationship between $\tilde{\mu}_d$ and conditional variance of the equity claim return, but the effect is very weak, somewhat similar to the weak empirical relationship in Table 2.8. Subjective and objective conditional variance of dividend strip returns are constant over time. However, higher subjectively expected growth implies higher weight of riskier longer-horizon dividend strips in the equity claim's value and hence higher variance of the return on the equity claim.

sion is run at quarterly frequency with the sum of four-quarter log excess returns from $t+1$ to $t+4$ as dependent variable and $p-d$ or $\tilde{\mu}_d$ as predictor. The blue line in the upper half of each plot in Figure 2.3 shows the average path across all simulations of this in-sample cumulative squared errors difference. The upward slope of this line and the fact that it ends up, on the right-hand side, above zero, both with $p-d$ as a predictor (top) and $\tilde{\mu}_d$ (bottom), indicates that the in-sample R^2 is greater than zero.

To assess the out-of-sample performance with each 360-quarter sample, we use a recursively expanding window, starting out at 80 quarters, to estimate the sample mean and the predictive regression. We calculate the next period out-of-sample squared prediction error of the trailing sample mean as a forecaster minus the squared prediction error of the fitted predictive regression. We cumulate these squared error differences forward and we average the resulting paths across all simulations. The red lines in the lower half of each plot in Figure 2.3 show that this path is, on average, in negative territory, which means that the predictive regression forecast underperforms the trailing sample mean as a forecaster. That the slope is still negative on average towards the end of the sample period shows that even after having observed almost 90 years of data, the trailing sample mean is typically still a better forecaster.

Thus, even though there is true return predictability in this model under the econometrician's objective probability measure, this predictability is not exploitable in real-time for typical sample sizes. The data generated by the model is therefore consistent with the lack of out-of-sample predictability found empirically by Welch and Goyal (2007).

The out-of-sample exercise also demonstrates that it would not be easy for the agent within the model to recognize that the loss of memory and the resulting reliance on relatively recent experiences in estimating endowment growth rates is detrimental to forecast performance. In this sense, one can interpret our model as a near-rational model.

2.7 Conclusion

We have shown that learning with fading memory can reconcile asset prices and survey expectations in a highly tractable framework. In our model, asset prices are volatile because subjective growth expectations are time-varying. Risk premia are high because subjective long-run growth rate uncertainty is high. The model produces realistic asset price behavior in a simple setting with IID endowment growth and constant risk aversion. While objective expected excess returns are strongly counter-cyclical, subjective beliefs about stock market excess returns are slightly pro-cyclical. As a consequence, subjective expectations errors are predictable, as they are in the survey data. As predicted by the model, long-run weighted averages of past real per-capita payout growth or past real stock index returns are a good empirical predictor of excess returns and subjective expectations errors. Unlike in leading rational expectations explanations of return predictability, and consistent with the data, movements in objective expected excess returns in our model are not associated with movements in conditional market return volatility.

Because memory of past data fades away, subjective beliefs about long-run growth fluctuate perpetually in our model. That these belief fluctuations persist is plausible because it would be difficult for an agent to detect that the loss of memory is detrimental to her investment decisions. While returns generated by our model economy are predictable to an econometrician examining a sample ex post, standard out-of-sample tests show that they are not predictable in real time in typical sample sizes, consistent with the empirical lack of out-of-sample predictability. Overall, these results suggest that subjective belief dynamics could be central to asset pricing and that learning with fading memory can provide a unifying account of many asset pricing phenomena and the evidence on subjective beliefs about stock returns in investor surveys.

2.8 Tables and Figures

Table 2.1: Predicting Returns with Experienced Real Growth

Dependent variable is the log return of the CRSP value-weighted index in quarter $t+1$ in excess of the return on a 3-month T-bill. In Panel A, experienced payout growth denotes a long-run exponentially weighted average of overlapping quarterly observations of four-quarter per-capita repurchase-adjusted real dividend growth rates leading up to and including quarter t, constructed with weights implied by constant gain learning with quarterly gain $\nu = 0.018$. In Panel B, experienced returns are constructed analogously as an exponentially weighted average of quarterly log stock market index returns (S&P Composite before 1926; then CRSP value-weighted index). Inflation is measured as the average log CPI inflation rate during the four quarters $t-3$ to t; $p-d$ refers to the log price-dividend ratio of the CRSP value-weighted index at the end of quarter t. The table shows slope coefficient estimates, with bootstrap bias-adjusted coefficient estimates in brackets. Intercepts are not shown. Bootstrap p-values are shown in parentheses. The reported R^2 are based on bias-adjusted estimates.

	(1) 1927-2016	(2) 1927-2016	(3) 1927-2016	(4) 1946-2016	(5) 1946-2016
Panel A: Predicting returns with experienced real payout growth					
Experienced real payout growth	-5.79	-6.25	-5.82	-2.99	-1.29
[bias-adj. coeff.]	[-5.71]	[-6.14]	[-5.40]	[-2.82]	[-1.14]
(p-value)	(0.00)	(0.00)	(0.00)	(0.04)	(0.29)
Inflation		-0.73	-0.71	-1.60	-2.13
[bias-adj. coeff.]		[-0.75]	[-0.73]	[-1.68]	[-2.14]
(p-value)		(0.10)	(0.13)	(0.01)	(0.00)
$p-d$			-0.01		-0.04
[bias-adj. coeff.]			[0.01]		[-0.02]
(p-value)			(0.56)		(0.04)
Observations	360	360	360	284	284
R^2	0.033	0.037	0.034	0.027	0.044
Panel B: Predicting returns with experienced returns					
Experienced real returns	-2.36	-2.58	-2.22	-3.17	-2.38
[bias-adj. coeff.]	[-1.74]	[-1.79]	[-2.06]	[-2.21]	[-2.26]
(p-value)	(0.02)	(0.02)	(0.05)	(0.01)	(0.03)
Inflation		-0.60	-0.57	-2.47	-2.78
[bias-adj. coeff.]		[-0.57]	[-0.50]	[-2.42]	[-2.76]
(p-value)		(0.19)	(0.21)	(0.00)	(0.00)
$p-d$			-0.01		-0.03
[bias-adj. coeff.]			[0.00]		[-0.01]
(p-value)			(0.44)		(0.04)
Observations	360	360	360	284	284
R^2	0.015	0.017	0.016	0.045	0.061

Table 2.2: Survey Return Expectations and Experienced Real Growth

In Panel A, the dependent variable is the average subjective expected stock return of survey respondents in quarter t minus the one-year treasury yield at the end of quarter $t-1$, which we regress on experienced real payout growth or experienced real returns leading up to and including quarter $t-1$. Lagged one-year return refers to the return of the CRSP value-weighted index over the four quarters $t-4$ to $t-1$. In Panel B, the dependent variable is the expectation error, i.e., the realized return on the CRSP value-weighted index during quarters $t+1$ to $t+4$ minus the subjective expected return of survey respondents in quarter t. We use expectations data up to and including quarter 2016:4. Newey-West standard errors are reported in parentheses (12 lags in Panel A; 6 lags in Panel B).

	(1)	(2)	(3)	(4)	(5)
Panel A: Subjective expected excess returns					
Experienced real payout growth	0.31		0.37		
	(0.97)		(0.91)		
Experienced real returns				0.86	0.38
				(0.63)	(0.76)
Lagged one-year return		0.03	0.03		0.03
		(0.01)	(0.01)		(0.01)
Constant	0.05	0.05	0.05	0.04	0.05
	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)
Observations	125	125	125	125	125
Adj. R^2	-0.004	0.075	0.073	0.031	0.073
Panel B: Expectation error: Realized - subj. expected					
Experienced real payout growth	-12.34		-12.59		
	(6.75)		(6.80)		
Experienced real returns				-14.27	-15.47
				(6.48)	(7.08)
Lagged one-year return		-0.10	-0.12		0.07
		(0.14)	(0.15)		(0.13)
Constant	0.12	0.03	0.14	0.25	0.27
	(0.06)	(0.03)	(0.06)	(0.09)	(0.10)
Observations	125	125	125	125	125
Adj. R^2	0.055	0.002	0.060	0.106	0.102

Table 2.3: Long-Run Growth Expectations and Experienced Real Growth

The dependent variable is an aggregate of I/B/E/S analysts long-term earnings-per-share (EPS) growth forecast, deflated using CPI inflation expectations from the Survey of Professional forecasters from the quarter prior to the EPS forecast date. Each month, we collect the median forecasts for each individual stock, form a value-weighted average across stocks using each stock's monthly median forecasts of current fiscal year earnings as the weight, with the sample restricted to stocks with positive current year forecasted earnings. We then average the resulting monthly time-series observations within calendar quarters 1984:3 to 2016:4. We regress the dependent variable measured in quarter t, and expressed as a per-quarter growth rate, on experienced real payout growth or experienced real returns leading up to and including quarter $t-1$. The one-year EPS forecast is the expected growth rate from current fiscal year to next fiscal year earnings, averaged across stocks in the same way as the long-term forecast. Lagged one-year return refers to the return of the CRSP value-weighted index over the four quarters $t-4$ to $t-1$. Newey-West standard errors are reported in parentheses (24 lags).

	(1)	(2)	(3)	(4)	(5)	(6)
Experienced real payout growth	0.30	0.61	0.62			
	(0.15)	(0.18)	(0.20)			
Experienced real returns				0.55	0.52	0.64
				(0.14)	(0.13)	(0.16)
One-year EPS growth forecast		0.28	0.29		0.05	0.08
		(0.09)	(0.11)		(0.05)	(0.05)
Lagged one-year return			-0.00			-0.01
			(0.00)			(0.00)
Constant	0.02	0.02		0.01	0.01	0.01
	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)
Observations	130	130	130	130	130	130
Adjusted R^2	0.077	0.308	0.312	0.355	0.362	0.452

Table 2.4: Baseline Model Parameters

This table reports the parameters values we use in the baseline calibration of our model at a quarterly frequency. The gain parameter ν is fixed at the value that Malmendier and Nagel (2015) estimated from survey data on inflation expectations. For endowment process parameters and preferences, we set most at the same values as in Bansal et al. (2012) and Collin-Dufresne et al. (2016a).

Parameter	Symbol	Value	Source
Belief updating			
Gain	ν	0.018	MN (2016) (survey data)
Endowment process			
Leverage ratio	λ	3	CJL (2017)
Dividend cointegration parameter	α	0.001	
Mean consumption growth	μ	0.45%	CJL (2017)
Consumption growth volatility	σ	1.35%	CJL (2017)
Dividend growth volatility	σ_d	1%	
Preferences			
Risk aversion	γ	4	
EIS	ψ	1	
Time discount factor	δ	0.9967	BKY (2012)

Table 2.5: Unconditional Moments

The second column in this table presents the model population moments obtained as average across 1,000 simulations of the model for 50,000 periods plus a 2,000-period burn-in period to compute $\tilde{\mu}$, $\tilde{\mu}_d$ and $\tilde{\mu}_r$ at the start of each sample, using a diffuse prior ($\phi = 1$) and $\psi = 1$. The third column shows results for the model with $\psi = 1.5$ and an informative prior with $\phi = 0.99$. The first column shows the corresponding empirical moments from the data for the 1927 to 2016 period. For the empirical versions of $\tilde{\mu}_d$ and $\tilde{\mu}_r$, we use data from 1871 to 1926 as pre-sample information to calculate their values at the start of the sample in 1927. Consumption growth is calculated from quarterly real per-capita consumption expenditure on nondurables and services in chained 2012 dollars from NIPA for 1947-2016, annual NIPA data on nondurables and services expenditure from 1929 to 1947, and annual real per-capita consumption expenditure from Barro and Ursúa (2008) for 1926 to 1929. In both columns, returns are annualized as follows: The means of risky returns are multiplied by four and standard deviation multiplied by two. For the risk-free rate, $\tilde{\mu}$, and $\tilde{\mu}_d$ we multiply quarterly means and standard deviations by four. We estimate the empirical moments of Δc from four-quarter changes of quarterly log nondurables and services consumption and those of Δd from four-quarter changes in the log of repurchase-adjusted dividends on the CRSP value-weighted index. The simulated statistics for $p - d$ use a four-quarter trailing sum of dividends in the calculation of $p - d$, just like in the empirical version of $p - d$.

	Data 1927-2016	Model $\psi = 1$, $\phi = 1$	Model $\psi = 1.5$, $\phi = 0.99$
$\mathbb{E}(\Delta c)$	1.84	1.80	1.80
$\sigma(\Delta c)$	2.72	2.70	2.70
$\mathbb{E}(\Delta d)$	2.38	1.80	1.80
$\sigma(\Delta d)$	13.31	8.35	8.35
$\sigma(\tilde{\mu})$	-	0.51	0.51
$\rho(\tilde{\mu})$	-	0.98	0.98
$\sigma(\tilde{\mu}_d)$	1.32	1.55	1.55
$\rho(\tilde{\mu}_d)$	0.97	0.98	0.98
$\text{corr}(\tilde{\mu}, \tilde{\mu}_d)$	-	0.96	0.96
$\text{corr}(\tilde{\mu}, \tilde{\mu}_r)$	-	0.85	0.77
$\mathbb{E}(R_m - R_f)$	8.11	7.16	7.75
$\sigma(R_m - R_f)$	22.41	13.31	16.35
$SR(R_m - R_f)$	0.36	0.54	0.47
$\mathbb{E}(p - d)$	3.40	2.81	2.98
$\sigma(p - d)$	0.44	0.14	0.22
$\rho(p - d)$	0.97	0.91	0.94
$\mathbb{E}(r_f)$	0.67	1.64	0.61
$\sigma(r_f)$	2.47	0.51	0.34

Table 2.6: Predictive Regressions in Simulated Data

This table reports the mean return predictability regression coefficients and adj. R^2 across 10,000 simulations of the model for 360 quarters plus a 400-quarter burn-in period to compute $\tilde{\mu}_d$ and $\tilde{\mu}_r$ at the start of each simulated sample. The dependent variable is the log excess return on the equity claim. The predictors $\tilde{\mu}_d$ and $\tilde{\mu}_r$ are constructed as the exponentially-weighted average of experienced payout growth and experienced log returns, respectively, with gain parameter $\nu = 0.018$. Each block of rows represents regressions with a different (single) predictor variable. Columns (1) to (3) show results using a diffuse prior ($\phi = 1$) and $\psi = 1$. Columns (4) to (6) show the corresponding results for $\psi = 1.5$ and an informative prior with $\phi = 0.99$.

	$\psi = 1$, $\phi = 1$			$\psi = 1.5$, $\phi = 0.99$		
	1Q (1)	1Y (2)	5Y (3)	1Q (4)	1Y (5)	5Y (6)
$\tilde{\mu}_d$	-2.40	-9.24	-38.37	-3.07	-11.78	-48.84
R^2_{adj}	0.01	0.05	0.21	0.01	0.05	0.22
$\tilde{\mu}_r$	-1.51	-5.79	-23.77	-1.48	-5.64	-23.12
R^2_{adj}	0.01	0.03	0.15	0.01	0.03	0.14
$p - d$	-0.06	-0.22	-0.93	-0.05	-0.20	-0.82
R^2_{adj}	0.01	0.04	0.18	0.01	0.05	0.21

Table 2.7: Dynamics of Subjective Expectations of Excess Returns in Simulated Data

This table reports the mean estimates from regressing subjective expected excess returns and expectation errors on $\tilde{\mu}_d$, $\tilde{\mu}_r$, and 1-year lagged log returns from 10,000 simulations of 360 quarters with a 400-quarter burn-in period, using the model with a diffuse prior ($\phi = 1$) and $\psi = 1$. In Panel A, the dependent variable is $(\tilde{\mathbb{E}}_t[R_{m,t+1}])^4 - (R_{f,t})^4$, which we regress on experienced real returns leading up to and including quarter t, and/or lagged one-year log returns over the four quarters $t-3$ to t. In Panel B, the dependent variable is the expectation error, defined as $\prod_{i=1}^{4} R_{m,t+i} - (\tilde{\mathbb{E}}_t[R_{m,t+1}])^4$.

	(1)	(2)	(3)	(4)	(5)
Panel A: Subjective expected excess returns					
$\tilde{\mu}_d$	0.83		0.83		
$\tilde{\mu}_r$				0.60	0.62
$r_{t-3,t}$		0.01	0.00		0.00
R^2_{adj}	0.93	0.08	0.93	0.83	0.83
Panel B: Expectation error: Realized - subj. expected					
$\tilde{\mu}_d$	-10.75		-11.05		
$\tilde{\mu}_r$				-6.80	-7.24
$r_{t-3,t}$		-0.06	0.02		0.03
R^2_{adj}	0.05	0.01	0.06	0.04	0.04

Table 2.8: Predicting Market Return Variance

Dependent variable is the sum of squared daily log returns of the CRSP value-weighted index in quarter $t+1$. Experienced ("Exp.") real payout growth denotes a long-run exponentially weighted average of overlapping quarterly observations of four-quarter per-capita repurchase-adjusted real dividend growth rates leading up to and including quarter t, constructed with weights implied by constant gain learning with quarterly gain $\nu = 0.018$; $p - d$ refers to the log price-dividend ratio of the CRSP value-weighted index at the end of quarter t. For readability of the estimates, we use $(p - d)/100$ as predictor variable. Newey-West standard errors with six lags are shown in parentheses.

	(1) 1927-2016	(2) 1927-2016	(3) 1927-2016	(4) 1946-2016	(5) 1946-2016	(6) 1946-2016
Exp. payout growth	0.47			0.41		
	(0.41)			(0.35)		
$(p - d)/100$		-0.30	-0.30		0.41	0.40
		(0.42)	(0.40)		(0.12)	(0.12)
Real risk-free rate			0.21			-0.17
			(0.26)			(0.14)
Constant	0.00	0.02	0.02	0.00	-0.01	-0.01
	(0.00)	(0.02)	(0.01)	(0.00)	(0.00)	(0.00)
Observations	360	360	360	284	284	284
Adjusted R^2	0.012	0.008	0.016	0.015	0.033	0.044

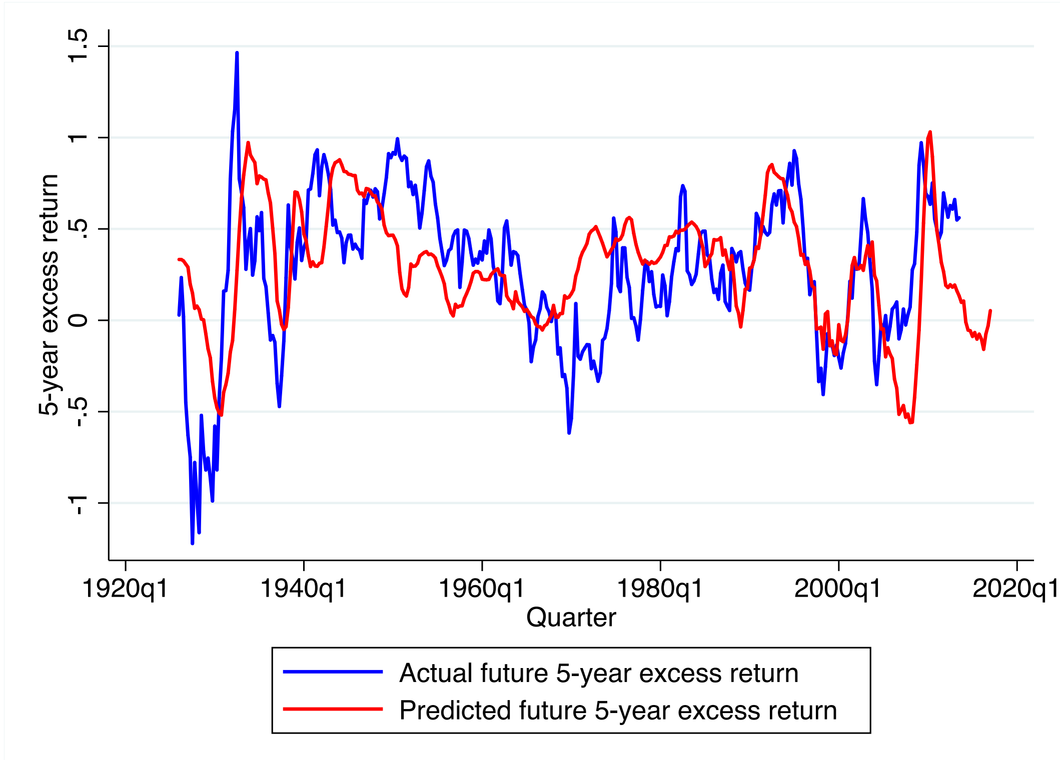

Figure 2.1: Predicted Five-year Excess Returns and Subsequent Actual Cumulative Five-year Excess Returns

Predicted returns are calculated based on bootstrap bias-adjusted coefficients from the predictive regression of log excess returns on experienced real payout growth shown in column (1) of Table 2.1, Panel A, applied to experienced payout growth in quarter t, and iterating using an AR(1) with AR coefficient $1 - \nu = 0.982$ to obtain 5-year forecasts. The actual cumulative five-year excess returns refers to the sum of excess log returns on the CRSP value-weighted index in quarters $t + 1$ to $t + 20$.

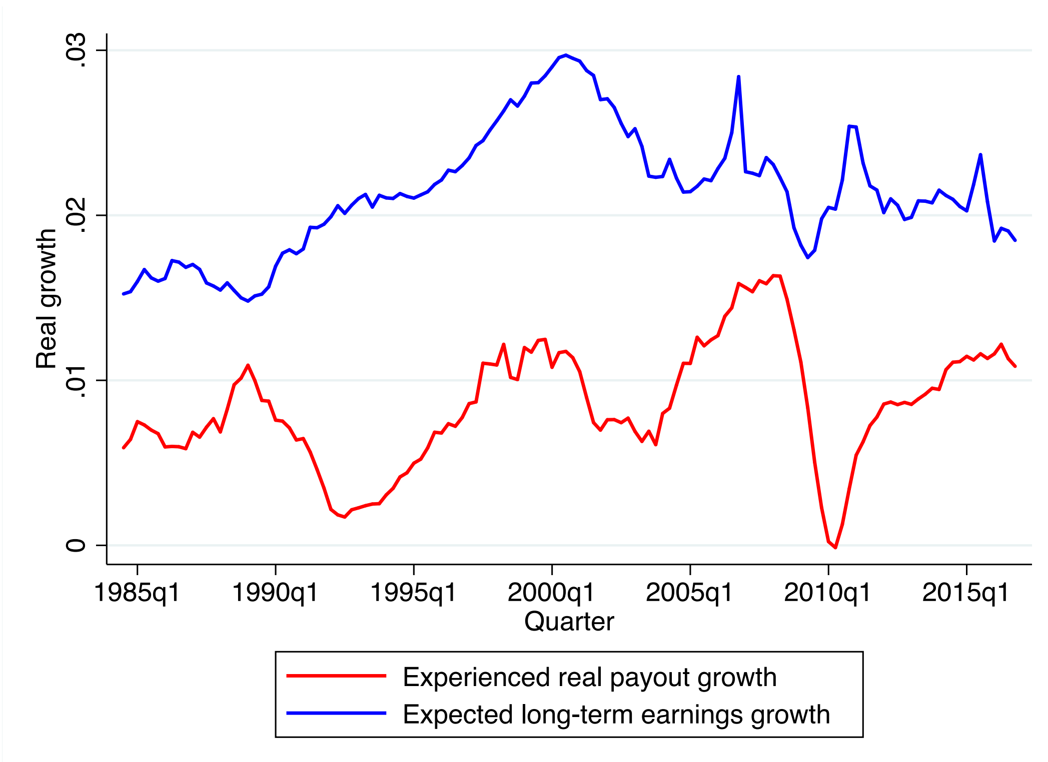

Figure 2.2: Experienced Real Payout Growth and Analysts' Long-run Earnings Growth Expectations

The figure shows the value-weighted I/B/E/S long-term earnings-per-share growth forecast, value-weighted across stocks based on forecasted current year total earnings.

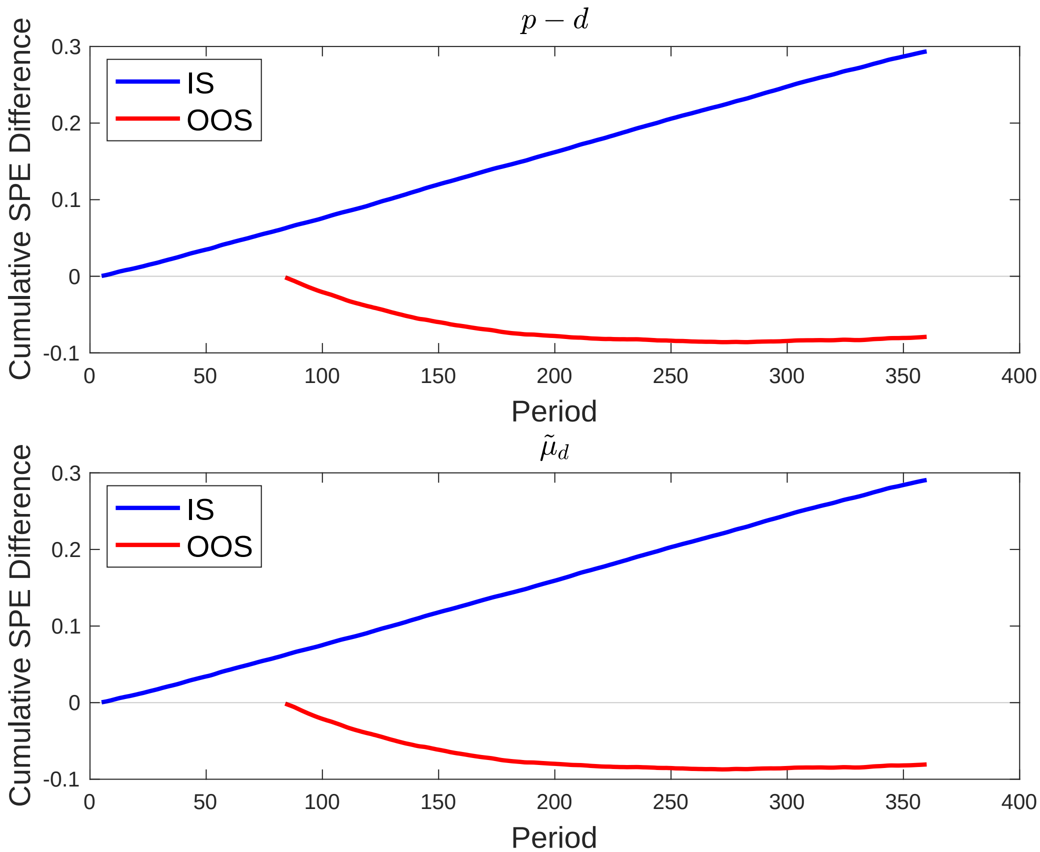

Figure 2.3: Out-of-sample Predictive Performance

In-sample and out-of-sample performance of predictors $p-d$ and $\tilde{\mu}_d$ from 10,000 simulations of 360 quarters with a 400-quarter burn-in period to compute $\tilde{\mu}_d$ at the start of each sample. The IS line plots the cumulative squared demeaned excess returns minus the cumulative squared full-sample regression residual. The OOS line plots the cumulative squared prediction errors of conditional mean minus the cumulative squared prediction errors of predictors. Both lines are the average path across simulations.

CHAPTER III

An Information-based Theory of the Profitability Anomaly

3.1 Introduction

Firms with high profitability earn substantially higher future stock returns than firms with low profitability, even after controlling for Fama and French (1993) three factors. This empirical pattern has been referred to as the "profitability" anomaly and has attracted a lot of attention in the literature (Fama and French 2015; Hou et al. 2014). The profitability anomaly yields a premium that is both statistically and economically large (Novy-Marx 2013), is robust to different measures of profitability (Ball et al. 2016), and it has low trading costs and a high strategy capacity (Novy-Marx and Velikov 2015).[1]

The strong empirical performance of the profitability anomaly has generated a lot of interest in understanding its underlying economic mechanism. In the rational expectations framework, it has been argured that profitable firms are intrinsictlly riskier than unprofitable firms. For example, Kogan and Papanikolaou (2013) and Kogan

[1]For example, Ball et al. (2016) document a three-factor-adjusted return spread of monthly 0.89% with a t-statistic of 8.48 between portfolios of the most and the least profitable firms. Novy-Marx and Velikov (2015) shows that the gross profitability strategy as in Novy-Marx (2013) has a capacity of around \$130 billions, which is among the highest of considered strategies.

et al. (2019) find that profitable firms are more exposed to investment-specific shocks and demand shocks. Deng (2019) argues that the difference in risk exposure is particularly large in highly competitive industries. Deviating from the rational expectations framework, explanations are proposed based on sentiment (Stambaugh et al. 2012; Lam et al. 2016), analyst bias (Bouchaud et al. 2019), and investors' inattention (Wang and Yu 2015).

In this paper, I contribute to this debate by documenting a new empirical pattern of the profitability anomaly. I show that the Fama-French three-factor (FF3F) alphas of the profitability anomaly exist only among firms with high information frictions (IF), proxied by young age, high forecast dispersion, high past return volatility, or high option-implied volatility. In the double-sorting portfolio analysis, the FF3 alphas of the long-short profitability portfolios are insignificant both economically and statistically in low-IF quintiles, but are as high as monthly 1.46% (t-statistics 4.06) in high-IF quintiles. This pattern persists with different measures of profitability (Novy-Marx 2013; Ball et al. 2016). Given that my sample tiltes towards firms that are either followed by stock analysts or issuing options, the pattern is also robust to excluding small and micro firms. In addition, I show that the a level factor captures most (about 60%) cross-sectional variation in proxies of information frictions. Using factor loadings, I construct a firm-level variable that captures the main variation, which I refer to as the "PIF" measure. In the double-sorting analysis, I show that the FF3F alphas increase monotonically with PIF quintiles for all three meaures of profitability. The highest FF3F alphas are obtained with the Ball et al. (2016) cash-based operating profitability anomaly in the highest PIF quintile with a monthly value of 1.38% (t-statistics 3.64). Also, I show that the stock analysts' forecast errors of firms in the highest PIF tercile do not converge to zero as forecast horizon shorterns, which is in contrast to firms in the lower PIF terciles.

Next, I investigate possible explanations that can account for this empirical pattern. I consider both empirically and theoretically motivated explanations. One possible explanation is that the profitability anomaly comes from the investors' sentiment, and firms with high IFs are also firms subject to high sentiment. Combined with short-sale constraints (Nagel 2005; Pontiff 2006), the profitability anomaly persists in high IF quintiles. I show that a range of measures of short-sale constraints and trading costs do not differ materially between long and short legs of the profitability anomaly in the high PIF quintile. Moreover, both long and short legs contribute significant FF3F alphas to the performance of the profitability anomaly in the high PIF quintile. Thus, the sentiment with short-sale constraints do not seem to be a main driving force of this empirical pattern. I also consider the explanation based on financial distress, following the argument in Avramov et al. (2009) and Avramov et al. (2013). I show that the emprical pattern is still present after dropping out stressed firms. Finally, I show that the investment-based model (Lin and Zhang 2013; Hou et al. 2014) and the sticky analyst model (Bouchaud et al. 2019) do not direcly speak to the empirical pattern documented here.

Finally, I provide an information-based model in the spirit of Admati (1985) to qualitatively reconcile the results. In the model, the investors are trying to learn about asset payoffs realized in the future from current price and a private signal in a Bayesian way. The IF is modeled as a high average posterior uncertainty across investors or a low average signal precision. Thus, the price of firms with high IFs reflect less information from payoff signals. In other words, among firms with high IFs, firms with high (low) payoffs are under-valued (over-valued) from an ex post point of view. Thus, for an econometrician who examines the data and use lagged payoff potential to measure the profitability of a firm, she will find the cross-sectional return predictability of the lagged profitability signal to be stronger among firms with higher IFs, consistent with the emprical pattern documented here.

My paper is related to the growing literature that studies the underlying economic mechanism of the profitability anomaly. As mentioned in the beginning, explanations in the rational expectations framework include Kogan and Papanikolaou (2013), Kogan et al. (2019), and Deng (2019). Behavioral explanations includeStambaugh et al. (2012), Lam et al. (2016)), Bouchaud et al. (2019)), and Wang and Yu (2015). The closest paper to mine is Wang and Yu (2015), who document that the ROE anomaly is much stronger among small, young, and volatile firms. They find that this strong performance mainly comes from the short legs, consistent with the idea that this overpricing is not corrected due to short-sale constraints. The distinction here is that I instead find both long and short legs contribute significant alphas to the strong performance of a wide range of profitability anomalies, suggesting this story may not be at play.

The rest of the paper is organized as follows. Section 3.2 describes the data. Section 3.3 documents the main empirical results that the profitability anomaly performs drastically among firms with high and low IFs. Section 3.4 explores possible explanations in the previous literature. Section 3.5 builds a simple information-based modelto qualitatively accoount for the empirical pattern. Section 3.6 concludes.

3.2 Data

The data come from Compustat, CRSP, I/B/E/S, and OptionMetrics. The sample period is from July 1963 to December 2016.

3.2.1 Equity price and accounting data

The sample includes all the non-financial ordinary common shares traded on NYSE, AMEX, and NASDAQ.

Delisting returns are taken from CRSP. If a delisting return is missing and the delisting is performance-related, I follow Shumway (1997) and Shumway and Warther (1999) to use a delisting return of -30% for stocks traded on NYSE and AMEX and a delisting return of -55% for stocks traded on NASDAQ.

I consider three measures of profitability from the literature (Novy-Marx 2013; Ball et al. 2016), namely the gross profitability (GP/A), the operating profitability (OP), and the cash-based operating profitability (C-OP). These measures are accounting-based and change at an annual frequency. I replicate all three anomalies in Table E.1 and confirm the findings of original papers.

I also construct other market-based and balance-sheet-based characteristics of stocks, which are summarized in Appendix D.1.

The Fama-French five factors are directly taken from Kenneth French's website.[2]

3.2.2 Proxies of information frictions

To construct proxies of information frictions, I follow Pastor and Veronesi (2003b) and Zhang (2006) to consider age, analyst forecast dispersion, and past return volatility. In addition, I use near-the-money call-option-implied volatility as a measure of forward-looking uncertainty. All proxies change at a monthly frequency except the age, which changes at an annual frequency. The detailed definitions of these variables are described below.

Age is defined as the number of years since a firm's first appearance in the CRSP or COMPUSTAT database, or the number of years since a firm's first appearance on CRSP, whichever is earlier, following Pastor and Veronesi (2003b).

The analyst forecast dispersion is constructed using analysts forecasts of earnings-per-share for individual firms from I/B/E/S database. I consider analyst forecasts

[2] `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`

of both the first fiscal year and long-term growth rate. The standard deviations are obtained from the unadjusted summary file to avoid rounding errors from adjustments. For standard deviations of fiscal-year-one forecasts, I scale them by the absolute value of mean forecasts (Diether et al. 2002). Because the long-term forecast is for growth rates, it is already comparable across firms and no further scaling is needed (Zhang 2006).

Past return volatility is calculated as the standard deviations of weekly excess returns in previous year, following Zhang (2006).

The near-the-money call-option-implied volatility is obtained from OptionMetrics. I apply the same option filters as in Christoffersen et al. (2017), which require options have days to maturity between 30 days to 365 days, moneyness between 0.9 to 1.1, and implied volatility between 5% and 150%. The near-the-money call-option-implied volatility for an individual firm is then weighted using open interest across options.[3]

3.3 Empirical Analysis

In this section, I document my main empirical findings. I will first show that the profitability anomaly remains a prominent anomaly. I will then show that the profitability anomaly performs distinctly between firms with high and low information frictions. This pattern is robust to different measures of profitability and information frictions.

3.3.1 Prominence of the profitability anomaly

The profitability anomaly still poses a challenge for the latest work-horse factor models include the FF5F model from Fama and French (2015) and the q-factor model

[3]The caveat here is that the volatility calculated is under the risk-neutral measure and thus contains a risk premium component.

from Hou et al. (2014). The alphas of quintile long-short portfolios from Column (1), (3), and (5) in Table 3.1 are controlled for the FF5F model. They remain economically and statistically significant across all three measures of profitability. A similar observation holds for the q-factor model. Column (2), (4), and (6) in Table 3.1 show that only the GP/A anomaly is able to be explained away by the q-factor model in terms of statistical significance. This is surprising given that these models explicitly incorporate a profitability factor.[4]

3.3.2 Initial portfolio-sorting analysis

This section investigates the performance of the profitability anomaly across firms with different information frictions by using a portfolio-sorting approach. I do an independent double-sorting with respect to proxies of profitability and information frictions (IF). Firms with non-missing characteristics are sorted into twenty-five portfolios based on the quintiles of the ranked values. I use NYSE breakpoints and value-weighted returns following Hou et al. (2018). Given the number of proxies, this approach yields fifteen combinations of double-sortings. Most combinations are balanced monthly due to the monthly-changing IF signals (except age).

Table 3.2 reports the excess returns of independently double-sorted portfolios. Across different combinations of PROF and IF, return spreads from long-short profitability portfolios that are both statistically and economically significant only emerge in groups of stocks with high IF. For example, the return spread from long-short gross profitability portfolios among firms with the lowest one-year analyst forecast dispersion (1YFD) is essentially 0% (t-statistics -0.01). But the return spread among firms with highest 1YFD is monthly 0.51% (t-statistics 2.22). Across all combinations, the highest return spread is monthly 1.27% (t-statistics 4.01) from long-short C-OP portfolios

[4]Hou et al. (2018) arrive at a similar conclusion that profitability anomalies remain significant to the q-factor model.

formed within firms with highest option-implied volatility.

Spreads in excess returns can come from different exposures to other factors. I next study the alphas of long-short portfolios after controlling for the Fama-French three factors. Table 3.3 shows that the economic and statistical significance of FF3F alphas both increase when moving from firms with low IF to those with high IF. For most combinations considered, the FF3F alphas are insignificant among firms with the lowest IF, while the most economically and statistically significant FF3F alphas concentrate among firms with high IF. For example, the highest FF3F alphas are obtained from long-short C-OP portfolios among firms with highest option-implied volatility, which are monthly 1.46% (t-statistics 4.06). In stark comparison, the FF3F alphas from portfolios sorting on the same profitability measure are merely monthly -0.17% (t-statistics -0.89) among firms with the lowest option-implied volatility.

3.3.3 Use principle components of IF proxies

For a parsimonious way to summarize the information contained in various IF proxies, I utilize the principal components (PCs) of these measures. To facilitate the extraction, I first perform a log transformation of the IF proxies. This is to ensure the skewness of these variables do not affect the linear projection used in the principal component analysis (PCA). The histograms from Figure E.1 show that the transformed variables are approximately normally distributed.

Figure E.2 shows that there is a strong level factor of the IF proxies I have considered.[5] The first PC, x_t, explains 56.3% of the variation in co-movement of IF proxies. To come up with a single measure for each individual firm in each month, I define a "PIF" measure as:

$$PIF_{t,i} = \sum_{k=1}^{5} e_k^1 IF_{t,k}, \tag{3.1}$$

[5] Here I pool all individual firms together to do PCA.

where e_k^1 denotes the first PC's loading on the k-th IF proxy in the pooling PCA, and $IF_{t,k}$ is the value of the firm's k-th IF proxy at time t. $\{PIF_{t,i}\}$ will serve as a firm-level sum-all measure of IF.

To have a non-missing PIF measure, the firm is required to: 1) followed by stock analysts and have both one-year ahead forecast and long-term growth forecast; 2) have outstanding stock options. This biases the sample towards big stocks, which further alleviates the concern that my result is driven by micro-cap stocks. Figure E.3 shows that the available stocks in the restricted sample still account for at least 65% and as high as 90% of total market capitalization over the sample period.

Table 3.4 reports the excess returns and FF3F alphas of the independently double-sorted portfolios on profitability measures and PIF. For C-OP, there is only statistically significant return spread among firms with the highest PIF; for GP/A and OP, there is no statistically significant return spread for any PIF quintile but the spreads increase almost monotonically with PIF. In terms of FF3F alphas, they all become statistically insignificant once the long-short portfolios are constructed among low-PIF firms. The highest alpha is still achieved from long-short C-OP portfolios in the fifth PIF quintile, which amounts to monthly monthly 1.38% (t-statistics 3.64). Moreover, the long leg of profitability strategy consistently contributes a sizable alpha.

For parsimony, I focus on using C-OP as the profitability measure, given that C-OP is the highest-alpha-generating anomaly out of the three (also see Table 3.1).

I visualize the economic magnitude of difference in anomaly performance across PIF quintiles. I plot the cumulative (sum of log) returns and FF3F alphas of these portfolios in Figure 3.1 and Figure 3.2, with a comparison to the original C-OP anomaly as in Ball et al. (2016). Figure 3.2 shows that the investors can earn much higher FF3F alphas by investing only among the firms with the highest IF and trading on the C-OP signal. In addition, the cumulative FF3F alphas of the long and short legs of the portfolios are plotted in Figure 3.3. Figure 3.3 shows that the cumulative alphas of both the long and short legs of the C-OP portfolios are highest among highest PIF quintiles.

3.3.4 Interpretation of IF proxies

The IF proxies I use can signal various aspects that differ across firms. For example, analyst forecast dispersion has been interpreted as a proxy of disagreement or information uncertainty (Diether et al. 2002; Patton and Timmermann 2010; Yu (2011)); option-implied volatility contains a risk premium component that reflects investors' risk attitude; age can relate to firm growth (Haltiwanger et al. 2013) or liquidity. In this section, I provide some evidence that support the interpretation that PIF can act as a proxy for information friction or difficulty of learning by examining how forecast errors evolve over time.

The forecast error is defined as

$$\text{Forecast Error} = \frac{\text{Mean}_{adj} - \text{Actual}_{adj}}{|\text{Actual}_{adj}|}. \tag{3.2}$$

I follow Diether et al. (2002) to combine information from unadjusted summary file and unadjusted detail history actuals from I/B/E/S and use CFACSHR from CRSP to adjust for share splits. I drop forecast errors with absolute values that are bigger than one to limit the effect of outliers, which amount to approximately 7% of total observations. I also drop stocks covered with less than three analysts. The forecast horizon is defined as the number of months between forecast issuing date and actual earnings announcement date (I/B/E/S variables "anndats" and "statpers"). Forecast errors are value-weighted by market capitalizations across firms for each forecast horizon.

Figure 3.4 plots the aggregate forecast errors for three PIF groups with 30th and 70th percentile as cutoff points. The figure yields the following observations. First, stock analysts on average are optimistic, consistent with previous literature (Bondt and Thaler 1990). Moreover, The magnitude of forecast errors are smaller for firms in the lower PIF tercile. For example, at the five-month horizon, average forecast errors of firms in the highest tercile are roughly twice of those of firms in the first and second tercile. Finally, the average forecast errors converge to zero as forecast horizon shortens for the lower PIF terciles, but the average forecast errors stay above zero for firms in the highest PIF tercile even for very short forecast horizons.

3.4 Possible Explanations

In this section, I examine whether existing explanations for the profitability anomaly from the literature can account for the empirical pattern that I have documented in the previous section. I consider both empirically and theoretically motivated explanations.

3.4.1 Short-sale constraints

Given that high PIF firms are young and volatile, one possibility is that PIF is a proxy for the short-sale constraints and my results indicate that the profitability anomaly represents mispricing (Stambaugh et al. 2012).

I first use instituitional ownership as a proxy for short-sale constraints (Nagel 2005). Table 3.5, Panel B, shows that there is no material difference in institutional ownership between long and short legs in the high PIF quintile. For example, in the highest PIF quintile, the average institutional ownership is 66% for low C-OP firms and 62% for high C-OP firms. I then consider transaction and holding costs as another source of friction (Pontiff 2006). Table 3.5, Panel C and D, show that there is also no material difference in idiosyncratic volatility or bid-ask spread between long and short legs in the high PIF quintile.[6] This indicates that short-sale constraints may not be the reason

[6] I follow Pontiff (2006) to use idiosyncratic volatility as a proxy for holding costs.

behind the high FF3F alphas of the profitability anomaly in the highest PIF quintile.

In addition, Table 3.4 shows that the long leg of the profitability anomaly, which is not subject to the short-sale constraints, contributes significantly to the FF3F alphas. For example, for the C-OP anomaly in the highest PIF quintile, the short leg contributes a monthly alpha of -0.78% (t-statistics -2.64) while the long leg contributes a comparable 0.59% (t-statistics 2.36).

3.4.2 Financial distress

Avramov et al. (2009) and Avramov et al. (2013) argue that the performance of many anomalies are derived from high credit risk firms with deteriorating conditions.

To test whether the PIF is just a proxy for financial distress and signals the high credit risk firms, I redo the portfolio sorting analysis after dropping the worst-rated firms.[7] Without the "worst-rated" firms, the remaining sample is considerably smaller judged by the number of firms and all three profitability anomalies become weaker in terms of FF3F alphas. To avoid portfolios that are too sparse, I use tercile-sorting on PIF with 30th and 70th percentiles as breakpoints.

Table 3.6 shows that the empirical pattern documented in the previous section persists. The OP and C-OP anomalies only have a statistically significant FF3F alpha–monthly 0.65% with t-statistics 2.23 and monthly 0.77% with t-statistics 2.67, respectively–in the highest PIF tercile, while the FF3F alphas of GP/A increase with PIF (but are all statistically insignificant).

3.4.3 Investment-based models

Lin and Zhang (2013) and Hou et al. (2014) motivate the inclusion of a profitability

[7]Following Avramov et al. (2009) and Avramov et al. (2013), I use S&P Long-Term Domestic Issuer Credit Rating from Compustat and label firms with a rating of BB+ or below as the "worst-rated" firms.

factor to capture the cross-section of stock returns in an investment-based framework. The expected stock return for firm i in their two-period setting, $\mathbb{E}_0 r_{i,1}$, satisfies:

$$\mathbb{E}_0 r_{i,1} = \frac{\mathbb{E}_0 \Pi_{i,1}}{1 + a(I_{i,0}/K_{i,0})} \tag{3.3}$$

where $\mathbb{E}_0 \Pi_{i,1}$ is the expected profitability, a is a constant that determines the adjustment cost, $I_{i,0}$ is the investment in the first period, and $K_{i,0}$ is the firm's capital in the first period.

To map to the empirical setting, previous work uses current profitability to proxy for the expected profitability. Thus, to be consistent with the empirical pattern documented in the previous section, it should be the case that current profitability only serves as a good proxy of expected profitability spread when PIF is high. This observation is not a priori clear in the standard investment-based framework.

3.4.4 Analyst bias

Bouchaud et al. (2019) considers a model of sticky expectation to explain the presence of the profitability anomaly. In their framework, the analysts have sticky expectation and underreact to information in past earnings. As a result, past profits predict future returns

$$cov(R_{t+1}, \pi_t) = (1 + m\rho)\frac{\rho}{1 - \lambda\rho^2}\lambda^2\sigma_u^2, \tag{3.4}$$

with $m = \frac{1-\lambda}{1+r-\rho}$. Here π_t is the current profit, ρ is the persistence of firm's profitability, σ_u is the standard deviation of the signal, r is the required rate of return from the investors, and a higher λ indicates a higher degree of stickiness of investors' expectation. This model shows that when the profitability signal is more volatile (a high σ_u), the signal has a stronger predicting power for future stock returns. Thus, if the model provides an explanation for the empirical pattern here, it should be the case that profitability is much more volatile in the highest PIF quintile. Still, the central focus of Bouchaud et al. (2019) is the stickiness λ but not σ_u.

3.5 A Simple Information-based Model

In the previous section, we see that most existing explanations–either empirically or theoretically motivated–do not directly speak to or fully account for the observation that the profitability anomaly only performs well in the high PIF quintiles.

In this section, I propose a simple information-based model in the spirit of Admati (1985) that can generate qualitative predictions that are consistent with the empirical patterns I have documented.

Environment The model has two periods. There is a continuum of investors indexed by $j \in [0, 1]$. The investors trade in period 1 and consume in period 2. Each investor can invest her initial wealth $W_{j,0}$ between a riskless asset and N risky assets. In period 2, the riskless asset pays r units and risky asset n pays f units of the single consumption good.[8] I assume that the stochastic payoffs $f \sim \mathcal{N}(0,\ \Sigma)$. The riskless asset is taken to be the numeraire and the price vector of the risky assets is denoted as p. The risky assets have stochastic supplies $\bar{x} + x$, where $\bar{x}$ is a constant vector and $x \sim \mathcal{N}(0,\ \sigma_x^2 \mathrm{I}_N)$.

Preference The investors have mean-variance preference with absolute risk aversion parameter ρ. They choose portfolio holdings q_j to maximize their utility:

$$U_j = \rho \tilde{\mathbb{E}}_j[W_j] - \frac{\rho^2}{2} \tilde{\mathrm{V}}_j[W_j] \tag{3.5}$$

subject to the budget constraint $W_j = rW_0 + q_j'(f - pr)$. For simplicity, I have assumed $W_{j,0} = W_0, \forall j$. $\tilde{\mathbb{E}}_j(\cdot)$ and $\tilde{\mathrm{V}}_j(\cdot)$ denote the mean and variance expected by investor j,

[8] Here r is exogenously given.

respectively.

Information The true mean and variance of asset payoffs are common knowledge to investors but the realization of those payoffs, f, are unknown in period 1. Each investor has a private signal in period 1 that is informative of future payoff realizations f, $\eta_j \sim \mathcal{N}(f, \Sigma_{\eta j})$.

Equilibrium Following Veldkamp (2011), with Bayesian belief updating and market clearing, a rational expectations equilibrium can be characterized by the investors' portfolio choices, $\{q_j\}$, and the price vector, p:

$$q_j = \frac{1}{\rho}\hat{\Sigma}_j^{-1}\left(\tilde{\mathbb{E}}_j[f] - pr\right), \quad \forall j \tag{3.6}$$

$$p = \frac{1}{r}(A + Bf + Cx), \tag{3.7}$$

where

$$\tilde{\mathbb{E}}_j[f] = \hat{\Sigma}_j(\Sigma_{\eta j}^{-1}\eta_j + \Sigma_p^{-1}\eta_p)\ ,\ \hat{\Sigma}_j^{-1} = \Sigma^{-1} + \Sigma_{\eta j}^{-1} + \Sigma_p^{-1}, \tag{3.8}$$

$$A = -\rho\bar{\Sigma}\bar{x}\ ,\ B = \mathrm{I}_N - \bar{\Sigma}\Sigma^{-1}\ ,\ \bar{\Sigma}^{-1} = \int \hat{\Sigma}_j^{-1} dj, \tag{3.9}$$

$$\eta_p = rp - A\ ,\ \Sigma_p = \mathrm{Var}(rp - A). \tag{3.10}$$

Econometrician's problem Now consider an outside econometrcian who simply examines the data ex post. She is able to observe the realized asset returns, $f - pr$, in period 2. To summarize the information in the cross-section of returns, she uses a lagged signal of realized payoffs f, $\zeta \sim (f, \Sigma_\zeta)$. This maps to the earlier empirical exercise of using a certain profitability signal to predict the cross-section of stock returns, as long as the profitability measure is informative of future cash flows from the stocks. It can be shown that

$$\beta \equiv \frac{\text{Cov}(\zeta, f - pr)}{V(\zeta)} = \bar{\Sigma}\Sigma^{-1}\frac{V(f)}{V(f) + \Sigma_\zeta} \tag{3.11}$$

$$= \left(\text{I}_N + \Sigma\bar{\Sigma}_\eta^{-1} + \frac{1}{\rho^2\sigma_x^2}\Sigma\bar{\Sigma}_\eta^{-1\prime}\bar{\Sigma}_\eta^{-1}\right)^{-1}\frac{V(f)}{V(f) + \Sigma_\zeta}, \tag{3.12}$$

where $\bar{\Sigma}_\eta^{-1} = \int \Sigma_{\eta j}^{-1} dj$ and $V(\cdot)$ Note that β represents the coefficient from regressing future realized returns on the lagged payoff signal.

Model's predictions Focusing on the coefficient β, this model yields several predictions. First, β increases in the payoff volatility under the econometrician's measure, $V(f)$. This echoes the prediction from Bouchaud et al. (2019) (see Section 3.4.4 for more details). Second, β increases in investors' average posterior uncertainty, $\bar{\Sigma}$, and decreases in their average signal precision, $\bar{\Sigma}_\eta^{-1}$. This is the key message from the model: for firms with higher uncertainty, the payoff signal predicts more strongly future asset returns. This is consistent with the empirical pattern documented in previous sections if we interpret PIF to be a proxy of information friction that makes the firms hard to learn.

3.6 Conclusion

This paper documents that the Fama-French three-factor (FF3F) alphas of the profitability anomaly exist only among firms with high information frictions (IF), proxied by young age, high forecast dispersion, high past return volatility, and/or high option-implied volatility. The FF3 alphas of the long-short profitability portfolios are insignificant both economically and statistically in low-IF quintiles, but are as high as monthly 1.46% (t-statistics 4.06) in high-IF quintiles. The results are robust to excluding micro-firms and using different measures of profitability. Short-sale constraints, financial distress, and standard investment-based framework do not fully account for or directly speak to this empirical pattern.

Instead of a risk-based explanation, I show that this empirical pattern is consistent with a noisy rational expectations equilibrium model in which investors learn about future firm payoffs from private signals and market prices. The model predicts that for firms with higher posterior uncertainty or lower signal precision, the econometrician should find a stronger relation between profitability measure and future returns.

3.7 Tables and Figures

Table 3.1: Capturing the Profitability Anomaly

This table reports alphas of profitability anomaly portfolios with respect to FF5F model from Fama and French (2015) and q-4 model from Hou et al. (2014). The measures of profitability include gross profitability (GP/A), operating profitability (OP), and cash-based operating profitability (C-OP). The definitions of these measures follow Novy-Marx (2013) and Ball et al. (2016), with a modification to GP/A where gross profit is deflated by the lagged book value of total assets. Portfolios are quintile-sorted on a given meausre of profitability with NYSE breakpoints. Portfolio returns are value-weighted by market capitalizations. Alphas are reported in monthly percentage points. Newey-West t-statistics with a lag of 12 months are reported in brackets. The sample period is from July 1963 to December 2016.

	GP/A		OP		C-OP	
Portfolios	FF5F (1)	q-4 (2)	FF5F (3)	q-4 (4)	FF5F (5)	q-4 (6)
Low	-0.10	-0.03	-0.18	-0.13	-0.26	-0.20
	[-1.76]	[-0.37]	[-2.50]	[-1.31]	[-3.62]	[-2.44]
2	-0.13	-0.05	-0.13	-0.10	-0.21	-0.17
	[-1.56]	[-0.55]	[-1.87]	[-1.06]	[-3.42]	[-2.16]
3	-0.05	0.03	-0.04	0.04	-0.06	-0.01
	[-0.88]	[0.39]	[-0.75]	[0.60]	[-1.13]	[-0.15]
4	0.00	0.04	-0.02	-0.00	-0.02	0.01
	[0.06]	[0.54]	[-0.38]	[-0.06]	[-0.37]	[0.22]
High	0.28	0.21	0.19	0.18	0.23	0.23
	[3.91]	[2.32]	[3.89]	[2.64]	[4.98]	[3.40]
High - Low	0.38	0.24	0.37	0.31	0.49	0.43
	[3.67]	[1.68]	[3.99]	[2.47]	[5.11]	[3.54]

Table 3.2: Excess Returns of Independently Double-Sorted Portfolios

This table reports the excess returns of anomaly portfolios that are independently double-sorted on PROF and IF. For IF proxies, I consider fiscal year 1 EPS forecast dispersion (1YFD), long-term growth EPS forecast dispersion (LTGFD), at-the-money option-implied volatility (OIV), lagged weekly return volatility (RetVol), and inverse of age (1/AGE). Portfolios are quintile sorted on each dimension. NYSE breakpoints are used and portfolio returns are value-weighted. Returns are reported in monthly percentage points. The sample starts from July 1963 to December 2016.

	GP/A						OP						C-OP					
Portfolios	Low	2	3	4	High	High - Low	Low	2	3	4	High	High - Low	Low	2	3	4	High	High - Low
Panel A: 1YFD																		
Low	0.70	0.44	0.71	0.63	0.76	0.06	0.73	0.59	0.60	0.70	0.73	-0.00	0.69	0.66	0.45	0.75	0.73	0.04
	[3.59]	[2.08]	[3.38]	[3.13]	[4.06]	[0.37]	[3.21]	[3.32]	[3.09]	[3.58]	[3.84]	[-0.01]	[3.11]	[3.33]	[2.43]	[3.87]	[3.88]	[0.22]
2	0.47	0.69	0.54	0.57	0.63	0.15	0.43	0.73	0.61	0.62	0.50	0.07	0.50	0.53	0.61	0.63	0.52	0.02
	[2.57]	[3.01]	[2.35]	[2.39]	[2.81]	[0.82]	[1.99]	[3.74]	[2.86]	[2.92]	[2.17]	[0.36]	[2.12]	[2.76]	[3.04]	[2.86]	[2.28]	[0.13]
3	0.48	0.60	0.63	0.62	0.74	0.26	0.50	0.66	0.63	0.64	0.62	0.12	0.35	0.63	0.62	0.57	0.68	0.33
	[2.37]	[2.55]	[2.66]	[2.27]	[2.75]	[1.22]	[2.14]	[3.16]	[2.76]	[2.71]	[2.31]	[0.62]	[1.39]	[3.10]	[2.86]	[2.42]	[2.58]	[1.80]
4	0.44	0.44	0.58	0.87	1.05	0.61	0.30	0.59	0.61	0.70	0.95	0.65	0.28	0.49	0.50	0.62	1.10	0.81
	[1.72]	[1.78]	[2.41]	[2.87]	[3.56]	[2.79]	[1.08]	[2.37]	[2.51]	[2.74]	[3.36]	[3.02]	[1.04]	[1.99]	[2.09]	[2.51]	[3.88]	[3.98]
High	0.47	0.52	0.65	0.46	0.49	0.02	0.31	0.63	0.67	0.47	0.82	0.51	0.28	0.58	0.75	0.72	0.93	0.65
	[1.58]	[1.82]	[2.22]	[1.32]	[1.33]	[0.08]	[1.04]	[2.17]	[2.25]	[1.52]	[2.36]	[2.22]	[0.93]	[2.10]	[2.52]	[2.37]	[2.71]	[2.76]
High - Low	-0.22	0.07	-0.06	-0.17	-0.27		-0.42	0.03	0.07	-0.24	0.10		-0.41	-0.08	0.30	-0.03	0.20	
	[-0.86]	[0.35]	[-0.27]	[-0.60]	[-0.84]		[-1.77]	[0.15]	[0.26]	[-0.99]	[0.32]		[-1.81]	[-0.38]	[1.27]	[-0.10]	[0.68]	
Panel B: LTGFD																		
Low	0.76	0.62	0.65	0.68	0.79	0.04	0.82	0.64	0.50	0.76	0.84	0.01	0.86	0.67	0.60	0.73	0.81	-0.06
	[3.78]	[2.75]	[2.88]	[3.14]	[3.78]	[0.18]	[3.65]	[3.39]	[2.41]	[3.71]	[4.14]	[0.07]	[3.63]	[3.54]	[3.09]	[3.49]	[4.05]	[-0.29]
2	0.58	0.78	0.65	0.71	1.02	0.44	0.54	0.79	0.79	0.87	0.83	0.28	0.60	0.53	0.75	0.86	0.84	0.25
	[2.83]	[3.56]	[2.83]	[3.18]	[4.75]	[2.32]	[2.34]	[3.70]	[3.80]	[3.76]	[4.01]	[1.52]	[2.48]	[2.56]	[3.65]	[3.82]	[3.96]	[1.31]
3	0.77	0.50	0.77	0.71	0.82	0.05	0.64	0.63	0.75	0.83	0.76	0.11	0.56	0.72	0.83	0.67	0.73	0.17
	[3.16]	[2.17]	[3.38]	[2.91]	[3.45]	[0.27]	[2.49]	[2.74]	[3.28]	[3.68]	[3.30]	[0.59]	[2.10]	[3.00]	[3.78]	[2.91]	[3.26]	[0.88]
4	0.56	0.49	0.79	0.78	0.93	0.38	0.48	0.71	0.61	0.80	0.96	0.48	0.26	0.68	0.70	0.70	0.95	0.69
	[2.18]	[1.90]	[2.98]	[2.56]	[3.09]	[1.55]	[1.80]	[2.81]	[2.37]	[3.06]	[3.19]	[2.07]	[0.94]	[2.61]	[2.84]	[2.56]	[3.18]	[3.25]
High	0.43	0.32	0.77	0.74	1.04	0.60	0.46	0.49	0.63	0.72	0.82	0.36	0.15	0.49	0.67	0.67	0.95	0.80
	[1.35]	[1.12]	[2.45]	[2.02]	[2.81]	[2.25]	[1.41]	[1.62]	[2.06]	[2.24]	[2.31]	[1.49]	[0.47]	[1.63]	[2.17]	[2.05]	[2.77]	[3.50]
High - Low	-0.32	-0.30	0.12	0.06	0.25		-0.36	-0.15	0.13	-0.04	-0.02		-0.71	-0.18	0.07	-0.06	0.15	
	[-1.20]	[-1.28]	[0.46]	[0.20]	[0.77]		[-1.39]	[-0.61]	[0.50]	[-0.14]	[-0.05]		[-2.76]	[-0.76]	[0.26]	[-0.23]	[0.49]	

Table 3.2: Excess Returns of Independently Double-Sorted Portfolios (Continued)

Portfolios	GP/A Low	GP/A 2	GP/A 3	GP/A 4	GP/A High	GP/A High - Low	OP Low	OP 2	OP 3	OP 4	OP High	OP High - Low	C-OP Low	C-OP 2	C-OP 3	C-OP 4	C-OP High	C-OP High - Low
Panel C: OIV																		
Low	0.64	0.51	0.62	0.34	0.58	-0.06	0.69	0.58	0.55	0.51	0.46	-0.23	0.71	0.67	0.57	0.45	0.52	-0.20
	[2.80]	[2.06]	[2.50]	[1.44]	[2.48]	[-0.29]	[2.69]	[2.68]	[2.46]	[2.19]	[2.00]	[-0.95]	[2.80]	[3.05]	[2.51]	[1.92]	[2.23]	[-0.78]
2	0.45	0.62	0.72	0.57	0.84	0.39	0.42	0.77	0.91	0.58	0.77	0.35	0.58	0.48	0.88	0.63	0.64	0.06
	[1.41]	[2.03]	[2.25]	[1.85]	[2.65]	[1.52]	[1.24]	[2.53]	[2.85]	[1.87]	[2.57]	[1.46]	[1.76]	[1.54]	[2.99]	[1.97]	[2.15]	[0.22]
3	0.64	0.56	0.68	0.44	1.39	0.76	0.35	0.67	0.82	1.06	0.80	0.45	0.32	0.57	0.78	0.91	0.90	0.57
	[1.69]	[1.50]	[1.81]	[1.13]	[3.85]	[2.38]	[0.91]	[1.81]	[2.18]	[2.98]	[2.09]	[1.57]	[0.82]	[1.63]	[2.22]	[2.52]	[2.34]	[1.93]
4	0.69	0.39	0.86	1.34	1.18	0.48	0.48	0.98	0.87	0.43	1.30	0.82	0.37	0.73	0.96	0.75	1.32	0.95
	[1.50]	[0.87]	[1.88]	[2.66]	[2.79]	[1.51]	[1.07]	[2.24]	[1.92]	[0.97]	[2.80]	[2.52]	[0.82]	[1.71]	[2.13]	[1.65]	[2.81]	[2.95]
High	0.07	0.46	0.63	1.16	1.08	1.02	0.08	0.84	0.51	0.69	1.09	1.01	-0.01	0.77	0.69	0.82	1.26	1.27
	[0.10]	[0.74]	[1.03]	[1.81]	[1.85]	[2.83]	[0.12]	[1.40]	[0.84]	[1.12]	[1.73]	[2.96]	[-0.01]	[1.23]	[1.20]	[1.26]	[2.03]	[4.01]
High - Low	-0.58	-0.05	0.01	0.82	0.50		-0.61	0.26	-0.04	0.18	0.63		-0.72	0.10	0.12	0.37	0.74	
	[-0.96]	[-0.10]	[0.01]	[1.35]	[0.90]		[-1.05]	[0.48]	[-0.07]	[0.33]	[1.04]		[-1.27]	[0.17]	[0.24]	[0.63]	[1.24]	
Panel D: RetVol																		
Low	0.46	0.49	0.67	0.46	0.65	0.19	0.50	0.44	0.50	0.62	0.53	0.03	0.60	0.45	0.53	0.55	0.58	-0.02
	[3.09]	[3.28]	[4.07]	[2.78]	[4.08]	[1.31]	[3.10]	[3.09]	[3.62]	[3.93]	[3.46]	[0.25]	[3.55]	[3.12]	[3.62]	[3.75]	[3.75]	[-0.16]
2	0.48	0.54	0.60	0.56	0.67	0.18	0.41	0.59	0.57	0.65	0.67	0.26	0.42	0.51	0.50	0.70	0.70	0.27
	[2.51]	[2.79]	[3.18]	[2.90]	[3.61]	[1.20]	[2.02]	[3.11]	[3.02]	[3.50]	[3.72]	[1.76]	[1.99]	[2.66]	[2.65]	[3.75]	[3.88]	[1.87]
3	0.52	0.57	0.72	0.75	0.84	0.32	0.56	0.79	0.65	0.69	0.74	0.19	0.55	0.50	0.73	0.77	0.74	0.20
	[2.23]	[2.66]	[3.31]	[3.31]	[3.85]	[1.95]	[2.33]	[3.56]	[3.01]	[3.26]	[3.44]	[1.25]	[2.37]	[2.23]	[3.41]	[3.72]	[3.45]	[1.39]
4	0.61	0.39	0.54	0.64	0.75	0.14	0.48	0.55	0.62	0.51	0.67	0.19	0.46	0.37	0.68	0.47	0.74	0.28
	[2.26]	[1.53]	[2.14]	[2.31]	[2.85]	[0.73]	[1.80]	[2.17]	[2.41]	[2.02]	[2.53]	[1.08]	[1.74]	[1.48]	[2.72]	[1.85]	[2.77]	[1.76]
High	0.26	0.37	0.49	0.53	0.93	0.68	0.15	0.61	0.50	0.42	0.80	0.64	-0.06	0.67	0.59	0.58	0.94	1.01
	[0.74]	[1.13]	[1.46]	[1.59]	[2.76]	[3.80]	[0.45]	[1.85]	[1.55]	[1.33]	[2.37]	[3.80]	[-0.19]	[2.01]	[1.90]	[1.73]	[2.84]	[6.45]
High - Low	-0.20	-0.13	-0.18	0.07	0.28		-0.34	0.17	0.00	-0.19	0.27		-0.66	0.22	0.06	0.03	0.37	
	[-0.64]	[-0.48]	[-0.61]	[0.26]	[0.98]		[-1.13]	[0.60]	[0.01]	[-0.74]	[0.94]		[-2.28]	[0.73]	[0.23]	[0.09]	[1.31]	
Panel E: 1/AGE																		
Low	0.47	0.52	0.58	0.44	0.63	0.16	0.48	0.54	0.55	0.58	0.51	0.03	0.52	0.41	0.54	0.54	0.55	0.03
	[2.75]	[2.83]	[3.27]	[2.51]	[3.71]	[1.14]	[2.33]	[3.16]	[3.31]	[3.29]	[3.20]	[0.23]	[2.51]	[2.37]	[3.18]	[3.25]	[3.40]	[0.22]
2	0.47	0.52	0.56	0.48	0.52	0.05	0.53	0.44	0.48	0.56	0.48	-0.06	0.40	0.42	0.57	0.53	0.59	0.18
	[2.75]	[2.67]	[2.76]	[2.17]	[2.91]	[0.32]	[2.75]	[2.43]	[2.63]	[3.05]	[2.46]	[-0.38]	[1.89]	[2.34]	[3.20]	[2.94]	[3.18]	[1.18]
3	0.35	0.48	0.57	0.68	0.59	0.23	0.45	0.48	0.48	0.69	0.57	0.12	0.37	0.46	0.44	0.69	0.65	0.28
	[1.76]	[2.15]	[2.78]	[2.99]	[2.83]	[1.50]	[1.94]	[2.33]	[2.36]	[3.43]	[2.66]	[0.79]	[1.55]	[2.22]	[2.22]	[3.50]	[2.98]	[1.62]
4	0.26	0.46	0.74	0.71	0.89	0.63	0.23	0.60	0.63	0.62	0.82	0.59	0.20	0.49	0.65	0.67	0.82	0.61
	[1.09]	[2.17]	[3.13]	[2.84]	[3.66]	[3.69]	[0.88]	[2.78]	[2.87]	[2.72]	[3.43]	[3.63]	[0.78]	[2.20]	[2.97]	[3.09]	[3.48]	[4.18]
High	0.35	0.35	0.50	0.53	0.81	0.46	0.20	0.58	0.56	0.47	0.71	0.52	0.11	0.60	0.70	0.56	0.83	0.72
	[1.35]	[1.44]	[1.82]	[1.97]	[3.05]	[3.14]	[0.72]	[2.42]	[2.34]	[1.90]	[2.75]	[3.59]	[0.39]	[2.52]	[2.96]	[2.29]	[3.26]	[5.25]
High - Low	-0.12	-0.18	-0.09	0.09	0.18		-0.28	0.04	0.02	-0.11	0.20		-0.41	0.19	0.15	0.02	0.28	
	[-0.66]	[-1.13]	[-0.49]	[0.46]	[0.85]		[-1.44]	[0.28]	[0.10]	[-0.68]	[1.03]		[-2.07]	[1.16]	[0.95]	[0.10]	[1.49]	

Table 3.3: Alphas of Independently Double-Sorted Portfolios

This table reports the FF3F alphas of anomaly portfolios that are independently double-sorted on PROF and IF. For IF proxies, I consider fiscal year 1 EPS forecast dispersion (1YFD), long-term growth EPS forecast dispersion (LTGFD), at-the-money option-implied volatility (OIV), lagged weekly return volatility (RetVol), and inverse of age (1/AGE). Portfolios are quintile sorted on each dimension. NYSE breakpoints are used and portfolio returns are value-weighted. Alphas are reported in monthly percentage points. Newey-West standard errors with a lag of 12 months are reported in brackets. The sample starts from July 1963 to December 2016.

	GP/A						OP						C-OP					
Portfolios	Low	2	3	4	High	High - Low	Low	2	3	4	High	High - Low	Low	2	3	4	High	High - Low
Panel A: 1YFD																		
Low	0.30	-0.10	0.17	0.20	0.38	0.09	0.24	0.06	0.11	0.20	0.35	0.11	0.15	0.17	-0.03	0.26	0.36	0.21
	[1.97]	[-0.80]	[1.75]	[2.16]	[3.22]	[0.43]	[1.65]	[0.40]	[0.92]	[2.37]	[3.51]	[0.65]	[1.11]	[1.32]	[-0.25]	[2.82]	[3.57]	[1.31]
2	-0.02	0.09	-0.07	0.04	0.12	0.14	-0.15	0.17	0.04	0.02	0.02	0.17	-0.14	-0.03	0.08	0.05	0.03	0.17
	[-0.18]	[0.59]	[-0.52]	[0.40]	[1.13]	[0.90]	[-0.87]	[1.40]	[0.35]	[0.20]	[0.20]	[0.84]	[-0.84]	[-0.27]	[0.57]	[0.57]	[0.30]	[0.91]
3	-0.12	-0.09	0.03	0.06	0.19	0.31	-0.17	0.00	-0.06	0.06	0.11	0.28	-0.38	0.02	-0.07	-0.06	0.18	0.56
	[-0.89]	[-0.69]	[0.21]	[0.47]	[1.23]	[1.39]	[-1.11]	[0.01]	[-0.45]	[0.47]	[0.82]	[1.71]	[-2.96]	[0.19]	[-0.67]	[-0.51]	[1.43]	[2.85]
4	-0.31	-0.36	-0.06	0.28	0.39	0.70	-0.51	-0.26	-0.19	0.01	0.45	0.96	-0.58	-0.35	-0.25	-0.07	0.58	1.16
	[-2.27]	[-2.75]	[-0.49]	[1.47]	[2.72]	[3.30]	[-2.96]	[-2.49]	[-1.55]	[0.08]	[2.74]	[4.05]	[-5.04]	[-2.95]	[-2.70]	[-0.45]	[3.47]	[5.74]
High	-0.43	-0.38	-0.15	-0.30	-0.37	0.07	-0.64	-0.30	-0.21	-0.36	0.09	0.72	-0.68	-0.33	-0.20	-0.12	0.20	0.88
	[-2.56]	[-2.04]	[-1.04]	[-1.79]	[-1.42]	[0.21]	[-3.73]	[-1.86]	[-1.30]	[-1.90]	[0.41]	[2.94]	[-3.89]	[-2.00]	[-1.09]	[-0.65]	[0.92]	[3.45]
High - Low	-0.73	-0.27	-0.33	-0.50	-0.75		-0.87	-0.36	-0.32	-0.57	-0.26		-0.83	-0.49	-0.17	-0.38	-0.16	
	[-3.19]	[-1.13]	[-1.80]	[-2.43]	[-2.57]		[-4.02]	[-1.61]	[-1.55]	[-2.55]	[-1.11]		[-3.75]	[-2.12]	[-0.70]	[-1.77]	[-0.67]	
Panel B: LTGFD																		
Low	0.25	-0.08	0.04	0.16	0.30	0.05	0.27	0.07	-0.08	0.20	0.39	0.13	0.23	0.11	0.05	0.20	0.35	0.12
	[1.61]	[-0.65]	[0.23]	[1.14]	[2.15]	[0.26]	[1.68]	[0.51]	[-0.51]	[1.53]	[3.11]	[0.60]	[1.36]	[0.82]	[0.37]	[1.49]	[2.95]	[0.56]
2	-0.00	0.13	-0.00	0.13	0.52	0.52	-0.08	0.12	0.19	0.23	0.35	0.43	-0.06	-0.08	0.14	0.26	0.35	0.40
	[-0.01]	[0.97]	[-0.03]	[1.20]	[4.44]	[3.15]	[-0.55]	[0.99]	[1.65]	[2.49]	[2.83]	[2.60]	[-0.40]	[-0.60]	[1.06]	[2.83]	[2.70]	[2.65]
3	-0.02	-0.25	0.13	0.09	0.25	0.27	-0.16	-0.11	0.06	0.14	0.20	0.36	-0.26	-0.04	0.20	0.00	0.16	0.42
	[-0.12]	[-2.03]	[0.88]	[0.83]	[1.85]	[1.43]	[-1.04]	[-0.85]	[0.59]	[1.24]	[1.59]	[1.84]	[-1.61]	[-0.29]	[1.59]	[0.02]	[1.34]	[2.10]
4	-0.22	-0.30	0.17	0.20	0.37	0.59	-0.34	-0.05	-0.18	0.11	0.43	0.77	-0.60	-0.11	0.00	0.01	0.38	0.98
	[-1.58]	[-2.15]	[0.94]	[1.41]	[2.47]	[2.73]	[-2.19]	[-0.36]	[-1.25]	[0.78]	[2.96]	[4.02]	[-3.95]	[-0.72]	[0.00]	[0.10]	[2.50]	[5.01]
High	-0.36	-0.49	0.14	0.11	0.47	0.83	-0.34	-0.31	-0.17	0.01	0.30	0.64	-0.67	-0.32	-0.13	-0.02	0.40	1.07
	[-1.83]	[-2.87]	[0.80]	[0.72]	[2.55]	[3.20]	[-1.70]	[-1.84]	[-1.10]	[0.06]	[2.16]	[2.68]	[-3.32]	[-1.99]	[-0.75]	[-0.11]	[2.94]	[4.40]
High - Low	-0.61	-0.41	0.11	-0.05	0.17		-0.61	-0.38	-0.08	-0.19	-0.09		-0.90	-0.43	-0.18	-0.22	0.05	
	[-2.42]	[-1.89]	[0.48]	[-0.22]	[0.71]		[-2.23]	[-1.58]	[-0.35]	[-0.81]	[-0.46]		[-3.22]	[-1.93]	[-0.78]	[-0.83]	[0.28]	

Table 3.3: Alphas of Independently Double-Sorted Portfolios (Continued)

Portfolios	GP/A Low	GP/A 2	GP/A 3	GP/A 4	GP/A High	GP/A High - Low	OP Low	OP 2	OP 3	OP 4	OP High	OP High - Low	C-OP Low	C-OP 2	C-OP 3	C-OP 4	C-OP High	C-OP High - Low
Panel C: OIV																		
Low	0.32	0.04	0.21	0.02	0.28	-0.04	0.31	0.25	0.19	0.08	0.16	-0.15	0.39	0.33	0.20	0.03	0.22	-0.17
	[2.31]	[0.31]	[1.40]	[0.17]	[1.97]	[-0.23]	[1.96]	[1.75]	[1.26]	[0.54]	[1.34]	[-0.90]	[2.17]	[2.33]	[1.35]	[0.24]	[1.97]	[-0.89]
2	-0.14	-0.02	0.11	0.06	0.31	0.45	-0.20	0.13	0.28	-0.05	0.31	0.51	-0.05	-0.16	0.27	0.03	0.18	0.23
	[-0.72]	[-0.09]	[0.60]	[0.37]	[2.13]	[1.66]	[-1.02]	[0.80]	[2.30]	[-0.33]	[2.39]	[1.90]	[-0.29]	[-0.84]	[2.01]	[0.23]	[1.30]	[1.11]
3	-0.13	-0.20	-0.06	-0.16	0.78	0.91	-0.45	-0.12	0.04	0.40	0.24	0.69	-0.46	-0.19	0.04	0.22	0.34	0.80
	[-0.53]	[-1.02]	[-0.33]	[-0.62]	[3.46]	[2.62]	[-2.16]	[-0.85]	[0.18]	[2.51]	[1.25]	[2.46]	[-2.27]	[-1.13]	[0.21]	[1.28]	[1.78]	[3.03]
4	-0.20	-0.48	0.03	0.60	0.44	0.63	-0.44	0.10	-0.01	-0.37	0.61	1.05	-0.55	-0.17	0.04	-0.03	0.63	1.18
	[-0.63]	[-2.00]	[0.11]	[2.33]	[1.91]	[1.89]	[-1.72]	[0.64]	[-0.04]	[-1.44]	[2.40]	[3.24]	[-2.41]	[-0.96]	[0.17]	[-0.13]	[2.47]	[3.77]
High	-1.00	-0.69	-0.38	0.19	0.22	1.22	-1.01	-0.26	-0.58	-0.30	0.18	1.19	-1.12	-0.35	-0.38	-0.13	0.34	1.46
	[-3.22]	[-1.98]	[-1.45]	[0.68]	[0.65]	[2.98]	[-2.59]	[-0.81]	[-2.19]	[-0.95]	[0.58]	[2.66]	[-4.54]	[-1.06]	[-1.83]	[-0.29]	[1.26]	[4.06]
High - Low	-1.32	-0.73	-0.59	0.17	-0.06		-1.33	-0.51	-0.76	-0.38	0.02		-1.51	-0.68	-0.58	-0.16	0.12	
	[-3.67]	[-2.03]	[-1.84]	[0.49]	[-0.15]		[-3.06]	[-1.39]	[-2.36]	[-1.01]	[0.05]		[-4.60]	[-1.85]	[-2.11]	[-0.33]	[0.35]	
Panel D: RetVol																		
Low	0.02	0.05	0.18	0.07	0.35	0.33	0.02	0.00	0.10	0.16	0.18	0.17	0.13	0.03	0.08	0.14	0.23	0.10
	[0.20]	[0.59]	[1.98]	[0.60]	[2.89]	[2.13]	[0.15]	[0.03]	[1.41]	[1.80]	[1.81]	[1.04]	[1.01]	[0.24]	[0.98]	[1.52]	[2.19]	[0.68]
2	-0.09	-0.07	0.01	0.12	0.29	0.38	-0.23	-0.01	-0.03	0.10	0.31	0.54	-0.21	-0.08	-0.08	0.16	0.32	0.54
	[-0.94]	[-0.67]	[0.13]	[1.16]	[3.18]	[2.39]	[-1.95]	[-0.09]	[-0.37]	[1.23]	[3.83]	[3.28]	[-1.78]	[-0.83]	[-0.75]	[2.34]	[4.00]	[3.39]
3	-0.20	-0.09	0.10	0.22	0.34	0.54	-0.17	0.03	-0.02	0.07	0.31	0.48	-0.15	-0.25	0.07	0.16	0.30	0.45
	[-1.65]	[-0.69]	[0.89]	[2.17]	[3.04]	[2.99]	[-1.30]	[0.24]	[-0.25]	[0.62]	[3.09]	[3.19]	[-1.34]	[-2.13]	[0.68]	[1.69]	[2.88]	[3.27]
4	-0.23	-0.33	-0.21	0.01	0.16	0.39	-0.40	-0.26	-0.14	-0.14	0.12	0.52	-0.39	-0.41	-0.10	-0.20	0.18	0.57
	[-1.33]	[-2.40]	[-1.74]	[0.11]	[1.31]	[1.83]	[-2.59]	[-2.28]	[-1.15]	[-1.10]	[1.16]	[2.90]	[-2.75]	[-3.56]	[-0.90]	[-1.51]	[1.61]	[3.18]
High	-0.60	-0.52	-0.30	-0.22	0.24	0.84	-0.76	-0.31	-0.33	-0.33	0.18	0.95	-0.98	-0.24	-0.25	-0.16	0.29	1.27
	[-3.28]	[-3.00]	[-1.94]	[-1.36]	[1.44]	[4.28]	[-4.04]	[-1.98]	[-2.03]	[-2.07]	[1.01]	[4.95]	[-6.58]	[-1.59]	[-1.64]	[-0.83]	[1.65]	[6.86]
High - Low	-0.62	-0.57	-0.48	-0.29	-0.11		-0.78	-0.31	-0.43	-0.50	-0.00		-1.10	-0.27	-0.33	-0.30	0.06	
	[-2.56]	[-2.89]	[-2.31]	[-1.24]	[-0.47]		[-3.08]	[-1.68]	[-2.27]	[-2.29]	[-0.00]		[-5.29]	[-1.30]	[-1.72]	[-1.21]	[0.26]	
Panel E: 1/AGE																		
Low	-0.10	-0.06	0.06	0.07	0.32	0.42	-0.20	-0.04	0.04	0.06	0.19	0.39	-0.10	-0.14	-0.01	0.06	0.22	0.32
	[-1.37]	[-0.57]	[0.78]	[0.88]	[2.61]	[2.61]	[-1.99]	[-0.51]	[0.55]	[0.69]	[2.34]	[2.66]	[-0.87]	[-1.87]	[-0.15]	[0.82]	[2.65]	[2.03]
2	-0.12	-0.07	-0.05	-0.02	0.20	0.32	-0.13	-0.16	-0.13	0.06	0.11	0.25	-0.30	-0.20	-0.02	0.05	0.19	0.49
	[-1.33]	[-0.62]	[-0.49]	[-0.22]	[1.94]	[2.32]	[-1.26]	[-1.79]	[-1.14]	[0.64]	[1.44]	[1.66]	[-2.65]	[-1.98]	[-0.25]	[0.51]	[2.55]	[3.48]
3	-0.23	-0.11	0.01	0.19	0.16	0.39	-0.28	-0.13	-0.03	0.18	0.15	0.43	-0.38	-0.16	-0.03	0.14	0.23	0.61
	[-2.51]	[-1.01]	[0.14]	[1.90]	[1.49]	[2.75]	[-2.30]	[-1.21]	[-0.33]	[2.33]	[1.37]	[2.80]	[-3.24]	[-1.93]	[-0.26]	[1.92]	[1.95]	[3.46]
4	-0.35	-0.18	0.17	0.19	0.49	0.84	-0.51	-0.04	-0.03	0.07	0.44	0.94	-0.54	-0.15	0.01	0.15	0.43	0.97
	[-2.61]	[-1.80]	[1.44]	[1.74]	[3.81]	[4.14]	[-3.99]	[-0.41]	[-0.40]	[0.74]	[4.33]	[6.00]	[-4.58]	[-1.55]	[0.16]	[1.46]	[4.90]	[6.43]
High	-0.24	-0.25	-0.09	0.01	0.34	0.58	-0.43	-0.10	-0.03	-0.12	0.30	0.72	-0.57	-0.02	0.11	0.04	0.39	0.96
	[-1.51]	[-2.07]	[-0.84]	[0.11]	[3.01]	[3.45]	[-2.46]	[-0.86]	[-0.28]	[-1.07]	[3.18]	[4.33]	[-3.91]	[-0.18]	[1.12]	[0.37]	[3.90]	[6.39]
High - Low	-0.13	-0.20	-0.16	-0.06	0.02		-0.23	-0.05	-0.07	-0.17	0.11		-0.47	0.13	0.12	-0.02	0.17	
	[-0.92]	[-1.60]	[-1.18]	[-0.45]	[0.15]		[-1.34]	[-0.51]	[-0.64]	[-1.37]	[0.78]		[-2.67]	[1.19]	[0.99]	[-0.13]	[1.16]	

Table 3.4: Portfolios Independently Double-Sorted on PROF and PIF

This table reports the excess returns and FF3F alphas of anomaly portfolios that are independently double-sorted on PROF and PIF. PIF is defined as in Equation (3.1) using the first PC of the log of IF proxies. Portfolios are quintile sorted on each dimension. NYSE breakpoints are used and portfolio returns are value-weighted. Returns are reported in monthly percentage points. The sample starts from February 1996 to December 2016.

	GP/A						OP						C-OP					
Portfolios	Low	2	3	4	High	High - Low	Low	2	3	4	High	High - Low	Low	2	3	4	High	High - Low
Panel A: Excess Returns																		
Low	0.50	0.48	0.54	0.39	0.71	0.21	0.57	0.53	0.49	0.71	0.53	-0.03	0.86	0.48	0.52	0.52	0.60	-0.27
	[1.95]	[1.69]	[2.07]	[1.64]	[2.93]	[0.82]	[1.96]	[2.27]	[2.03]	[2.76]	[2.25]	[-0.12]	[3.04]	[1.91]	[2.20]	[1.99]	[2.46]	[-1.04]
2	0.57	0.51	0.72	0.82	0.83	0.27	0.42	0.81	0.57	0.68	0.82	0.40	0.36	0.61	0.71	0.64	0.77	0.40
	[1.77]	[1.79]	[2.55]	[2.56]	[2.73]	[0.98]	[1.32]	[2.70]	[1.87]	[2.25]	[3.00]	[1.49]	[1.14]	[2.05]	[2.56]	[2.11]	[2.64]	[1.42]
3	0.68	0.51	0.69	0.71	1.10	0.42	0.49	0.71	0.81	0.56	0.85	0.35	0.66	0.66	0.67	0.80	0.98	0.32
	[1.94]	[1.57]	[2.06]	[1.73]	[2.92]	[1.30]	[1.37]	[1.99]	[2.48]	[1.63]	[2.21]	[1.01]	[1.86]	[1.96]	[1.88]	[2.38]	[2.45]	[0.88]
4	0.65	0.33	0.74	1.06	0.95	0.30	0.57	0.29	0.67	0.89	0.96	0.39	0.27	0.60	0.71	0.70	1.01	0.74
	[1.67]	[0.85]	[1.75]	[2.05]	[2.09]	[0.79]	[1.51]	[0.72]	[1.62]	[2.17]	[1.99]	[0.97]	[0.67]	[1.59]	[1.82]	[1.70]	[2.04]	[2.02]
High	0.25	0.54	1.13	1.08	0.93	0.68	0.53	0.86	0.59	0.91	0.97	0.44	0.15	0.60	0.90	0.58	1.27	1.12
	[0.49]	[1.05]	[2.05]	[1.82]	[1.63]	[1.85]	[0.99]	[1.63]	[0.99]	[1.73]	[1.72]	[1.26]	[0.30]	[1.21]	[1.67]	[0.95]	[2.31]	[3.17]
High - Low	-0.25	0.06	0.59	0.69	0.22		-0.04	0.34	0.10	0.20	0.44		-0.71	0.12	0.38	0.06	0.67	
	[-0.55]	[0.14]	[1.18]	[1.23]	[0.40]		[-0.09]	[0.70]	[0.18]	[0.43]	[0.80]		[-1.53]	[0.26]	[0.73]	[0.10]	[1.26]	
Panel B: FF3F Alphas																		
Low	0.20	-0.04	0.11	0.07	0.43	0.22	0.25	0.17	0.11	0.28	0.25	0.00	0.52	0.10	0.15	0.14	0.30	-0.23
	[1.00]	[-0.20]	[0.65]	[0.53]	[2.47]	[0.97]	[1.21]	[0.94]	[0.54]	[1.77]	[1.78]	[0.01]	[2.12]	[0.56]	[0.87]	[0.82]	[2.01]	[-1.03]
2	-0.03	-0.04	0.27	0.28	0.34	0.37	-0.11	0.21	0.05	0.10	0.37	0.48	-0.18	0.03	0.23	0.07	0.31	0.49
	[-0.16]	[-0.27]	[1.66]	[1.48]	[2.18]	[1.38]	[-0.50]	[1.24]	[0.32]	[0.72]	[2.73]	[1.75]	[-0.68]	[0.18]	[1.68]	[0.50]	[2.19]	[1.53]
3	0.00	-0.14	0.11	0.16	0.52	0.52	-0.19	0.02	0.13	-0.08	0.39	0.58	-0.07	-0.00	-0.02	0.18	0.49	0.56
	[0.02]	[-0.83]	[0.65]	[0.82]	[1.94]	[1.39]	[-0.78]	[0.10]	[0.81]	[-0.55]	[1.90]	[1.66]	[-0.29]	[-0.01]	[-0.12]	[1.01]	[2.22]	[1.52]
4	-0.13	-0.40	0.03	0.43	0.33	0.46	-0.23	-0.48	-0.12	0.18	0.40	0.63	-0.56	-0.16	-0.02	-0.01	0.41	0.98
	[-0.66]	[-1.84]	[0.12]	[1.92]	[1.26]	[1.28]	[-1.37]	[-1.96]	[-0.63]	[0.71]	[2.02]	[2.46]	[-4.25]	[-0.71]	[-0.08]	[-0.04]	[1.99]	[3.87]
High	-0.61	-0.32	0.40	0.36	0.18	0.79	-0.35	0.01	-0.32	0.21	0.29	0.64	-0.78	-0.25	0.05	-0.12	0.59	1.38
	[-2.09]	[-0.97]	[1.18]	[1.50]	[0.51]	[2.21]	[-1.06]	[0.05]	[-1.13]	[0.70]	[1.10]	[1.84]	[-2.64]	[-0.93]	[0.18]	[-0.26]	[2.36]	[3.64]
High - Low	-0.81	-0.28	0.29	0.29	-0.25		-0.60	-0.16	-0.43	-0.07	0.04		-1.31	-0.35	-0.10	-0.25	0.29	
	[-2.23]	[-0.63]	[0.73]	[0.96]	[-0.55]		[-1.46]	[-0.50]	[-1.26]	[-0.24]	[0.11]		[-3.52]	[-1.16]	[-0.31]	[-0.49]	[1.00]	

Table 3.5: Characteristics of Portfolios Independently Double-Sorted on C-OP and PIF

This table reports the characteristics of anomaly portfolios that are independently double-sorted on C-OP and PIF. PIF is defined as in Equation (3.1) using the first PC of the log of IF proxies. Portfolios are quintile sorted on each dimension. NYSE breakpoints are used and characteristics are value-weighted. Market capitalization are in billions. The sample starts from February 1996 to December 2016.

	C-OP				
PIF	Low	2	3	4	High
Panel A.1: Market Capitalization					
Low	380	295	475	588	890
2	188	247	377	463	622
3	176	212	249	333	658
4	153	178	173	277	834
High	225	168	162	221	565
Panel B: Institutional Ownership					
Low	0.54	0.57	0.56	0.60	0.58
2	0.64	0.64	0.61	0.62	0.63
3	0.66	0.65	0.66	0.65	0.60
4	0.68	0.68	0.66	0.67	0.57
High	0.66	0.66	0.67	0.66	0.62
Panel C: Idiosyncratic Volatility					
Low	0.010	0.011	0.011	0.011	0.010
2	0.013	0.013	0.012	0.012	0.012
3	0.014	0.014	0.014	0.014	0.014
4	0.017	0.016	0.016	0.016	0.016
High	0.021	0.020	0.020	0.020	0.019
Panel D: Bid-Ask Spread					
Low	0.021	0.021	0.021	0.021	0.020
2	0.025	0.024	0.022	0.023	0.024
3	0.026	0.026	0.025	0.026	0.026
4	0.031	0.029	0.030	0.030	0.030
High	0.037	0.035	0.036	0.037	0.036

Table 3.6: Portfolios Independently Double-Sorted on PROF and PIF without Worst-Rated Firms

This table reports the excess returns and FF3F alphas of anomaly portfolios that are independently double-sorted on PROF and PIF. PIF is defined as in Equation (3.1) using the first PC of the log of IF proxies. The sample excludes firms with a S&P Long-Term Domestic Issuer Credit Rating of BB+ or below. Portfolios are quintile sorted on each dimension. NYSE breakpoints are used and portfolio returns are value-weighted. Returns are reported in monthly percentage points. The sample starts from February 1996 to December 2016.

	GP/A						OP						C-OP					
Portfolios	Low	2	3	4	High	High - Low	Low	2	3	4	High	High - Low	Low	2	3	4	High	High - Low
Panel A: Excess Returns																		
Low	0.70	0.62	0.69	0.57	0.55	-0.15	0.52	0.71	0.58	0.52	0.65	0.13	0.75	0.50	0.61	0.64	0.57	-0.18
	[2.68]	[1.75]	[2.16]	[2.15]	[2.06]	[-0.58]	[1.84]	[2.40]	[1.85]	[1.90]	[2.62]	[0.50]	[2.51]	[1.77]	[1.84]	[2.36]	[2.32]	[-0.60]
2	0.71	0.50	0.66	0.79	0.95	0.24	0.59	0.93	0.64	0.58	0.81	0.22	0.69	0.65	0.65	0.66	0.81	0.12
	[2.21]	[1.70]	[2.20]	[1.94]	[3.08]	[0.82]	[1.89]	[2.99]	[2.09]	[1.79]	[2.52]	[0.74]	[2.15]	[2.21]	[2.20]	[1.94]	[2.37]	[0.42]
High	0.49	0.24	0.83	1.06	0.93	0.44	0.54	0.63	0.48	0.77	1.11	0.57	0.50	0.32	0.82	0.10	1.14	0.65
	[1.28]	[0.62]	[1.97]	[2.10]	[1.78]	[1.01]	[1.48]	[1.58]	[1.12]	[1.92]	[2.28]	[1.50]	[1.26]	[0.82]	[1.94]	[0.23]	[2.35]	[1.90]
High - Low	-0.21	-0.38	0.14	0.50	0.38		0.02	-0.08	-0.10	0.25	0.46		-0.26	-0.18	0.21	-0.54	0.57	
	[-0.69]	[-1.34]	[0.37]	[1.15]	[0.75]		[0.08]	[-0.22]	[-0.27]	[0.74]	[1.00]		[-0.77]	[-0.61]	[0.55]	[-1.42]	[1.24]	
Panel B: FF3F Alphas																		
Low	0.27	-0.01	0.18	0.21	0.22	-0.05	0.07	0.22	0.15	0.08	0.32	0.25	0.29	0.04	0.07	0.24	0.24	-0.06
	[1.29]	[-0.06]	[0.87]	[1.39]	[1.25]	[-0.19]	[0.29]	[1.11]	[0.66]	[0.44]	[2.26]	[0.90]	[1.22]	[0.20]	[0.35]	[1.13]	[1.67]	[-0.21]
2	0.10	-0.01	0.15	0.25	0.47	0.37	0.02	0.37	0.08	-0.01	0.36	0.34	0.08	0.12	0.10	0.10	0.34	0.26
	[0.51]	[-0.05]	[0.92]	[1.20]	[2.57]	[1.52]	[0.12]	[2.30]	[0.44]	[-0.05]	[2.14]	[1.27]	[0.50]	[0.83]	[0.64]	[0.50]	[2.19]	[1.18]
High	-0.20	-0.40	0.22	0.44	0.41	0.60	-0.07	-0.07	-0.29	0.17	0.59	0.65	-0.17	-0.46	0.20	-0.46	0.60	0.77
	[-1.05]	[-1.95]	[0.90]	[1.83]	[1.32]	[1.55]	[-0.34]	[-0.31]	[-1.00]	[0.63]	[2.63]	[2.23]	[-1.06]	[-1.99]	[0.87]	[-1.58]	[2.43]	[2.67]
High - Low	-0.47	-0.38	0.04	0.23	0.19		-0.14	-0.29	-0.44	0.09	0.27		-0.47	-0.50	0.13	-0.70	0.36	
	[-1.74]	[-1.59]	[0.13]	[0.94]	[0.47]		[-0.52]	[-0.91]	[-1.51]	[0.32]	[0.95]		[-1.82]	[-1.99]	[0.50]	[-1.98]	[1.24]	

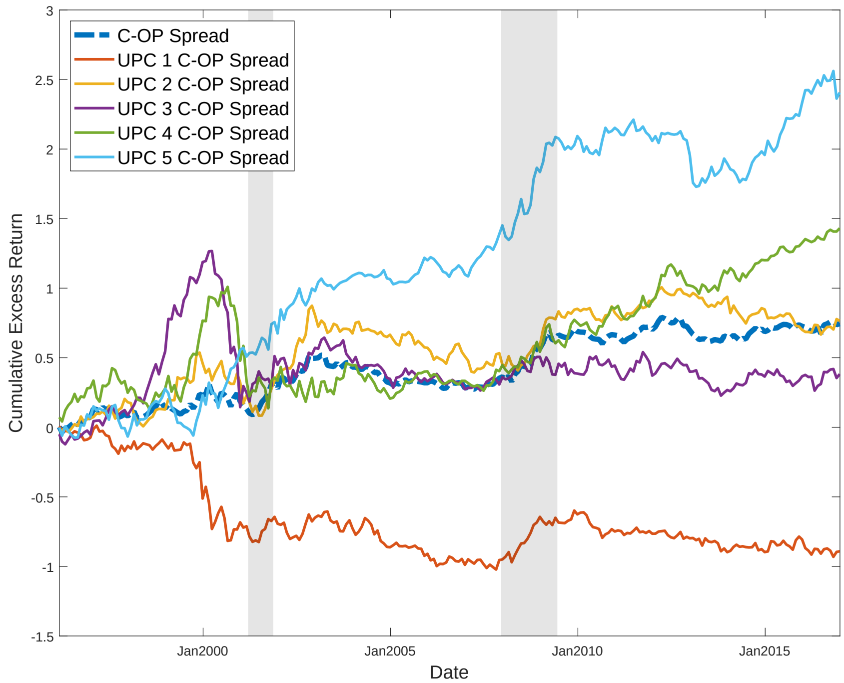

Figure 3.1: Cumulative Excess Returns of C-OP Portfolios Across PIF Quintiles

This figure plots the cumulative (sum of log) returns of C-OP long-short portfolios across PIF quintiles. The sample starts from February 1996 to December 2016.

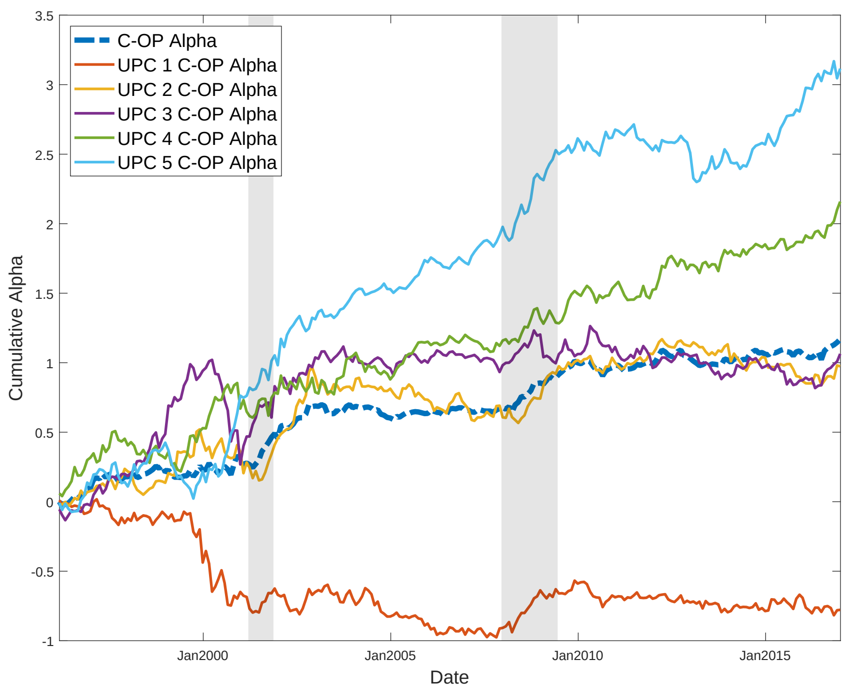

Figure 3.2: Cumulative FF3F Alphas of C-OP Portfolios Across PIF Quintiles

This figure plots the cumulative (sum of log) FF3F alphas of C-OP long-short portfolios across PIF quintiles. The FF3F loadings are estimated with full sample. The sample starts from February 1996 to December 2016.

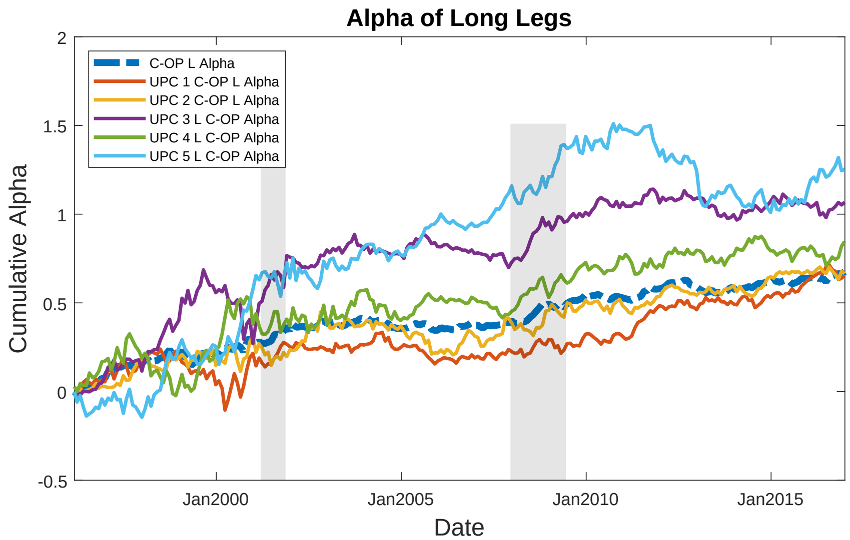

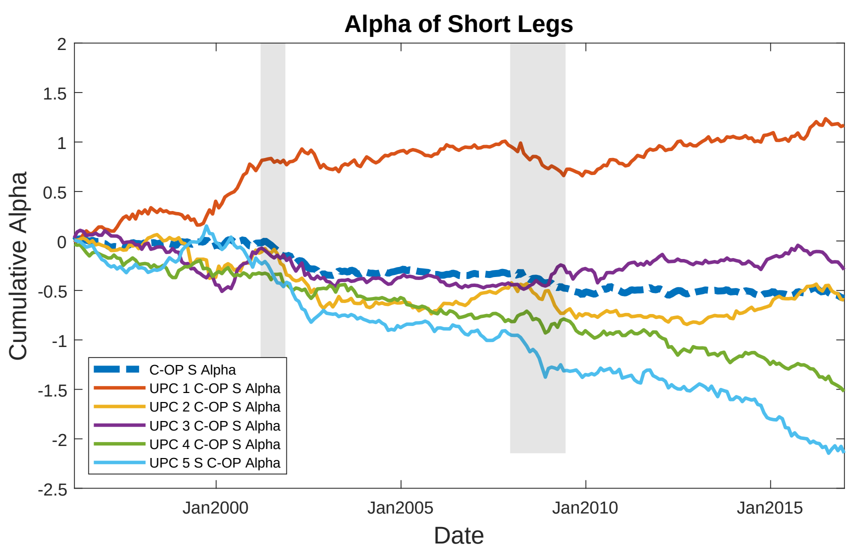

Figure 3.3: Cumulative FF3F Alphas of C-OP Long/Short Legs Across PIF Quintiles

This figure plots the cumulative (sum of log) FF3F alphas of long/short legs of C-OP long-short portfolios across PIF quintiles. The FF3F loadings are estimated with full sample. The sample starts from February 1996 to December 2016.

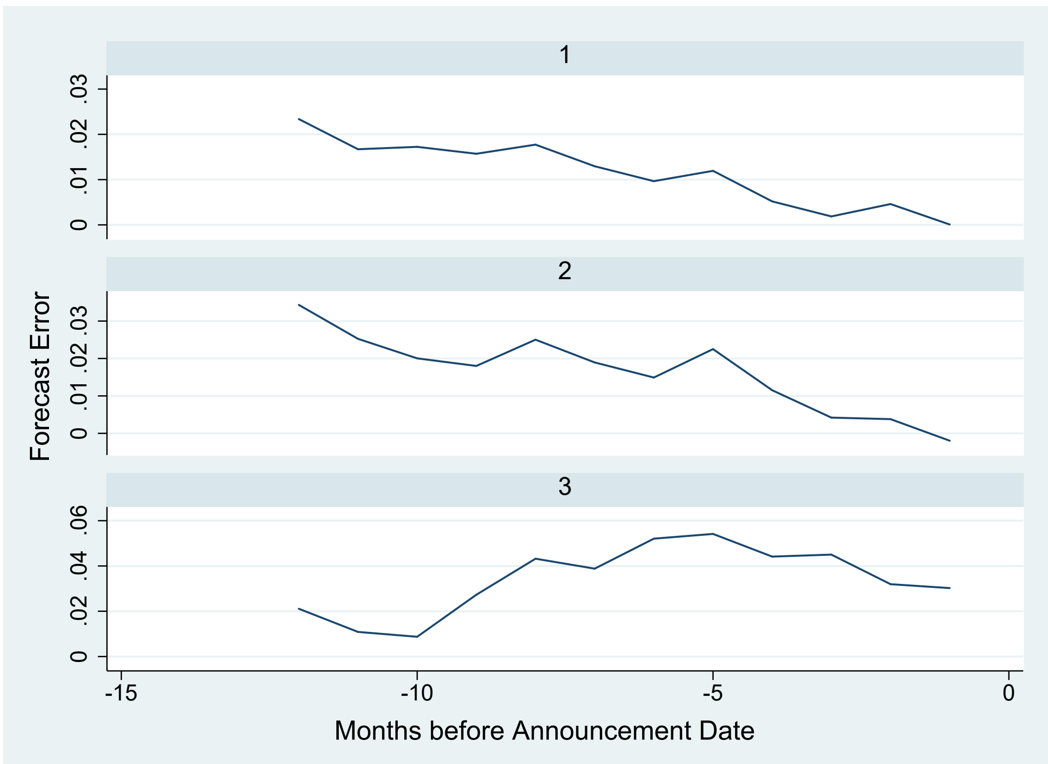

Figure 3.4: Evolution of Forecast Errors Across Uncertainty Groups

This figure plots how forecast errors, defined as share adjusted forecasts minus share adjusted actual earnings scaled by absolute value of share adjusted actual earnings, evolves before announcement date. The forecast errors at each horizon are market-capitalization weighted. Firms are allocated into each tercile based on PIF using 30th and 70th cutoff points. The maximum forecast horizon is 12 months.

APPENDICES

APPENDIX A

Expectation Formation: Data Construction, Derivations, and Proofs

A.1 Data Sources

A.1.1 Surveys

Blue Chip Financial Forecasts. The Blue Chip Financial Forecasts (BCFF) is a monthly survey of about 45 professional forecasters from leading financial institutions, beginning in November 1982. The survey is typically published on the first day of the month and based on responses collected during the last week of the previous month. The participants are asked to provide forecasts of a given variable in the following quarters, up to six quarters ahead. The survey covers interest rate forecasts of the federal fund rate, 3-month/6-month/1-year Treasury bills, and 2-year/5-year/10-year/30-year Treasury notes. The survey also includes forecasts of real/nominal GDP growth and CPI.

For long-term forecasts of the same variables, BCFF has provided the results of a semi-annual long-range consensus survey since October 1983. Before December 1996, the long-range surveys were typically published in March and October (except for 1985, which were published in March and November). The publishing months switched to June and December beginning in December 1996 (one exception is the January 2003 long-range survey). Forecasts are made for the annual averages of selected variables for the following five years. The survey has also provided a 5-year average forecast for the next 6 to 10 years (or the next 7 to 11 years) since 1986.

Blue Chip Economic Indicator. The Blue Chip Economic Indicators (BCEI) is a monthly survey of about 50 professional forecasters from leading financial institutions, beginning in August 1976. The survey is typically published on the tenth day of the month and based on responses collected during the first week of the same month. The survey includes forecasts for real GDP growth, CPI, and 3-month Treasury bill rates.

BCEI has published a semi-annual long-range survey in March and October consistently since October 1983. Long-range surveys from BCEI and BCFF have the same forecast horizons.

Survey of Professional Forecasters. The Survey of Professional Forecasters (SPF) is a quarterly survey of professional forecasters maintained by the Philadelphia Fed that covers a wide range of macroeconomic variables and selective yields (3-month Treasury bills and 10-year Treasury bonds). The survey published in a given quarter is based on questionnaires sent to the panelists during the end of the first month and collected no later than the middle month. The documentation available on the Philadelphia Fed's website provides more details of the survey.[1]

Goldsmith-Nagan Survey. The Goldsmith-Nagan (GN) survey is a quarterly survey of a selected panel of approximately 50 market professionals who have subscribed to the *Goldsmith-Nagan Bond and Money Market Letter*. The survey is conducted late in the last month of each quarter for forecasts of a set of interest rates on the last business day of the coming quarter and of the quarter following. Friedman (1980) and Froot (1989) provide more details on the survey. The data is kindly provided by Kenneth Froot.

A.1.2 Interest rates

Zero-coupon Treasury yields. Monthly nominal zero-coupon Treasury yields are from Liu and Wu (2019), with maturities ranging from 1 month to 30 years at one-month intervals. The data is available from the authors' website.[2]

Daily nominal zero-coupon Treasury yields are from Gürkaynak et al. (2007), available from the Federal Reserve Board.

[1] `https://www.philadelphiafed.org/-/media/research-and-data/real-time-center/survey-of-professional-forecasters/spf-documentation.pdf?la=en`

[2] `https://sites.google.com/view/jingcynthiawu/yield-data`

Daily real zero-coupon yields calculated from TIPS are from Gürkaynak et al. (2010), available from the Federal Reserve Board.

For nominal zero-coupon yields before 1972, I use data from McCulloch and Kwon (1993) beginning in 1952.

Fama bond portfolio returns. Monthly Fama bond portfolio returns for different yield maturities are obtained from CRSP.

1-month risk-free rate. Monthly 1-month risk-free rates are obtained from CRSP.

3-month Treasury bill rate. I use the monthly secondary market rates since January 1954 from the H.15 Selected Interest Rates, available from the Federal Reserve Board.

A.1.3 Macro variables

Inflation. Inflation is constructed from the monthly core Consumer Price Index (CPI) from January 1957 to December 2018, available from the Bureau of Labor Statistics. For January 1947 to December 1956, I use the all-items CPI, also obtained from the Bureau of Labor Statistics. For inflation before 1947, I use the CPI from Shiller (2005).

Real GDP. The quarterly vintages of real GNP/GDP (ROUTPUT) are obtained from the *Real-Time Data Set for Macroeconomists*, available from the Philadelphia Fed for 1965Q4 to 2018Q4.

Employment growth. The monthly vintages of nonfarm payroll employment (EMPLOY) are from the *Real-Time Data Set for Macroeconomists*, available from the Philadelphia Fed for December 1964 to December 2018.

A.2 Yield Decompositions

This section provides derivations regarding yield decompositions. Basic notation:

- y_t^n: yield on a nominal zero-coupon bond with a unit payoff upon maturity at $t+n$.
- p_t^n: log price of a nominal zero-coupon bond defined as above. $p_t^n = -ny_t^n$.
- i_t: nominal short rate. $i_t = y_t^1$.
- s_t^n: term spread as the difference between an n-maturity yield and the short rate. $s_t^n = y_t^n - i_t$.

- $f_t^{n,m}$: forward rate at time t on an n-maturity bond, m periods into the future (the bond matures at $t+n+m$). $f_t^{n,m} = \frac{1}{n}[(m+n)y_t^{m+n} - my_t^m]$.

I define an m-holding-period return and the excess return on an n-maturity bond realized at time $t+m$, respectively, as

$$r_{t,t+m}^n = p_{t+m}^{n-m} - p_t^n = ny_t^n - (n-m)y_{t+m}^{n-m}, \tag{A.1}$$

$$rx_{t,t+m}^n = p_{t+m}^{n-m} - p_t^n + p_t^m = ny_t^n - (n-m)y_{t+m}^{n-m} - my_t^m. \tag{A.2}$$

For one-period returns and excess returns, I use the simplified notations r_{t+1}^n and rx_{t+1}^n.

Solving forward with $m=1$ yields

$$\begin{aligned} p_t^n = -ny_t^n &= -\sum_{k=0}^{n-1} r_{t+k+1}^{n-k} \\ &= -\sum_{k=0}^{n-1} \left(rx_{t+k+1}^{n-k} + i_{t+k} \right). \end{aligned} \tag{A.3}$$

Taking a generic expectation, $\widehat{\mathbb{E}}_t$, of both sides yields

$$\begin{aligned} y_t^n &= \frac{1}{n}\widehat{\mathbb{E}}_t\left(\sum_{k=0}^{n-1} i_{t+k}\right) + \frac{1}{n}\widehat{\mathbb{E}}_t\left(\sum_{k=0}^{n-1} rx_{t+k+1}^{n-k}\right) \\ &= \frac{1}{n}\widehat{\mathbb{E}}_t\left(\sum_{k=0}^{n-1} i_{t+k}\right) + \frac{1}{n}\widehat{\mathbb{E}}_t\left(\sum_{k=1}^{n} rx_{t+n-k+1}^{k}\right) \\ &\equiv \frac{1}{n}\widehat{\mathbb{E}}_t\left(\sum_{k=0}^{n-1} i_{t+k}\right) + tp_t^n. \end{aligned} \tag{A.4}$$

The last term is typically referred to as the "term premium".

A.3 Model Solutions

A.3.1 Kalman filtering

A.3.1.1 Real endowment growth

The agent at time t subjectively perceives

$$g_t = \tilde{\mu}_t + (\tilde{C}_t - \tilde{C}_{t-1}) + \tilde{\xi}_t, \qquad \tilde{\xi}_t \sim \mathcal{N}(0, \tilde{\sigma}_\xi^2), \tag{A.5}$$

$$\tilde{\mu}_t = \phi_\mu \tilde{\mu}_{t-1} + (1-\phi_\mu)\mu + \tilde{\omega}_t, \qquad \tilde{\omega}_t \sim \mathcal{N}(0, \tilde{\sigma}_\omega^2), \tag{A.6}$$

$$\tilde{C}_t = \tilde{\rho}_c \tilde{C}_{t-1} + \tilde{\zeta}_t, \qquad \tilde{\zeta}_t \sim \mathcal{N}(0, \tilde{\sigma}_\zeta^2). \tag{A.7}$$

The standard state-space representation with lagged variables in the measurement equation is

$$g_t - \mu = H' B_t + \xi_t, \tag{A.8}$$

$$B_{t+1} = F B_t + \psi_{t+1}, \tag{A.9}$$

where $H = (1, 1, -1)'$, and

$$B_t = \begin{pmatrix} \tilde{\mu}_t - \mu \\ \tilde{C}_t \\ \tilde{C}_{t-1} \end{pmatrix}, \quad F = \begin{pmatrix} \phi_\mu & 0 & 0 \\ 0 & \tilde{\rho}_c & 0 \\ 0 & 1 & 0 \end{pmatrix}, \quad \psi_t = \begin{pmatrix} \tilde{\omega}_t \\ \tilde{\zeta}_t \\ 0 \end{pmatrix}. \tag{A.10}$$

I denote the conditional expectation, given an infinite history of observations $\mathcal{H}_t \equiv \{g_{t-k}\}_{k=0}^{\infty}$, as

$$\widehat{B}_{t+s|t} \equiv \tilde{\mathbb{E}}[B_{t+s}|\mathcal{H}_t]. \tag{A.11}$$

Following Hamilton (1994), the conditional minimum variance estimate of B_{t+1} in the steady-state evolves as

$$\widehat{B}_{t+1|t} = F\widehat{B}_{t|t-1} + K(g_t - \mu - H'\widehat{B}_{t|t-1}), \tag{A.12}$$

where K is the Kalman gain matrix. The steady-state MSE matrix, $P \equiv \tilde{\mathbb{E}}[(B_{t+1} - \widehat{B}_{t+1|t})(B_{t+1} - \widehat{B}_{t+1|t})']$, and the Kalman gain matrix, $K \equiv (\nu_\mu\ ,\ \nu_{c,1}\ ,\ \nu_{c,2})'$, are obtained from the following equations:

$$P = F\left[P - PH(H'PH + \tilde{\sigma}_\xi^2)^{-1}H'P\right]F' + Q, \tag{A.13}$$

$$K = FPH(H'PH + \tilde{\sigma}_\xi^2)^{-1}, \tag{A.14}$$

where

$$Q = \begin{pmatrix} \tilde{\sigma}_\omega^2 & 0 & 0 \\ 0 & \tilde{\sigma}_\zeta^2 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{A.15}$$

Note that Equation (A.14) implies $\nu_{c,2} = \nu_{c,1}/\tilde{\rho}_c$ and that it can be shown that $\widehat{C}_{t+1|t} = \tilde{\rho}_c\widehat{C}_{t|t}$. Denoting $\widehat{\mu}_t \equiv \widehat{\mu}_{t+1|t}$ and $\widehat{c}_t \equiv \widehat{C}_{t+1|t} - \widehat{C}_{t|t}$, Equation (A.12) can be written as

$$\widehat{\mu}_t = \phi_\mu\widehat{\mu}_{t-1} + (1-\phi_\mu)\mu + \nu_\mu\big(g_t - \widehat{\mu}_{t-1} - \widehat{c}_{t-1}\big), \tag{A.16}$$

$$\widehat{c}_t = \tilde{\rho}_c\widehat{c}_{t-1} - \nu_c\big(g_t - \widehat{\mu}_{t-1} - \widehat{c}_{t-1}\big), \tag{A.17}$$

where

$$\nu_c = \nu_{c,2} - \nu_{c,1}. \tag{A.18}$$

The predictive distribution of g_{t+1} is

$$g_{t+1}|\mathcal{H}_t \sim \mathcal{N}\left(\widehat{\mu}_t + \widehat{c}_t\ ,\ \tilde{\sigma}_g^2\right), \tag{A.19}$$

where

$$\tilde{\sigma}_g^2 = H'PH + \tilde{\sigma}_\xi^2. \tag{A.20}$$

A.3.1.2 Inflation

The equations for the subjective dynamics of inflation are:

$$\pi_t = \tilde{\tau}_t + (\tilde{K}_t - \tilde{K}_{t-1}) + \tilde{\eta}_t, \qquad \tilde{\eta}_t \sim \mathcal{N}(0, \tilde{\sigma}_\eta^2), \tag{A.21}$$

$$\tilde{\tau}_t = \phi_\tau\tilde{\tau}_{t-1} + (1-\phi_\tau)\tau + \tilde{\varepsilon}_t, \quad \tilde{\varepsilon}_t \sim \mathcal{N}(0, \tilde{\sigma}_\varepsilon^2), \tag{A.22}$$

$$\tilde{K}_t = \tilde{\rho}_\kappa\tilde{K}_{t-1} + \tilde{\iota}_t, \qquad \tilde{\iota}_t \sim \mathcal{N}(0, \tilde{\sigma}_\iota^2). \tag{A.23}$$

Similarly, the steady-state MSE matrix, P_π, and Kalman gain matrix, K_π, are obtained from the following equations

$$P_\pi = F_\pi \left[P_\pi - P_\pi H_\pi (H'_\pi P_\pi H_\pi + \tilde{\sigma}^2_\eta)^{-1} H'_\pi P_\pi \right] F'_\pi + Q_\pi, \tag{A.24}$$

$$K_\pi = F_\pi P_\pi H_\pi (H'_\pi P_\pi H_\pi + \tilde{\sigma}^2_\eta)^{-1}, \tag{A.25}$$

where

$$F_\pi = \begin{pmatrix} \phi_\tau & 0 & 0 \\ 0 & \tilde{\rho}_\kappa & 0 \\ 0 & 1 & 0 \end{pmatrix}, \quad Q_\pi = \begin{pmatrix} \tilde{\sigma}^2_\varepsilon & 0 & 0 \\ 0 & \tilde{\sigma}^2_\iota & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{A.26}$$

Denote $K_\pi \equiv (\nu_\tau \ , \ \nu_{\kappa,1} \ , \ \nu_{\kappa,2})'$. Given an infinite history of observations $\mathcal{I}_t = \{\pi_{t-k}\}_{k=0}^{\infty}$, the optimal forecasts $\widehat{\tau}_t \equiv \tilde{\mathbb{E}}_t[\tau_{t+1}|\mathcal{I}_t]$ and $\widehat{\kappa}_t \equiv \tilde{\mathbb{E}}[K_{t+1}|\mathcal{I}_t] - \tilde{\mathbb{E}}[K_t|\mathcal{I}_t]$ are:

$$\widehat{\tau}_t = \phi_\tau \widehat{\tau}_{t-1} + (1-\phi_\tau)\tau + \nu_\tau \big(x_t - \widehat{\tau}_{t-1} - \widehat{\kappa}_{t-1} \big), \tag{A.27}$$

$$\widehat{\kappa}_t = \tilde{\rho}_\kappa \widehat{\kappa}_{t-1} - \nu_\kappa \big(x_t - \widehat{\tau}_{t-1} - \widehat{\kappa}_{t-1} \big), \tag{A.28}$$

where

$$\nu_\kappa = \nu_{\kappa,2} - \nu_{\kappa,1}. \tag{A.29}$$

The predictive distribution of π_{t+1} is

$$\pi_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\widehat{\tau}_t + \widehat{\kappa}_t \ , \ \tilde{\sigma}^2_\pi \right), \tag{A.30}$$

where

$$\tilde{\sigma}^2_\pi = H'_\pi P_\pi H_\pi + \tilde{\sigma}^2_\eta. \tag{A.31}$$

A.3.2 Stochastic discount factor

The Epstein-Zin real log SDF can be written as

$$m_{t+1} = \theta \log \delta - \frac{\theta}{\psi} g_{t+1} + (\theta - 1) r_{w,t+1} \tag{A.32}$$

with $\theta = (1-\gamma)/(1-1/\psi)$.

I log-linearize the return on wealth as

$$r_{w,t+1} = k_0 + k_1 z_{t+1} - z_t + g_{t+1}, \tag{A.33}$$

where $z_t \equiv \log((W_t - X_t)/X_t)$. I conjecture that z_t is linear in the state variable $\widehat{\mu}_t$

and $\widehat{c}_t$:

$$z_t = A_0 + A_1 \widehat{\mu}_t + A_2 \widehat{c}_t. \tag{A.34}$$

By applying the subjective pricing equation to $r_{w,t+1}$, it can be shown that

$$A_0 = \frac{\log \delta + k_0 + k_1 A_1 (1-\phi_\mu)\mu + \frac{1}{2}\theta(1 - 1/\psi + k_1 A_1 \nu_\mu - k_1 A_2 \nu_c)^2 \tilde{\sigma}_g^2}{1 - k_1}, \tag{A.35}$$

$$A_1 = \frac{1 - 1/\psi}{1 - k_1 \phi_\mu} \quad , \quad A_2 = \frac{1 - 1/\psi}{1 - k_1 \tilde{\rho}_c} \ , \tag{A.36}$$

where

$$\bar{z} = A_0 + A_1 \mu \ , \ k_1 = \frac{\exp(\bar{z})}{1 + \exp(\bar{z})} \ , \ k_0 = \log(1 + e^{\bar{z}}) - \bar{z} k_1. \tag{A.37}$$

Thus, the real log SDF is

$$m_{t+1} = \tilde{\mu}_m - \frac{1}{\psi}\left(\widehat{\mu}_t + \widehat{c}_t\right) - \xi \tilde{\sigma}_g \tilde{\delta}_{t+1}, \tag{A.38}$$

where $\{\tilde{\delta}\}$ are i.i.d. standard normal shocks under the subjective measure and

$$\Lambda = 1 - 1/\psi + k_1 A_1 \nu_\mu - k_1 A_2 \nu_c, \tag{A.39}$$

$$\tilde{\mu}_m = \log \delta - \frac{1}{2}\theta(\theta - 1)\Lambda^2 \tilde{\sigma}_g^2, \tag{A.40}$$

$$\xi = (1 - \theta)\Lambda + \frac{1}{\psi}. \tag{A.41}$$

The nominal log SDF is

$$m_{t+1}^{\$} = m_{t+1} - \pi_{t+1}. \tag{A.42}$$

The real short rate is

$$i_t \equiv y_t^1 = -\tilde{\mu}_m + \frac{1}{\psi}\left(\widehat{\mu}_t + \widehat{c}_t\right) - \frac{1}{2}\xi^2 \tilde{\sigma}_g^2. \tag{A.43}$$

and the nominal short rate is

$$i_t^{\$} = -\tilde{\mu}_m + \frac{1}{\psi}\left(\widehat{\mu}_t + \widehat{c}_t\right) + \left(\widehat{\tau}_t + \widehat{\kappa}_t\right) - \frac{1}{2}\xi^2 \tilde{\sigma}_g^2 - \frac{1}{2}\tilde{\sigma}_\pi^2. \tag{A.44}$$

A.3.3 Bond pricing

A.3.3.1 Real bonds

I conjecture that the log real bond prices are affine functions of the state variables $(\widehat{\mu}_t, \widehat{c}_t)$:

$$p_t^n = a_n + b_n \widehat{\mu}_t + c_n \widehat{c}_t. \tag{A.45}$$

Through the recursive subjective Euler equations,

$$p_t^n = \log \tilde{\mathbb{E}}_t[\exp(m_{t+1} + p_{t+1}^{n-1})], \tag{A.46}$$

the coefficients are obtained as

$$a_n = n\tilde{\mu}_m - \left(n - \frac{1-\phi_\mu^n}{1-\phi_\mu}\right)\frac{\mu}{\psi} + \frac{1}{2}\sum_{k=1}^{n}(-\xi + b_{k-1}\nu_\mu - c_{k-1}\nu_c)^2 \tilde{\sigma}_g^2, \tag{A.47}$$

$$b_n = -\frac{1-\phi_\mu^n}{1-\phi_\mu}\frac{1}{\psi} \quad , \quad c_n = -\frac{1-\tilde{\rho}_c^n}{1-\tilde{\rho}_c}\frac{1}{\psi}. \tag{A.48}$$

The real yields are obtained as

$$y_t^n \equiv -\frac{1}{n}p_t^n = -\frac{a_n}{n} + \frac{1-\phi_\mu^n}{n(1-\phi_\mu)}\frac{1}{\psi}\widehat{\mu}_t + \frac{1-\tilde{\rho}_c^n}{n(1-\tilde{\rho}_c)}\frac{1}{\psi}\widehat{c}_t. \tag{A.49}$$

The real yield spread is

$$\begin{aligned} s_t^n \equiv y_t^n - y_t^1 = & \frac{1}{2}\xi^2\tilde{\sigma}_g^2 - \frac{1}{2n}\sum_{k=1}^{n}(-\xi + b_{k-1}\nu_\mu - c_{k-1}\nu_c)^2\tilde{\sigma}_g^2 \\ & + \left[1 - \frac{1-\phi_\mu^n}{n(1-\phi_\mu)}\right]\frac{1}{\psi}(\mu - \widehat{\mu}_t) - \left[1 - \frac{1-\tilde{\rho}_c^n}{n(1-\tilde{\rho}_c)}\right]\frac{1}{\psi}\widehat{c}_t. \end{aligned} \tag{A.50}$$

By taking n to infinity, we have

$$\begin{aligned} \lim_{n\to\infty} s_t^n = & \frac{1}{2}\left(2\psi\xi + \frac{\nu_\mu}{1-\phi_\mu} - \frac{\nu_c}{1-\tilde{\rho}_c}\right)\left(\frac{\nu_c}{1-\tilde{\rho}_c} - \frac{\nu_\mu}{1-\phi_\mu}\right)\frac{\tilde{\sigma}_g^2}{\psi^2} \\ & + \frac{1}{\psi}(\mu - \widehat{\mu}_t - \widehat{c}_t). \end{aligned} \tag{A.51}$$

A.3.3.2 Nominal bonds

Similarly, I conjecture that the log nominal bond prices are affine functions of the state variables $(\widehat{\mu}_t, \widehat{c}_t^*, \widehat{\tau}_t, \widehat{\kappa}_t^*)$; that is:

$$p_t^{\$,n} = a_n^{\$} + b_n^{\$}\widehat{\mu}_t + c_n^{\$}\widehat{c}_t + d_n^{\$}\widehat{\tau}_t + e_n^{\$}\widehat{\kappa}_t. \tag{A.52}$$

The coefficients are obtained from recursive subjective Euler equations as

$$\begin{aligned} a_n^{\$} = n\tilde{\mu}_m - \left(n - \frac{1-\phi_\mu^n}{1-\phi_\mu}\right)\frac{\mu}{\psi} + \frac{1}{2}\sum_{k=1}^{n}(-\xi + b_{k-1}^{\$}\nu_\mu - c_{k-1}^{\$}\nu_c)^2\tilde{\sigma}_g^2 \\ - \left(n - \frac{1-\phi_\tau^n}{1-\phi_\tau}\right)\tau + \frac{1}{2}\sum_{k=1}^{n}(-1 + d_{k-1}^{\$}\nu_\tau - e_{k-1}^{\$}\nu_\kappa)^2\tilde{\sigma}_\pi^2, \end{aligned} \tag{A.53}$$

$$b_n^{\$} = b_n = -\frac{1-\phi_\mu^n}{1-\phi_\mu}\frac{1}{\psi} \quad , \quad c_n^{\$} = c_n = -\frac{1-\tilde{\rho}_c^n}{1-\tilde{\rho}_c}\frac{1}{\psi} \ , \tag{A.54}$$

$$d_n^{\$} = -\frac{1-\phi_\tau^n}{1-\phi_\tau} \quad , \quad e_n^{\$} = -\frac{1-\tilde{\rho}_\kappa^n}{1-\tilde{\rho}_\kappa} \ . \tag{A.55}$$

The nominal yields are obtained as

$$y_t^{\$,n} \equiv -\frac{1}{n}p_t^{\$,n} = -\frac{a_n^{\$}}{n} + \frac{1-\phi_\mu^n}{n(1-\phi_\mu)}\frac{1}{\psi}\widehat{\mu}_t + \frac{1-\tilde{\rho}_c^n}{n(1-\tilde{\rho}_c)}\frac{1}{\psi}\widehat{c}_t + \frac{1-\phi_\tau^n}{n(1-\phi_\tau)}\widehat{\tau}_t + \frac{1-\tilde{\rho}_\kappa^n}{n(1-\tilde{\rho}_\kappa)}\widehat{\kappa}_t. \tag{A.56}$$

The nominal yield spread is

$$\begin{aligned} s_t^{\$,n} \equiv y_t^{\$,n} - y_t^{\$,1} = s_t^n + \frac{1}{2}\tilde{\sigma}_\pi^2 - \frac{1}{2n}\sum_{k=1}^{n}(-1 + d_{k-1}^{\$}\nu_\tau - e_{k-1}^{\$}\nu_\kappa)^2\tilde{\sigma}_\pi^2 \\ + \left[1 - \frac{1-\phi_\tau^n}{n(1-\phi_\tau)}\right](\tau - \widehat{\tau}_t) - \left[1 - \frac{1-\tilde{\rho}_\kappa^n}{n(1-\tilde{\rho}_\kappa)}\right]\widehat{\kappa}_t. \end{aligned} \tag{A.57}$$

By taking n to infinity, we have

$$\begin{aligned} \lim_{n\to\infty} s_t^{\$,n} = \lim_{n\to\infty} s_t^n + \tau - \widehat{\tau}_t - \widehat{\kappa}_t \\ + \frac{1}{2}\left(2 + \frac{\nu_\tau}{1-\phi_\tau} - \frac{\nu_\kappa}{1-\tilde{\rho}_\kappa}\right)\left(\frac{\nu_\kappa}{1-\tilde{\rho}_\kappa} - \frac{\nu_\tau}{1-\phi_\tau}\right)\tilde{\sigma}_\pi^2. \end{aligned} \tag{A.58}$$

A.3.4 Forecasts and forecast errors

A.3.4.1 Forecasts of future endowment growth and inflation

Based on the filtering results, the subjective forecasts of future endowment growth and inflation are:

$$\tilde{\mathbb{E}}_t \mu_{t+j} = \phi_\mu^{j-1} \widehat{\mu}_t + (1 - \phi_\mu^{j-1})\mu \ , \ \tilde{\mathbb{E}}_t c_{t+j} = \tilde{\rho}_c^{j-1} \widehat{c}_t, \tag{A.59}$$

$$\tilde{\mathbb{E}}_t g_{t+j} = \phi_\mu^{j-1} \widehat{\mu}_t + (1 - \phi_\mu^{j-1})\mu + \tilde{\rho}_c^{j-1} \widehat{c}_t, \tag{A.60}$$

and

$$\tilde{\mathbb{E}}_t \tau_{t+j} = \phi_\tau^{j-1} \widehat{\tau}_t + (1 - \phi_\tau^{j-1})\tau \ , \ \tilde{\mathbb{E}}_t \kappa_{t+j} = \tilde{\rho}_\kappa^{j-1} \widehat{\kappa}_t, \tag{A.61}$$

$$\tilde{\mathbb{E}}_t \pi_{t+j} = \phi_\tau^{j-1} \widehat{\tau}_t + (1 - \phi_\tau^{j-1})\tau + \tilde{\rho}_\kappa^{j-1} \widehat{\kappa}_t. \tag{A.62}$$

In comparison, an econometrician who has full knowledge sees

$$\mathbb{E}_t g_{t+j} = \mu + \rho_c^{j-1} c_t^e \ , \ \mathbb{E}_t \pi_{t+j} = \tau + \rho_\kappa^{j-1} \kappa_t^e, \tag{A.63}$$

where $c_t^e = (\rho_c - 1)c_t$ and $\kappa_t^e = (\rho_\kappa - 1)\kappa_t$.

A.3.4.2 Forecasts of future state variables

Under the subjective measure,

$$\tilde{\mathbb{E}}_t \widehat{\mu}_{t+j} = \phi_\mu^{j} \widehat{\mu}_t + (1 - \phi_\mu^{j})\mu \ , \ \tilde{\mathbb{E}}_t \widehat{c}_{t+j} = \tilde{\rho}_c^{j} \widehat{c}_t, \tag{A.64}$$

$$\tilde{\mathbb{E}}_t \widehat{\tau}_{t+j} = \phi_\tau^{j} \widehat{\tau}_t + (1 - \phi_\tau^{j})\tau \ , \ \tilde{\mathbb{E}}_t \widehat{\kappa}_{t+j} = \tilde{\rho}_\kappa^{j} \widehat{\kappa}_t. \tag{A.65}$$

Under the objective measure, the expressions are derived from iterations with the following forms:

$$\mathbb{E}_t \widehat{\mu}_{t+j} = a_{g,j}\mu + b_{g,j} c_t^e + c_{g,j} \widehat{\mu}_t + d_{g,j} \widehat{c}_t, \tag{A.66}$$

$$\mathbb{E}_t \widehat{c}_{t+j} = e_{g,j}\mu + f_{g,j} c_t^e + g_{g,j} \widehat{\mu}_t + h_{g,j} \widehat{c}_t. \tag{A.67}$$

Similarly, we have

$$\mathbb{E}_t\widehat{\tau}_{t+j} = a_{\pi,j}\tau + b_{\pi,j}\kappa_t^e + c_{\pi,j}\widehat{\tau}_t + d_{\pi,j}\widehat{\kappa}_t, \tag{A.68}$$

$$\mathbb{E}_t\widehat{\kappa}_{t+j} = e_{\pi,j}\tau + f_{\pi,j}\kappa_t^e + g_{\pi,j}\widehat{\tau}_t + h_{\pi,j}\widehat{\kappa}_t. \tag{A.69}$$

Here, I only provide formulas for the special case $j = 1$:

$$\mathbb{E}_t\widehat{\mu}_{t+1} = (\phi_\mu - \nu_\mu)\widehat{\mu}_t + (1 - \phi_\mu + \nu_\mu)\mu + \nu_\mu c_t^e - \nu_\mu\widehat{c}_t, \tag{A.70}$$

$$\mathbb{E}_t\widehat{c}_{t+1} = (\tilde{\rho}_c + \nu_c)\widehat{c}_t - \nu_c c_t^e - \nu_c(\mu - \widehat{\mu}_t), \tag{A.71}$$

and similarly for inflation. The online supplementary appendix provides more details.

A.3.4.3 Short-rate forecast errors

I focus on forecast errors of nominal short rates since they are directly observable. For one-period-ahead, by combining Equation (A.44) and results from Section A.3.4.2, we have

$$\begin{aligned}\mathbb{E}_t i_{t+1}^{\$} - \tilde{\mathbb{E}}_t i_{t+1}^{\$} = & \frac{1}{\psi}(\nu_\mu - \nu_c)\left[\mu + c_t^e - \widehat{\mu}_t - \widehat{c}_t\right] \\ & + (\nu_\tau - \nu_\kappa)\left[\tau + \kappa_t^e - \widehat{\tau}_t - \widehat{\kappa}_t\right].\end{aligned} \tag{A.72}$$

For relatively long-term horizons, the general expressions are

$$\begin{aligned}\mathbb{E}_t i_{t+j}^{\$} - \tilde{\mathbb{E}}_t i_{t+j}^{\$} = & \frac{1}{\psi}\left(a_{g,j}^*\mu + b_{g,j}^* c_t^e + c_{g,j}^*\widehat{\mu}_t + d_{g,j}^*\widehat{c}_t\right) \\ & + \left(a_{\pi,j}^*\tau + b_{\pi,j}^*\kappa_t^e + c_{\pi,j}^*\widehat{\tau}_t + d_{\pi,j}^*\widehat{\kappa}_t\right),\end{aligned} \tag{A.73}$$

with

$$a_{g,j}^* = a_{g,j} + e_{g,j} - (1 - \phi_\mu^j)\ ,\ b_{g,j}^* = b_{g,j} + f_{g,j}, \tag{A.74}$$

$$c_{g,j}^* = c_{g,j} + g_{g,j} - \phi_\mu^j\ ,\ d_{g,j}^* = d_{g,j} + h_{g,j} - \tilde{\rho}_c^j, \tag{A.75}$$

and similarly for inflation. Since simple closed-form expressions are not available, I rely on numerical calculations. For very long-term horizons, the forecast errors converge to zero.

A.3.5 Objective and subjective risk premia

I focus on risk premia on nominal bonds. The realized one-period log return on a n-maturity zero-coupon bond is

$$\begin{aligned} r_{t+1}^{n} \equiv p_{t+1}^{\$,n-1} - p_{t}^{\$,n} = a_{n-1}^{\$} - a_{n}^{\$} + b_{n-1}^{\$}(1-\phi_{\mu})\mu + d_{n-1}^{\$}(1-\phi_{\tau})\tau \\ + \frac{1}{\psi}\left(\widehat{\mu}_t + \widehat{c}_t\right) + \left(b_{n-1}^{\$}\nu_{\mu} - c_{n-1}^{\$}\nu_{c}\right)\left(g_{t+1} - \widehat{\mu}_t - \widehat{c}_t\right) \\ + \left(\widehat{\tau}_t + \widehat{\kappa}_t\right) + \left(d_{n-1}^{\$}\nu_{\tau} - e_{n-1}^{\$}\nu_{\kappa}\right)\left(\pi_{t+1} - \widehat{\tau}_t - \widehat{\kappa}_t\right). \end{aligned} \tag{A.76}$$

Thus, the subjective risk premium is

$$\tilde{\mathbb{E}}_t r_{t+1}^{n} - i_t^{\$} + \frac{1}{2}\tilde{\text{var}}_t(r_{t+1}^{n}) = \xi(b_{n-1}^{\$}\nu_{\mu} - c_{n-1}^{\$}\nu_{c})\tilde{\sigma}_g^2 + (d_{n-1}^{\$}\nu_{\tau} - e_{n-1}^{\$}\nu_{\kappa})\tilde{\sigma}_{\pi}^2, \tag{A.77}$$

which is a maturity-specific constant.

Under the objective expectation, the risk premium is

$$\begin{aligned} \mathbb{E}_t r_{t+1}^{n} - i_t^{\$} + \frac{1}{2}\text{var}_t(r_{t+1}^{n}) = & A_n + \left(b_{n-1}^{\$}\nu_{\mu} - c_{n-1}^{\$}\nu_{c}\right)\left[\mu + c_t^{e} - \widehat{\mu}_t - \widehat{c}_t\right] \\ & + B_n + \left(d_{n-1}^{\$}\nu_{\tau} - e_{n-1}^{\$}\nu_{\kappa}\right)\left[\tau + \kappa_t^{e} - \widehat{\tau}_t - \widehat{\kappa}_t\right], \end{aligned} \tag{A.78}$$

where

$$A_n = \frac{1}{2}\xi^2\tilde{\sigma}_g^2 + \frac{1}{2}(b_{n-1}^{\$}\nu_{\mu} - c_{n-1}^{\$}\nu_{c})^2(\sigma_{\xi}^2 + \sigma_{\zeta}^2) - \frac{1}{2}(-\xi + b_{n-1}^{\$}\nu_{\mu} - c_{n-1}^{\$}\nu_{c})^2\tilde{\sigma}_g^2, \tag{A.79}$$

$$B_n = \frac{1}{2}\tilde{\sigma}_{\pi}^2 + \frac{1}{2}(d_{n-1}^{\$}\nu_{\tau} - e_{n-1}^{\$}\nu_{\kappa})^2(\sigma_{\eta}^2 + \sigma_{\iota}^2) - \frac{1}{2}(-1 + d_{n-1}^{\$}\nu_{\tau} - e_{n-1}^{\$}\nu_{\kappa})^2\tilde{\sigma}_{\pi}^2. \tag{A.80}$$

A.3.6 Proofs of theoretical results

Proof of Theorem I.1: Under the objective expectation, from Equation (A.16) and (A.17), we have

$$\mathbb{E}[\widehat{\mu}_t] = \mu\ ,\ \mathbb{E}[\widehat{c}_t] = 0. \tag{A.81}$$

The latter equation holds because the unconditional mean of C_t is zero. Thus, Equation (A.50) implies

$$\mathbb{E}[ns_t^{n}] - \mathbb{E}[(n-1)s_t^{n-1}] = \frac{1}{2}\xi^2\tilde{\sigma}_g^2 - \frac{1}{2}(-\xi + b_{n-1}\nu_{\mu} - c_{n-1}\nu_{c})^2\tilde{\sigma}_g^2 > 0, \quad \forall n > 1, \tag{A.82}$$

provided $(-\xi + b_{n-1}\nu_\mu - c_{n-1}\nu_c)^2 < \xi^2$. Thus, if

$$(-\xi + b_{n-1}\nu_\mu - c_{n-1}\nu_c)^2 < \xi^2, \quad \forall n > 1, \tag{A.83}$$

or

$$0 < b_{n-1}\nu_\mu - c_{n-1}\nu_c < 2\xi, \quad \forall n > 1, \tag{A.84}$$

we have

$$\mathbb{E}[ns_t^n] > \mathbb{E}[(n-1)s_t^{n-1}] > \cdots > \mathbb{E}[s_t^1] = 0, \tag{A.85}$$

which implies $\mathbb{E}[s_t^n] > 0$. In addition, Equation (A.51) shows that

$$\lim_{n\to\infty} \mathbb{E}[s_t^n] = \frac{1}{2}\left(2\psi\xi + \frac{\nu_\mu}{1-\phi_\mu} - \frac{\nu_c}{1-\tilde{\rho}_c}\right)\left(\frac{\nu_c}{1-\tilde{\rho}_c} - \frac{\nu_\mu}{1-\phi_\mu}\right)\frac{\tilde{\sigma}_g^2}{\psi^2} > 0 \tag{A.86}$$

if and only if

$$0 < \frac{\nu_c}{1-\tilde{\rho}_c} - \frac{\nu_\mu}{1-\phi_\mu} < 2\psi\xi. \tag{A.87}$$

Proof of Corollary I.2: From Equations (A.50) and (A.57), we have

$$\mathbb{E}[s_t^{\$,n} - s_t^n] = \frac{1}{2}\tilde{\sigma}_\pi^2 - \frac{1}{2n}\sum_{k=1}^{n}(-1 + d_{k-1}^{\$}\nu_\tau - e_{k-1}^{\$}\nu_\kappa)^2\tilde{\sigma}_\pi^2. \tag{A.88}$$

Thus, if

$$(-1 + d_{n-1}^{\$}\nu_\tau - e_{n-1}^{\$}\nu_\kappa)^2 < 1, \quad \forall n > 1, \tag{A.89}$$

or

$$0 < d_{n-1}^{\$}\nu_\tau - e_{n-1}^{\$}\nu_\kappa < 2, \quad \forall n > 1, \tag{A.90}$$

we have

$$\mathbb{E}[n(s_t^{\$,n} - s_t^n)] > \mathbb{E}[(n-1)(s_t^{\$,n-1} - s_t^{n-1})] > \cdots > \mathbb{E}[s_t^{\$,1} - s_t^1] = 0, \tag{A.91}$$

which implies $\mathbb{E}[s_t^{\$,n}] > \mathbb{E}[s_t^n]$. In addition, Equation (A.58) and (A.51) imply that

$$\lim_{n\to\infty} \mathbb{E}[s_t^{\$,n}] > \lim_{n\to\infty} \mathbb{E}[s_t^n] \tag{A.92}$$

if and only if

$$0 < \frac{\nu_\kappa}{1-\tilde{\rho}_\kappa} - \frac{\nu_\tau}{1-\phi_\tau} < 2. \tag{A.93}$$

A.4 Interpolation and Bootstrapping of BCFF Yield Forecast Data

Here, I describe in detail how I construct month-end zero-coupon yield forecasts of constant forecast horizons using BCFF data.

A.4.1 Treatment of forecasts for the current quarter

Denote the current-quarter forecasts made in the first, second, and third month of a calendar quarter as F_1^0, F_2^0, and F_3^0.

For F_2^0 and F_3^0, the forecasters already have knowledge of yield realizations in the previous month(s) at the time of making forecasts. To arrive at true forward-looking forecasts, I need to adjust the *quarter-average* forecasts to account for past realizations.

For F_2^0, the forecasters have knowledge of the first-three-week average of yields in the first month, $\bar{y}_{1,3week}$ (BCFF also publishes this information in the middle month.) I calculate the forward-looking measure as $\frac{3F_2^0 - \bar{y}_{1,3week}}{2}$ and treat it as a forecast of the end of the second month. The effective forecast horizon is one month.

For F_3^0, the forecasters have knowledge of yield averages in the first month, $\bar{y}_1$, and the first-three-week in the second month, $\bar{y}_{2,3week}$ (BCFF also publishes this information in the middle month.) I calculate the forward-looking measure as $3F_3^0 - \bar{y}_{2,3week} - \bar{y}_1$ and treat it as a forecast of the end of the third month. The effective forecast horizon is one month.

For F_1^0, no adjustment for past realizations is needed. I treat it as a forecast of the end of the second month. The effective forecast horizon is two months.

A.4.2 For constant forecast horizons

I follow Kim and Orphanides (2012) by treating the quarter-average forecasts as forecasts made for the middle month and interpolating between mid-quarter points to get constant-horizon forecasts. For example, I assume that the 1-quarter-ahead forecast *published* in January is made for May and has an effective 5-month forecast horizon.[3] Thus, the effective forecast horizons of forecasts made in the first month range from 5 to 17 months. I then interpolate along adjacent horizons to get constant-horizon forecasts with 3-month to 18-month horizons using a smooth spline. Forecasts

[3]Chun (2010) looks at quarterly average and treats the effective forecast horizon in that example as three months, which is the distance between the beginning of the publication month and the beginning of the forecast quarter.

of 18-month horizons are only interpolated after January 1997, when 5-quarter-ahead forecasts (in BCFF terms) are available.

A.4.3 Bootstrap to get zero-coupon yield forecasts

BCFF only provides par yield forecasts for Treasury securities with maturities above one year. To get implied zero-coupon yield forecasts, I first fit a Nelson-Siegel model for 6-month, 1-year, 2-year, 5-year, and 10-year Treasury yields. I exclude yield forecasts of 3-month Treasury bills and 30-year Treasury bonds because the Nielson-Siegel model has difficulty fitting the very short end of the observed yield curve (Gürkaynak et al. 2007). Due to the limited availability of yield maturities, I follow Diebold and Li (2006) in applying a simple regression method instead of a numerical optimization approach. I choose 2 years as the maturity at which the loading on the medium-term factor achieves its maximum, since 2-year forecasts are directly observable. The choice of this 3-factor model seems reasonable given that the first two principle components summarize about 99.8% of the cross-sectional variation in yield forecasts across all forecast horizons.[4] The zero-coupon yield forecasts are then bootstrapped from fitted par yield forecasts (see Equations 2 and 7 in Gürkaynak et al. 2007).

A.5 Objective Risk Premia Construction

To construct the objective risk premia, I use the fitted values from predictive regressions of future bond excess returns on the cycle factor from Cieslak and Povala (2015) and the macro factor from Ludvigson and Ng (2009).

For the cycle factor, I project yields with maturities of 1, 2, 5, 7, 10, and 15 years on the trend inflation defined as in Cieslak and Povala (2015).[5] The regression residuals are then used to construct the short cycle, $c_t^{(1)}$, and average cycle, $\bar{c}_t$. The cycle factor is obtained by regressing the maturity-weighted average excess returns, $\overline{rx}_{t+1} \equiv \frac{1}{14}\sum_{k=2}^{15} rx_{t+1}^{k}/k$, on lagged cycle variables

$$\overline{rx}_{t+1} = \gamma_0 + \gamma_1 \bar{c}_t + \gamma_2 c_t^{(1)} + \varepsilon_{t+1} \tag{A.94}$$

[4]For example, the first (second) principal component explains 96.4% (3.5%) of the cross-sectional variation in my sample for 1-year-ahead forecasts from December 1987 to December 2018.

[5]Cieslak and Povala (2015) use 20-year bond instead of 15-year bond. I choose 15-year bond here because the 20-year zero-coupon yield is only available starting from July 1981.

and

$$\widehat{cf}_t = \widehat{\gamma}_0 + \widehat{\gamma}_1 \bar{c}_t + \widehat{\gamma}_2 c_t^{(1)}. \tag{A.95}$$

For the macro factor, I use a similar approach and regress the average excess returns on the principle components from Ludvigson and Ng (2009):

$$\overline{rx}_{t+1} = \beta_0 + \beta_1 \widehat{F}_{1t} + \beta_2 \widehat{F}_{1t}^3 + \beta_3 \widehat{F}_{2t} + \beta_4 \widehat{F}_{3t} + \beta_5 \widehat{F}_{4t} + \beta_6 \widehat{F}_{8t} + u_{t+1} \tag{A.96}$$

and

$$F6_t = \widehat{\beta}_0 + \widehat{\beta}_1 \widehat{F}_{1t} + \widehat{\beta}_2 \widehat{F}_{1t}^3 + \widehat{\beta}_3 \widehat{F}_{2t} + \widehat{\beta}_4 \widehat{F}_{3t} + \widehat{\beta}_5 \widehat{F}_{4t} + \widehat{\beta}_6 \widehat{F}_{8t}. \tag{A.97}$$

Finally, the objective risk premia are constructed by running predictive regressions of individual bonds' excess returns on these two factors and using the fitted values

$$rx_{t+1}^n = \alpha_1 \widehat{cf}_t + \alpha_2 F6_t + \eta_{t+1} \tag{A.98}$$

and

$$\mathbb{E}_t rx_{t+1}^n = \widehat{\alpha}_1 \widehat{cf}_t + \widehat{\alpha}_2 F6_t. \tag{A.99}$$

Table B.1 reports the performance of these factors.

A.6 Bootstrap for the Expectation Hypothesis Test

The bootstrap is carried out as follows: First, the m-period ahead-forecasts on n-maturity yields under the null are calculated, using Equation (18), as

$$\tilde{\mathbb{E}}_t^{null} y_{t+m}^n = y_t^m + \frac{n+m}{n}(y_t^{n+m} - y_t^m) + \hat{c}_{n+m,m}, \tag{A.100}$$

where $\hat{c}_{n+m,m}$ is the sample mean of scaled subjective risk premia. The forecast residuals under the null are thus defined as

$$e_t^{m,n} \equiv \tilde{\mathbb{E}}_t y_{t+m}^n - \tilde{\mathbb{E}}_t^{null} y_{t+m}^n. \tag{A.101}$$

Next, I follow Crump and Gospodinov (2019) in using a nonparametric method to construct bootstrapped yields. The main idea of their approach is to represent forward short rates from period $t+n-1$ to $t+n$, $f_t^{1,n-1}$, as the sum of "difference returns,"

dr_t^n:

$$f_t^{1,n-1} \equiv ny_t^n - (n-1)y_t^{n-1} = f_{t-N+n}^N - \sum_{k=0}^{N-n-1} dr_{t-k}^{n+1+k}, \tag{A.102}$$

where

$$dr_t^n \equiv r_t^n - r_t^{n-1}. \tag{A.103}$$

The bootstrapped yields are re-constructed by resampling the following matrix of forward rates and difference returns with a block bootstrap and conditioning on the first observation:

$$Z = \begin{pmatrix} f_2^N - f_1^N & dr_2^2 & dr_2^3 & \cdots & dr_2^N & x_{12} & \cdots & x_{K2} \\ f_3^N - f_2^N & dr_3^2 & dr_3^3 & \cdots & dr_3^N & x_{13} & \cdots & x_{K3} \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ f_t^N - f_{t-1}^N & dr_t^2 & dr_t^3 & \cdots & dr_t^N & x_{1t} & \cdots & x_{Kt} \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ f_T^N - f_{T-1}^N & dr_T^2 & dr_T^3 & \cdots & dr_T^N & x_{1T} & \cdots & x_{KT} \end{pmatrix}, \tag{A.104}$$

where $\{x_{kt}\}$ are additional variables that can be resampled at the same time to preserve the cross-section correlation structure.[6] More details can be found in Crump and Gospodinov (2019). I choose $\{x_{kt}\} = \{e_t^{m,n} : m \in \mathfrak{M}, n \in \mathfrak{N}\}$ where $\mathfrak{M}$ and $\mathfrak{N}$ are sets of forecast horizons and yield maturities considered, respectively.

Given a resampled matrix Z^*, I can use the reconstructed yields, $\{y_t^{n,*}\}$, and forecast residuals, $\{e_t^{m,n,*}\}$, to calculate bootstrapped forecasts as

$$\tilde{\mathbb{E}}_t^* y_{t+m}^n = y_t^{m,*} + \frac{n+m}{n}(y_t^{n+m,*} - y_t^{m,*}) + e_t^{m,n,*}. \tag{A.105}$$

Finally, I re-run the regressions, as in Equation (1.6), and record the Newey-West t-statistics with respect to the null hypothesis $\beta = 1$ in the bootstrapped samples. The bootstrapped p-values are obtained by comparing the sample t-statistic to the distribution of bootstrapped t-statistics.

[6]The block length is determined as in Politis and White (2004). I thank Andrew Patton for making the Matlab code available online.

A.7 Proxies for Long-term Macroeconomic Expectations

Proxies of long-term macroeconomic expectations are constructed as the exponentially weighted average of past data realizations

$$g_t^{LT} = \nu_g g_t + (1-\nu_g) g_{t-1}^{LT} = \nu_g \sum_{k=0}^{\infty} (1-\nu_g)^k g_{t-k}, \tag{A.106}$$

$$\pi_t^{LT} = \nu_\pi \pi_t + (1-\nu_\pi) \pi_{t-1}^{LT} = \nu_\pi \sum_{k=0}^{\infty} (1-\nu_\pi)^k \pi_{t-k}, \tag{A.107}$$

for real GDP growth and inflation, respectively. To avoid an arbitrary truncation of the sum, I use a very long history of past observations on real GDP growth and inflation, up to the point at which the weights become close to negligible for initial observations. From 1947 onwards, I use quarterly vintages of first-release values from the Philadelphia Fed for real GDP growth and lagged log changes in core CPI for inflation.[7] Before 1947, I use real GDP growth from the Maddison Project Database (Bolt et al. 2018), available at the University of Groningen, and historical CPI from Shiller (2005).[8] The gain parameters are estimated to match long-term survey forecasts from 1983 to 2018 at $\nu_g = 0.025$ and $\nu_\pi = 0.045$ in quarterly updating terms. These values are close to estimates in the previous literature that match 1-year-ahead forecasts (for a review, see Cieslak and Povala 2015). Figure B.2 plots the fitted proxies for long-term expectations with a comparison to survey data. Regressions show that g^{LT} (π^{LT}) explains the long-term average forecasts of real GDP growth (inflation) with a coefficient of 0.96 (0.85) and R^2 of 0.57 (0.87), indicating a high informativeness.

A.8 Bootstrap for Predictive Regressions

I use a bootstrap approach that follows Greenwood and Vayanos (2014). The bootstrapped samples are constructed by re-drawing dependent variables and predictors with a circular block bootstrap. The optimal block length is determined as in Politis and White (2004). The sample Newey-West t-statistic is calculated with respect

[7]SPF (a) sends survey questionnaires after the Bureau of Economic Analysis releases the advance estimate to the public and (b) collects questionnaires before the second releases. More information can be found in the documentation of SPF.

[8]CPI is not subject to revisions (Croushore and Stark 2001). Real GDP growth from the Maddison Project Database is estimated based on historical benchmarks to reflect historical living standards (Bolt et al. 2018). Thus, the effect of look-ahead bias is unknown but arguably not large.

to the null that slope coefficient β_1 is zero. Since the bootstrapped samples are not simulated under the null, the bootstrapped Newey-West t-statistics are calculated with respect to the null that the bootstrapped estimate β_1^* is equal to the sample estimate $\hat{\beta}_1$ (MacKinnon 2009). The bootstrapped p-values are obtained by comparing the sample Newey-West t-statistic to the distribution of bootstrapped t-statistics.

A.9 An AR(1) Approximation of AR(2) Processes

Assume x_t follows an AR(2) process

$$(1 - \rho_{x,1} L - \rho_{x,2} L^2) x_t = \zeta_t, \tag{A.108}$$

where L is the lag operator and ζ_t is an i.i.d. shock. If $\rho_{x,1}$ and $\rho_{x,2}$ are such that

$$\rho_{x,1} \approx 2\rho_x \ , \ \rho_{x,2} \approx -\rho_x^2 \tag{A.109}$$

for some ρ_x, we can write Equation (A.108) approximately as

$$(1 - \rho_x L)^2 x_t = \zeta_t, \tag{A.110}$$

or

$$(1 - \rho_x L) x_t = \zeta_t^*, \tag{A.111}$$

with $\zeta_t = \zeta_t^* - \rho_x \zeta_{t-1}^*$. It is easy to see that $\{\zeta_t^*\}$ cannot have a zero auto-correlation. However, I make the AR(1) approximation that $\rho_x \approx \frac{\rho_{x,1}}{2}$ by ignoring the auto-correlation structure.

APPENDIX B

Expectation Formation: Additional Results

B.1 Additional Results

Table B.1: Constructing Objective Risk Premia

This table reports the results of constructing objective risk premia using the cycle factor, $\widehat{cf}_t$, as in Cieslak and Povala (2015), and the macro factor, $F6_t$, as in Ludvigson and Ng (2009). Panel A reports factor estimation results of regressing maturity-weighted average excess returns, $\overline{rx}_{t+1} \equiv \frac{1}{14}\sum_{k=2}^{15} rx_{t+1}^k/k$, on cycle variables and macroeconomic principle components. Panel B reports the predictive regressions of individual bond returns on the constructed factors. Newey-West standard errors with 18-month lags are reported in the parentheses. The sample period is from November 1971 to December 2018.

Panel A: Factor estimation, $\overline{rx}_{t+1} = \boldsymbol{\gamma}^T \boldsymbol{f_t} + \varepsilon_{t+1}$

	Cycle		Macro					
	γ_1	γ_2	β_1	β_2	β_3	β_4	β_5	β_6
coeff.	1.31	-0.70	0.94	-0.49	1.03	0.45	-0.27	0.07
(s.e.)	(0.15)	(0.09)	(0.26)	(0.11)	(0.31)	(0.16)	(0.33)	(0.42)
Adj. R^2	0.46		0.18					

Panel B: Individual bond regression, $\frac{1}{n} rx_{t+1}^n = \alpha_1 \widehat{cf}_t + \alpha_2 F6_t + \eta_{t+1}$

	$n=2$			$n=5$			$n=10$		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
$\widehat{cf}_t$	0.73		0.59	1.03		0.85	1.02		0.89
	(0.12)		(0.12)	(0.12)		(0.12)	(0.11)		(0.11)
$F6_t$		0.86	0.51		1.12	0.61		1.01	0.43
		(0.18)	(0.13)		(0.17)	(0.12)		(0.16)	(0.13)
Adj. R^2	0.29	0.16	0.37	0.41	0.20	0.48	0.48	0.18	0.51

Table B.2: Short-Rate Process

This table reports the AR(1) estimation results of the 1-month zero-coupon yields from Liu and Wu(2019)

$$y^1_{t+1} = \mu + \rho y^1_t + \sigma \varepsilon_{t+1}.$$

The bias-adjusted coefficients are calculated from sample estimates $\hat{\rho}$, $\hat{\mu}$, and $\hat{\sigma}$ as

$$\rho = (\hat{\rho} + 1/T)/(1 - 3/T),$$

$$\mu = \hat{\mu}(1 - \rho)/(1 - \hat{\rho}),$$

$$\sigma = \hat{\sigma}[(1 - \rho^2)/(1 - \hat{\rho}^2)]^{\frac{1}{2}}.$$

The sample period is from January 1972 to December 2018.

	1972M1 - 2018M12
$\hat{\rho}$	0.9860
Bias-adj ρ	0.9930
$\hat{\mu}$	0.0638
Bias-adj μ	0.0317
$\hat{\sigma}$	0.5857
Bias-adj σ	0.4137

Table B.3: Trend and Cycle in Short-Rate Expectations

Panel A reports results of regressing 3-month Treasury bill (tbill) rate forecasts at different horizons on long-term average forecasts (forecasts for 7- to 11-year-ahead averages) and nowcasts (forecasts for the current quarter). Newey-West t-statistics with 6-quarter lags are reported in brackets. The sample period is from 1983Q3 to 2018Q4, for which long-term average forecasts are available.
Panel B reports results of 1-year and long-term average forecasts of tbill rates on short- and long-term forecasts of real GDP growth (g) and inflation (π). Newey-West t-statistics with a 6-quarter lag are reported in brackets. The sample period is from 1980Q4 to 2018Q4 and 1983Q3 to 2018Q4 for 1-year and long-term average forecasts, respectively.

Panel A: Trend and cycle in tbill rate forecasts

	(1) 1Q	(2) 2Q	(3) 4Q	(4) 2Y	(5) 5Y
LT avg.	-0.06	-0.04	0.16	0.60	0.88
	[-3.03]	[-0.93]	[1.85]	[6.50]	[13.59]
Nowcast	1.01	0.99	0.85	0.48	0.06
	[97.37]	[43.69]	[17.09]	[6.85]	[2.69]
Obs.	142	142	142	142	142
Adj. R^2	0.99	0.99	0.98	0.94	0.98

Panel B: Links to macroeconomic expectations

	1-year tbill			LT avg. tbill		
	(1)	(2)	(3)	(1)	(2)	(3)
LT avg. g	3.64		2.87	1.30		1.15
	[5.44]		[4.20]	[4.43]		[3.38]
LT avg. π	1.37		1.23	0.78		0.75
	[18.03]		[13.12]	[8.98]		[8.89]
Nowcast g		0.30	0.21		0.17	0.05
		[1.55]	[3.48]		[1.08]	[0.79]
Nowcast π		1.16	0.23		0.57	0.06
		[8.33]	[2.91]		[4.32]	[0.83]
Obs.	153	153	153	142	142	142
Adj. R^2	0.89	0.56	0.91	0.88	0.36	0.89

Table B.4: Constructing Objective Real GDP Growth Expectations

This table reports the results of 1-year-ahead predictive regressions for real GDP growth

$$g_{t+1} = \gamma_0 + \gamma_1 g_t + \gamma_2 s_t^{\$,n} + \varepsilon_{t+1}.$$

g_t is the year-on-year real GDP growth and $s_t^{\$,n}$ is the difference between 7-year and 3-month zero-coupon yields. The regressions are performed at quarterly frequency from 1964Q1 to 2018Q4. Newey-West t-statistics with 6-quarter lags are reported in brackets.

	(1)	(2)	(3)
Yield spread	0.68		0.74
	[2.98]		[3.42]
Lagged obs.		0.15	0.21
		[1.61]	[2.04]
Obs.	216	216	216
Adj. R^2	0.14	0.02	0.18

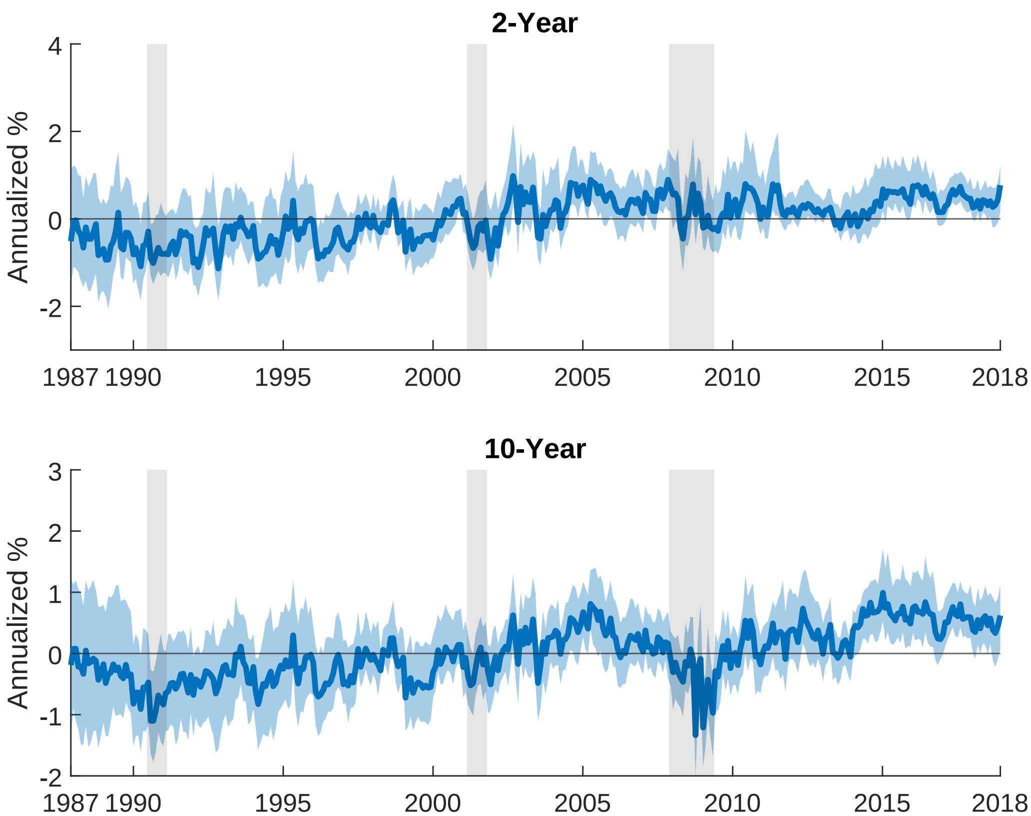

Figure B.1: Forecasts Implied by Constant Subjective Risk Premia

The solid blue lines plots the difference between real-time survey forecasts of n-maturity yields from BCFF, $\tilde{\mathbb{E}}_t y_{t+1}^n$, and forecasts implied by constant subjective risk premia:

$$\tilde{\mathbb{E}}_t^* y_{t+1}^n \equiv \frac{(n+1)y_t^{n+1} - y_t^1 - \tilde{\mathbb{E}}_t^* rx_{t+1}^{n+1}}{n}.$$

Here, constant subjective risk premia on an n-maturity zero-coupon bond, $\tilde{\mathbb{E}}_t^* rx_{t+1}^n$, are sample averages of $\tilde{\mathbb{E}}_t rx_{t+1}^n$. The blue shaded areas are bounded by the difference between top-/bottom-10 average forecasts (also from BCFF) and $\tilde{\mathbb{E}}_t^* y_{t+1}^n$. The upper panel uses 2-year yield and the bottom uses 10-year yield. The sample period is from December 1987 to December 2018.

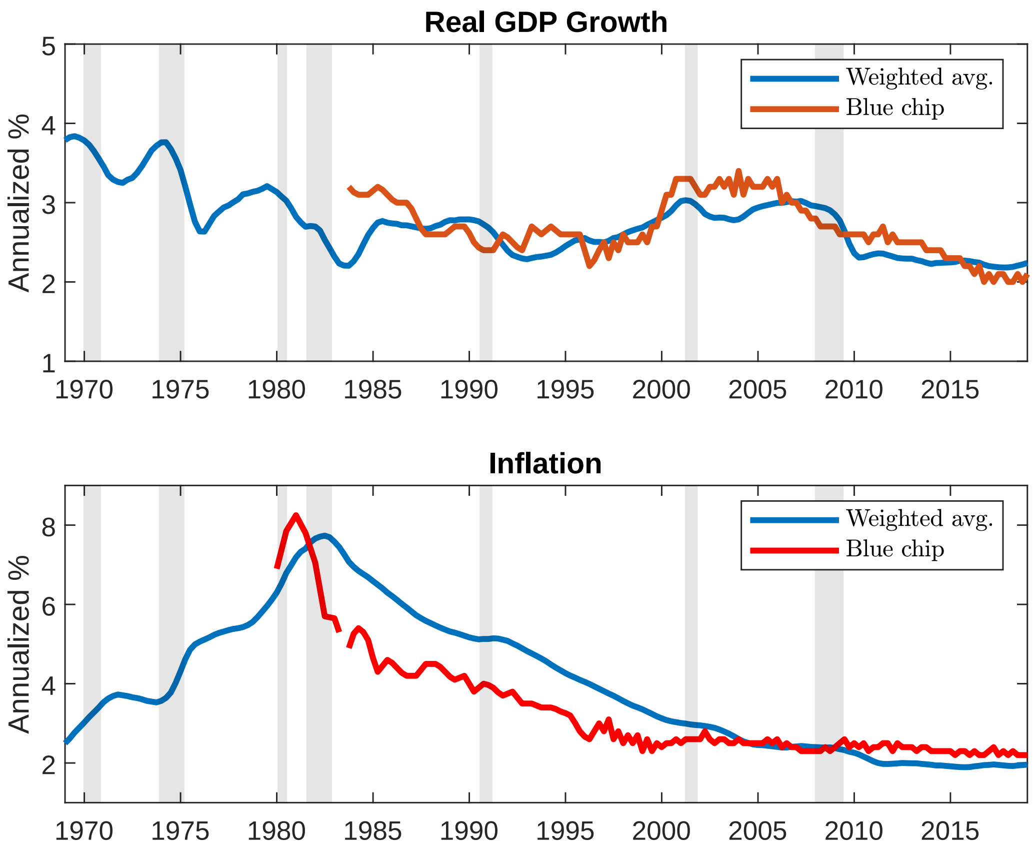

Figure B.2: Proxies for Long-term Macroeconomic Expectations

This figure compares the long-term average macroeconomic forecasts from blue chip surveys and proxies constructed as the exponentially-weighted average of past data realizations

$$g_t^{LT} = \nu_g g_t + (1-\nu_g) g_{t-1}^{LT} = \nu_g \sum_{k=0}^{\infty} (1-\nu_g)^k g_{t-k},$$

$$\pi_t^{LT} = \nu_\pi \pi_t + (1-\nu_\pi) \pi_{t-1}^{LT} = \nu_\pi \sum_{k=0}^{\infty} (1-\nu_\pi)^k \pi_{t-k}.$$

Here g and π denote real GDP growth and inflation, respectively. To avoid an arbitrary truncation of the sum, I use a very long history of past observations on real GDP growth and inflation, up to the point at which the weights become close to negligible for initial observations. Descriptions of data used are given in Appendix A.7. The gain parameters are estimated to match long-term survey forecasts from 1983 to 2018 at $\nu_g = 0.025$ and $\nu_\pi = 0.045$ in quarterly updating terms.

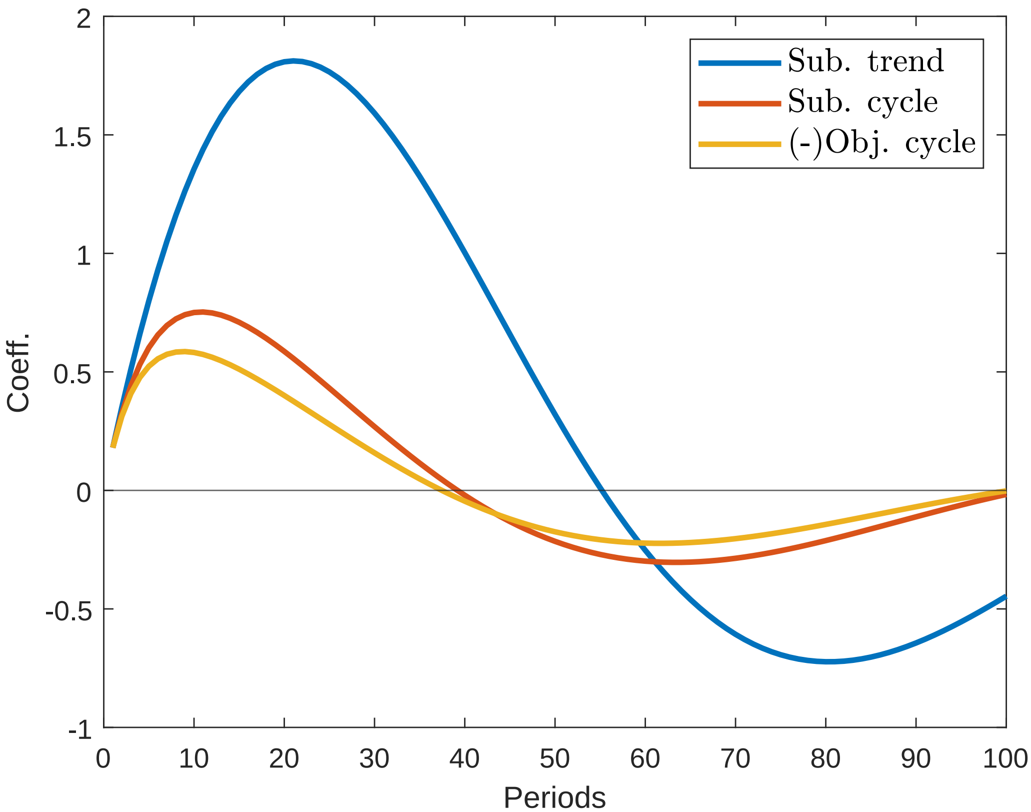

Figure B.3: Function Coefficients of Model-implied Short-rate Expectation Wedges

This figure plots the coefficients of $\widehat{\mu}_t$ $(c^*_{g,j})$, $\widehat{c}_t$ $(d^*_{g,j})$, and the negative of c^e_t $(-b^*_{g,j})$ in the following function of model-implied short-rate expectation wedges at different horizons:

$$\mathbb{E}_t i^{\$}_{t+j} - \tilde{\mathbb{E}}_t i^{\$}_{t+j} = \frac{1}{\psi}\left(a^*_{g,j}\mu + b^*_{g,j}c^e_t + c^*_{g,j}\widehat{\mu}_t + d^*_{g,j}\widehat{c}_t\right) + \left(a^*_{\pi,j}\tau + b^*_{\pi,j}\kappa^e_t + c^*_{\pi,j}\widehat{\tau}_t + d^*_{\pi,j}\widehat{\kappa}_t\right).$$

The parameter values are

$$\phi_\mu = 0.96\ ,\ \tilde{\rho}_c = 0.82\ ,\ \rho_c = 0.75,\ ,\ \nu_\mu = 0.02\ ,\ \nu_c = 0.2.$$

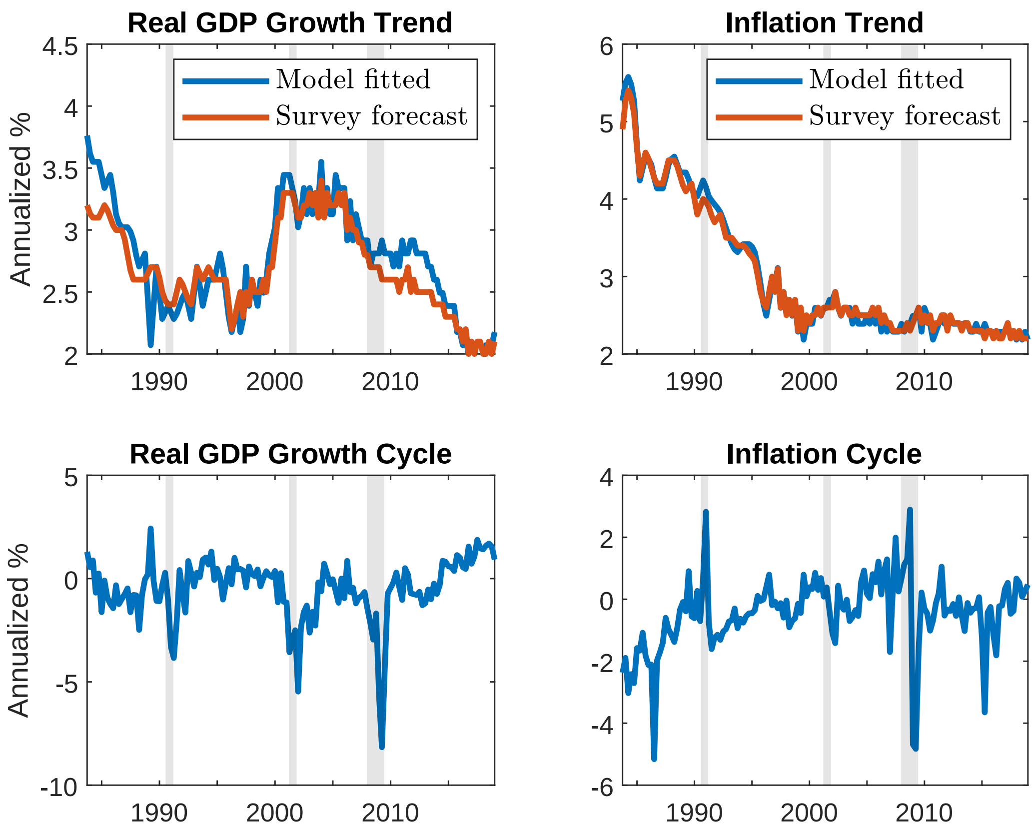

Figure B.4: Model-implied Trend and Cycle Expectations

This figure plots the model-implied trend and cycle expectations in real GDP growth and inflation. For the top two panels, I plot model-implied long-term average forecasts, with a comparison to observed data. For the bottom two panels, I plot $\widehat{c}_t$ and $\widehat{\kappa}_t$. The sample period is from 1983Q3 to 2018Q4.

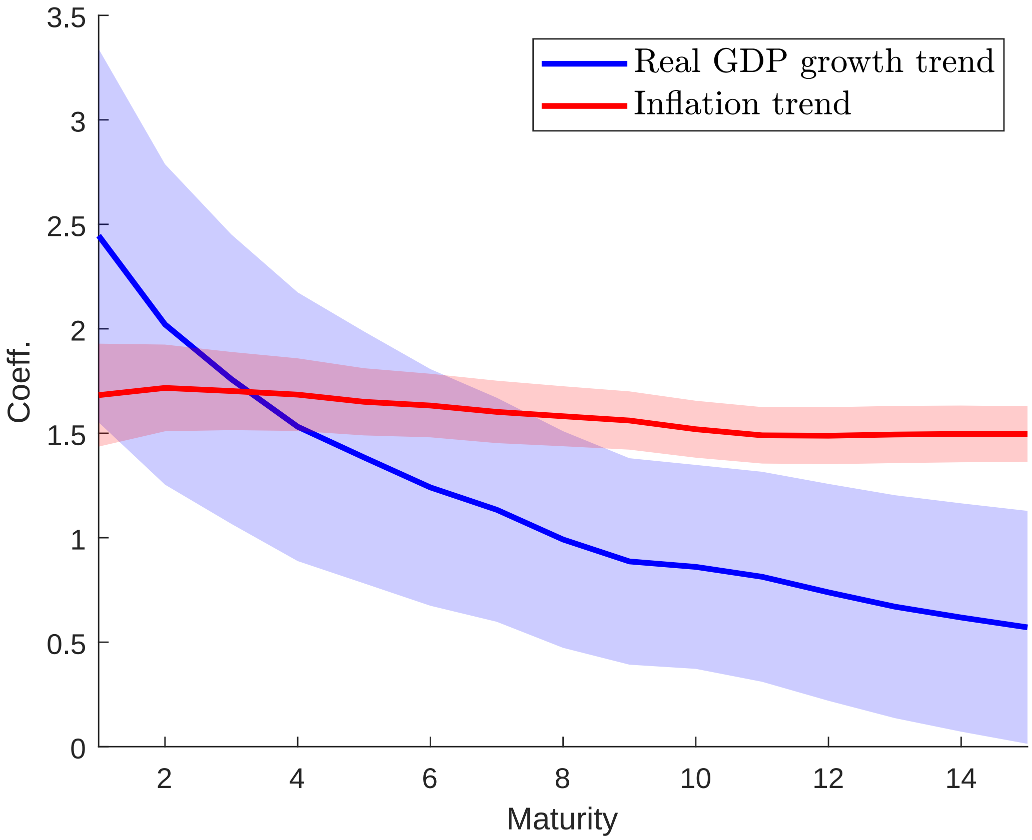

Figure B.5: Nominal Yields' Loading on Trend Expectations

This figure plots the loadings of nominal yields on trend expectations of real GDP growth and inflation. The trend expectations are proxied by g^{LT} and π^{LT}. I consider nominal yields with maturities from 1 to 15 years. The shaded areas are bounded by $\pm$ 2 standard errors of the coefficient estimates. The sample period is from 1971Q4 to 2018Q4, with yield data sampled in the middle month of a calendar quarter.

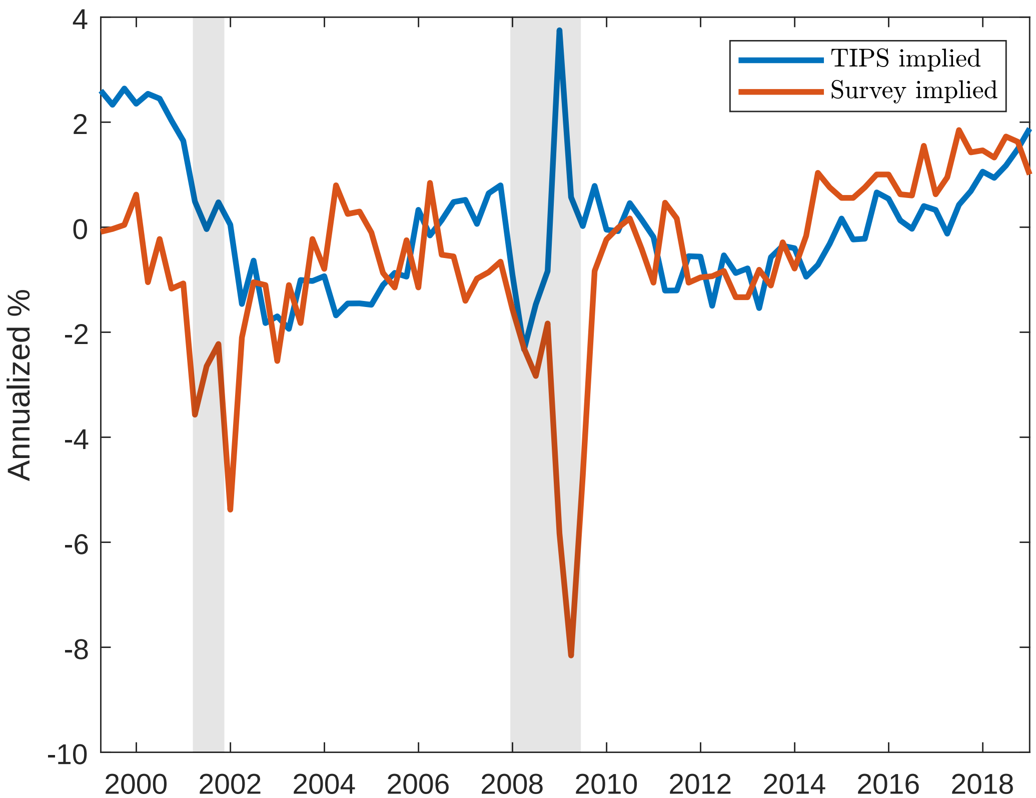

Figure B.6: Comparison of Implied Real Growth Cycle Expectation

This figure plots the real GDP growth cycle expectations implied by TIPS and survey forecasts. The blue line plots the residuals from regressing 2-year TIPS yield on g^{LT}. The red line plots the survey-implied cycle expectations from Section 1.6.

APPENDIX C

Fading Memory: Data Construction, Derivations, and Proofs

C.1 Data

C.1.1 Financial market data

Population. We obtain quarterly data on the population of the United States from 1947 onwards from the FRED database at the Federal Reserve Bank of St. Louis. For 1871 to 1946, we obtain annual population numbers from the Historical Statistics of the United States, available at `https://hsus.cambridge.org/`.

Stock index returns. We calculate quarterly stock index returns from 1871 to 1925 using data from Shiller (2005) back to 1871. From 1926 onwards, we use quarterly returns on the value-weighted CRSP index.

Inflation. We use the Consumer Price Index (CPI) series (and the pre-cursors of the official CPI) in Shiller's data set to deflate returns and the payout growth series that we describe next.

Payouts. From 1926 onwards, we calculate quarterly aggregate dividends using the lagged total market value of the CRSP value-weighted index and the difference between quarterly returns with and without dividends. Furthermore, working with the monthly CRSP individual stock files and following Bansal et al. (2005), we use reductions in the shares outstanding (after adjusting for stock splits, stock dividends, etc. using the CRSP share adjustment factors) as reported by CRSP to calculate stock repurchases. To eliminate the effect of data errors (there are instances where the shares outstanding drop by a huge amount and jump back up a few months later), we drop observations where the shares outstanding fall by more than 10 percent within one month. Repurchases account for an economically significant share of payouts only from the 1980s onwards. We aggregate the sum of dividends and repurchases across firm within each month and then at the aggregate level within each quarter. Dividing

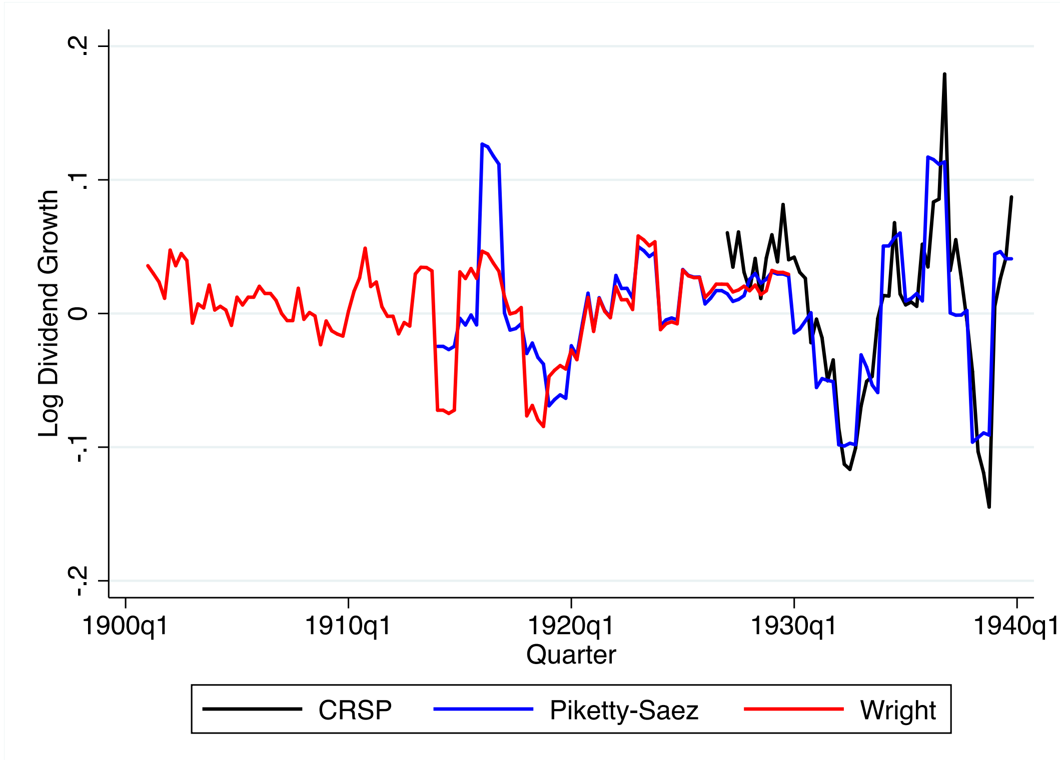

Figure C.1: Dividend Growth from Three Data Sources

by the size population at each point in time, we obtain per-capita payouts. To avoid seasonality effects, we compute growth rates as the four-quarter change in log per-capita aggregate payouts, divided by four.

Prior to 1926, we use annual data on aggregate household dividend receipts from tax data in Piketty et al. (2017) for the period 1913 to 1926. While this data source covers only the portion of dividends received by households, its advantage is that it is based on high-quality administrative data. As long as the share of total aggregate dividends received by households does not change much from year to year, the growth rates calculated from this data set should approximate well the growth rates in total aggregate dividends. Figure C.1 suggests that this is the case: in years when the Piketty-Saez data overlaps with CRSP, the per-capita growth rates obtained from the two data sets are closely aligned.

For the period from 1900 to 1913, we use a series of annual aggregate corporate non-farm non-financial dividends from Wright (2004). Figure C.1 shows that the per-capita growth rates calculated from the Wright data are, with a few exceptions, very close to those from the Piketty-Saez data. For the period from 1871 to 1900, we use real per-capita GDP growth rates from Barro and Ursúa (2008) (which are in turn based on Balke and Gordon (1989)) as proxy for Δd from 1871 to 1900.

Bond and bill yields. To calculate subjectively expected excess return from survey expectations, we use the one-year constant maturity Treasury yield, obtained from FRED database at the Federal Reserve Bank of St. Louis. To calculate quarterly stock market index excess returns back to 1926, we use the three-month T-Bill yield series from Nagel (2016), extended until the end of 2016 with 3-month T-bill yields from the FRED database (where we convert the reported discount yields into effective annual yields).

C.1.2 Survey data on return expectations

Three surveys provide us with direct measures of percentage expected stock market returns over a one-year horizon: UBS/Gallup (1998-2007, monthly); Vanguard Research Initiative (VRI) survey of Ameriks et al. (2016) (2014, one survey); and Surveys of Lease et al. (1974) and Lewellen et al. (1977) (one survey per year in 1972 and 1973). In part of the sample, the UBS/Gallup survey respondents report only the return they expect on their own portfolio. We impute market return expectations by regressing expected market returns on own portfolio expectations using the part of the sample where both are available and using the fitted value from this regression when the market return expectation is not reported. The VRI survey asks about expected growth in the Dow Jones Industrial Average (DJIA). Since the DJIA is a price index, we add to the price growth expectation the dividend yield of the CRSP value-weighted index at the time of the survey. Figure C.2 shows the time series of expectations from these surveys.

In the next step, we use the data from the Michigan Survey of Consumers (MSC). The MSC elicits the perceived probability that an investment in a diversified stock fund would increase in value over a one-year horizon. For comparability with the other surveys above, which are all based on surveys of people that hold stocks, we restrict the sample to respondents that report to hold stocks (as the MSC does for the aggregate stock market beliefs series that they publish on their website). To impute percentage expectations, we regress the percentage expectations from the UBS/Gallup and VRI surveys on the MSC probability. The red line in Figure C.2 shows the resulting fitted value. In the periods when the series overlaps with the UBS/Gallup and VRI samples, the fit is very good, indicating that the simple imputation procedure delivers reasonable results. In time periods when the UBS/Gallup and VRI surveys are not available, we use this fitted value.

Finally, we bring in data from the Conference Board (1986-2016, monthly) and Roper surveys (1974-1997, one survey per year). These surveys elicit respondents simple categorical beliefs about whether the "stock prices" will likely increase, decrease, or stay the same (or whether they are undecided, which we include in the "same" category). We construct the ratio of the proportion of those who respond with "increase" to the sum of the proportions of "decrease" and "same." We then regress the expected return series that we obtained from the surveys above on this ratio. More precisely, since the Conference Board and Roper surveys ask about stock price increases, we subtract the current dividend yield of the CRSP value weighted index from the dependent

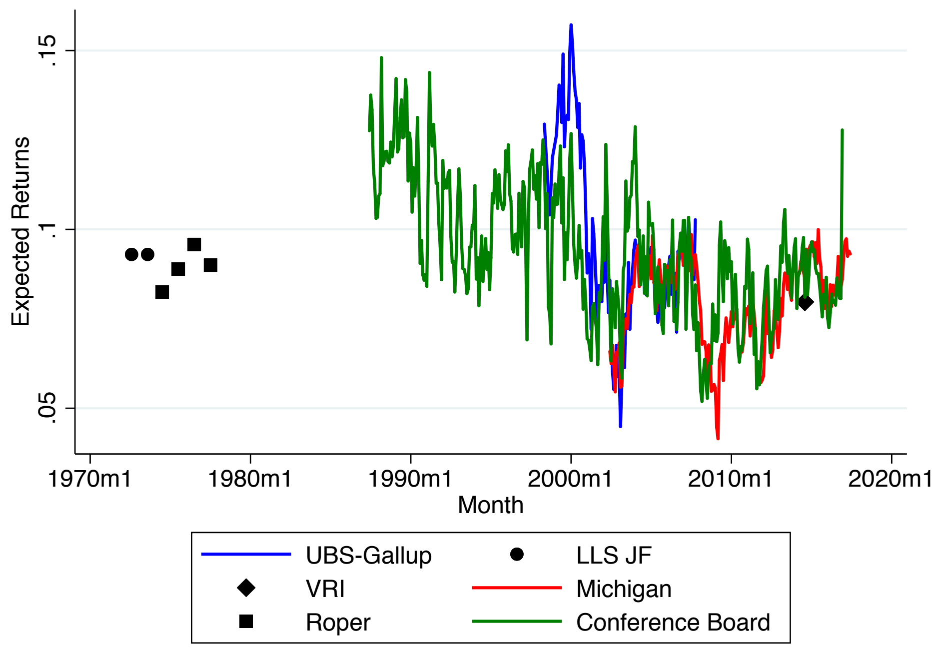

Figure C.2: Expected Return Imputation

.

variable in this regression and we add it back to the fitted value. The green line in Figure C.2 shows the fitted value from this regression for the Conference Board series and the four squares show the fitted value for the Roper surveys. Except for a relatively short period around the year 2000, the fitted series tracks the expected returns from UBS/Gallup, MSC, and VRI very well. In time periods when the UBS/Gallup, MSC and VRI surveys are not available, we use this fitted value.

C.1.3 Analysts long-term earnings expectations

We collect the median of analysts' stock-level long-term median earnings-per-share (EPS) growth forecasts from the I/B/E/S Unadjusted US Summary Statistics database, focusing on U.S. firms with earnings forecasts in U.S. dollars. The data is available at a monthly frequency starting in December 1981, but shares outstanding information that we need to aggregate across stocks becomes available only in September 1984 for all but 10% of stocks covered by I/B/E/S. We therefore start the sample in September 1984. We also collect the median forecasts of current fiscal year EPS. We use those to calculate forecasted current year total earnings by multiplying forecasted EPS with shares outstanding. We aggregate across stocks by forming a value-weighted average each quarter using total forecasted current fiscal year earnings as the weight. Stocks with negative values of this weight variable are excluded. We then aggregate the monthly long-term EPS growth forecasts across months within each calendar quarter.

We also collect the median forecasts for next fiscal year EPS. We calculate the growth rate from current year to next-year earnings. We then aggregate the data, across stocks, then within quarter, in the same way as described above for the long-term forecasts.

We convert these nominal expected growth rates to real expected growth rates by deflating with inflation forecasts from the Survey of Professional Forecasters (SPF), which we obtain from the Federal Reserve Bank of Philadelphia.[1] We use the one-year and ten-year CPI median inflation forecasts in the quarter prior to the quarter in which the analyst forecast is made. The 10-year forecast is available in the SPF only starting in quarter 1991:4. Prior to this date we use the extended series constructed by the Philadelphia Fed from the Blue Chip Economic Indicators and the Livingston survey that is available on the same website. In quarters when the extended series has missing values, we substitute the value from the previous quarter.

Long-term earnings growth rates in the I/B/E/S data are meant to represent an expected growth rate over the next three to five years, or a "full business cycle" (see, e.g., Sharpe (2002)). To deflate with an expected inflation rate that approximates this forecast horizon, we subtract the average of the one-year and ten-year inflation forecast from the long-term earnings forecast. We subtract the one-year inflation forecast from the one-year earnings growth forecast. As a last step, we convert the real expected earnings growth series to quarterly frequency by dividing the annualized numbers by four.

C.2 Bootstrap Simulations for Predictive Regressions

Our bootstrap simulations closely follow those in Kothari and Shanken (1997), but extended to multiple predictor variables. We start by estimating AR(1) processes for the predictor variables and we add $1/T$ to the slope coefficient to perform first-order bias-adjustment (and we adjust the intercept accordingly). We also estimate the predictive regression for returns by OLS and record the residuals.

Using these bias-adjusted coefficients from the estimated AR(1) for the predictors, we then simulate a VAR(1) with a diagonal coefficient matrix, where the innovations are the bootstrapped residuals from the estimated AR(1). As in Kothari and Shanken (1997), we condition on the first observation of the predictor time series. We preserve contemporaneous correlations of the innovations by drawing vectors of residuals for the different predictors.

Based on the simulated predictor series, we then also simulate two return series by combining the predictor time-series with bootstrapped residuals from the predictive regression. For the first return series, we set the predictive regression slope coefficients equal to the OLS predictive regression estimate, i.e., we simulate under the alternative. For the second series, we set the predictive regression slope coefficients equal to zero,

[1] Available at `https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/historical-data/inflation-forecasts`

i.e., in this case we are simulating under the null hypothesis of no predictability.

We use the described approach to simulate 10,000 bootstrap samples of predictors and the two returns series. We then run the predictive regressions on the bootstrap samples and record the regression coefficients and t-statistics. We obtain the predictive regression bias-adjustment by comparing the mean slope coefficients from the bootstrap samples with the first return series (alternative) to the OLS estimate. We obtain the p-values by comparing the sample predictive regression t-statistic to the quantiles of the distribution of the t-statistic in the bootstrap regressions with the second return series (null).

C.3 Properties of the Predictive Distribution

We describe now the properties of $\tilde{\varepsilon}_{t+j}$, $j = 1, 2, \ldots$ under the time-t predictive distribution.

We first show that the subjective conditional variance of $\tilde{\varepsilon}_{t+j}$ is decreasing in the forecast horizon. First note the perceived consumption growth process has the following autocovariance structure

$$\tilde{\mathrm{cov}}_t(\Delta c_{t+i}, \Delta c_{t+j}) = \nu\sigma^2, \quad j > i \geq 1, \tag{C.1}$$

which arises from the agent's uncertainty about μ. From the definition of $\tilde{\varepsilon}$, we obtain

$$\begin{aligned}
\tilde{\varepsilon}_{t+i} &= \frac{\Delta c_{t+i} - \tilde{\mu}_{t+i-1}}{\sqrt{1+\nu}\sigma}. \\
&= \frac{\Delta c_{t+i} - \nu\Delta c_{t+i-1} - (1-\nu)\tilde{\mu}_{t+i-2}}{\sqrt{1+\nu}\sigma}. \\
&= \frac{\Delta c_{t+i} - \nu\sum_{j=1}^{i-1}(1-\nu)^{j-1}\Delta c_{t+i-j} - (1-\nu)^{i-1}\tilde{\mu}_t}{\sqrt{1+\nu}\sigma}.
\end{aligned} \tag{C.2}$$

Because of the constant autocovariance structure of perceived consumption growth,

$$\begin{aligned}
\tilde{\mathrm{var}}_t(\tilde{\varepsilon}_{t+i-1}) &= \tilde{\mathrm{var}}_t\left(\frac{\Delta c_{t+i-1} - \nu\sum_{j=1}^{i-2}(1-\nu)^{j-1}\Delta c_{t+i-1-j}}{\sqrt{1+\nu}\sigma}\right) \\
&= \tilde{\mathrm{var}}_t\left(\frac{\Delta c_{t+i} - \nu\sum_{j=1}^{i-2}(1-\nu)^{j-1}\Delta c_{t+i-j}}{\sqrt{1+\nu}\sigma}\right),
\end{aligned} \tag{C.3}$$

and

$$\begin{aligned}
\tilde{\mathrm{var}}_t(\tilde{\varepsilon}_{t+i}) &= \tilde{\mathrm{var}}_t(\tilde{\varepsilon}_{t+i-1}) + \nu^2(1-\nu)^{2i-4} - \frac{2\nu^2}{1+\nu}(1-\nu)^{2i-4} \\
&= \tilde{\mathrm{var}}_t(\tilde{\varepsilon}_{t+i-1}) - \frac{1-\nu}{1+\nu}\nu^2(1-\nu)^{2i-4}, \qquad i \geq 2.
\end{aligned} \tag{C.4}$$

This leads to

$$\tilde{\text{var}}_t(\tilde{\varepsilon}_{t+i}) = 1 - \frac{1-\nu}{1+\nu}\nu^2 \frac{1-(1-\nu)^{2i-2}}{1-(1-\nu)^2}, \tag{C.5}$$

which decreases over time and converges to $1 - \frac{1-\nu}{1+\nu}\frac{\nu}{2-\nu}$.

Using these results, we can calculate the time-t perception of the autocovariance of future $\tilde{\varepsilon}_{t+i}$,

$$\begin{aligned}
&\tilde{\text{cov}}_t(\tilde{\varepsilon}_{t+i}, \tilde{\varepsilon}_{t+i+1}) \\
=&\tilde{\text{cov}}_t\left(\frac{\Delta c_{t+i} - \nu\sum_{j=1}^{i-1}(1-\nu)^{j-1}\Delta c_{t+i-j}}{\sqrt{1+\nu}\sigma}, \frac{\Delta c_{t+i+1} - \nu\sum_{j=1}^{i-1}(1-\nu)^{j-1}\Delta c_{t+i+1-j}}{\sqrt{1+\nu}\sigma}\right) \\
=& -\frac{1}{1+\nu} + (1-\nu)\tilde{\text{var}}_t(\tilde{\varepsilon}_{t+i}) \\
=& -\frac{\nu^2}{1+\nu} - \frac{(1-\nu)^2}{1+\nu}\nu^2\frac{1-(1-\nu)^{2i-2}}{1-(1-\nu)^2} < 0.
\end{aligned} \tag{C.6}$$

i.e., it is negative, as we claimed in the main text.

C.4 Kalman Filtering Interpretation

Here we show that our model can be mapped into a full memory model that is equivalent in terms of the relevant subjective belief dynamics and asset prices. In this equivalent version of the model, the agent perceives a latent AR(1) trend growth rate and she uses the Kalman filter to optimally track this latent trend, while objectively the trend growth rate is constant. The agent uses full memory and the information structure is a filtration and it is Markovian. In the agent's subjective view, past data gradually loses relevance for forecasting not because of fading memory but because it is perceived as irrelevant given the perceived stochastic drift over time in the trend growth rate.

C.4.1 Diffuse prior

Suppose the agent at time t perceives the law of motion

$$\Delta c_t = \mu_t + \xi_t, \qquad \xi_t \sim \mathcal{N}(0, \sigma_\xi^2), \tag{C.7}$$

$$\mu_{t+1} = \mu_t + \zeta_{t+1}, \qquad \zeta_{t+1} \sim \mathcal{N}(0, \sigma_\zeta^2), \tag{C.8}$$

where the agent knows σ_ξ^2 and σ_ζ^2, but not μ_t. With diffuse prior and an infinite history, H_t, of observed data on Δc, the predictive distribution can be obtained from the steady-state Kalman filter (see, e.g., Hamilton (1994)) as

$$\mu_{t+1}|H_t \sim \mathcal{N}(\hat{\mu}_{t+1|t}, \omega^2 + \sigma_\zeta^2), \tag{C.9}$$

where the optimal forecast $\hat{\mu}_{t+1|t} \equiv \hat{E}(\mu_{t+1}|H_t)$ evolves as

$$\hat{\mu}_{t+1|t} = \hat{\mu}_{t|t-1} + K(\Delta c_t - \hat{\mu}_{t|t-1}), \tag{C.10}$$

with

$$K = \frac{\omega^2 + \sigma_\zeta^2}{\omega^2 + \sigma_\zeta^2 + \sigma_\xi^2}, \tag{C.11}$$

and

$$\omega^2 = K\sigma_\xi^2. \tag{C.12}$$

Thus the predictive distribution of Δc_{t+1} at time t is

$$\Delta c_{t+1} \sim \mathcal{N}(\hat{\mu}_{t+1|t}, \omega^2 + \sigma_\zeta^2 + \sigma_\xi^2). \tag{C.13}$$

To map into our fading memory setup, we choose

$$K = \nu, \quad \sigma_\xi^2 = (1-\nu^2)\sigma^2, \quad \sigma_\zeta^2 = \nu^2(1+\nu)\sigma^2. \tag{C.14}$$

The time-t predictive distribution of Δc_{t+1} then is exactly the same as in our fading memory setting. The time-t predictive distribution of Δc_{t+j} for $j > 1$ is different from the fading memory setting, though, because here the agent perceives $\tilde{\mu}_t$ as a martingale and the predictive distribution inherits these martingale dynamics, while in our fading memory setting the predictive distribution converges to a stationary one at long horizons. For pricing, however, this difference in the perceived distribution for $j > 1$ does not matter, because under resale valuation, pricing is based on a chain of valuations of one-period ahead payoffs from selling the asset. Thus, pricing in this perceived stochastic trend setting here is the same as in our fading memory setting.

C.4.2 Informative prior

Suppose the agent at time t perceives the law of motion

$$\Delta c_t = \mu_t + \xi_t, \qquad \xi_t \sim \mathcal{N}(0, \sigma_\xi^2), \tag{C.15}$$

$$\mu_{t+1} = (1-h)\mu + h\mu_t + \zeta_{t+1}, \qquad \zeta_{t+1} \sim \mathcal{N}(0, \sigma_\zeta^2), \tag{C.16}$$

where $0 \leq h < 1$ and the value of h is known to the agent. Steady-state Kalman filter updating yields optimal forecasts of the state as

$$\hat{\mu}_{t+1|t} = (1-h)\mu + h\hat{\mu}_{t|t-1} + K(\Delta c_t - \hat{\mu}_{t|t-1}), \tag{C.17}$$

with

$$K = h\frac{\sigma_\zeta^2 + h^2\omega^2}{\sigma_\zeta^2 + \sigma_\xi^2 + h^2\omega^2}, \tag{C.18}$$

$$\omega^2 = K\sigma_\xi^2/h. \tag{C.19}$$

Iterating yields

$$\hat{\mu}_{t+1|t} = \frac{1-h}{1-h+K}\mu + K\sum_{j=0}^{\infty}(h-K)^j\Delta c_{t-j}. \tag{C.20}$$

The predictive distributions are

$$\mu_{t+1}|H_t \sim \mathcal{N}(\hat{\mu}_{t+1|t}, h^2\omega^2 + \sigma_\zeta^2), \tag{C.21}$$

and

$$\Delta c_{t+1}|H_t \sim \mathcal{N}(\hat{\mu}_{t+1|t}, h^2\omega^2 + \sigma_\zeta^2 + \sigma_\xi^2). \tag{C.22}$$

We can map this into our fading memory setup with informative prior by choosing

$$K = \phi\nu, \quad h = 1-\nu+\phi\nu, \quad \sigma_\xi^2 = \frac{1-\nu}{1-\nu+\phi\nu}(1+\phi\nu)\sigma^2 \tag{C.23}$$

to obtain equivalence in terms of the relevant subjective belief dynamics and asset prices.

C.5 Model Solution for $\psi = 1$

C.5.1 Stochastic discount factor

Following Hansen et al. (2008), we start with value function iteration

$$v_t = \frac{\delta}{1-\gamma}\log\tilde{\mathbb{E}}_t[e^{(1-\gamma)(v_{t+1}+\Delta c_{t+1})}], \tag{C.24}$$

where $v_t = \log(V_t/C_t)$ and V_t is the continuation value. We conjecture the solution to be linear in the state variable, i.e.

$$v_t = \mu_v + U_v\tilde{\mu}_t. \tag{C.25}$$

Plugging in the conjectured solution we get

$$U_v = \frac{\delta}{1-\delta}, \tag{C.26}$$

and

$$\mu_v = \frac{1}{2}(1-\gamma)U_v(\nu U_v + 1)^2(1+\nu)\sigma^2. \tag{C.27}$$

We obtain the log SDF

$$
\begin{aligned}
m_{t+1|t} &= \log\left(\delta \frac{C_t}{C_{t+1}} \frac{(V_{t+1})^{1-\gamma}}{\tilde{\mathbb{E}}_t[(V_{t+1})^{1-\gamma}]}\right) \\
&= \log\delta - \Delta c_{t+1} + (1-\gamma)\log(V_{t+1}) - \log\tilde{\mathbb{E}}_t[(V_{t+1})^{1-\gamma}] \\
&= \log\delta - \Delta c_{t+1} + (1-\gamma)(v_{t+1} + c_{t+1}) - \log\tilde{\mathbb{E}}_t(e^{(1-\gamma)(v_{t+1}+c_{t+1})}) \\
&= \tilde{\mu}_m - \tilde{\mu}_t - \xi\sigma\tilde{\varepsilon}_{t+1}, \qquad \text{(C.28)}
\end{aligned}
$$

where

$$
\tilde{\mu}_m = \log\delta - \frac{1}{2}(1-\gamma)^2(\nu U_v + 1)^2(1+\nu)\sigma^2, \qquad \text{(C.29)}
$$

$$
\xi = [1-(1-\gamma)(\nu U_v + 1)]\sqrt{1+\nu}. \qquad \text{(C.30)}
$$

C.5.2 Consumption claim valuation

Let $\zeta \equiv W_t/C_t$. The return on the consumption claim is

$$
R_{W,t+1} \equiv \frac{W_{t+1}}{W_t - C_t} = \frac{C_{t+1}}{C_t}\frac{\zeta}{\zeta - 1}, \qquad \text{(C.31)}
$$

and in logs,

$$
\begin{aligned}
r_{w,t+1} &= \Delta c_{t+1} + \log(\zeta/(\zeta-1)) \\
&= \tilde{\mu}_t + \sqrt{1+\nu}\sigma\tilde{\varepsilon}_{t+1} + \log(\zeta/(\zeta-1)). \qquad \text{(C.32)}
\end{aligned}
$$

Plugging the return on the consumption claim into the Euler equation and taking logs,

$$
\begin{aligned}
\log(\zeta/(\zeta-1)) &= -\tilde{\mu}_m + \xi\sqrt{1+\nu}\sigma^2 - \frac{1}{2}(1+\nu)\sigma^2 - \frac{1}{2}\sigma^2\xi^2 \\
&= -\log\delta, \qquad \text{(C.33)}
\end{aligned}
$$

which we can solve for the wealth-consumption ratio

$$
\zeta = \frac{1}{1-\delta}. \qquad \text{(C.34)}
$$

That the consumption-wealth ratio is constant can also be seen by valuing consumption strips. Denoting with w_t^1 the log of the component of time-t wealth that derives from the one-period ahead endowment flow, we have

$$\begin{aligned} w_t^1 - c_t &= \log \tilde{\mathbb{E}}_t \left[M_{t+1|t} \frac{C_{t+1}}{C_t} \right] \\ &= \log \tilde{\mathbb{E}}_t \left[\exp(\tilde{\mu}_m + (\sqrt{1+\nu} - \xi)\sigma \tilde{\varepsilon}_{t+1} \right] \\ &= \tilde{\mu}_m + \frac{1}{2}(\sqrt{1+\nu} - \xi)^2 \sigma^2, \end{aligned} \tag{C.35}$$

i.e., $w_t^1 - c_t$ is constant. It does not vary with $\tilde{\mu}_t$ because, going from the first to the second line, $-\tilde{\mu}_t$ in $m_{t+1|t}$ cancels with $\tilde{\mu}_t$ in Δc_{t+1}. Working through the valuation equation backwards in time, we obtain the price of an n-period consumption strip

$$w_t^n - c_t = n\tilde{\mu}_m + \frac{n}{2}(\sqrt{1+\nu} - \xi)^2 \sigma^2. \tag{C.36}$$

Plugging in the solutions for $\tilde{\mu}_m$ and ξ from the previous subsection, we get

$$w_t^n - c_t = n \log \delta. \tag{C.37}$$

Summing the value of consumption strips at all horizons strips yields the consumption-wealth ratio in (C.34).

C.5.3 Dividend strip valuation

By analyzing dividend strips that are claims to single dividends in the future, we can transparently analyze the conditions needed for a finite price. The price of the n-period dividend strip is

$$P_t^n \equiv \tilde{\mathbb{E}}_t[M_{t+1}|_t \tilde{\mathbb{E}}_{t+1}[\cdots \tilde{\mathbb{E}}_{t+n-1}[M_{t+n}|_{t+n-1} D_{t+n}]]]. \tag{C.38}$$

As we discussed earlier, when we evaluate these expectations, we do so by iterating backwards from the payoff at $t+n$, evaluating one conditional expectation at a time without relying on the Law of Iterated Expectations (LIE).

Taking logs and evaluating (C.38), we obtain

$$p_t^n - d_t = [1 - (1-\alpha)^n] \left(c_t - d_t + \mu_{dc} + \frac{\lambda - 1}{\alpha} \tilde{\mu}_t \right) + n\tilde{\mu}_m + \frac{1}{2}(A_n \sigma^2 + B_n \sigma_d^2), \tag{C.39}$$

where

$$A_n = \sum_{k=0}^{n-1} \left\{ \sqrt{1+\nu} \left[\nu(\lambda - 1)\frac{1-(1-\alpha)^k}{\alpha} + (\lambda-1)(1-\alpha)^k + 1 \right] - \xi \right\}^2, \tag{C.40}$$

and

$$B_n = \frac{1-(1-\alpha)^{2n}}{1-(1-\alpha)^2}. \tag{C.41}$$

For very large n, approximately,

$$A_n \approx n\left[\sqrt{1+\nu}\left(1+\nu\frac{\lambda-1}{\alpha}\right)-\xi\right]^2, \tag{C.42}$$

and B_n, which does not grow with n, becomes very small relative to A_n. Thus for the price to be well-defined, we need the terms that grow with n in (C.39) to be (weakly) negative. Using (2.34), we see that this requires

$$\tilde{\mu}_m + \frac{1}{2}\left[\sqrt{1+\nu}\left(1+\nu\frac{\lambda-1}{\alpha}\right)-\xi\right]^2\sigma^2 \leq 0. \tag{C.43}$$

In our calibration, we will work with a value for α that satisfies this condition.

As equation (2.33) shows, the dividend strip log price-dividend ratio is increasing in $\tilde{\mu}_t$ if $\lambda > 1$, i.e., if the dividend claim is levered. With $\lambda = 1$ the effect of a higher expected growth rate of dividends would be offset by a higher risk-free rate, just like it is for the consumption claim. With leverage, the effect of higher expected dividend growth is stronger than the risk-free rate effect. Since dividends and consumption are co-integrated and hence have the same growth rate in the long-run, the reason why the agent can expect the dividend and consumption growth rates to differ for a substantial period of time may not be immediately obvious. When the agent revises upward her posterior mean $\tilde{\mu}_t$, then her expectation of the dividend growth rate in the next period gets revised upward by $\lambda - 1$ times the revision in $\tilde{\mu}_t$. She expects that over the near future dividend growth will exceed consumption growth, leading to a rise in $d - c$. Eventually, the higher log dividend-consumption ratio will—through the cointegration relationship in (2.31)—generate enough negative offset to bring the dividend growth rate back down to $\tilde{\mu}_t$. Thus, she expects the process to settle down with similar mean growth rates, but at a higher $d-c$.[2] Is it economically plausible that the agent perceives higher unconditional mean economic growth to be associated with a higher dividend-consumption ratio? It is impossible to answer this question within a model with exogenous endowment flow. In Appendix C.7 we present calculations based on a Ramsey-Cass-Koopmans model with endogenous investment and production that suggests that a positive relationship between growth and the capital income to consumption ratio is indeed plausible.

To get the expected returns of dividend strips, we start with (2.33) to compute returns. For the one-period claim, we get

$$r^1_{t+1} = \lambda\Delta c_{t+1} - (\lambda-1)\tilde{\mu}_t - \tilde{\mu}_m - \frac{1}{2}\left(\sqrt{1+\nu}\lambda - \xi\right)^2\sigma^2. \tag{C.44}$$

[2] A similar mechanism is at work in Collin-Dufresne et al. (2016a), but there dividends and consumption are not cointegrated, and so $d - c$ can grow without bound, but the unconditional mean growth rate of consumption (and hence dividends) is truncated.

Subtracting $r_{f,t} = -\tilde{\mu}_m + \tilde{\mu}_t - \frac{1}{2}\xi^2\sigma^2$ yields

$$r^1_{t+1} - r_{f,t} = \lambda(\Delta c_{t+1} - \tilde{\mu}_t) - \frac{1}{2}\left(\sqrt{1+\nu}\lambda - \xi\right)^2 \sigma^2 + \frac{1}{2}\xi\sigma^2. \tag{C.45}$$

The subjective conditional variance of r^1_{t+1} is $(1+\nu)\lambda^2\sigma^2$, and so, after taking subjective expectations of (C.45), we obtain

$$\log \tilde{\mathbb{E}}_t[R^1_{t+1}] - r_{f,t} = \lambda\xi\sqrt{1+\nu}\sigma^2. \tag{C.46}$$

The objective conditional variance of r^1_{t+1} is only $\lambda^2\sigma^2$, and so taking objective expectations of (C.45) yields,

$$\log \mathbb{E}_t[R^1_{t+1}] - r_{f,t} = \lambda\xi\sqrt{1+\nu}\sigma^2 - \frac{1}{2}\nu\lambda^2\sigma^2 + \lambda(\mu - \tilde{\mu}_t). \tag{C.47}$$

For the infinite-horizon claim, again starting from (2.33), we get

$$r^\infty_{t+1} = \Delta c_{t+1} + \frac{\lambda-1}{\alpha}(\tilde{\mu}_{t+1} - \tilde{\mu}_t) - \tilde{\mu}_m - \frac{1}{2}\left[\sqrt{1+\nu}\left(1+\nu\frac{\lambda-1}{\alpha}\right) - \xi\right]^2 \sigma^2, \tag{C.48}$$

and, after subtracting the risk-free rate,

$$r^\infty_{t+1} - r_{f,t} = \Delta c_{t+1} + \frac{\lambda-1}{\alpha}(\tilde{\mu}_{t+1} - \tilde{\mu}_t) - \tilde{\mu}_t - \frac{1}{2}\left[\sqrt{1+\nu}\left(1+\nu\frac{\lambda-1}{\alpha}\right) - \xi\right]^2 \sigma^2 + \frac{1}{2}\xi^2\sigma^2. \tag{C.49}$$

The subjective conditional variance of r^1_{t+1} is $(1+\nu)\left(1+\nu\frac{\lambda-1}{\alpha}\right)^2\sigma^2$ and therefore, after taking subjective expectations of (C.49), we obtain

$$\log \tilde{\mathbb{E}}_t[R^\infty_{t+1}] - r_{f,t} = \left[1+\nu\frac{\lambda-1}{\alpha}\right]\xi\sqrt{1+\nu}\sigma^2. \tag{C.50}$$

The objective conditional variance of r^∞_{t+1} is only $\left(1+\nu\frac{\lambda-1}{\alpha}\right)^2\sigma^2$, and so taking objective expectations of (C.49) yields

$$\begin{aligned}\log \mathbb{E}_t[R^\infty_{t+1}] - r_{f,t} = &\left[1+\nu\frac{\lambda-1}{\alpha}\right]\xi\sqrt{1+\nu}\sigma^2 - \frac{1}{2}\nu\left(1+\nu\frac{\lambda-1}{\alpha}\right)^2\sigma^2 \\ &+ \left(1+\nu\frac{\lambda-1}{\alpha}\right)(\mu - \tilde{\mu}_t).\end{aligned} \tag{C.51}$$

C.5.4 Numerical solution

We use the analytical solutions for dividend strip prices to numerically compute the price, P_t, of the equity claim to the whole stream of dividends. For $n > J$ and some big enough J, equation (2.33) implies that

$$P_t^n \approx C_t e^{\mu_{dc}+\frac{1}{2}\frac{1}{1-(1-\alpha)^2}\sigma_d^2+\frac{\lambda-1}{\alpha}\tilde{\mu}_t} \exp(n\tilde{\mu}_m + \frac{1}{2}A_n\sigma^2), \quad n > J, \tag{C.52}$$

where we approximate

$$A_n \approx A_J + (n-J)[\sqrt{1+\nu}(\nu\frac{\lambda-1}{\alpha}+1)-\xi]^2, \quad n > J. \tag{C.53}$$

We can show that

$$P_t \approx \left(\sum_{n=1}^{J} P_t^n\right) + C_t V_J \exp\left(\mu_{dc} + \frac{1}{2}\frac{1}{1-(1-\alpha)^2}\sigma_d^2 + \frac{\lambda-1}{\alpha}\tilde{\mu}_t\right), \tag{C.54}$$

with

$$V_J = \frac{\exp\left((J+1)\tilde{\mu}_m + \frac{1}{2}A_J\sigma^2 + \frac{1}{2}[\sqrt{1+\nu}(\nu\frac{\lambda-1}{\alpha}+1)-\xi]^2\sigma^2\right)}{1-\exp\left(\tilde{\mu}_m + \frac{1}{2}[\sqrt{1+\nu}(\nu\frac{\lambda-1}{\alpha}+1)-\xi]^2\sigma^2\right)}. \tag{C.55}$$

We implement this by choosing a J big enough so that the value of P_t we obtain is not sensitive anymore to further changes in J. In our calibration, this requires $J \approx 7,000$.

We further use numerical methods to solve for the subjective equity premium in the $\psi = 1$ case. We follow the approach of Pohl et al. (2018). When $\psi = 1$, the wealth-consumption ratio is a constant

$$\log\frac{W_t - C_t}{C_t} = \log\frac{\delta}{1-\delta}, \tag{C.56}$$

and we only need to solve for the log price-dividend ratio. The log P/D ratio should be a function of both $\tilde{\mu}$ and $d_t - c_t$, i.e.

$$\log\frac{P_t}{D_t} = H(\tilde{\mu}_t, d_t - c_t). \tag{C.57}$$

In this case, because there are two state variables, the basis functions are now

$$\psi_{ij}(\tilde{\mu}, d_t - c_t) \equiv \Lambda_i(\tilde{\mu})\Lambda_j(d_t - c_t), \tag{C.58}$$

where Λ_i denotes the Chebyshev polynomials. We will approximate the log P/D ratio as

$$\hat{H}(\tilde{\mu}, d_t - c_t; \boldsymbol{\beta}_m) = \sum_{i=0}^{n_1-1}\sum_{j=0}^{n_2-1} \beta_{m,ij}\psi_{ij}(\tilde{\mu}, d_t - c_t). \tag{C.59}$$

Rewrite the subjective Euler equation

$$\tilde{\mathbb{E}}_t[M_{t+1}R_{m,t+1}] = 1 \tag{C.60}$$

as

$$\begin{aligned}
0 &= I(\tilde{\mu}_t, d_t - c_t) \\
&\equiv \tilde{\mathbb{E}}_t[e^{\tilde{\mu}_m - \tilde{\mu}_t - \xi\sigma\epsilon_{t+1} + \Delta d_{t+1}} \frac{e^{H(\tilde{\mu}_{t+1}, d_{t+1} - c_{t+1})} + 1}{e^{H(\tilde{\mu}_t, d_t - c_t)}}] - 1 \\
&= e^{\tilde{\mu}_m + (\lambda-1)\tilde{\mu}_t - \alpha(d_t - c_t - \mu_{dc})}\tilde{\mathbb{E}}_t[e^{(\lambda\sqrt{1+\nu} - \xi)\sigma\tilde{\varepsilon}_{t+1} + \sigma_d\eta_{t+1}} \frac{e^{H(\tilde{\mu}_{t+1}, d_{t+1} - c_{t+1})} + 1}{e^{H(\tilde{\mu}_t, d_t - c_t)}}] - 1. \quad \text{(C.61)}
\end{aligned}$$

We evaluate the function $I(\tilde{\mu}_t, d_t - c_t)$ on the two-dimensional grid of $\tilde{\mu}_t$ and $d_t - c_t$ and use the two-dimensional Gaussian quadrature approach to calculate the expectation part as an integral. Following Pohl et al. (2018), the numerical solution is implemented by the "fmincon" solver with the SQP algorithm in Matlab. We minimize a constant subject to the nonlinear constraints implied by Equation (C.61). We choose the degree of approximation, i.e., n_1 and n_2, such that the log P/D ratio computed using the projection method is closest to the analytically computed log P/D ratio as in Equation (C.54) in terms of the RMSE,

$$RMSE_{pd} = \sqrt{\frac{1}{t}\sum_{j=1}^{t}(pd_j^{Analytical} - pd_j^{Projection})^2}, \tag{C.62}$$

where $pd_t^{Analytical}$ is calculated from dividend strip prices as in (C.54). We explore different combinations of n_1 and n_2 up to a maximal degree of 8 and we choose the combination that minimizes $RMSE_{pd}$. Table C.1 summarizes the parameter choices for this numerical procedure.

After we obtain the coefficients for $H(\tilde{\mu}_t, d_t - c_t)$, we can calculate the subjective equity return as

$$\begin{aligned}
\tilde{\mathbb{E}}_t[R_{m,t+1}] &= \tilde{\mathbb{E}}_t[e^{\Delta d_{t+1}} \frac{e^{H(\tilde{\mu}_{t+1}, d_{t+1} - c_{t+1})} + 1}{e^{H(\tilde{\mu}_t, d_t - c_t)}}] \\
&= e^{\lambda\tilde{\mu}_t - \alpha(d_t - c_t - \mu_{dc})}\tilde{\mathbb{E}}_t[e^{\lambda\sqrt{1+\nu}\sigma\tilde{\varepsilon}_{t+1} + \sigma_d\eta_{t+1}} \frac{e^{H(\tilde{\mu}_{t+1}, d_{t+1} - c_{t+1})} + 1}{e^{H(\tilde{\mu}_t, d_t - c_t)}}]. \quad \text{(C.63)}
\end{aligned}$$

C.6 Model Solution for $\psi \neq 1$

C.6.1 Existence

Hansen and Scheinkman (2012) provide sufficient conditions for existence of equilibrium in a Markovian setting with Epstein-Zin preferences. As we show in Appendix

Table C.1: Parameter Values for Projection Method with $\psi = 1$

Parameter	Symbol	Value
Grid for $\tilde{\mu}_t$	$[\mu - 4\sigma(\tilde{\mu}), \mu + 4\sigma(\tilde{\mu})]$	[0.00 , 0.01]
Grid for $d_t - c_t$	$[E(dc) - 4\sigma(dc), E(dc) + 4\sigma(dc)]$	[-5.5 , -3.9]
# of points for Gaussian-Quad sampling	v	200
Degrees of approximation for $\tilde{\mu}_t$	n_1	3
Degrees of approximation for $d_t - c_t$	n_2	3
	Mean	0.008
$RMSE_{pd}$ across simulations	Median	0.008
	Max	0.038

C.4, our fading memory model can be mapped into an equivalent full memory model in which the information structure is a filtration and Markovian. This allows us to use the results in Hansen and Scheinkman (2012) to derive parameter restrictions sufficient to ensure existence of equilibrium.

In our equivalent "Kalman filtering" economy, we only have one state variable, which is $\{\hat{\mu}_{t+1|t}\}$ as in Equation (C.17). This equation also shows that Δc_{t+1} can be written as a function of $\hat{\mu}_{t+2|t+1}$ and $\hat{\mu}_{t+1|t}$. In addition, given the Markov property of $\{\hat{\mu}_{t+1|t}\}$, Assumption 1 in Hansen and Scheinkman (2012) is satisfied.

To distinguish the perceived time-t predictive distribution for Δc_{t+j} in this "Kalman filtering" economy here from the fading memory economy, we denote the subjective expectation here as $\tilde{\mathbb{E}}^*$. The Perron-Frobenius eigenvalue equation of interest is

$$\mathcal{T}v(x) = \exp(\eta)v(x), \quad v(\cdot) > 0, \tag{C.64}$$

where

$$\mathcal{T}f(x) = \tilde{\mathbb{E}}^*\left[f(\hat{\mu}_{t+2|t+1})\exp\left[(1-\gamma)\Delta c_{t+1}\right]|\hat{\mu}_{t+1|t} = x\right]. \tag{C.65}$$

Another random variable of interest is

$$N_{t+1} = \frac{e^{(1-\gamma)\Delta c_{t+1}}v(\hat{\mu}_{t+2|t+1})}{\exp(\eta)v(\hat{\mu}_{t+1|t})}. \tag{C.66}$$

Hansen and Scheinkman (2012) show that solutions exist for our model if the following additional assumptions are met:

Assumption 1.

$$\log\delta + \frac{\eta}{\theta} < 0. \tag{C.67}$$

Assumption 2.

$$\lim_{t\to\infty} \tilde{\mathbb{E}}^*[N_{t+1}v(\hat{\mu}_{t+2|t+1})^{-1/\theta}|\hat{\mu}_{1|0} = x] < \infty. \tag{C.68}$$

Assumption 3.

$$\lim_{t\to\infty} \tilde{\mathbb{E}}^*[N_{t+1}v(\hat{\mu}_{t+2|t+1})^{-1}|\hat{\mu}_{1|0} = x] < \infty. \tag{C.69}$$

To derive the explicit expressions of constraints in our model, we first calculate the Perron-Frobenius eigenvalue function v. We can show that (one of) the solution is

$$v(x) = \exp\left(\frac{1-\gamma}{\nu - \phi\nu} x\right), \tag{C.70}$$

$$\eta = (1-\gamma)\mu + \frac{1}{2}\left(\frac{1-\gamma}{1-\phi}\right)^2 (1+\phi\nu)\sigma^2. \tag{C.71}$$

With some algebra, both Assumption 2 and Assumption 3 can be reduced to the form

$$\lim_{t\to\infty} \tilde{\mathbb{E}}^*\Big[\exp\left(k_1 \Delta c_{t+1} + k_2 \hat{\mu}_{t+1|t}\right)\Big|\hat{\mu}_{1|0} = x\Big] < \infty, \tag{C.72}$$

for some corresponding pairs of constants (k_1, k_2).

With the Wold representation, we have

$$\hat{\mu}_{t+1|t} = [1-(h-K)L]^{-1} K\Delta c_t, \tag{C.73}$$

$$\Delta c_{t+1} = [1-(h-K)L]^{-1} K\Delta c_t + \tau_{t+1}, \tag{C.74}$$

$$\Delta c_t = \left[1 + (1-hL)^{-1}KL\right]\tau_t, \tag{C.75}$$

where $\{\tau_{t-j}\}$ are uncorrelated with variance $(1+\phi\nu)\sigma^2$. It suffices to show that

$$\lim_{t\to\infty} \tilde{\mathbb{E}}^*\Big[\exp\left(k\,[1-(h-K)L]^{-1} K \left[1 + (1-hL)^{-1}KL\right]\tau_t\right)\Big] < \infty, \tag{C.76}$$

or

$$\lim_{t\to\infty} \tilde{\mathbb{E}}^*\Big[\exp\left(kK(1-hL)^{-1}\tau_t\right)\Big] < \infty. \tag{C.77}$$

As long as $h < 1$, we have

$$\lim_{t\to\infty} \tilde{\mathbb{E}}^*\Big[\exp\left(kK(1-hL)^{-1}\tau_t\right)\Big] \tag{C.78}$$

$$= \lim_{t\to\infty} \tilde{\mathbb{E}}^*\Big[\exp\Big(kK\sum_{j=0}^{\infty} h^j L^j \tau_t\Big)\Big] \tag{C.79}$$

$$= \exp\left(\frac{1}{2}k^2K^2\sum_{j=0}^{\infty} h^{2j}(1+\phi\nu)\sigma^2\right) < \infty. \tag{C.80}$$

Finally, Assumption 1 translates to

$$\log\delta + \left(1 - \frac{1}{\psi}\right)\left[\mu + \frac{1}{2}\frac{1-\gamma}{(1-\phi)^2}(1+\phi\nu)\sigma^2\right] < 0, \tag{C.81}$$

and this is the only parameter constraint we apply to our model to ensure existence of equilibrium.

C.6.2 Log-linearized solution

We solve the model for $\psi \neq 1$ using log-linearization along similar lines as, e.g., in Beeler and Campbell (2012). We can write the Epstein-Zin log SDF as

$$m_{t+1} = \theta\log\delta - \frac{\theta}{\psi}\Delta c_{t+1} + (\theta - 1)r_{w,t+1}, \tag{C.82}$$

with $\theta = \frac{1-\gamma}{1-1/\psi}$.

We log-linearize the return on wealth and the return on the equity claim as

$$r_{w,t+1} = k_0 + k_1 z_{t+1} - z_t + \Delta c_{t+1}, \tag{C.83}$$
$$r_{m,t+1} = k_{d,0} + k_{d,1} z_{d,t+1} - z_{d,t} + \Delta d_{t+1}, \tag{C.84}$$

where $z_t \equiv \log((W_t - C_t)/C_t)$ and $z_{d,t} \equiv \log(P_t/D_t)$. We then conjecture z_t and $z_{d,t}$ are linear in the state variables

$$z_t = A_0 + A_1\phi\tilde{\mu}_t, \tag{C.85}$$
$$z_{d,t} = A_{d,0} + A_{d,1}\phi\tilde{\mu}_t + A_{d,2}(d_t - c_t). \tag{C.86}$$

By applying the subjective pricing equation to both $r_{w,t+1}$ and $r_{m,t+1}$, we can show that

$$A_0 = \frac{\log\delta + k_0 + (1-\phi)(1 - k_1 + \nu k_1)\mu A_1 + \frac{1}{2}(1 - 1/\psi + \phi\nu k_1 A_1)^2\theta(1+\phi\nu)\sigma^2}{1 - k_1}, \tag{C.87}$$

$$A_1 = \frac{1 - 1/\psi}{1 - (1 - \nu + \phi\nu)k_1}, \tag{C.88}$$

and

$$A_{d,0} = \frac{\theta \log \delta + (\theta-1)k_0 + (\theta-1)(k_1-1)A_0 + k_{d,0} + (1-\phi)(1-k_{d,1}+\nu k_{d,1})\mu A_{d,1}}{1-k_{d,1}} \tag{C.89}$$

$$+ \frac{(\theta-1)k_1\nu(1-\phi)\mu A_1 + (k_{d,1}A_{d,2}+1)\alpha\mu_{dc} + \frac{1}{2}(k_{d,1}A_{d,2}+1)^2\sigma_d^2}{1-k_{d,1}} \tag{C.90}$$

$$+ \frac{\frac{1}{2}[\theta - 1 - \theta/\psi + \lambda + (\lambda-1)k_{d,1}A_{d,2} + \phi\nu(k_{d,1}A_{d,1} + (\theta-1)k_1A_1)]^2(1+\phi\nu)\sigma^2}{1-k_{d,1}}, \tag{C.91}$$

$$A_{d,1} = \frac{\lambda - 1/\psi + (\lambda-1)k_{d,1}A_{d,2}}{1-(1-\nu+\phi\nu)k_{d,1}}, \tag{C.92}$$

$$A_{d,2} = \frac{\alpha}{(1-\alpha)k_{d,1}-1}. \tag{C.93}$$

Applying the subjective pricing equation to the risk-free payoff, we obtain

$$r_{f,t} = -\log\delta + \frac{1}{\psi}\left[\phi\tilde{\mu}_t + (1-\phi)\mu\right] \tag{C.94}$$

$$+ \frac{1}{2}(1-\gamma)\left(\frac{1}{\psi}-\gamma\right)\left(1+\frac{k_1\phi\nu}{1-(1-\nu+\phi\nu)k_1}\right)^2(1+\phi\nu)\sigma^2 \tag{C.95}$$

$$- \frac{1}{2}\left[\left(\frac{1}{\psi}-\gamma\right)\left(1+\frac{k_1\phi\nu}{1-(1-\nu+\phi\nu)k_1}\right)-\frac{1}{\psi}\right]^2(1+\phi\nu)\sigma^2. \tag{C.96}$$

We solve for the log-linearization coefficients by iterating on

$$\bar{z} = A_0 + A_1\phi\mu, \tag{C.97}$$

$$k_1 = \frac{\exp(\bar{z})}{1+\exp(\bar{z})}, \tag{C.98}$$

$$k_0 = \log(1+e^{\bar{z}}) - \bar{z}k_1, \tag{C.99}$$

and

$$\bar{z}_d = A_{d,0} + A_{d,1}\phi\mu + A_{d,2}\mathbb{E}[d_t - c_t], \tag{C.100}$$

$$k_{d,1} = \frac{\exp(\bar{z}_d)}{1+\exp(\bar{z}_d)}, \tag{C.101}$$

$$k_{d,0} = \log(1+e^{\bar{z}_d}) - \bar{z}_d k_{d,1}, \tag{C.102}$$

until we reach fixed points of k_0, k_1, $k_{d,0}$, and $k_{d,1}$, determined by a difference of less than 10^{-6}.

We then simulate the model and construct $\tilde{\mu}_t$. To calculate the objective risk premium, we directly use the return on equity claim as in (C.84) and the risk-free rate as in (C.96). To calculate subjective expected return, we take the subjective expectation of Equation (C.84) to yield

$$\tilde{\mathbb{E}}_t[r_{m,t+1}] = \frac{\phi}{\psi}\tilde{\mu}_t + \tilde{A}_{d,0}, \tag{C.103}$$

where

$$\begin{aligned}\tilde{A}_{d,0} &= (k_{d,1} - 1)A_{d,0} + k_{d,0} + (k_{d,1}A_{d,2} + 1)\alpha\mu_{dc} &\text{(C.104)}\\ &+ (1-\phi)\mu\left[k_{d,1}A_{d,1}\phi\nu + (\lambda - 1)k_{d,1}A_{d,2} + \lambda\right]. &\text{(C.105)}\end{aligned}$$

As a result, with log-linearization, the model implies a constant subjective premium.

C.7 Capital Income to Consumption Ratio in Ramsey-Cass-Koopmans Model

To illustrate plausible properties of the dividend-consumption ratio in a model in which investment and production is endogenous, this section presents a calculation of the capital income to consumption ratio in the Ramsey-Cass-Koopmans model with Cobb-Douglas technology based on Barro and i Martin (2004), Chapter 2. For ease of comparison, we use their notations here: IES $1/\theta$, interest rate r, capital and consumption in efficiency units $\hat{k}$ and $\hat{c}$, productivity growth rate x, population growth rate n, time discount rate ρ, capital share α and depreciation rate δ. Consistent with our baseline calibration, we set $1/\theta = 1$.

Solving their equations (2.24) and (2.25) for $\hat{c}$ and $\hat{k}$ with the left-hand side equal to zero in the steady state, and using the property of Cobb-Douglas technology that $f'(\hat{k})/\alpha = f(\hat{k})/\hat{k}$, we use this solution to calculate the ratio of capital income to consumption as

$$r\hat{k} = (x + \rho\hat{k})/\hat{c}. \tag{C.106}$$

Taking the derivative with respect to the productivity growth rate x, we obtain a positive derivative if

$$n - \rho < \left(\frac{1}{\alpha} - 1\right)\delta. \tag{C.107}$$

With an annual population growth rate of $n = 0.01$, and ρ slightly bigger than 1% as in our calibration, the left-hand side is negative and so (since $0 \leq \alpha \leq 1$ and $\delta > 0$) this inequality always holds. With typical values of $\delta = 0.05$ and $\alpha = 0.4$, the inequality holds unless population growth rates are implausibly high (more than 8% for $\rho = 0.01$) or the time discount rate implausibly low.

APPENDIX D

Profitability Anomaly: Data Construction

D.1 Construction of Characteristics

This section describes the construction of market-based and balance-sheet-based characteristics of stocks. For balance-sheet-based characteristics, I assume the information in fiscal year t-1 is known in June of year t.

Market-based

1. Following Ang et al. (2006), idiosyncratic volatility (Ivol) is estimated using daily log excess return from Fama-French three-factor model using one-month window (22 trading days) and requiring at least 17 trading days of non-missing data.
2. Bid-ask spread is the monthly average of daily $2 \times |Bid - Ask|/|Bid + Ask|$.
3. Dollar volume is the trailing 6-month average of share trading volume times prices.

Balance-sheet–based

Number of institutional investors and percentage of institutional ownership are obtained from Thompson Reuters 13F Holdings database which covers from January 1980 to December 2016.

The construction of following characteristics follows Freyberger et al. (forthcoming) and Green et al. (2017):

1. Book equity (BE) is shareholder equity (SH) plus deferred taxes (DT), minus preferred stocks (PS). Data from Davis et al. (2000) is supplemented, which is available from Kenneth French's website.[1]

[1] `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`

2. Market equity (ME) is common shares outstanding (SHROUT) times price (PRC) from CRSP.

3. Book-to-market ratio is book equity (BE) at the end of fiscal year $t-1$ dividend by market equity (ME) at the end of December in year $t-1$ to avoid a short momentum position.

4. Shareholder equity (SH) is as given in Compustat (SEQ) if available, or else common equity plus the carrying value of preferred stock (CEQ + PS) if available, or else total assets minus total liabilities (AT - LT).

5. Deferred taxes (DT) is deferred taxes and investment tax credits (TXDITC) if available, or else deferred taxes and/or investment tax credit (TXDB and/or ITCB).

6. Preferred stock is redemption value (PSTKRV) if available, or else liquidating value (PSTKL) if available, or else carrying value (PSTK).

7. Fiscal year-end market equity (MVE_F) is common shares outstanding (CSHO) times fiscal year-end close price (PRCC_F).

8. Cash dividend yield is defined as dividends (DVT) dividend by Fiscal year-end market equity (MVE_F).

9. Repurchase (REP) is positive change in Treasury stock (TSTKC) when the firm is not using retirement method, or positive difference between purchase of common and preferred stocks (PRSTKC) and sale of common and preferred stocks (SSTK), following Fama and French (2001).

10. Leverage (LEV) is the ratio of long term debt and debt in current liabilities (DLTT and DLC) to shareholders' equity (SH), long term debt and debt in current liabilities.

APPENDIX E

Profitability Anomaly: Additional Results

E.1 Additional Tables and Figures

Table E.1: Replicating the Profitability Anomalies

This table reports the excess returns, alphas, and characteristics of profitability anomaly portfolios. Characteristics include the average profitability measure (Prof.), portfolio size (ME, in billions of dollars), average firm size (Size, in millions of dollars), and average number of firms. Portfolios are quintile sorted on profitability proxies using NYSE breakpoints and portfolio returns are value-weighted. Excess returns and alphas are reported in monthly percentage points. Newey-West standard errors with a lag of 12 months are reported in brackets. The sample is from July 1963 to December 2016.

	FF3F alphas and factor loadings					Portfolio Characteristics			
Portfolios	r^e	α	MKT	SMB	HML	Prof.	ME	Size	n
Panel A: Gross Profitability									
Low	0.42	-0.15	0.92	-0.00	0.26	0.11	725	1143	582
		[-2.23]	[57.96]	[-0.15]	[6.51]				
2	0.50	-0.08	1.01	-0.07	0.23	0.22	831	1775	496
		[-1.09]	[35.92]	[-1.93]	[5.32]				
3	0.60	0.05	0.98	0.04	0.10	0.34	943	1809	538
		[0.87]	[47.18]	[1.31]	[2.17]				
4	0.50	0.06	0.99	-0.01	-0.17	0.49	1263	1981	620
		[1.09]	[49.39]	[-0.49]	[-4.44]				
High	0.62	0.28	0.94	-0.04	-0.36	0.81	1434	1740	773
		[4.11]	[42.43]	[-1.37]	[-9.67]				
High - Low	0.20	0.43	0.02	-0.04	-0.62				
		[3.98]	[0.66]	[-0.70]	[-10.14]				
Panel B: Operating Profitability									
Low	0.40	-0.28	1.08	0.19	0.21	0.06	514	512	936
		[-3.19]	[44.26]	[5.57]	[4.59]				
2	0.53	-0.06	0.93	0.03	0.30	0.12	568	1258	457
		[-1.09]	[56.63]	[0.68]	[5.51]				
3	0.58	0.04	0.95	-0.06	0.19	0.16	820	1807	469
		[0.76]	[44.72]	[-1.68]	[3.92]				
4	0.60	0.09	0.98	-0.05	0.07	0.21	1109	2261	489
		[1.89]	[65.56]	[-2.16]	[1.76]				
High	0.53	0.18	0.97	-0.06	-0.35	0.37	2185	3137	657
		[3.98]	[70.02]	[-3.57]	[-20.27]				
High - Low	0.14	0.46	-0.11	-0.26	-0.56				
		[4.43]	[-3.59]	[-6.24]	[-11.93]				
Panel C: Cash-based Operating Profitability									
Low	0.29	-0.37	1.12	0.27	0.06	0.02	513	483	963
		[-4.73]	[39.12]	[5.44]	[1.47]				
2	0.41	-0.17	0.97	0.06	0.19	0.10	622	1275	477
		[-3.22]	[54.46]	[2.20]	[4.36]				
3	0.52	-0.00	0.93	-0.06	0.19	0.14	758	1713	462
		[-0.08]	[51.13]	[-1.82]	[4.44]				
4	0.55	0.08	0.97	-0.08	0.09	0.18	1150	2313	480
		[1.59]	[70.65]	[-3.53]	[3.26]				
High	0.63	0.23	0.96	-0.08	-0.32	0.31	2152	3233	617
		[5.38]	[73.87]	[-4.41]	[-19.33]				
High - Low	0.33	0.64	-0.18	-0.36	-0.34				
	[3.22]	[6.24]	[-4.87]	[-6.19]	[-7.54]				

Table E.2: Examining Anomalies in Samples with Non-missing IF Proxies

This table examines the performance of profitability anomalies in samples for which IF proxies (1YFD and OIV) are non-missing. I report the FF3F alphas, FF5F alphas, and q-4 alphas. I also report characteristics of the available sample, including number of available firm-month observations, average firm size, average total market capitalization. These characteristics are all reported as ratios of those of full sample for comparison. The definitions of these variables follow Novy-Marx (2013) and Ball et al. (2016), with a modification to GP/A which is deflated by lagged book value of total assets. Portfolios are quintile sorted on profitability proxies using NYSE breakpoints and portfolio returns are value-weighted. Alphas are reported in monthly percentage points. Newey-West standard errors with a lag of 12 months are reported in brackets.

	GP/A			OP			C-OP				Sample Characteristics			
Proxies	FF3F	FF5F	q-4	FF3F	FF5F	q-4	FF3F	FF5F	q-4	Year	n	Size	MKTME	T
Panel A: Full Sample														
Low	-0.15	-0.10	-0.03	-0.28	-0.18	-0.13	-0.36	-0.26	-0.20					
	[-2.23]	[-1.76]	[-0.37]	[-3.19]	[-2.50]	[-1.31]	[-4.44]	[-3.62]	[-2.44]					
High	0.28	0.28	0.21	0.18	0.19	0.18	0.23	0.23	0.23					
	[4.11]	[3.91]	[2.32]	[3.98]	[3.89]	[2.64]	[5.38]	[4.98]	[3.40]					
High - Low	0.43	0.38	0.24	0.46	0.37	0.31	0.59	0.49	0.43					
	[3.98]	[3.67]	[1.68]	[4.43]	[3.99]	[2.47]	[5.96]	[5.11]	[3.54]					
Panel B: Fiscal Year 1 EPS Forecast Dispersion														
Low	-0.11	-0.03	-0.02	-0.27	-0.17	-0.13	-0.35	-0.24	-0.19	1976	0.16	4.50	0.76	197602 –201612
	[-1.36]	[-0.51]	[-0.18]	[-2.78]	[-2.05]	[-1.27]	[-4.03]	[-3.48]	[-2.24]	1986	0.44	2.15	0.96	
High	0.22	0.21	0.21	0.14	0.15	0.17	0.20	0.19	0.22	1996	0.55	1.77	0.97	
	[2.83]	[2.42]	[2.09]	[2.28]	[2.28]	[2.00]	[3.52]	[3.31]	[2.81]	2006	0.67	1.46	0.98	
High - Low	0.33	0.25	0.22	0.41	0.31	0.29	0.55	0.43	0.40	2016	0.76	1.30	0.99	
	[2.86]	[2.13]	[1.49]	[3.68]	[2.99]	[2.20]	[5.24]	[4.51]	[3.14]					
Panel C: At-the-money Option Implied Volatility														
Low	-0.11	-0.13	-0.10	-0.24	-0.19	-0.18	-0.32	-0.21	-0.22	1976	/	/	/	199604–201612
	[-0.95]	[-1.15]	[-0.82]	[-1.84]	[-1.65]	[-1.41]	[-2.60]	[-2.14]	[-1.85]	1986	/	/	/	
High	0.22	0.23	0.26	0.17	0.21	0.29	0.20	0.22	0.31	1996	0.28	3.21	0.89	
	[2.47]	[2.09]	[2.10]	[2.01]	[2.45]	[2.56]	[2.76]	[2.87]	[3.05]	2006	0.50	1.89	0.95	
High - Low	0.33	0.36	0.37	0.40	0.40	0.48	0.53	0.43	0.54	2016	0.70	1.41	0.99	
	[2.09]	[2.04]	[1.83]	[2.77]	[2.74]	[2.79]	[3.65]	[3.20]	[2.99]					

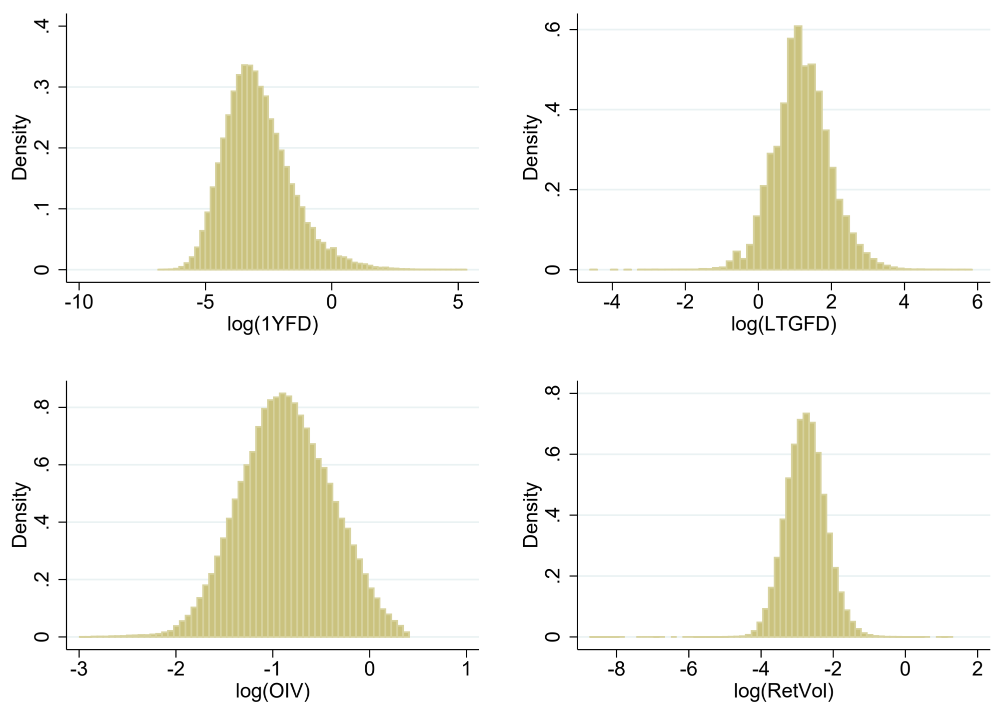

Figure E.1: Principle Component Analysis of IF Proxies

This figure plots the histograms of log-transformed fiscal year 1 EPS forecast dispersion (1YFD), long-term growth EPS forecast dispersion (LTGFD), at-the-money option-implied volatility (OIV), and lagged weekly return volatility (RetVol).

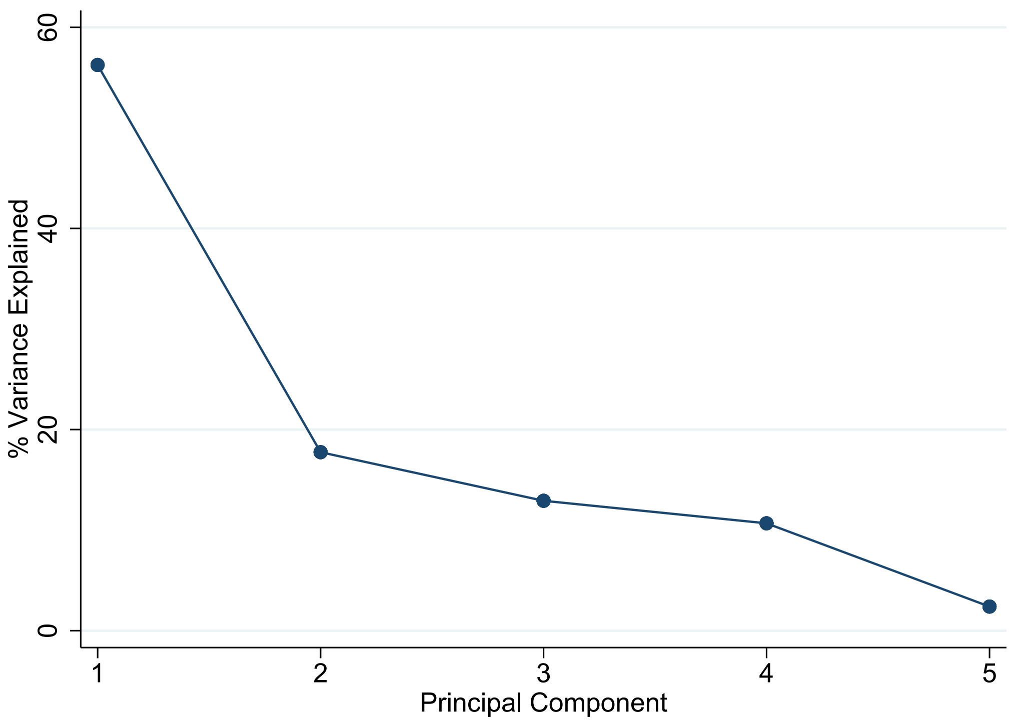

Figure E.2: Principle Component Analysis of Uncertainty Proxies

This figure plots the percentage of variance explained by each principle component of log-transformed fiscal year 1 EPS forecast dispersion, long-term growth EPS forecast dispersion, at-the-money option-implied volatility, lagged weekly return volatility, and age.

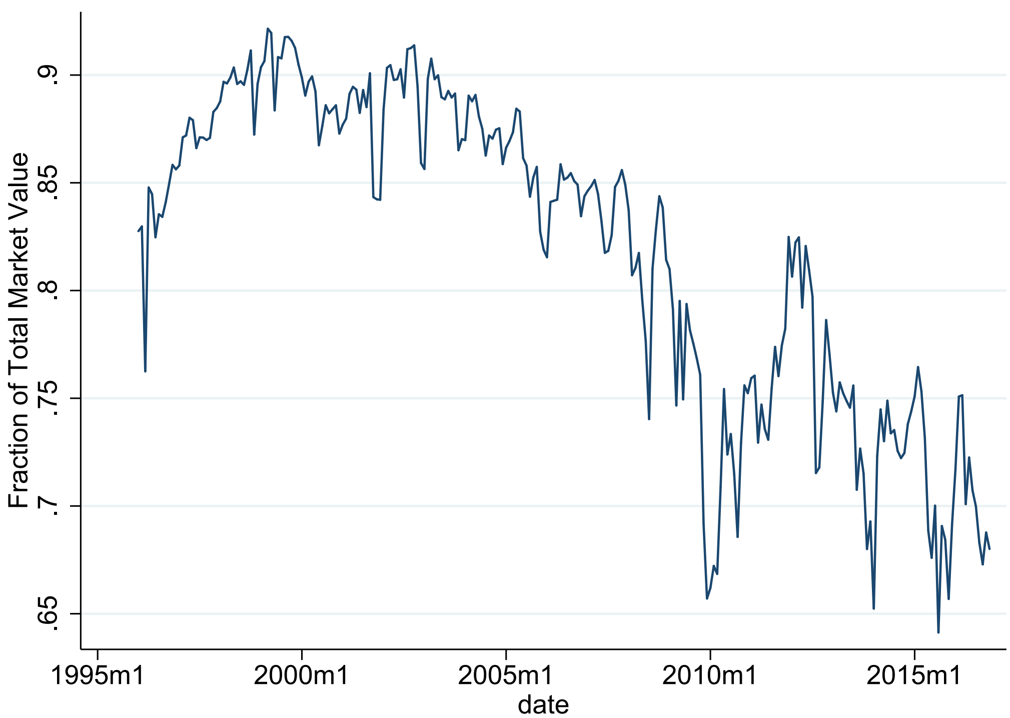

Figure E.3: Market Capitalization of Stocks with Non-missing PIF

This figure plots the total market capitalization of stocks with non-missing PIF measure as a fraction of total market capitalization. The sample starts from February 1996 to December 2016.

BIBLIOGRAPHY

BIBLIOGRAPHY

Adam, Klaus, Albert Marcet, and Johannes Beutel, 2017, Stock price booms and expected capital gains, *American Economic Review* 107, 2352–2408.

Admati, Anat R., 1985, A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets, *Econometrica* 53, 629–657.

Ameriks, John, Gábor Kézdi, Minjoon Lee, and Matthew D Shapiro, 2016, Heterogeneity in Expectations, Risk Tolerance, and Household Stock Shares, *University of Michigan Working Paper* .

Amromin, Gene, and Steven A. Sharpe, 2014, From the Horse's Mouth: Economic Conditions and Investor Expectations of Risk and Return, *Management Science* 60, 845–866.

Andonov, Aleksandar, and Joshua D. Rauh, 2019, The Return Expectations of Institutional Investors, *Stanford University Graduate School of Business Research Paper No. 18-5* .

Ang, Andrew, Geert Bekaert, and Min Wei, 2007, Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?, *Journal of Monetary Economics* 54, 1163–1212.

Ang, Andrew, Geert Bekaert, and Min Wei, 2008, The Term Structure of Real Rates and Expected Inflation, *The Journal of Finance* 63, 797–849.

Ang, Andrew, Robert j. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The Cross-Section of Volatility and Expected Returns, *The Journal of Finance* 61, 259–299.

Avramov, Doron, Tarun Chordia, Gergana Jostova, and Alexander Philipov, 2009, Credit ratings and the cross-section of stock returns, *Journal of Financial Markets* 12, 469 – 499.

Avramov, Doron, Tarun Chordia, Gergana Jostova, and Alexander Philipov, 2013, Anomalies and financial distress, *Journal of Financial Economics* 108, 139 – 159.

Azeredo da Silveira, Rava, and Michael Woodford, 2019, Noisy Memory and Over-Reaction to News, *AEA Papers and Proceedings* 109, 557–61.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, 2016, Measuring Economic Policy Uncertainty, *The Quarterly Journal of Economics* 131, 1593–1636.

Balke, Nathan S., and Robert J. Gordon, 1989, The Estimation of Prewar Gross National Product: Methodology and New Evidence, *Journal of Political Economy* 97, 38–92.

Ball, Ray, Joseph Gerakos, Juhani T. Linnainmaa, and Valeri Nikolaev, 2016, Accruals, cash flows, and operating profitability in the cross section of stock returns, *Journal of Financial Economics* 121, 28 – 45.

Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad, 2005, Consumption, Dividends, and the Cross Section of Equity Returns, *The Journal of Finance* 60, 1639–1672.

Bansal, Ravi, A. Ronald Gallant, and George Tauchen, 2007, Rational Pessimism, Rational Exuberance, and Asset Pricing Models, *The Review of Economic Studies* 74, 1005–1033.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, An Empirical Evaluation of the Long-Run Risks Model for Asset Prices, *Critical Finance Review* 1, 183–221.

Bansal, Ravi, and Ivan Shaliastovich, 2013, A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets, *The Review of Financial Studies* 26, 1–33.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, *The Journal of Finance* 59, 1481–1509.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2015, X-capm: An extrapolative capital asset pricing model, *Journal of Financial Economics* 115, 1 – 24.

Barberis, Nicholas, Ming Huang, and Tano Santos, 2001, Prospect Theory and Asset Prices, *The Quarterly Journal of Economics* 116, 1–53.

Barro, Robert J., and Xavier Sala i Martin, 2004, *Economic Growth*, second edition (MIT Press, Cambridge, MA).

Barro, Robert J., and José F. Ursúa, 2008, Macroeconomic Crises since 1870, *Brookings Papers on Economic Activity* 255–335.

Bauer, Michael D., and Glenn D. Rudebusch, 2019, Interest Rates Under Falling Stars, *Federal Reserve Bank of San Francisco Working Paper 2017-16* .

Beeler, Jason, and John Y. Campbell, 2012, The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment, *Critical Finance Review* 1, 141–182.

Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 1997, On Biases in Tests of the Expectations Hypothesis of the Term Structure of Interest Rates, *Journal of Financial Economics* 44, 309–348.

Berardi, Michele, and Jaqueson K. Galimberti, 2017, Empirical calibration of adaptive learning, *Journal of Economic Behavior & Organization* 144, 219 – 237.

Bernanke, Ben S., 2004, The Great Moderation, *Remarks at the meetings of the Eastern Economic Association* .

Bolt, Jutta, Robert Inklaar, Herman de Jong, and Jan Luiten van Zanden, 2018, Rebasing "Maddison": New Income Comparisons and the Shape of Long-run Economic Development, *Maddison Project Working paper 10* .

Bondt, Werner F. M. De, and Richard H. Thaler, 1990, Do Security Analysts Overreact?, *The American Economic Review* 80, 52–57.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2019a, Diagnostic expectations and stock returns, *The Journal of Finance* 74, 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2019b, Memory, Attention, and Choice, *Working Paper* .

Bouchaud, Jean-philippe, Philipp Krüger, Augustin Landier, and David Thesmar, 2019, Sticky Expectations and the Profitability Anomaly, *The Journal of Finance* 74, 639–674.

Branch, William A., and George W. Evans, 2006, A Simple Recursive Forecasting Model, *Economics Letters* 91, 158–166.

Buraschi, Andrea, Ilaria Piatti, and Paul Whelan, 2019, Subjective Bond Risk Premia and Belief Aggregation, *Saïd Business School Working Paper 2019-06* .

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

Campbell, John Y., 2000, Asset Pricing at the Millennium, *The Journal of Finance* 55, 1515–1567.

Campbell, John Y., and John Ammer, 1993, What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns, *The Journal of Finance* 48, 3–37.

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.

Campbell, John Y., and Robert J. Shiller, 1988, The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors, *The Review of Financial Studies* 1, 195–228.

Campbell, John Y., and Robert J. Shiller, 1991, Yield Spreads and Interest Rate Movements: A Bird's Eye View, *The Review of Economic Studies* 58, 495–514.

Chernov, Mikhail, and Philippe Mueller, 2012, The term structure of inflation expectations, *Journal of Financial Economics* 106, 367 – 394.

Choi, Hoyong, Philippe Mueller, and Andrea Vedolin, 2017, Bond Variance Risk Premiums, *Review of Finance* 21, 987–1022.

Choi, James J, and Thomas Mertens, 2019, Extrapolative Expectations and the Equity Premium, *Yale University Working Paper* .

Choi, Jung Ho, Alon Kalay, and Gil Sadka, 2016, Earnings news, expected earnings, and aggregate stock returns, *Journal of Financial Markets* 29, 110 – 143.

Christoffersen, Peter, Mathieu Fournier, and Kris Jacobs, 2017, The Factor Structure in Equity Options, *The Review of Financial Studies* 31, 595–637.

Chun, Albert Lee, 2010, Expectations, Bond Yields, and Monetary Policy, *The Review of Financial Studies* 24, 208–247.

Cieslak, Anna, 2018, Short-Rate Expectations and Unexpected Returns in Treasury Bonds, *The Review of Financial Studies* 31, 3265–3306.

Cieslak, Anna, and Pavol Povala, 2015, Expected Returns in Treasury Bonds, *The Review of Financial Studies* 28, 2859–2901.

Clark, Peter K., 1987, The Cyclical Component of U. S. Economic Activity, *The Quarterly Journal of Economics* 102, 797–814.

Clements, Michael P., David F. Hendry, Siem Jan Koopman, and Marius Ooms, 2012, Forecasting Economic Time Series Using Unobserved Components Time Series Models.

Cochrane, John H., 2015, Comments on "Robust Bond Risk Premia" by Michael Bauer and Jim Hamilton, Comments, Federal Reserve Bank of San Francisco.

Cochrane, John H., and Monika Piazzesi, 2005, Bond Risk Premia, *American Economic Review* 95, 138–160.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer, 2016a, Asset Pricing When 'This Time Is Different', *The Review of Financial Studies* 30, 505–535.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer, 2016b, Parameter Learning in General Equilibrium: The Asset Pricing Implications, *American Economic Review* 106, 664–98.

Croushore, Dean, and Tom Stark, 2001, A Real-time Data Set for Macroeconomists, *Journal of Econometrics* 105, 111–130, Forecasting and empirical methods in finance and macroeconomics.

Crump, Richard K., and Nikolay Gospodinov, 2019, Deconstructing the Yield Curve, Staff Reports 884, Federal Reserve Bank of New York.

Davis, James L., Eugene F. Fama, and Kenneth R. French, 2000, Characteristics, Covariances, and Average Returns: 1929 to 1997, *The Journal of Finance* 55, 389–406.

Deng, Yao, 2019, Product Market Competition and the Profitability Premium, *Working Paper* .

Diebold, Francis X., and Canlin Li, 2006, Forecasting the Term Structure of Government Bond Yields, *Journal of Econometrics* 130, 337–364.

Diether, Karl B., Christopher J. Malloy, and Anna Scherbina, 2002, Differences of Opinion and the Cross Section of Stock Returns, *The Journal of Finance* 57, 2113–2141.

Duffee, Gregory R., 2002, Term Premia and Interest Rate Forecasts in Affine Models, *The Journal of Finance* 57, 405–443.

Duffee, Gregory R., 2018, Expected Inflation and Other Determinants of Treasury Yields, *The Journal of Finance* 73, 2139–2180.

D'Amico, Stefania, Don H. Kim, and Min Wei, 2018, Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices, *Journal of Financial and Quantitative Analysis* 53, 395–436.

Ehling, Paul, Alessandro Graniero, and Christian Heyerdahl-Larsen, 2017, Asset Prices and Portfolio Choice with Learning from Experience, *The Review of Economic Studies* 85, 1752–1780.

Epstein, Larry G., and Stanley E. Zin, 1989, Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework, *Econometrica* 57, 937–969.

Evans, George W., and Seppo Honkapohja, 2001, *Learning and Expectations in Macroeconomics* (Princeton University Press, Princeton, NJ).

Fama, Eugene F., and Robert R. Bliss, 1987, The Information in Long-Maturity Forward Rates, *The American Economic Review* 77, 680–692.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3 – 56.

Fama, Eugene F, and Kenneth R French, 2001, Disappearing dividends: Changing firm characteristics or lower propensity to pay?, *Journal of Financial Economics* 60, 3 – 43.

Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1 – 22.

Freyberger, Joachim, Andreas Neuhierl, and Michael Weber, forthcoming, Dissecting Characteristics Nonparametrically, *The Review of Financial Studies* .

Friedman, Benjamin M., 1980, Survey Evidence on the 'Rationality' of Interest Rate Expectations, *Journal of Monetary Economics* 6, 453–465.

Froot, Kenneth A., 1989, New Hope for the Expectations Hypothesis of the Term Structure of Interest Rates, *The Journal of Finance* 44, 283–305.

Gabaix, Xavier, 2012, Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance, *The Quarterly Journal of Economics* 127, 645–700.

Giacoletti, Marco, Kristoffer Laursen, and Kenneth J. Singleton, 2018, Learning and Risk Premiums in an Arbitrage-free Term Structure Model, *Stanford Graduate School of Business Working Paper No. 3670* .

Grant, Angelia L., and Joshua C.C. Chan, 2017, A Bayesian Model Comparison for Trend-Cycle Decompositions of Output, *Journal of Money, Credit and Banking* 49, 525–552.

Green, Jeremiah, John R. M. Hand, and X. Frank Zhang, 2017, The Characteristics that Provide Independent Information about Average U.S. Monthly Stock Returns, *The Review of Financial Studies* 30, 4389–4436.

Greenwood, Robin, and Stefan Nagel, 2009, Inexperienced investors and bubbles, *Journal of Financial Economics* 93, 239 – 258.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of Returns and Expected Returns, *The Review of Financial Studies* 27, 714–746.

Greenwood, Robin, and Dimitri Vayanos, 2014, Bond Supply and Excess Bond Returns, *The Review of Financial Studies* 27, 663–713.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, The U.S. Treasury Yield Curve: 1961 to the Present, *Journal of Monetary Economics* 54, 2291–2304.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2010, The TIPS Yield Curve and Inflation Compensation, *American Economic Journal: Macroeconomics* 2, 70–92.

Haltiwanger, John, Ron S. Jarmin, and Javier Miranda, 2013, Who Creates Jobs? Small versus Large versus Young, *The Review of Economics and Statistics* 95, 347–361.

Hamilton, James Douglas, 1994, *Time Series Analysis* (Princeton Univ. Press, Princeton, NJ).

Hansen, Lars Peter, 2007, Beliefs, Doubts and Learning: Valuing Macroeconomic Risk, *American Economic Review* 97, 1–30.

Hansen, Lars Peter, John C. Heaton, and Nan Li, 2008, Consumption Strikes Back? Measuring Long-Run Risk, *Journal of Political Economy* 116, 260–302.

Hansen, Lars Peter, and José A. Scheinkman, 2012, Recursive Utility in a Markov Environment with Stochastic Growth, *Proceedings of the National Academy of Sciences* 109, 11967–11972.

Harvey, Andrew, 1985, Trends and Cycles in Macroeconomic Time Series, *Journal of Business & Economic Statistics* 3, 216–27.

Hasler, Michael, Mariana Khapko, and Roberto Marfè, 2019, Rational Learning and Term Structures, *Working Paper* .

Hirshleifer, David, Jun Li, and Jianfeng Yu, 2015, Asset pricing in production economies with extrapolative expectations, *Journal of Monetary Economics* 76, 87 – 106.

Hou, Kewei, Chen Xue, and Lu Zhang, 2014, Digesting Anomalies: An Investment Approach, *The Review of Financial Studies* 28, 650–705.

Hou, Kewei, Chen Xue, and Lu Zhang, 2018, Replicating Anomalies, *The Review of Financial Studies* hhy131.

Jin, Lawrence J, and Pengfei Sui, 2019, Asset Pricing with Return Extrapolation, *Working Paper* .

Jørgensen, Kasper, 2018, How Learning from Macroeconomic Experiences Shapes the Yield Curve, *Working Paper* .

Kahana, Michael J, and Jessica A Wachter, 2019, A Retrieved-Context Theory of Financial Decisions, *University of Pennsylvania Working Paper* .

Kim, Don H., and Athanasios Orphanides, 2012, Term Structure Estimation with Survey Data on Interest Rate Forecasts, *Journal of Financial and Quantitative Analysis* 47, 241—-272.

Kogan, Leonid, Jun Li, and Harold H. Zhang, 2019, A Production-based Economic Explanation for the Gross Profitability Premium, *Working Paper* .

Kogan, Leonid, and Dimitris Papanikolaou, 2013, Firm Characteristics and Stock Returns: The Role of Investment-Specific Shocks, *The Review of Financial Studies* 26, 2718–2759.

Kothari, S.P., Jonathan Lewellen, and Jerold B. Warner, 2006, Stock returns, aggregate earnings surprises, and behavioral finance, *Journal of Financial Economics* 79, 537 – 568.

Kothari, S.P., and Jay Shanken, 1997, Book-to-market, dividend yield, and expected market returns: A time-series analysis, *Journal of Financial Economics* 44, 169 – 203.

Kozicki, Sharon, and P.A. Tinsley, 2001, Shifting Endpoints in the Term Structure of Interest Rates, *Journal of Monetary Economics* 47, 613–652.

La Porta, Rafael, 1996, Expectations and the Cross-Section of Stock Returns, *The Journal of Finance* 51, 1715–1742.

Lam, F.Y. Eric C., Shujing Wang, and K.C. John Wei, 2016, The Profitability Premium: Macroeconomic Risks or Expectation Errors?, *Working Paper* .

Lease, Ronald C., Wilbur G. Lewellen, and Gary G. Schlarbaum, 1974, The individual investor: Attributes and attitudes, *The Journal of Finance* 29, 413–433.

Lettau, Martin, and Sydney C. Ludvigson, 2010, Chapter 11 - measuring and modeling variation in the risk-return trade-off, in YACINE AÏT-SAHALIA, and LARS PETER HANSEN, eds., *Handbook of Financial Econometrics: Tools and Techniques*, volume 1 of *Handbooks in Finance*, 617 – 690 (North-Holland, San Diego).

Lewellen, Jonathan, and Jay Shanken, 2002, Learning, Asset-Pricing Tests, and Market Efficiency, *The Journal of Finance* 57, 1113–1145.

Lewellen, Wilbur G., Ronald C. Lease, and Gary G. Schlarbaum, 1977, Patterns of Investment Strategy and Behavior Among Individual Investors, *The Journal of Business* 50, 296–333.

Lin, Xiaoji, and Lu Zhang, 2013, The investment manifesto, *Journal of Monetary Economics* 60, 351 – 366.

Liu, Yan, and Jing Cynthia Wu, 2019, Reconstructing the Yield Curve, *Working Paper* .

Ludvigson, Sydney C., and Serena Ng, 2009, Macro Factors in Bond Risk Premia, *The Review of Financial Studies* 22, 5027–5067.

MacKinnon, James G., 2009, *Bootstrap Hypothesis Testing*, chapter 6, 183–213 (John Wiley & Sons, Ltd).

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?, *The Quarterly Journal of Economics* 126, 373–416.

Malmendier, Ulrike, and Stefan Nagel, 2015, Learning from Inflation Experiences, *The Quarterly Journal of Economics* 131, 53–87.

Malmendier, Ulrike, Stefan Nagel, and Yan Zhen, 2019a, The Making of Hawks and Doves, *University of Chicago Working Paper* .

Malmendier, Ulrike, Demian Pouzo, and Victoria Vanasco, 2019b, Investor experiences and financial market dynamics, *Journal of Financial Economics* .

Mangel, Marc, 1990, Dynamic information in uncertain and changing worlds, *Journal of Theoretical Biology* 146, 317 – 332.

McCulloch, J. Huston, and Heon-Chul Kwon, 1993, U.S. Term Structure Data, 1947-1991.

Nagel, Stefan, 2005, Short sales, institutional investors and the cross-section of stock returns, *Journal of Financial Economics* 78, 277 – 309.

Nagel, Stefan, 2016, The Liquidity Premium of Near-Money Assets, *The Quarterly Journal of Economics* 131, 1927–1971.

Nagel, Stefan, and Zhengyang Xu, 2019, Asset Pricing with Fading Memory, Working Paper 26255, National Bureau of Economic Research.

Nakov, Anton, and Galo Nuño, 2015, Learning from experience in the stock market, *Journal of Economic Dynamics and Control* 52, 224 – 239.

Novy-Marx, Robert, 2013, The other side of value: The gross profitability premium, *Journal of Financial Economics* 108, 1 – 28.

Novy-Marx, Robert, and Mihail Velikov, 2015, A Taxonomy of Anomalies and Their Trading Costs, *The Review of Financial Studies* 29, 104–147.

Orphanides, Athanasios, and Min Wei, 2012, Evolving Macroeconomic Perceptions and the Term Structure of Interest Rates, *Journal of Economic Dynamics and Control* 36, 239–254.

Orphanides, Athanasios, and John C. Williams, 2005, The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations, *Journal of Economic Dynamics and Control* 29, 1927–1950, Expectations, learning, and monetary policy.

Pastor, Lubos, and Pietro Veronesi, 2003a, Stock Valuation and Learning about Profitability, *The Journal of Finance* 58, 1749–1789.

Pastor, Lubos, and Pietro Veronesi, 2003b, Stock Valuation and Learning about Profitability, *The Journal of Finance* 58, 1749–1789.

Pastor, Lubos, and Pietro Veronesi, 2006, Was there a nasdaq bubble in the late 1990s?, *Journal of Financial Economics* 81, 61 – 100.

Patton, Andrew, and Allan Timmermann, 2010, Why Do Forecasters Disagree? Lessons from the Term Structure of Cross-sectional Dispersion, *Journal of Monetary Economics* 57, 803–820.

Piazzesi, Monika, Juliana Salomao, and Martin Schneider, 2015, Trend and Cycle in Bond Premia, *Working Paper* .

Piazzesi, Monika, and Martin Schneider, 2006, Equilibrium Yield Curves, *NBER Macroeconomics Annual* 21, 389–472.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman, 2017, Distributional National Accounts: Methods and Estimates for the United States, *The Quarterly Journal of Economics* 133, 553–609.

Pohl, Walter, Karl Schmedders, and Ole Wilms, 2018, Higher Order Effects in Asset Pricing Models with Long-Run Risks, *The Journal of Finance* 73, 1061–1111.

Politis, Dimitris N., and Halbert White, 2004, Automatic Block-Length Selection for the Dependent Bootstrap, *Econometric Reviews* 23, 53–70.

Pontiff, Jeffrey, 2006, Costly arbitrage and the myth of idiosyncratic risk, *Journal of Accounting and Economics* 42, 35 – 52, Conference Issue on Implications of Changing Financial Reporting Standards.

Reuters, Thomson, 2009, Methodology for estimates.

Rubinstein, Mark, 1974, An Aggregation Theorem for Securities Markets, *Journal of Financial Economics* 1, 225–244.

Sadka, Gil, and Ronnie Sadka, 2009, Predictability and the earnings–returns relation, *Journal of Financial Economics* 94, 87 – 106.

Schraeder, Stefanie, 2015, Information Processing and Non-Bayesian Learning in Financial Markets, *Review of Finance* 20, 823–853.

Sharpe, Steven A., 2002, Reexamining Stock Valuation and Inflation: The Implications Of Analysts' Earnings Forecasts, *The Review of Economics and Statistics* 84, 632–648.

Shiller, Robert J., 2005, *Irrational Exuberance* (Princeton University Press).

Shumway, Tyler, 1997, The Delisting Bias in CRSP Data, *The Journal of Finance* 52, 327–340.

Shumway, Tyler, and Vincent A. Warther, 1999, The Delisting Bias in CRSP's Nasdaq Data and Its Implications for the Size Effect, *The Journal of Finance* 54, 2361–2379.

Stambaugh, Robert F., 1999, Predictive Regressions, *Journal of Financial Economics* 54, 375–421.

Stambaugh, Robert F., Jianfeng Yu, and Yu Yuan, 2012, The short of it: Investor sentiment and anomalies, *Journal of Financial Economics* 104, 288 – 302, Special Issue on Investor Sentiment.

Stock, James H., and Mark W. Watson, 2003, Forecasting Output and Inflation: The Role of Asset Prices, *Journal of Economic Literature* 41, 788–829.

Stock, James H., and Mark W. Watson, 2007, Why Has U.S. Inflation Become Harder to Forecast?, *Journal of Money, Credit and Banking* 39, 3–33.

Timmermann, Allan G., 1993, How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices, *The Quarterly Journal of Economics* 108, 1135–1145.

Veldkamp, Laura L., 2011, *Information Choice in Macroeconomics and Finance* (Princeton University Press).

Vissing-Jorgensen, Annette, 2004, Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions, in *NBER Macroeconomics Annual 2003, Volume 18*, 139–208 (National Bureau of Economic Research, Inc).

Wachter, Jessica A., 2006, A Consumption-based Model of the Term Structure of Interest Rates, *Journal of Financial Economics* 79, 365–399.

Wang, Chen, 2019, Under- and Over-Reaction in Yield Curve Expectations, *Working Paper* .

Wang, Huijun, and Jianfeng Yu, 2015, Dissecting the Profitability Premium, *Working Paper* .

Welch, Ivo, and Amit Goyal, 2007, A Comprehensive Look at The Empirical Performance of Equity Premium Prediction, *The Review of Financial Studies* 21, 1455–1508.

Wright, Stephen, 2004, Measures of Stock Market Value and Returns for the U.S. Nonfinancial Corporate Sector, 1900–2002, *Review of Income and Wealth* 50, 561–584.

Yu, Jialin, 2011, Disagreement and return predictability of stock portfolios, *Journal of Financial Economics* 99, 162 – 183.

Zhang, X. Frank, 2006, Information Uncertainty and Stock Returns, *The Journal of Finance* 61, 105–137.

Zhao, Guihai, 2019, Learning, Equilibrium Trend, Cycle, and Spread in Bond Yields, *Bank of Canada Working Paper* .