Asymmetries in Risk Premia, Macroeconomic Uncertainty and Business Cycles*

Christoph Görtz
University of Birmingham

Mallory Yeromonahos
University of Birmingham

February 2022

Abstract

A large literature suggests that the expected equity risk premium is countercyclical. Using a variety of different measures for this risk premium, we document that it also exhibits growth asymmetry, i.e. the risk premium rises sharply in recessions and declines much more gradually during the following recoveries. We show that a model with recursive preferences, in which agents cannot perfectly observe the state of current productivity, can generate the observed asymmetry in the risk premium. Key for this result are endogenous fluctuations in uncertainty which induce procyclical variations in agent's nowcast accuracy. In addition to matching moments of the risk premium, the model is also successful in generating the growth asymmetry in macroeconomic aggregates observed in the data, and in matching the cyclical relation between quantities and the risk premium.

Keywords: Risk Premium, Business cycles, Bayesian Learning, Asymmetry, Uncertainty, Nowcasting.

JEL Classification: E2, E3, G1.

*We thank Juan F. Rubio-Ramirez, two anonymous referees, Leonardo Melosi, Cédric Tille, Kaushik Mitra, Christoph Thoenissen and Christoph Himmels for extremely helpful comments and suggestions. All remaining errors are our own. Görtz (corresponding author): University of Birmingham, Department of Economics, B152TT Birmingham, UK. Email: c.g.gortz@bham.ac.uk. Yeromonahos: University of Birmingham, Department of Economics,, B152TT Birmingham, UK.

1 Introduction

The most recent U.S. investment and housing booms that ended abruptly in 2001 and 2007 have been associated with highly optimistic beliefs about profitability. In the former case, beliefs about profitability were linked to information technology and in the latter to house price gains. Both booms were associated with times of low uncertainty and saw long hikes in stock markets. Adjustments of beliefs about profitability resulted in sharp recessions, heightened uncertainty, and strong corrections in stock markets (see e.g. Beaudry and Portier (2006) and Shiller (2007)). In this situation, investors were less willing to bear financial risk and, for a given level of stock market risk, they required a higher compensation to hold stocks instead of a risk free short-term asset. Indeed, this equity risk premium increased sharply at the brink of both recessions, very much in contrast to the slow and gradual decline that could be observed during the preceding booms. This growth asymmetry shows in positively skewed distributions of risk premium growth — skewness is 0.64 and 1.00, respectively over these two business cycles.[1]

There is a large body of work in the finance literature on risk premia, which for example provides substantial empirical evidence that the equity risk premium varies over time and is countercyclical.[2] A recent and growing literature jointly studies both the behavior of risk premia and macroeconomic dynamics. We contribute to this body of work which has not considered the important asymmetric feature of risk premia. We first document the degree of asymmetry in the data and then develop a structural model with endogenous countercyclical variations in uncertainty that is consistent with this feature in the data. To the best of our knowledge this is the first paper that tackles this issue.

We start by computing statistics on growth asymmetry for a variety of expected equity risk premium measures that have been found relevant in the literature. We document for the post-WWII U.S. economy that growth rates of all risk premium measures exhibit positive skewness. In particular, we construct risk premium measures using models based on the

[1] These skewness statistics over the two business cycles are significant with p-values of 0.065 and 0.017, respectively. The positively skewed distribution implies that positive changes in risk premia are more extreme than negative changes.

[2] See e.g. Fama and French (1989), Ferson and Harvey (1991), and Bekaert and Harvey (1995).

historical mean of realized stock market returns in excess of Treasury bond yields, and models based on predictive time series regressions of equity returns on selected fundamentals. We further employ direct risk premium measures based on responses of Chief Financial Officers recorded in the Duke CFO Global Business Outlook Survey. The broad support for growth asymmetry is remarkable given the substantial diversity in assumptions underlying the various employed risk premium measures.

We then develop a structural model that is consistent with the empirically observed growth asymmetry in the risk premium. We build on the framework in Van Nieuwerburgh and Veldkamp (2006) who introduce a Bayesian learning mechanism into an otherwise standard RBC model. We extend the framework with capital adjustment costs and preferences of the Epstein and Zin (1989) type. This extension allows disentangling relative risk aversion and the elasticity of intertemporal substitution and gives rise to a risk premium. Our empirical measures for expected risk premia incorporate information about future returns, so that their value can differ from realized ex-post risk premia. In the model, deviations from fundamentals are possible, because agents need to form nowcasts about the state of total factor productivity (TFP). Nowcasting is required as the otherwise standard Cobb-Douglas type production function includes an additive noise term and neither TFP nor the noise term can be observed separately at the time decisions about production inputs are made.[3]

The key mechanism to generate growth asymmetry in the risk premium are endogenous changes in the degree of uncertainty about the state of productivity, which in turn induce procyclical variations in agent's nowcast precision. TFP follows a two-state Markov process and agents employ a Bayesian learning technology to form a nowcast about the state of current productivity. As in Van Nieuwerburgh and Veldkamp (2006), intensified use of production inputs amplifies the signal on the state of productivity relative to the noise and results in endogenous procyclical variations of the signal-to-noise ratio. When production inputs are high, agent's nowcasts are relatively accurate as uncertainty about the state of productivity is low. The peak of the cycle, can be observed relatively precisely. The following

[3]This setup is consistent with the nowcasting process of statistical agencies documented in Faust et al. (2005). They describe the preliminary nowcast to be the sum of the final GDP announcement and an additive noise term.

period is accompanied with a substantial decline in nowcast precision which leads to a sharp increase in the risk premium. The reduced use of production inputs during the following recession leads to a lower signal-to-noise ratio and a situation of heightened uncertainty about the state of productivity. For this reason, agent's nowcasting accuracy increases only slowly during the following recovery. The associated gradual decline in uncertainty comes along with a slow and gradual decline in the risk premium.

The model is calibrated to match nowcast precision in the Survey of Professional Forecasters (SPF). Overall, it is successful in generating the observed growth asymmetry in the risk premium and in matching the risk premium's relation with macroeconomic activity in the data. The model captures the empirically observed positive skewness in the risk premium, while a framework without the endogenous variation in nowcast accuracy does not imply a skewness substantially different from zero. The model's ability to generate growth asymmetry rests on the procyclicality of nowcast precision and the associated variations in uncertainty. This mechanism finds strong support in the data. Firstly, the median absolute nowcast error for real GDP growth from the SPF varies countercyclically and is notably heightened when the economy contracts. We also find this absolute median nowcast error is particularly high at times when the risk premium rises strongly. Secondly, we employ the dispersion in nowcasts for GDP growth from the SPF as a proxy for uncertainty to provide further corroborative evidence for the model mechanism. We report uncertainty varies countercyclically — in line with evidence in the literature, see e.g. Bloom (2014) — and is notably heightened at times when the economy contracts and when risk premia are high. Endogenous procyclical variations in agent's nowcasting precision are crucial also for the model's ability to generate the well known stylized business cycle fact of negatively skewed growth rates in macroeconomic aggregates — i.e. expansions in economic activity are long and gradual while recessions are sharp and short. In addition to the skewness statistics, it is notable that our model is also successful in matching the countercyclical movements of risk premia observed in the data.

Our paper is related to several strands of the literature. There is a large body of work in the empirical finance literature on risk premia. Yet, existing theoretical work in finance has mostly been confined to endowment economies that do not consider feedback between time-varying risk premia and macroeconomic aggregates.[4] On the other hand, most standard macroeconomic models do not include a meaningful role for the risk premium. Our work links to a growing literature that jointly studies the behavior of macroeconomic aggregates and risk premia in bond or equity markets (e.g. Jermann (1998) and Kaltenbrunner and Lochstoer (2010)). Gilchrist and Zakrajšek (2012) empirically document a close link between increases in the excess bond premium and a deterioration of macroeconomic conditions. Gourio (2013) develops a macroeconomic framework driven by variations in disaster risk that reproduces key features of corporate bond risk premia — such as their countercyclicality — and studies their implications for business cycles. Campbell et al. (2019) show how macroeconomic dynamics drive risk premia in bond and equity markets and Corradi et al. (2013) find that the level and volatility of fluctuations in the stock market are largely explained by business cycle factors. Bekaert et al. (2009) highlight the role of uncertainty for the countercyclical volatility of asset returns. We contribute to this literature by explaining the growth asymmetry in risk premia and macroeconomic aggregates.[5]

Our paper is also related to work that highlights the importance of beliefs about current and future TFP for fluctuations in macroeconomic and financial aggregates (e.g. Beaudry and Portier (2006), Barsky and Sims (2011), Pinter et al. (2017), Cascaldi-Garcia and Vukotic (2019)). Görtz et al. (2021) document a close link between changes in expectations about future TFP, stock prices and risk premia. Risk premia incorporate expectations about future stock market returns and as such, they can differ from ex-post realizations. In our model, this can be the case as agents need to form nowcasts to learn about the current state of productivity. Milani (2011) highlights the relevance of expectations and learning for output fluctuations. He relaxes the rational expectations assumption to allow for agent's learning in a New Keynesian framework and estimates the model using forecast data from the SPF.

[4] Jermann (1998) and Lettau and Uhlig (2000) stress that many asset pricing models which are successful in endowment economies do not generalize well to production economies.

[5] The implications of uncertainty on financial markets have also been examined through a different lens in a tangential branch of the finance literature. Papers such as e.g. Segal et al. (2015) explore how macroeconomic uncertainty can cause asymmetries in the cross sectional distributions of stock returns. However, this literature does not study the time series properties of the equity risk premium.

Enders et al. (2017) compute GDP nowcast errors based on the SPF and show that these are sizable and play a non-negligible role, accounting for up to 15% of output fluctuations. Cascaldi-Garcia and Galvao (2020) disentangle the effects of beliefs about future technology and uncertainty shocks. Johri and Karimzada (2020) document that fluctuations in firms' learning from production activity is an important source for fluctuation in macroeconomic aggregates over the business cycle.

While the above literature typically does not consider asymmetries, in our framework agents have to solve a signal extraction problem with time varying parameters to explain growth asymmetries in the data. In this respect, our work links closely to a literature that considers an asymmetric speed of learning and time variation in uncertainty (e.g. Veldkamp (2005), Boldrin and Levine (2001), Fajgelbaum et al. (2017)). It is important to note that learning at a constant speed is not sufficient to resemble the empirically observed asymmetry — the beginning of an expansion and a contraction would be observed with the same precision which triggers in the model a response of risk premia of equal size (in absolute terms). The key for the model to generate asymmetry in risk premia is that the speed of learning varies procyclically. Our mechanism for endogenous variations in the signal-to-noise ratio is closely related to Van Nieuwerburgh and Veldkamp (2006) and Ordoñez (2013) who employ it to explain steepness asymmetry in macroeconomic aggregates observed at business cycle frequencies. A similar mechanism is used in Saijo (2017), where agents learn about the efficiency of investments in an environment where uncertainty varies endogenously and has adverse effects on economic activity. Common across this literature is that endogenous variations in uncertainty imply state dependencies in the strength of agents' responses to shocks. While also our model relies on such a type of mechanism, the studies above use it to explain empirical facts related to macroeconomic aggregates. We add to this literature by studying asymmetries in risk premia.[6]

A procyclical speed of learning is not the only model mechanism that can facilitate

[6] In principle, a business cycle or asset pricing model that gives rise to positive countercyclical risk premia and features the discussed type of procyclical variations in the speed of learning, can generate the empirically documented growth asymmetries. To the best of our knowledge, our paper is the first to explicitly focus on this empirical fact. An advantage of our model is that it can generate both, the empirically observed asymmetry in risk premia and macroeconomic aggregates.

matching the growth asymmetry of risk premia in the data. In principle, this would be possible for any mechanism that generates countercyclical risk premia, and that implies a state dependent response of risk premia so that these rise faster at the onset of a recession than they decline during the following expansion.[7] This holds for example for models with occasionally binding constraints of the type in He and Krishnamurthy (2013). He and Krishnamurthy (2013) focus on the sharp increase in risk premia during asset market crisis. In their model, the level of financial intermediaries' equity capital is important in driving growth asymmetries in risk premia: when equity capital is particularly low, as for example during an asset market crisis, losses within the financial sector have significant effects on risk premia and will trigger a surge in the latter. Yet, when equity capital is high, as during an expansion, losses have no effect on risk premia and they decline more gradually. While the particular focus of He and Krishnamurthy (2013) is on times of severe frictions in financial markets and substantial reductions in bank equity, this mechanism would be less prevalent for the business cycles that do not coincide with a financial crisis. We instead rely on procyclical variations in agent's learning, which we find — consistent with the literature on varying signal-to-noise ratios discussed above — to be a feature over the entire sample.

The remainder of the paper is structured as follows. Section 2 provides an overview about the data. In Section 3 we provide details on the estimation of risk premium measures and document their growth asymmetry. Section 4 describes the model and Section 5 the calibration and computational details. Section 6 discusses the model mechanism that gives rise to asymmetries and results from simulations. Section 7 concludes.

2 Data

We construct measures for U.S. risk premia over a horizon from 1957Q3 to 2019Q2. For comparability with the existing literature, we follow the common practice and use the S&P

[7] The former is for example a feature in the model by Routledge and Zin (2010) with asymmetric preferences due to generalized disappointment aversion — they imply a higher mean risk premium during contractions. However, on their own, these preferences cannot deliver the latter as they imply a change in risk premia that is of equal size (in absolute terms) at the onset of an expansion and at the beginning of a contraction.

500 as a measure for equity prices and treasury yields for the risk-free rate (see e.g. Graham and Harvey (2007)). Quarterly time series for the S&P 500 index are from Robert Shiller's website. For most risk premium measures we consider an investment horizon of one quarter since it has become standard in the literature to estimate risk premia at this horizon.[8] A survey based measure can only be considered at an annual investment horizon due to data availability. Consistent with the respective horizon, we either use the 3-Month Treasury Bill rate (TB3MS) or the 1-Year Treasury Constant Maturity rate (DGS1) as measures for the risk-free rate which are obtained from the Board of Governors of the Federal Reserve System.

For the fundamentals in the regression based method to estimate risk premia, we use the cyclically adjusted price-earning Ratio (CAPE) available from Robert Shiller's website. As an alternative fundamental, we compute the cyclically adjusted price-dividend ratio (CAPD) based on data from the same source. Consistent with Shiller's cyclical adjustment to the price earnings-ratio, we compute the CAPD as the current real price of equity divided by the average of dividends over the previous ten years. The real price of equity is defined as the S&P 500 index deflated with the CPI.

The U.S. Bureau of Economic Analysis provides time series for real gross domestic product (GDPC1), real gross private domestic investment (GPDIC1), and real personal consumption expenditures (PCECC96). These series are quarterly, seasonally adjusted, and in billions of chained 2012 Dollars. Hours worked by all persons in the non-farm business sector (HOANBS) is available from the US Bureau of Labor Statistics. This source also provides a time series of civilian non-institutional population (CNP16OV) used to express the above macroeconomic aggregates in per-capita terms.

3 Empirical Evidence on Risk Premia

In this section, we estimate risk premia using a variety of models that have been found relevant in the literature. We then document that all measures for risk premia exhibit growth asymmetry.

[8] See for example Goyal and Welch (2008), Lettau and Ludvigson (2001a), Lettau and Ludvigson (2001b).

The equity risk premium is the compensation required to make agents indifferent at the margin between investing in a risky market portfolio and a risk-free bond. Formally, the equity risk premium at time t over investment horizon k, $ERP_{t,t+k}$, is defined as the difference between the expected return on equity, $R^e_{t,t+k}$, and the risk-free rate, $R^f_{t,t+k}$, over horizon k,

$$E_t[ERP_{t,t+k}] = E_t[R^e_{t,t+k}] - R^f_{t,t+k}. \tag{1}$$

The term $R^f_{t,t+k}$, as it is risk-free, is known at time t, while the future expected performance of the stock market is not. Investors can only observe with certainty the past returns of the stock market up to time t, and can use the information available to form expectations.

To compute the risk premium in equation (1), a variety of methods have been suggested in the literature. Duarte and Rosa (2015) provide an extensive overview about the most widely used models and classify these in five categories. We will estimate risk premia based on three models for investment horizon $k = 1$ and one model for investment horizon $k = 4$ which, according to Duarte and Rosa (2015)'s classification, are part of three of these categories. The first category comprises models based on the historical mean of realized equity premia, the second includes models that employ time series regressions and the third is based on survey data. Models in these three categories have the advantage that they rely on a minimum of assumptions, and importantly, allow us to compute long time series for risk premia. The other two methods classified by Duarte and Rosa (2015) are undoubtedly very useful in other circumstances, but have substantial drawbacks for our purposes.[9] We now provide a brief overview over the models we employ to compute risk premia.

[9]Models based on cross-sectional regressions (see e.g. Adrian et al. (2013)) impose tight restrictions on the estimation of risk premia and results are heavily dependent on the portfolios, state variables and risk factors used (Harvey et al. (2016)). While models in our three considered categories use information in real time where investors don't have information sets that include future realizations, this method uses full-sample regression estimates which is particularly problematic in our context with a focus on asymmetries. Risk premium estimates based on dividend discount models (see e.g. Damodaran (2019)) require additional strong assumptions, for example on the computation of future expected dividends and a discount rate for these dividends.

3.1 Historical Mean of Realized Returns

This method is the most straightforward of all approaches to compute the future risk premium from time t to $t+k$. Following Goyal and Welch (2008), it is simply the historical mean of realized stock market returns in excess of the risk-free rate over H periods preceding time t. This can be formalized as

$$ERP_{t,t+k} = \frac{1}{H}\sum_{h=0}^{H}(R^e_{t-k-h,t-h} - R^f_{t-k-h,t-h}).$$

We specify $H = t - k$ as in Goyal and Welch (2008) who use systematically all the available historical data since the beginning of the sample.

The validity of this method relies on the assumption about consistent behavior between past and future. This means the mean of excess returns should either be constant or very slow moving to avoid a systematic bias in the estimates. We verify that there is no trend in realized excess stock market returns using the augmented Dickey-Fuller test (for details see Appendix A.1).

3.2 Time series regressions

This method is based on the idea to utilize the relationship between time series of economic variables and stock market returns to predict future equity returns from a linear regression. One can then subtract the contemporaneous risk-free rate to recover an estimate of the risk premium, as in Fama and French (1988), Fama and French (2002) and Campbell and Thompson (2008). We estimate the following predictive regression

$$R^e_{t,t+k} = \alpha + \beta \cdot fundamental_t + \varepsilon_t, \tag{2}$$

where $fundamental_t$ represents a variable that theory and practice have found likely to drive future excess stock returns. This method links as directly as possible to equation (1) by computing the equity risk premium

$$E_t[ERP_{t,t+k}] = \hat{\alpha} + \hat{\beta} \cdot fundamental_t - R^f_{t,t+k}, \tag{3}$$

based on the estimates $\hat{\alpha}$ and $\hat{\beta}$ for α and β. Generally the literature relies on a single fundamental in this regression, as using several variables at once has been found to reduce model's out-of-sample accuracy. The fundamental used is typically a valuation ratio such as the price-dividend ratio or the price-earning ratio. These valuation ratios are known to be negatively correlated with future stock returns since the works of Rozeff (1984), Campbell and Shiller (1988a), and Campbell and Shiller (1988b).

We compute risk premia from two different models based on the above time series regressions and follow the detailed methodology in Campbell and Thompson (2008). The models differ in the variable used as fundamental, where we either employ Shiller's cyclically adjusted price-earning ratio (CAPE) or the cyclically adjusted price-dividend ratio (CAPD). For each quarter t in our sample, we estimate parameters in equation (2) based on a sample up to time $t-1$. The risk premium is then constructed according to equation (3) using an out-of-sample forecast. To estimate α and β we use a sample that begins 20 years prior to 1957Q3.[10] We further implement the two restrictions suggested by Campbell and Thompson (2008), i.e. $\hat{\beta}$ must have the sign predicted by theory, otherwise it is replaced by zero, and the predicted risk premium must be positive, otherwise the historical mean is used as a predictor instead. Out-of-sample forecasts are produced for each quarter t from 1957Q3 to 2019Q2.

3.3 Survey based risk premium measures

The third method we consider to derive a measure for the risk premium is based on survey data. The Duke CFO Global Business Outlook Survey is the longest ongoing survey about the expected equity return (conducted quarterly since 2000Q2) in the United States.[11]

[10] Our results are robust also to using a sample beginning in 1881Q1, when both fundamentals are first available.

[11] Every quarter, on average about 350 Chief Financial Officers from a sample of representative US firms respond to the following question: "The current annual yield on a 10-year Treasury bond is x%. Please

Graham and Harvey (2018) then recover the 10-year ahead expected risk premium by subtracting the known risk-free Treasury bond annual yield to the median forecast of future S&P 500 annual returns. Since 2004Q1, the survey also includes a question on the expected return of the S&P 500 over the next year. We use the responses to this question to compute, analogously to Graham and Harvey (2018), the expected risk premium for an investment horizon of one year. Responses and questions based upon which we could construct risk premia with an investment horizon of one quarter are not available in this survey.

3.4 Asymmetries in Risk Premia

In this section, we show skewness statistics for the growth rate of the risk premium measures described above. In particular, we report results based on a model that relies on the historical mean of realized returns, results based on two time series regression models (using either the cyclically adjusted price-dividend or the price-earning ratio as fundamental), and results based on survey evidence.

Throughout the paper, the skewness measure that we employ as a baseline is defined as follows: define m_r as the r^{th} moment about the mean $\bar{x}$, where n is the number of observations in the sample. Given $m_r = \frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^r$, the coefficient of skewness is then defined as $m_3 {m_2}^{-3/2}$.[12] The first three columns of Table 1 summarize results based on an investment horizon of one quarter ($k = 1$) for the historical average method and the two approaches based on time series regressions. All three risk premium measures exhibit growth asymmetry which manifests in positively skewed distributions. Over our sample (1957Q3 - 2019Q2) the growth in the risk premium based on the historical average method and the CAPD and CAPE time series regression models have a skewness of 0.12, 1.12 and 0.25 respectively.[13] All three statistics are highly significant. The positive skewness implies that the risk premium exhibits growth asymmetry: it declines gradually and rises much more

complete the following: Over the next 10 years, I expect the average annual S&P 500 return will be: ...%". Here x% is replaced by the the actual yield on a 10-year Treasury bond at the time of the survey. A corresponding question is asked for a one-year investment horizon.

[12] In Appendix A.2 we show that all risk premium measures also exhibit a positive skewness when we apply the quantile-based skewness measure suggested by Ghysels et al. (2016).

[13] For comparison, the skewness of the GDP growth rate over this sample is -0.52.

sharply. We attach most weight to the historical average measure as, in comparison to the two time series methods, it relies on minimal assumptions and there is widespread agreement in the empirical literature that it is the method with the highest accuracy (see e.g. Welch and Goyal (2003), Simin (2008) and Welch and Goyal (2008)).[14]

Table 1: Skewness statistics for growth in different measures of risk premia

Historical average	Time series (fundamental = CAPE)	Time series (fundamental = CAPD)	Survey
0.12 (0.029)	0.25 (0.004)	1.12 (0.000)	0.8 (0.012)

Notes. The sample is 1957Q3-2019Q2 for the historical mean and time series methods. Survey data are available only from 2004Q1 to 2019Q2. "Historical average", refers to the expectations obtained using the historical average method. "Time series (fundamental = CAPE)" and "Time series (fundamental = CAPD)" refer to the expectations obtained using the time series regression method, using the CAPE and CAPD ratios respectively as fundamentals. "Survey" refers to a risk premium measure based on the Duke CFO Global Business Outlook Survey. The investment horizon is one year for the survey measure and one quarter for all other measures. Skewness statistics are calculated from the first difference of the logarithm of the risk premium. P-values, in parenthesis, are based on Bai and Ng (2005) test statistics.

The historical average method and the CAPE and CAPD time series methods have very similar means, 0.06, 0.07 and 0.06, respectively. Given the similarity in the underlying methodology, the CAPE and CAPD time series methods also have a relatively similar standard deviation (0.33 and 0.34), while this statistic is substantially lower for the historical average method (0.03). All three risk premium estimates are positively correlated. The correlation of the historical average measure with the CAPE and CAPD time series measures is 0.67 and 0.33, respectively. The correlation between the two time series measures is 0.30. These findings are consistent with Duarte and Rosa (2015) who document similarities and differences in first and second moments across risk premium estimates. To supplement the discussion on moments in the data, we show in Appendix A.2 also the time series of risk premia estimated at the one quarter investment horizon. It is reassuring that despite any differences between these time series, the results from Table 1 indicate they all exhibit a positive skewness.

[14]We confirm this for our risk premium estimates. We calculate the excess Root Mean Square Error (RMSE), as in Welch and Goyal (2008), as the difference between the RMSE of the historical average method and the RMSE of an alternative method (both at the one-quarter investment horizon). The excess RMSE of the CAPE and CAPD methods are -0.21 and -0.42, which indicates higher accuracy of the historical average method.

We also briefly discuss the survey-based measure, which, for data availablility reasons, is constructed with an investment horizon of one year, over the time sample 2014Q1-2019Q2.[15] Although it is calculated on a shorter sample and with a different investment horizon, this survey-based measure brings interesting corrobative evidence as it is based on a very different methodology. The growth in the risk premium based on the survey method exhibits a skewness of 0.80 (shown in the fourth column of Table 1) and hence also provides evidence for steepness asymmetry in the risk premium.

The quantitative differences between skewness statistics in Table 1 are not surprising — and consistent with findings for first and second moments in the literature (see e.g. Duarte and Rosa (2015)) — in light of the substantial diversity in assumptions and the underlying methodologies to derive the risk premium measures. Given this, it is striking that all considered measures support the notion of positive skewness. Overall, this section provides broad evidence that the risk premium — measured in a variety of ways and at different investment horizons — exhibits growth asymmetries: declines are long and gradual and rises are sharp and short.

3.5 Cyclical variations in nowcast precision and uncertainty

We have documented in the section above that risk premia exhibit growth asymmetry. In this section, we consider empirical evidence that can be informative for mechanisms in structural models to resemble this growth asymmetry. In particular, we provide evidence on procyclical variations in nowcast precision and contercyclical variations in uncertainty.

We employ the median absolute nowcast error for GDP growth from the Survey of Professional Forecasters (SPF) as a measure for nowcast accuracy and hence the speed of learning. Table 2 shows results from the regression $y_t = \alpha + \beta x_t + \varepsilon$, where y_t is the median absolute nowcast error for GDP growth and x_t is either GDP growth or a dummy indicating a contraction of the economy. We find a positive relationship between GDP growth and nowcast

[15] The Duke CFO Global Business Outlook Survey does not include a question that corresponds to the one quarter investment horizon. Based on this survey, Graham and Harvey (2018) also provide a risk premium measure for a 10 year investment horizon though. Skewness for growth in this measure is 0.151 over a 2000Q2-2019Q2 sample.

accuracy and we document that nowcast errors are particularly large during contractions. Our results based on nowcasting accuracy are consistent with findings in the literature on procyclical forecast precision. See for example evidence in Jaimovich and Rebelo (2009) based on the Livingston Survey, or in Van Nieuwerburgh and Veldkamp (2006), who consider an exercise in the spirit of the one above, based on forecast data from the SPF. Table 2 also provides evidence on the relationship between nowcast accuracy and growth of expected risk premia. We regress the median absolute nowcast error for GDP growth on a dummy for large increases in risk premia — indicating a quarter-on-quarter growth rate of expected risk premia of at least to 2%.[16] The result of this regression indicates that surges in risk premia coincide with times of a slow speed of learning.

Table 2: Nowcast accuracy, economic contractions and surges in risk premia

	Negative GDP growth dummy	GDP growth rate	Positive risk premium growth dummy
$\hat{\alpha}$	1.436 (0.000)	2.160 (0.000)	0.753 (0.034)
$\hat{\beta}$	2.537 (0.000)	-0.177 (0.000)	1.638 (0.000)
Adjusted R^2	0.23	0.09	0.02

Results of the time series regression $y_t = \alpha + \beta x_t + \varepsilon$ where y_t is the median absolute nowcast error for real GDP growth from the Survey of Professional Forecasters (SPF), and x_t is either the quarter-on-quarter growth rate of real GDP, or a dummy variable equal to one when the quarter-on-quarter growth of real GDP is negative, or a dummy variable equal to one when the quarter-on-quarter growth rate of risk premium exceeds 2%. The sample is limited to 1968Q4-2019Q2 by the availability of the SPF. P-values are reported in parenthesis.

Next, we turn to regressions where we employ the dispersion of nowcasts for GDP growth from the SPF as dependent variable. Dispersion is defined as the difference between the 75th and the 25th percentile of the projections for quarter-on-quarter growth. Disagreement of private sector expectations, as reported in the SPF, are a widely used proxy for uncertainty.[17] Considering GDP growth as independent variable, Table 3 reveals a negative link between output growth and uncertainty. A regression with a dummy — indicating times when the economy contracts — as independent variable corroborates this finding, reporting a signifi-

[16] This is a conservative classification for a period to exhibit a large increase in the risk premium. The dummy is one in 71 out of a total of 203 quarters. Results in Tables 2 and 3 are robust also when we apply a tighter threshold that includes surges in risk premia above about 4%. This implies the dummy is unity in 25 periods which is the same number of quarters covered by the dummy indicating a contraction in GDP.

[17] See e.g. Bachmann et al. (2013).

cant positive relationship between contractions and nowcast dispersion. Our results on the adverse link between uncertainty and economic activity are consistent with findings in the literature (see e.g. Bachmann et al. (2013), Bloom (2014)). Using the dummy for strong surges in risk premia as independent variable shows that risk premia are heightened at times of high uncertainty. Investors tend to be more uncertain about the current and future state of the economy during economic contractions which requires compensation through higher risk premia. Our results are consistent with evidence in Corradi et al. (2013) who report the volatility of risk premia to be strongly countercyclicyal and with Baker et al. (2012) who use firm level data to document that uncertainty raises stock price volatility.

The evidence provided in Tables 2 and 3 suggest a link between uncertainty, strong risk premium growth, the state of the business cycle and variations in the speed of learning Informed by this evidence, we introduce in the next section a structural model that can explain the empirically observed asymmetry in risk premia due to endogenous procyclical variations in nowcasting precision.

Table 3: Uncertainty, economic contractions and surges in risk premia

	Negative GDP growth dummy	GDP growth rate	Positive risk premium growth dummy
$\hat{\alpha}$	1.267 (0.000)	1.570 (0.000)	1.352 (0.000)
$\hat{\beta}$	1.139 (0.000)	-0.069 (0.000)	0.369 (0.034)
Adjusted R^2	0.19	0.05	0.02

Results of the time series regression $y_t = \alpha + \beta x_t + \varepsilon$ where y_t is the dispersion of individual nowcasts for real GDP growth from the Survey of Professional Forecasters (SPF), and x_t is either the quarter-on-quarter growth rate of real GDP, or a dummy variable equal to one when the quarter-on-quarter growth of real GDP is negative, or a dummy variable equal to one when the quarter-on-quarter growth rate of risk premium exceeds 2%. The dispersion of nowcasts is measured as the difference between the 75th percentile and the 25th percentile of the nowcasts for quarter-on-quarter GDP growth nowcasts, expressed in annualized percentage points. The sample is limited to 1968Q4-2019Q2 by the availability of the SPF. P-values are reported in parenthesis.

4 The model

The core of our model is a representative agent real business cycle (RBC) model which is extended with two key mechanisms. Firstly, households have recursive preferences of the Epstein and Zin (1989) type. It is well known that standard RBC models with Arrow-Pratt preferences and a reasonable degree of relative risk aversion (RRA) fail to account for the existence of risk premia. This is due to the fact that the intertemporal elasticity of substitution (EIS) and the RRA are reciprocal of each other. A small EIS of the magnitude necessary to justify meaningful risk premia necessarily leads to an excessively large RRA. Recursive preferences separate the RRA and the EIS. Secondly, agents cannot directly observe productivity. Instead, they receive a noisy signal about previous period's productivity and use a Bayesian learning technology to form nowcasts. Agent's varying speed of learning over the business cycle is the key to match empirically observed asymmetries in risk premia and macroeconomic variables. The model builds on the framework in Van Nieuwerburgh and Veldkamp (2006) who introduce the Bayesian learning mechanism into an otherwise standard RBC model. We extend their model by introducing Epstein and Zin (1989) preferences and investment adjustment cost which gives rise to meaningful risk premia and at the same time retains the model tractable.

4.1 Production and technology

The economy comprises of a continuum of perfectly competitive identical firms with unit mass. Firms use the following Cobb-Douglas production function to produce output, y_t,

$$y_t = A_t k_t^{\alpha} l_t^{1-\alpha} + \nu_t, \qquad 0 < \alpha < 1, \tag{4}$$

by employing capital, k_t, and labour, l_t. Output further depends on a productivity shock, A_t, and an additive noise shock, ν_t. This production function is based on Van Nieuwerburgh and Veldkamp (2006) and is consistent with the notion in Faust et al. (2005) who characterize the preliminary GDP announcement of statistical agencies as the sum of a final GDP announcement and a noise term. The productivity shock takes the form of a Markov process with two states, high and low $A_t = \{A_t^H, A_t^L\} \; \forall \, t$, and a standard deviation σ_A. The Markov chain is ergodic and has a symmetric transition matrix, Π, to ensure any asymmetry in the resulting model dynamics is endogenous. An asymmetric transition matrix could in itself lead to asymmetric business cycles, even without endogenous uncertainty and learning. The noise shock is independent and identically normal distributed with zero mean and standard deviation σ_ν.

The assumptions about agent's information set are such that — even though they know the underlying shock processes — they cannot separately observe the productivity and noise shock. Further, agents make decisions about production inputs before they know the level of output since both shocks are realized only at the end of each period. To make an informed decision about production inputs, agents use a Bayesian learning technology to infer the level of current period's productivity based on their noisy observation of output in the previous period.[18]

Both, firms as well as households have the same belief about current productivity since all agents have the same information set and have access to the same Bayesian updating technology. In the following sections, we will discuss optimal decision making of firms and households, given their beliefs about productivity, and show how these agents employ the Bayesian learning technology to update their beliefs.

4.2 Firms

Firms enter the period with knowledge about their capital stock. They use Bayesian updating, to be described in detail below, to form a belief about productivity at the beginning of the period. Given this information, firms decide about labor demand and investment, where the latter determines next period's capital stock. Firms own the capital stock, rather than rent it from households, but issue shares and pay out dividends.

At the beginning of the period, after firms have formed a belief about productivity, they expect cash flow, $\tilde{f}_t$, to be

$$\tilde{f}_t = \tilde{A}_t k_t^\alpha {l_t^d}^{1-\alpha} + \tilde{\nu}_t - w_t l_t^d - i_t,$$

[18]These timing assumptions are consistent with nowcasting in public policy institutions. Bok et al. (2017) document that the New York Fed Staff Nowcast for GDP on the last quarter is only observable at about the beginning of the next quarter. They also describe that nowcasts and forecasts are based on surveys and limited number of reporting units, i.e. they filter.

where $\tilde{A}_t$ denotes the beliefs about productivity and $\tilde{\nu}_t$ the belief about the noise. Following the discussion in the section above, agent's expectation about the noise, $\tilde{\nu}_t$, is zero. w_t denotes the real wage, l_t^d stands for labor demand, and i_t for investment. In general, notation $\tilde{x}_t$ indicates agent's belief about a particular variable. This belief is formed at the beginning of the current period, t, given the information set at the beginning of the current period, $\mathcal{I}_t$, such that $\tilde{x}_t = \mathbb{E}_t[x_t \mid \mathcal{I}_t]$. Then, x_t denotes the realization of this variable at the end of period t.

Firms have to respect their investment financing constraint

$$i_t = \tilde{y}_t - w_t l_t^d - \tilde{d}_t s_t^s + p_t(s_{t+1}^s - s_t^s), \tag{5}$$

where the difference between s_{t+1}^s and s_t^s represents the supplied number of shares to be traded at price p_t between firms and households. Expected dividends, $\tilde{d}_t$, communicated to the households at the beginning of the period, are given by

$$\tilde{d}_t = \frac{\tilde{y}_t - w_t l_t^d - i_t + p_t(s_{t+1}^s - s_t^s)}{s_t^s}. \tag{6}$$

Actual dividends are paid out at the end of the period and will absorb the effects of incorrect beliefs and balance out the investment financing constraint. Note that realized cash flow,

$$f_t = A_t k_t^{\alpha} l_t^{d\,1-\alpha} + \nu_t - w_t l_t^d - i_t,$$

will differ from expected cash flow most of the time as they include realized productivity as well as the realization of the noise term.

The law of motion for capital is

$$k_{t+1} = \left[(1-\delta) + \Phi\left(\frac{i_t}{k_t}\right)\right] k_t, \tag{7}$$

where δ is the depreciation rate and the capital adjustment cost function $\Phi\left(\frac{i_t}{k_t}\right)$ is positive and concave. The concavity implies that large changes in the investment ratio are more expensive than gradual adjustments. As in Hayashi (1982) and Jermann (1998), the adjustment cost has the functional form

$$\Phi\left(\frac{i_t}{k_t}\right) = \frac{a_1}{1-\chi}\left(\frac{i_t}{k_t}\right)^{1-\chi} + a_2, \qquad \chi > 1,$$

where χ is the elasticity of the investment ratio with respect to Tobin's q and parameters a_1 and a_2 ensure costs are zero in the steady state. The use of these capital adjustment costs allows us to derive the expression for the return on equity shown as shown in Appendix B.2.

Firms maximize their value, which is equivalent to the sum of discounted expected cash flow

$$\max_{l_t^d, i_t, k_{t+1}} \mathbb{E}_t \left[\sum_{j=0}^{+\infty} m_{t,t+j} \left(\tilde{A}_{t+j} k_{t+j}^{\alpha} l^{d\,1-\alpha}_{t+j} - w_{t+j} l_{t+j}^d - i_{t+j} \right) \middle| \mathcal{I}_t \right], \tag{8}$$

where $m_{t,t+j}$ is the household's discount factor to be specified in the next section. We maximize equation (8) with respect to i_t, k_{t+1} and l_t^d subject to equation (7) and the constraints $i_t \geq 0$, $k_t \geq 0$ to obtain the first order conditions

$$q_t = \frac{1}{\Phi'(i_t/k_t)}, \tag{9}$$

$$q_t = \mathbb{E}_t \left\{ m_{t,t+1} \left[\tilde{A}_{t+1} l^{d\,1-\alpha}_{t+1} \alpha k_{t+1}^{\alpha-1} - \frac{i_{t+1}}{k_{t+1}} + q_{t+1} \left(1 - \delta + \Phi\left(\frac{i_{t+1}}{k_{t+1}}\right)\right)\right] \middle| \mathcal{I}_t \right\}, \tag{10}$$

$$w_t = (1-\alpha)\tilde{A}_t l_t^{d\,-\alpha} k_t^{\alpha}, \tag{11}$$

where q_t denotes the Lagrange multiplier and can be interpreted as Tobin's q. Equation (9) determines the real price of investment and equation (10) determines optimal investment. The labor supply function (11) states that the real wage is equal to the expected marginal productivity of labour, since the actual marginal productivity is unobservable.

4.3 Households

There is a continuum of identical households with unit mass. At the beginning of each period, households decide how much labor to supply and how many shares to buy. Based on their expected cash flow, firms also inform households on the amount of dividends, $\tilde{d}_t$, they expect to pay. When firms observe their realized cash flow at the end of the period, they pay dividends, d_t, which may differ from the expected dividends. At this point, households update their views about their income which they subsequently use for consumption, c_t. In other words, consumption expenditures absorb any unexpected realizations due to incorrect beliefs to satisfy the households' budget at the end of the period. At the beginning of the period, the households' expected budget constraint is

$$\tilde{c}_t + p_t(s^d_{t+1} - s^d_t) = w_t l^s_t + \tilde{d}_t s^d_t, \tag{12}$$

where $\tilde{c}_t$ is the expected consumption level, labor supply is l^s_t, and the difference between s^d_{t+1} and s^d_t represents the demand for the number of new shares.[19]

Households have preferences as in Epstein and Zin (1989) so that recursive utility is a CES aggregate of their period utility function and a certainty equivalent for next period utility,

$$U_t = \left[(1-\beta)u_t(\tilde{c}_t, l^s_t)^{1-\frac{1}{\psi}} + \beta\left(\mathbb{E}_t\left[U^{1-\gamma}_{t+1} \mid \mathcal{I}_t\right]\right)^{\frac{1-\frac{1}{\psi}}{1-\gamma}}\right]^{\frac{1}{1-\frac{1}{\psi}}}, \tag{13}$$

where $\psi > 1$ is the elasticity of inter-temporal substitution, $\gamma \in [0, +\infty) \setminus \{1\}$ is the relative risk aversion, and $\beta \in (0,1)$ is the discount factor. Period utility takes the form

$$u_t(\tilde{c}_t, l^s_t) = \tilde{c}^\kappa_t (1 - l^s_t)^{1-\kappa}, \tag{14}$$

with $\kappa \in (0,1)$ which controls labor supply.

Households maximize equation (13) subject to (12) and to the interiority conditions $\tilde{c}_t \geq 0$, $c_t \geq 0$ and $0 \leq l^s_t \leq 1$. We obtain the following labor supply function from the household's maximization problem (details are provided in Appendix B.1)

$$\tilde{c}_t = \frac{\kappa}{1-\kappa}(1 - l^s_t)w_t, \tag{15}$$

which provides an intratemporal link between labor supply, the real wage and the beginning

[19] Note that the end of period budget constraint would abstract from the tildes in equation (12). It would be exactly valid because the residuals — consumption or dividends — adjust.

of period belief about consumption. Combining the first order condition with respect to s_{t+1}^d and the Envelope Theorem for s_t^d (shown in Appendix B.1) we obtain the Lucas equation[20]

$$1 = \mathbb{E}_t \left[m_{t,t+1} \frac{\tilde{d}_{t+1} + p_{t+1}}{p_t} \,\middle|\, \mathcal{I}_t \right], \tag{16}$$

where

$$m_{t,t+1} = \beta \left(\frac{U_{t+1}^{1-\gamma}}{\mathbb{E}_t \left[U_{t+1}^{1-\gamma} \mid \mathcal{I}_t \right]} \right)^{1-\frac{1}{\theta}} \left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t} \right)^{\frac{\kappa(1-\gamma)}{\theta}-1} \left(\frac{1-l_{t+1}^s}{1-l_t^s} \right)^{\frac{(1-\gamma)(1-\kappa)}{\theta}}, \tag{17}$$

is the stochastic discount factor, using $\theta := (1-\gamma)/(1-\frac{1}{\psi})$. The risk-free rate between period t and $t+1$ is thus defined as

$$R_{t,t+1}^f = \frac{1}{\mathbb{E}_t[m_{t+1,t} \mid \mathcal{I}_t]}, \tag{18}$$

and the expected return on equity between period t and $t+1$ is

$$\mathbb{E}_t[R_{t,t+1}^e \mid \mathcal{I}_t] = \mathbb{E}_t \left[\frac{\tilde{d}_{t+1} + p_{t+1}}{p_t} \,\middle|\, \mathcal{I}_t \right]. \tag{19}$$

Then the expected risk premium is given by

$$\mathbb{E}_t[ERP_{t,t+1} \mid \mathcal{I}_t] = \mathbb{E}_t[R_{t,t+1}^e \mid \mathcal{I}_t] - R_{t,t+1}^f.$$

Note that, as for example in Heer and Maußner (2012), we do not explicitly take into account equation (5) in the maximization programme of the representative firm. This is because irrespective of the choice of labour and investment, it is always possible to find a combination of dividends and number of shares that satisfies equation (5). Since we also do not impose a specific dividend policy, we cannot directly compute the return on equity based on dividends and share prices. However, as shown in Appendix B.2, we can recover

[20] As explained above, $\tilde{d}$ denotes beginning-of-period expected dividends. We preserve the tilde notation despite the presence of the expectation around the right-hand side to avoid any confusion with the realised end-of-period dividends. For this reason, we apply the notation with tilde for dividends and consumption, where appropriate, throughout this section.

the expected return on equity to be

$$\mathbb{E}_t\left[\frac{\tilde{d}_{t+1}+p_{t+1}}{p_t}\,\middle|\,\mathcal{I}_t\right]=\mathbb{E}_t\left[\frac{q_{t+1}k_{t+2}+y_{t+1}-w_{t+1}l_{t+1}-i_{t+1}}{q_t k_{t+1}}\,\middle|\,\mathcal{I}_t\right]. \tag{20}$$

Variables on the right hand side of this equation can be recovered using household's and firm's programmes, given expectations about future productivity. We will discuss in the next section how these expectations about productivity can be formed. Using the right hand side of equation (20) to compute the expected return on equity has the advantage that it limits the state space of the dynamic programming problem and thereby keeps our computational problem tractable.

4.4 Bayesian learning

We now turn to a description of the Bayesian learning mechanism which agents use to form a belief about current technology, $\tilde{A}_t$. Information set $\mathcal{I}_t$ contains all information available to the agents at the beginning of period t

$$\mathcal{I}_t := \{y^{t-1}, c^{t-1}, d^{t-1}, p^t, w^t, l^{d^t}, l^{s^t}, i^t, k^t, s^{d^{t+1}}, s^{s^{t+1}}\},$$

where x^t denotes the history of variable x up to time t. The technology and the noise shocks are never individually observed, but agents have information about their underlying processes.[21] This includes the transition matrix, Π, which consists of the probabilities of a state change as detailed in Section 4.1.

Agents use the following Bayesian filter to forecast A_t given information set $\mathcal{I}_t$

$$P(A_{t-1}=A^H \mid \mathcal{I}_t) = \frac{\phi(y_{t-1}\mid A^H,\mathcal{I}_{t-1})P(A_{t-1}=A^H\mid\mathcal{I}_{t-1})}{\phi(y_{t-1}\mid A^H,\mathcal{I}_{t-1})P(A_{t-1}=A^H\mid\mathcal{I}_{t-1})+\phi(y_{t-1}\mid A^L,\mathcal{I}_{t-1})P(A_{t-1}=A^L\mid\mathcal{I}_{t-1})}, \tag{21}$$

[21] This assumption is consistent with the fact that TFP is not directly observed in the data and available estimates, e.g. the one provided by Fernald (2019), is an imperfectly cleansed version of the Solow residual which includes measurement error (Kimball et al. (2006)). TFP measures have been subject to substantial revisions and Kurmann and Sims (2019) for example document that revisions in the Fernald (2019) estimate have substantially affected the time series properties of the TFP series across different vintages.

$$[P(A_t = A^H \mid \mathcal{I}_t), P(A_t = A^L \mid \mathcal{I}_t)] = [P(A_{t-1} = A^H \mid \mathcal{I}_t), P(A_{t-1} = A^L \mid \mathcal{I}_t)]\mathbf{\Pi}. \tag{22}$$

This filter comprises a Bayesian updating formula, equation (21), and an adjustment for the possibility of a state change, equation (22), where ϕ is the normal probability density function. In equation (21) Bayes' law gives the posterior probability at time t for productivity to be in a high state in the previous period. The reciprocal posterior probability for a low state, $P(A_{t-1} = A^L|\mathcal{I}_t)$, is obtained analogously. Then, agents adjust for the possibility of a state change from period $t-1$ to t using equation (22) by multiplying the vector of posterior probabilities with the transition matrix, to obtain a prior belief about the current state of productivity. Agents can subsequently form a belief about the productivity level in the current period by multiplying the vector of priors with the vector of productivity states

$$\tilde{A}_t = [P(A_t = A^H \mid \mathcal{I}_t), P(A_t = A^L \mid \mathcal{I}_t)][A^H, A^L]'. \tag{23}$$

Note that for agents to compute the risk-free rate and the return on equity they need to form expectations about several variables in period $t+1$. To do so, they need to estimate the probability that productivity will be in the high or low state in $t+1$, given their beliefs about the current state of productivity, $P(A_t = A^H \mid \mathcal{I}_t)$ and $P(A_t = A^L \mid \mathcal{I}_t)$. Then they multiply the vector of prior probabilities with the transition matrix,

$$[P(A_{t+1} = A^H \mid \mathcal{I}_t), P(A_{t+1} = A^L \mid \mathcal{I}_t)] = [P(A_t = A^H \mid \mathcal{I}_t), P(A_t = A^L \mid \mathcal{I}_t)]\mathbf{\Pi}, \tag{24}$$

so that agents employ part of the learning technology analogously to the case described above.

4.5 Equilibrium and social planner problem

Equilibrium. At the end of each period the equilibrium in the decentralized economy presented above is a sequence of quantities $\{c_t, l_t^s, l_t^d, i_t, d_t, k_t, s_t^d, s_t^d\}_{t=0}^{\infty}$ and prices $\{w_t, p_t\}_{t=0}^{\infty}$, given k_0, s_0, and A_0, such that the problem of firms is solved, the problem of households is solved, the markets for goods, labour and firm's shares clear

$$y_t = i_t + c_t, \qquad l_t^s = l_t^d = l_t, \qquad s_t^s = s_t^d = s_t.$$

Social Planner Problem. The decentralized economy has a social planner analogue which can be solved in a recursive fashion. At the beginning of each period the planner maximizes the utility of the representative household, equation (13), subject to the capital accumulation constraint (7), the aggregate resource constraint $\tilde{y}_t = i_t + \tilde{c}_t$ (which is the combination of the households' budget constraint (12) and the firms' investment financing constraint (5)), and the interiority conditions $c_t \geq 0$, $\tilde{c}_t \geq 0$, $0 \leq l_t \leq 1$, $k_t \geq 0$ and $i_t \geq 0$.

The benevolent planner enters the period with knowledge about two state variables: the capital stock, k_t, and a belief about current period's productivity, $\tilde{A}_t$. The belief is established by using the Bayesian updating mechanism in equations (21)-(23). Given these state variables, the planner chooses hours worked, l_t, and investment, i_t, which then implies beliefs for the levels of output, $\tilde{y}_t$, and consumption, $\tilde{c}_t$. The planner uses this information together with the technology (24) to derive the risk-free rate, the expected return on equity and subsequently the risk premium. Then, the actual productivity shock A_t and the noise ν_t are realized, but not observed separately. The realization of these shocks implies that the planner can observe the actual level of output, y_t, which will typically differ from the belief about output, $\tilde{y}_t$. Subsequently, actual consumption, c_t is realized as a residual.

Formally, the planner solves the following Bellman equation, where V denotes the value function:

$$V(k_t, \tilde{A}_t) = \max_{l_t, i_t, k_{t+1}} \left[(1-\beta)(\tilde{c}_t^{\kappa}(1-l_t)^{1-\kappa})^{\frac{1-\gamma}{\theta}} + \beta(\mathbb{E}_t V^{1-\gamma}(k_{t+1}, \tilde{A}_{t+1} \,\middle|\, \mathcal{I}_t))^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}$$

$$\text{s.t. } k_{t+1} = \left(1 - \delta + \Phi\left(\frac{i_t}{k_t}\right)\right) k_t,$$

$$\tilde{c}_t = \tilde{y}_t - i_t,$$

$$i_t \geq 0, \quad k_t \geq 0, \quad \tilde{c}_t \geq 0, \quad c_t \geq 0, \quad 0 \leq l_t \leq 1, \quad \text{and } k_0,\ A_0 \text{ given},$$

where

$$\Phi\left(\frac{i_t}{k_t}\right) = \frac{b_1}{1-\kappa}\left(\frac{i_t}{k_t}\right)^{1-\kappa} + b_2 \qquad \text{and} \qquad \tilde{y}_t = \tilde{A}_t k_t^{\alpha} l_t^{1-\alpha} + \tilde{\nu},$$

and the updating rules (21)-(24) are taken as given.

The social planner equilibrium is achievable in the decentralized economy since the planner uses information that is available to all agents at no cost, the constraints and first order conditions of the planner are consistent with those of the agents, technology is convex and the preferences are insatiable.[22]

5 Calibration and computation

5.1 Calibration

Table 4 summarizes the parameter values used to calibrate the model. Consistent with the empirical sections above the model is calibrated at quarterly frequency. Several values are standard in the literature. We calibrate the share of capital in production, $\alpha = 0.36$, the discount factor, $\beta = 0.98$, and the capital depreciation rate, $\delta = 0.025$ (see e.g. Kydland and Prescott (1982)). We set the steady state labour supply to 1/3 which then implies $\kappa = 0.37$ for equation (15) to hold in steady state. The capital adjustment cost parameter is set to $\chi = 4$, consistent with the value in Jermann (1998). The two parameters related to the adjustment costs, a_1 and a_2, ensure zero capital adjustment costs in steady state and can be expressed as functions of other parameters (derivations of their functional forms are shown in Appendix B.3). Based on Caldara et al. (2012), we calibrate the degree of relative risk aversion, γ, to be 5 as our benchmark, which is also in line with the value used in Gourio (2012). Empirical evidence on the degree of relative risk aversion is scarce. For robustness

[22]It is important to note that the formulation with a social planning economy rules out agent's active experimentation. In our setup there is no feedback between actions and beliefs and learning is passive. This is a common assumption in the literature, see e.g. Van Nieuwerburgh and Veldkamp (2006), Ordoñez (2013) and Saijo (2017). Active learning would invalidate the Welfare Theorems in the social planning economy and hence there would be no decentralized counterpart to the planner's equilibrium. The passive learning is reflected in the planner's recursive problem above: the state variables, including beliefs about productivity, are determined before optimal production decisions are made, after which subsequently beliefs are updated again. This process can be repeated until beliefs about productivity coincide with its actual realization.

we verify $\gamma \in \{1, 10\}$ which does not significantly alter our results.[23]

Given the parameters above, we calibrate the elasticity of inter-temporal substitution, ψ, to be 0.01, so that the mean risk premium in the model matches the equivalent moment in the data. This calibration is also consistent with the empirical estimates in Yogo (2004) and Gomes and Paz (2011) for the elasticity of inter-temporal substitution. We use the risk premium based on the historical average measure as a benchmark to calibrate our model as it relies on a minimum of assumptions while at the same time we observe a long time series. We consider this measure at an investment horizon of one quarter, which is consistent with the setup of our model. Importantly, as detailed in Section 3.4, it dominates the other quarterly risk premium measures (and also all measures at the annual investment horizon) in terms of accuracy. While the model is calibrated to match the level of the mean risk premium in the data (0.063 vs. 0.069), it is reassuring that given the above parameters the model also delivers levels for the expected return on equity (0.103 vs. 0.088) and the risk-free rate (0.042 vs. 0.045) that are comparable to their data equivalents.[24, 25]

The model's learning technology relies on three parameters that require calibrating. We set the states of the two-state Markov chain to be $A^H = 1 + 0.032$ and $A^L = 1 - 0.032$ so that the standard deviation of the technology process, σ_A, is consistent with the findings in Cooley and Prescott (1995) and Fernald (2019) based on estimates of Solow residuals. In fact, σ_A conforms exactly to the corresponding moment in Fernald's quarterly series for our sample. Note that the distance between the two states matters for the volatility of the process, but since we evaluate deviations from the steady state, the absolute level of the technology process is not important. Let p_{ij} denote the probability for a change from state $i = \{H, L\}$ to state $j = \{H, L\}$, then the ergodicity of the Markov chain implies $p_{ij} \in (0, 1)$ and $p_{iH} + p_{iL} = 1$. In combination with the symmetry assumption on the transition matrix

[23] These results are available upon request.

[24] The difference between statistics for the expected return on equity and the risk-free rate do not exactly match the ones provided for the risk premium. The reason is that the risk premium is computed for every period before the average is taken across all simulations.

[25] Even though we haven't explicitly targeted these moments, the model also generates risk premia with a standard deviation (0.01, standard error 0.00) and kurtosis (3.86, standard error 0.03) that are consistent with the moments observed in the data (0.01 and 4.60). It is further reassuring that the the correlation between the risk premium and the risk free rate in the model (-0.39, standard error 0.01) is consistent with the empirical counterpart (-0.28). We used the 3-month Treasury Bill Rate as measure for the risk free rate.

this implies $p_{LH} = p_{HL}$ and $p_{HH} = p_{LL} = (1 - p_{LH})$. Hence, the autocorrelation for technology can be pinned down by the probability of a state change, p_{LH}, which we set to 0.05. This implies an auto-correlation for productivity of 0.95, which is consistent with the estimate in Cooley and Prescott (1995), and gives an autocorrelation for output in the model (0.93) that is in line with the corresponding statistic in the data (0.84). Finally, we calibrate the standard deviation of the noise shock to be $\sigma_\nu = 0.01$, so that our model matches the negative correlation between the median absolute nowcast error for GDP growth and real GDP growth in the data.[26] The variance of the noise shock affects the signal-to-noise ratio and thereby determines the speed of learning. If the volatility of the noise shock is too large, it becomes impossible to extract any information from the signal received. If the volatility of the noise shock is too small, it becomes straightforward to infer real productivity and learning is trivial. Our value for σ_ν is between these extreme cases so that learning is neither impossible nor trivial.

5.2 Computational details

We solve the model using Value Function Iteration. Epstein and Zin (1989) show that a version of the contraction mapping theorem still holds with recursive preferences. The algorithm requires the choice of two grids, for hours and for capital. We use 1000 grid points for capital and 500 grid points for hours. The upper and lower bounds of the grids are equal to 125% and 75% of the respective steady state values of the variables. These values ensure that the choices of the representative agent are not constrained by the boundaries, while maintaining a high grid density for precision of the solution. During simulations we do not visit the grid points at the boundaries of the state space. Consumption and the belief about consumption do not require a specific grid, as their values can be recovered using the grids for capital and labour. We use the policy functions to simulate 500 time series of 248 quarters after 50 periods are discarded. This is consistent with the length of the time horizon in the

[26] The nowcast data is from the Survey of Professional Forecasters (SPF). The SPF provides quarterly nowcasts over a horizon 1968Q4-2019Q2. The nowcasts are on GNP growth, up to 1991Q4, and GDP growth, from 1992Q1. Throughout the paper we compute nowcast errors using the corresponding series for GNP and GDP growth from the Bureau of Economic Analysis.

Table 4: Calibrated Parameters

Description	Parameter	Value
Income share of capital	α	0.36
Discount factor	β	0.98
Depreciation rate of capital	δ	0.025
Probability of state change in transition matrix	p_{LH}	0.05
Standard deviation of productivity shock	σ_A	0.032
Standard deviation of noise shock	σ_ν	0.01
Relative risk aversion	γ	5
Elasticity of inter-temporal substitution	ψ	0.01
Capital adjustment cost parameter	χ	4
Period utility parameter	κ	0.37

empirical sections above.

6 Asymmetries in the model

We have documented in Section 3.4 that there is substantial growth asymmetry in risk premia and Section 3.5 provides empirical evidence on procyclical variations in nowcasting precision. In this section, we show that our model can resemble this asymmetry in risk premia due to endogenous variations in agent's nowcasting precision about productivity.[27] The key for this mechanism is the formulation for output (4), which consists of the product of TFP and the function of production inputs, as well as the additive noise term. Agents employ output realized at the end of the previous period in the Bayesian learning technology to infer the current state of productivity. When production inputs are low, agents learn slowly about productivity because the noise variance is relatively large in comparison to the variance of the signal. A recession is hence a time of high uncertainty and low nowcast accuracy. During a recovery, intensified use of production inputs amplifies changes in technology so

[27] Enders et al. (2017) find that productivity shocks have a statistically and economically significant impact on nowcast errors and report evidence for Granger causality. They also investigate potential links between a variety of other non-technology shocks and nowcast errors, but cannot find significant effects of such shocks on nowcast accuracy.

that the variance of the signal increases. Given our assumption of a constant noise variance this implies a rising signal-to-noise ratio during a recovery. This decline in uncertainty raises nowcast precision so that agent's speed of learning increases with output and the risk premium declines gradually. At the peak, a situation of low uncertainty and high output, a decline in the state of productivity can be observed relatively precisely. The result is a strong negative adjustment in production inputs, an increase in uncertainty, and a sharp rise in the risk premium. Hence, procyclical fluctuations in the signal-to-noise ratio lead to endogenous variations in nowcasting accuracy which generate asymmetries in the risk premium and the other macroeconomic aggregates.[28] The empirical evidence in Section 3.5 corroborates this model mechanism.[29]

We now evaluate the model's ability to resemble the risk premium's growth asymmetry observed in the data. We report moments for the risk premium based on the historical average method at the one-quarter investment horizon as this measure has been employed to calibrate the model. Table 5 reports a selection of moments for the risk premium in the data (Panel A) and implied by the model (Panel B). Second moments are computed based on cyclical components that are extracted using the methodology in Hamilton (2018) with the improvements suggested in Quast and Wolters (2020).[30] The appropriate transformation to detect growth asymmetry, as shown in Sichel (1993), is by computing the skewness from log first-differences. The model matches the countercyclicality of the risk premium and the autocorrelation observed in the data rather well. Also the risk premium's volatility relative to output volatility is reasonably close to the statistic reported in the data. Most notable however is that the model is able to generate positive skewness in the risk premium (0.156)

[28] In Section 4.1 we assumed the variance of the noise to be constant. This assumption has been made for simplicity to keep the computational problem tractable. In principle, we can relax this assumption so that the noise variance can even rise when the use of production inputs increases. As long as it rises at a rate less than $k_t^{\alpha} l_t^{1-\alpha}$, this still guarantees a procyclical signal-to-noise ratio.

[29] Further supportive evidence that learning is faster at the peak than during recessions is provided by the following statistic: considering all business cycle peaks in our sample (starting in 1968Q4, the first period for which the SPF is available), on average the absolute median nowcast error for quarterly GDP growth from the SPF is 0.71 percentage points. One period after the peaks, on average, this compares to an absolute median nowcast error of 1.29 percentage points which is considerably higher. Evidence based on model simulations is consistent with this empirical evidence in that the median absolute nowcast error for output growth is substantially smaller at the peak than in the period after the peak.

[30] Appendix C shows that the results in Table 5 are robust to using a one-sided HP-filter.

that comes close to the one observed in the data (0.122). The positive skewness implies that increases in the risk premium are larger than decreases. Together with the observed negative correlation with output, this is consistent with the mechanism outlined above and the empirical evidence in Section 3.5: the risk premium declines gradually during a recovery and increases sharply when a recession occurs.

It is interesting to contrast this result with statistics based on a model without the learning mechanism. The difference to the baseline model is that the state of productivity is revealed at the beginning of the period. The corresponding moments are shown in Table 5, Panel D. While the baseline model can generate the empirically observed positive skewness in the risk premium rather well, the model without the learning mechanism fails to generate this asymmetry. Skewness in this model is not substantially different from zero; in fact it is slightly negative (-0.054). Concerning the risk premium statistics, this is the main difference to the baseline model with learning. The model without learning nevertheless implies a countercyclical risk premium, correctly ranks the risk premium to be more volatile than output, and generates a positive autocorrelation, albeit the latter is somewhat weaker than in the data.

It is important to note that learning in itself is not sufficient to generate the asymmetry in the risk premium. The key mechanism to generate a positive skewness is that learning varies procyclically: at the peak, when agents learn fast, a decline in the state of productivity will be observed relatively precisely which triggers a sharp rise in the risk premium. Instead, the onset of a recovery, when the signal-to-noise ratio is low, can be observed much less precisely which results in a more gradual decline in the risk premium. Under the assumption of a constant speed of learning, these changes in the risk premium would be of equal size (in absolute value). To confirm this intuition in our model, we consider a version with learning, but where learning does not vary over the business cycle. In particular, in this version of the model we hold the noise-to-signal ratio constant and equal to the median noise-to-signal ratio in the baseline model. It is evident from Panel C that, in line with the discussion above, learning with a constant signal-to-noise ratio is not sufficient to generate the empirically observed skewness. This model variant delivers a skewness for the risk premium which is close to zero.

We also report in Table 5 moments for macroeconomic aggregates. The model without learning (Panel D) corresponds to a relatively standard Real Business Cycle framework. As such, and in line with the existing literature (see e.g. Hansen (1985)), it matches the corresponding volatilities and correlations in the data reasonably well. The same holds for the models with learning (Panels B and C). Consumption is less volatile and investment more volatile than output.[31] The capital adjustment costs in our model have a long tradition in the literature since they have been found helpful for business cycle models to match features in the data. In fact, frictionless adjustment of the capital stock has been found a major weakness of such frameworks as it leads to too smooth consumption and overly volatile investment (see e.g. Jermann (1998)). A well known issue of real business cycle models is the low volatility of hours worked which is larger than the volatility of output in the data.[32] Our baseline model is able to replicate the negative skewness of macroeconomic aggregates observed in the data. The growth asymmetry in macroeconomic aggregates implies that increases are long and gradual and declines are short and sharp. This is a well documented feature of business cycles. It is consistent with the evidence in Van Nieuwerburgh and Veldkamp (2006), Görtz and Tsoukalas (2013) and Ordoñez (2013) who employ signal extraction processes similar to ours to generate asymmetries in macroeconomic aggregates. Görtz and Tsoukalas (2013) report asymmetry in these variables is a salient feature of business cycles across G7 countries and Ordoñez (2013) show it is stronger even in countries with less developed financial sectors. Given the above discussion, it is not surprising that the model without learning fails also to generate the growth asymmetry in macroeconomic aggregates. Skewness of output and consumption is close to zero. Investment and hours are only slightly negatively skewed, of about the size of one standard error, and by far not as much as in the data. It is

[31] To make learning non-trivial the variance of the noise shock needs to be large enough to disguise the true technology state. This however implies an unrealistically low autocorrelation and high volatility of output. We follow Van Nieuwerburgh and Veldkamp (2006) to resolve this conflict between learning and output volatility and report all moments for the model's output based on a filtered series given public information available at the end or the period, i.e. the persistent component of end-of-period output $\hat{y}_t = \mathbb{E}_t[A_t \mid \mathcal{I}_{t+1}]k_t^{\alpha}l_t^{1-\alpha}$. They show that $\hat{y}$ can be interpreted as revised data which is typically collected by data agencies who would like to report $y_t - \nu_t$ but don't observe the noise. For the national income accounts to balance also consumption must be filtered so that $\hat{y}_t = \hat{c}_t - i_t$.

[32] A remedy discussed in the literature can e.g. be to use the indivisible labor approach of Hansen (1985).

hence apparent that the learning mechanism is crucial to align the model outcomes with the empirically observed asymmetry in the risk premium and the macroeconomic aggregates.

7 Conclusion

The expected risk premium on equity is the expected excess return above the risk-free rate that investors require as compensation for the higher uncertainty associated with risky assets. We estimate a variety of measures for the expected risk premium on equity for the post-WWII U.S. economy based on models that have been found relevant in the literature. We document these measures exhibit growth asymmetry in the sense that increases in the risk premium are sharp and short while declines are more gradual and long. We show this positive skewness is a salient feature of risk premium growth. A real business cycle model with Epstein-Zin preferences is consistent with this fact in the data. We demonstrate that, in this model, the key mechanism to generate growth asymmetry in risk premia are procyclical variations in nowcast accuracy due to endogenous changes in the degree of uncertainty about productivity. This mechanism finds support in the data using measures for uncertainty and nowcast precision from the Survey of Professional Forecasters. In addition to matching the growth asymmetry in risk premia, the model is also successful in generating the empirically observed countercyclicality of risk premia and the negative skewness in macroeconomic aggregates.

Table 5: Key moments of the risk premium and macroeconomic aggregates

	Relative std deviation	Correlation with output	1st order auto-cor.	Skewness
Panel A: U.S. data				
Risk premium	2.229	-0.479	0.768	0.122
Output	1.000	1.000	0.867	-0.523
Investment	4.836	0.885	0.866	-0.697
Hours	1.267	0.878	0.921	-0.965
Consumption	0.747	0.873	0.880	-0.672
Panel B: Baseline model (with learning and varying noise-to-signal ratio)				
Risk premium	2.972 (0.028)	-0.511 (0.006)	0.640 (0.007)	0.156 (0.028)
Output	1.000 (0.000)	1.000 (0.000)	0.932 (0.003)	-0.591 (0.058)
Investment	2.038 (0.007)	0.735 (0.005)	0.890 (0.004)	-0.461 (0.063)
Hours	0.209 (0.001)	0.696 (0.006)	0.859 (0.004)	-0.453 (0.067)
Consumption	0.944 (0.004)	0.911 (0.001)	0.825 (0.007)	-0.137 (0.040)
Panel C: Baseline model (with learning but constant noise-to-signal)				
Risk premium	2.852 (0.080)	-0.420 (0.025)	0.206 (0.022)	0.008 (0.015)
Output	1.000 (0.000)	1.000 (0.000)	0.820 (0.024)	-0.041 (0.046)
Investment	2.104 (0.023)	0.627 (0.021)	0.810 (0.021)	-0.066 (0.054)
Hours	0.191 (0.021)	0.660 (0.028)	0.791 (0.012)	-0.083 (0.076)
Consumption	0.894 (0.024)	0.874 (0.014)	0.781 (0.019)	0.027 (0.063)
Panel D: Model without learning				
Risk premium	3.676 (0.057)	-0.456 (0.006)	0.185 (0.007)	-0.054 (0.018)
Output	1.000 (0.000)	1.000 (0.000)	0.921 (0.004)	-0.029 (0.036)
Investment	1.992 (0.005)	0.983 (0.001)	0.908 (0.004)	-0.059 (0.049)
Hours	0.205 (0.001)	0.933 (0.003)	0.883 (0.000)	-0.071 (0.054)
Consumption	0.725 (0.002)	0.989 (0.000)	0.923 (0.003)	0.042 (0.046)

Values reported in parentheses are standard errors. The sample in panel A is 1957Q3 - 2019Q2. Statistics shown for the risk premium in Panel A are based on the historical average measure with one quarter investment horizon. The models in panels B, C and D are simulated 500 times over 298 periods after which the first 50 periods are discarded. Second moments are calculated based on percentage deviations from trend. The trend is computed using the Hamilton (2018) filter with the improvements suggested in Quast and Wolters (2020). Skewness is calculated from log first-differenced series.

References

Adrian, T., Crump, R. K., and Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1):110–139.

Altug, S. and Labadie, P. (2008). *Asset pricing for dynamic economies*. Cambridge University Press, UK.

Bachmann, R., Elstner, S., and Sims, E. R. (2013). Uncertainty and economic activity: evidence from business survey data. *American Economic Journal: Macroeconomics*, 5(2):217–249.

Bai, J. and Ng, S. (2005). Tests for skewness, kurtosis, and normality for time series data. *Journal of business and economic statistics*, 23(1):49–60.

Baker, S. R., Bloom, N., and Davis, S. J. (2012). Has economic policy uncertainty hampered the recovery? *Chicago Booth Research Paper No. 12-06*.

Barsky, R. B. and Sims, E. R. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58(3):273–289.

Beaudry, P. and Portier, F. (2006). Stock prices, news, and economic fluctuations. *American Economic Review*, 96(4):1293–1308.

Bekaert, G., Engstrom, E., and Xing, Y. (2009). Risk, uncertainty, and asset prices. *Journal of Financial Economics*, 91(1):59–82.

Bekaert, G. and Harvey, C. R. (1995). Time-varying world market integration. *Journal of Finance*, 50(2):403–445.

Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives*, 28(2):153–177.

Bok, B., Caratelli, D., Giannone, D., Sbordone, A., and Tambalotti, A. (2017). Macroeconomic nowcasting and forecasting with big data. *Federal Reserve Bank of New York Staff Report No. 830*.

Boldrin, M. and Levine, D. K. (2001). Growth cycles and market crashes. *Journal of Economic Theory*, 96(1-2):13–40.

Caldara, D., Fernández-Villaverde, J., Rubio-Ramírez, J. F., and Yao, W. (2012). Computing dsge models with recursive preferences and stochastic volatility. *Review of Economic Dynamics*, 15(2):188–206.

Campbell, J. Y., Pflueger, C., and Viceira, L. M. (2019). Macroeconomic drivers of bond and equity risks. *Harvard Business School Finance Working Paper No. 14-031.*

Campbell, J. Y. and Shiller, R. J. (1988a). The dividend-price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies*, 1(3):195–228.

Campbell, J. Y. and Shiller, R. J. (1988b). Stock prices, earnings and expected dividends. *The Journal of Finance*, 43(3):661–676.

Campbell, J. Y. and Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies*, 21(4):1509–1531.

Cascaldi-Garcia, D. and Galvao, A. B. (2020). News and Uncertainty Shocks. *Journal of Money, Credit and Banking*, forthcoming.

Cascaldi-Garcia, D. and Vukotic, M. (2019). Patent-based news shocks. *The Warwick Economics Research Paper Series 1225.*

Cooley, T. F. and Prescott, E. (1995). *Frontiers of business cycle research.* Princeton University Press, Princeton, NJ.

Corradi, V., Distaso, W., and Mele, A. (2013). Macroeconomic determinants of stock volatility and volatility premiums. *Journal of Monetary Economics*, 60(2):203–220.

Damodaran, A. (2019). Equity risk premiums (erp): determinants, estimation and implications, the 2019 edition. *New York University - Stern School of Business working paper.*

Davidson, R. and MacKinnon, J. G. (2000). Bootstrap test: how many bootstraps? *Econometric Reviews*, 19(1):55–68.

Duarte, F. and Rosa, C. (2015). The equity risk premium: A review of models. *Economic Policy Review*, 21(2):39–58.

Enders, Z., Kleemann, M., and Müller, G. (2017). Growth expectations, undue optimism, and short-run fluctuations. *Bundesbank Discussion Paper No. 11/2017*.

Epstein, L. G. and Zin, S. E. (1989). Substutition, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969.

Fajgelbaum, P. D., Schaal, E., and Taschereau-Dumouchel, M. (2017). Uncertainty traps. *The Quarterly Journal of Economics*, 132(4):1641–1692.

Fama, E. F. and French, K. R. (1988). Dividend yields and expected stock returns. *Journal of Financial Economics*, 22(1):3–25.

Fama, E. F. and French, K. R. (1989). Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics*, 25(1):23–49.

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

Faust, J., Rogers, J., and Wright, J. (2005). News and noise in g-7 gdp announcements. *Journal of Money, Credit and Banking*, 37(3):403–419.

Fernald, J. G. (2012 - 2019). A quarterly, utilization-adjusted series on total factor productivity. *FRBSF working paper series 2012-19*.

Ferson, W. E. and Harvey, C. R. (1991). The variation of economic risk premiums. *Journal of Political Economy*, 99(2):385–416.

Ghysels, E., Plazzi, A., and Valkanov, R. (2016). Why invest in emerging markets? the role of conditional return asymmetry. *The Journal of Finance*, 71(5):2145–2192.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 109(4):1692–1720.

Gomes, F. A. R. and Paz, L. S. (2011). Narrow replication of yogo (2004) estimating the elasticity of intertemporal substitution when instruments are weak. *Journal of Applied Econometrics*, 26(7):1215–1216.

Görtz, C. and Tsoukalas, J. D. (2013). Learning, capital-embodied technology and aggregate fluctuation. *Review of Economic Dynamics*, 16(4):708–723.

Görtz, C., Tsoukalas, J. D., and Zanetti, F. (2021). News shocks under financial frictions. *American Economic Journal: Macroeconomics*, forthcoming.

Gourio, F. (2012). Disaster risk and business cycles. *The American Economic Review*, 102(6):2734–2766.

Gourio, F. (2013). Credit risk and disaster risk. *American Economic Journal: Macroeconomics*, 5(3):1–34.

Goyal, A. and Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 21(4):1455–1508.

Graham, J. R. and Harvey, C. (2007). The equity risk premium in january 2007: Evidence from the global cfo outlook survey. *Financial Risk Management*, 4:46–61.

Graham, J. R. and Harvey, C. (2018). The equity risk premium in 2018. *Duke University - Fuqa School of Business working paper*.

Hamilton, J. D. (2018). Why You Should Never Use the Hodrick-Prescott Filter. *The Review of Economics and Statistics*, 100(5):831–843.

Hansen, G. D. (1985). Indivisible labour and the business cycle. *Journal of Monetary Economics*, 16(3):309–328.

Harvey, C. R., Liu, Y., and Zhu, H. (2016). ...and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–69.

Hayashi, F. (1982). Tobin's marginal and and average q: a neoclassical interpretation. *Econometrica*, 50(1):213–224.

He, Z. and Krishnamurthy, A. (2013). Intermediary Asset Pricing. *American Economic Review*, 103(2):732–770.

Heer, B. and Maußner, A. (2012). Asset returns, the business cycle and the labor market. *German Economic Review*, 14(3):372–397.

Hodrick, R. J. (2020). An exploration of trend-cycle decomposition methodologies in simulated data. Working Paper 26750, National Bureau of Economic Research.

Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.

Jermann, U. J. (1998). Asset pricing in production economies. *Journal of Monetary Economics*, 41(2):257–275.

Johri, A. and Karimzada, M. (2020). Learning Efficiency Shocks, Knowledge Capital and the Business Cycle: A Bayesian Evaluation. *Canadian Journal of Economics*, forthcoming.

Kaltenbrunner, G. and Lochstoer, L. A. (2010). Long-run risk through consumption smoothing. *Review of Financial Studies*, 23(8):3190–3225.

Kimball, M. S., Fernald, J. G., and Basu, S. (2006). Are Technology Improvements Contractionary? *American Economic Review*, 96(5):1418–1448.

Kurmann, A. and Sims, E. (2019). Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks. *Review of Economics and Statistics*, forthcoming.

Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370.

Lettau, M. and Ludvigson, S. (2001a). Consumption, aggregate wealth, and expected stock returns. *Journal of Finance*, 56(3):815–850.

Lettau, M. and Ludvigson, S. (2001b). Resurrecting the (c)capm: A cross-sectional test when risk premia are time-varying. *Journal of Political Economy*, 109(6):1238–1288.

Lettau, M. and Uhlig, H. (2000). Can habit formation be reconciled with business cycle facts? *Review of Economic Dynamics*, 3(1):79–99.

Milani, F. (2011). Expectation shocks and learning as drivers of the business cycle. *Economic Journal*, 121(552):379–399.

Ordoñez, G. (2013). The asymmetric effects of financial frictions. *Journal of Political Economy*, 121(5):844–896.

Pinter, G., Theodoridis, K., and Yates, T. (2017). Risk news shocks and the business cycle. Bank of England working papers 483, Bank of England.

Quast, J. and Wolters, M. H. (2020). Reliable real-time output gap estimates based on a modified hamilton filter. *Journal of Business & Economic Statistics*, 0(0):1–17.

Routledge, B. R. and Zin, S. E. (2010). Generalized disappointment aversion and asset prices. *The Journal of Finance*, 65(4):1303–1332.

Rozeff, M. S. (1984). Dividend yields are equity risk premiums. *Journal of Portfolio Management*, 11(1):68–75.

Saijo, H. (2017). The uncertainty multiplier and business cycles. *Journal of Economic Dynamics and Control*, 18:1–26.

Segal, G., Shaliastovich, I., and Yaron, A. (2015). Good and bad uncertainty: Macroeconomic and financial market implications. *Journal of Financial Economics*, 117(2):369 – 397.

Shiller, R. J. (2007). Understanding recent trends in house prices and home ownership. *NBER Working Paper No. 13553*.

Sichel, D. E. (1993). Business cycle asymmetry: A deeper look. *Economic Inquiry*, 31(2):224–237.

Simin, T. (2008). The poor predictive performance of asset pricing models. *The Journal of Financial and Quantitative Analysis*, 43(2):355–380.

Van Nieuwerburgh, S. and Veldkamp, L. (2006). Learning asymetries in real business cycles. *Journal of Monetary Economics*, 53(4):753–772.

Veldkamp, L. L. (2005). Slow boom, sudden crash. *Journal of Economic Theory*, 124(2):230–258.

Welch, I. and Goyal, A. (2003). Predicting the equity premium with dividend ratios. *Management Science*, 49(5):639–655.

Welch, I. and Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508.

Yogo, M. (2004). Estimating the elasticity of intertemporal substitution when instruments are weak. *The Review of Economics and Statistics*, 86(3):797–810.

Appendix

A Additional evidence on the estimated risk premia

A.1 ADF test results for realized risk premia

Table 6 shows that ADF tests overwhelmingly reject the null hypothesis of a unit root in realized ex-post risk premia at the quarterly horizon. These are computed as the ex-post difference between the stock market return and the risk-free rate. This test statistic validates the use of the historical mean method to compute the risk premium using the historical mean method.

Table 6: Quarterly risk premia

Test statistic	1% critical value	5% critical value	10% critical value
-13.401	-3.461	-2.880	-2.570
MacKinnon approximate p-value = 0.000			

A.2 Additional empirical evidence on risk premia

Alternative Skewness Measure. In this section, we provide evidence for the positive skewness of growth in risk premia based on the quantile-based measures suggested by Ghysels et al. (2016).

Ghysels et al. (2016) propose the following quantile-based measure of skewness denoted $SK_{INT}(r_{t,n})$ and defined as

$$SK_{INT}(r_{t,n}) = 6 \cdot RA_{INT}(r_{t,n}) \frac{\int_{0.5}^{1} q_\alpha(z) d\alpha}{\int_{0.5}^{1} q_\alpha^2(z) d\alpha},$$

with

$$RA_{INT}(r_{t,n}) = \frac{\int_{0.5}^{1} \{[q_\alpha(r_{t,n}) - q_{0.5}(r_{t,n})] - [q_{0.5}(r_{t,n}) - q_{1-\alpha}(r_{t,n})]\} d\alpha}{\int_{0.5}^{1} \{q_\alpha(r_{t,n}) - q_{1-\alpha}(r_{t,n})\} d\alpha},$$

where $r_{t,n}$ represents the growth rate in period t of the expected equity risk premium with an investment horizon of n quarters. The expression $q_\alpha(r_{t,n})$ is quantile α of the unconditional distribution of $r_{t,n}$, and $q_\alpha(z)$ is quantile α of $z \sim N(0,1)$.

Table 7 shows that this quantile-based statistic corroborates our findings in Section 3.4: all risk premium measures exhibit positive growth skewness.

Table 7: Ghysels et al. (2016) quantile-based skewness statistics for growth in different measures of risk premia (SK_{INT}).

Historical average	Time series (fundamental = CAPE)	Time series (fundamental = CAPD)	Survey
0.15 (0.049)	0.14 (0.072)	0.26 (0.000)	0.43 (0.008)

Notes. The sample is 1957Q3-2019Q2 for the historical mean and time series methods. Survey data are available only from 2004Q1 to 2019Q2. "Historical average", refers to the expectations obtained using the historical average method. "Time series (fundamental = CAPE)" and "Time series (fundamental = CAPD)" refer to the expectations obtained using the time series regression method, using the CAPE and CAPD ratios respectively as fundamentals. "Survey" refers to a risk premium measure based on the Duke CFO Global Business Outlook Survey. The investment horizon is one year for the survey measure and one quarter for all other measures. Skewness statistics are calculated from the first difference of the logarithm of the risk premium. P-values, in parenthesis, are derived from a truncated Monte Carlo significance test procedure as in Davidson and MacKinnon (2000).

Time Series of Risk Premia at One-Quarter Investment Horizon. Figure 1 shows the time series of the one-quarter investment horizon risk premium estimates. In particular, it shows risk premia based on the historical average method (grey line), the CAPD time series method (blue line) and the CAPE time series method (orange line). It is evident from the figure, and in line with the statistics reported in Section 3.4, that all three measures are positively correlated.

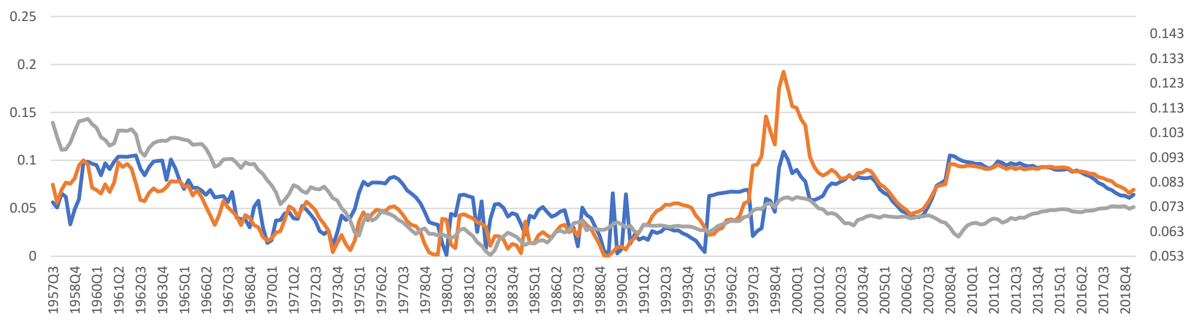

Figure 1: Time series of equity risk premium estimates. Historical average method (grey), CAPD time series method (blue) and CAPE time series method (orange) at one quarter investment horizons. The right y-axis refers to the measures using the historical average method and the left y-axis to the measures based on time series methods.

B Model derivations

B.1 Household's optimality conditions

The recursive structure of the utility function immediately implies the Bellman equation

$$F(s_t^d, \tilde{A}_t) = \max_{s_{t+1}^d, l_t^s, \tilde{c}_t} W\left(u(\tilde{c}_t, l_t^s), \mu_t\right),$$

where

$$W\left(u(\tilde{c}_t, l_t^s), \mu_t\right) = \left[(1-\beta)u(\tilde{c}_t, l_t^s)^{\frac{1-\gamma}{\theta}} + \beta\mu_t^{\frac{1-\gamma}{\theta}}\right]^{\frac{\theta}{1-\gamma}},$$

with

$$\mu_t = \left(\mathbb{E}_t\left[F_{t+1}^{1-\gamma} \mid \mathcal{I}_t\right]\right)^{\frac{1}{1-\gamma}}$$

and $\theta := \frac{1-\gamma}{1-\frac{1}{\psi}}$, subject to

$$\tilde{c}_t = w_t l_t^s + \tilde{d}_t s_t^d - p_t(s_{t+1}^d - s_t^d).$$

To ease notation in the following algebraic derivations, we use the simplified notations u_t to denote the period utility function $u(\tilde{c}_t, l_t)$, W_t to denote the CES aggregation $W\left(u(\tilde{c}_t, l_t^s), \mu_t\right)$, and F_t to denote the value function $F(s_t^d, \tilde{A}_t)$.

B.1.1 Derivation of the Lucas equation and the SDF

The first order condition with respect to s_{t+1} yields

$$\frac{\partial W_t}{\partial u_t}\frac{\partial u_t}{\partial \tilde{c}_t}\frac{\partial \tilde{c}_t}{\partial s_{t+1}^d} + \frac{\partial W_t}{\partial \mu_t}\frac{\partial \mu_t}{\partial \mathbb{E}_t[F_{t+1}^{1-\gamma} \mid \mathcal{I}_t]}\mathbb{E}_t\left[\frac{\partial F_{t+1}^{1-\gamma}}{\partial F_{t+1}}\frac{\partial F_{t+1}}{\partial s_{t+1}^d}\,\middle|\, \mathcal{I}_t\right] = 0.$$

The Envelope theorem for s_t yields

$$\frac{\partial F_t}{\partial s_t^d} = \frac{\partial W_t}{\partial u_t}\frac{\partial u_t}{\partial \tilde{c}_t}\frac{\partial \tilde{c}_t}{\partial s_t^d}.$$

Combining both conditions, we obtain

$$\frac{\partial W_t}{\partial u_t}\frac{\partial u_t}{\partial \tilde{c}_t}\frac{\partial \tilde{c}_t}{\partial s^d_{t+1}} + \frac{\partial W_t}{\partial \mu_t}\frac{\partial \mu_t}{\partial \mathbb{E}_t[F_{t+1}^{1-\gamma} \mid \mathcal{I}_t]}\mathbb{E}_t\left[\frac{\partial F_{t+1}^{1-\gamma}}{\partial F_{t+1}}\frac{\partial W_{t+1}}{\partial u_{t+1}}\frac{\partial u_{t+1}}{\partial \tilde{c}_{t+1}}\frac{\partial \tilde{c}_{t+1}}{\partial s^d_{t+1}} \,\middle|\, \mathcal{I}_t\right] = 0. \tag{B.1}$$

We can then recover

$$\frac{\partial W_t}{\partial u_t} = F_t^{1-\frac{1-\gamma}{\theta}}(1-\beta)u_t^{\frac{1-\gamma}{\theta}-1},$$

$$\frac{\partial W_t}{\partial \mu_t} = \beta F_t^{1-\frac{1-\gamma}{\theta}}\left(\mathbb{E}_t[F_{t+1}^{1-\gamma} \mid \mathcal{I}_t]\right)^{\frac{1}{\theta}-\frac{1}{1-\gamma}},$$

$$\frac{\partial \mu_t}{\partial \mathbb{E}_t[F_{t+1}^{1-\gamma} \mid \mathcal{I}_t]} = \frac{1}{1-\gamma}\mathbb{E}_t[F_{t+1}^{1-\gamma} \mid \mathcal{I}_t]^{\frac{1}{1-\gamma}-1},$$

$$\frac{\partial \tilde{c}_{t+1}}{\partial s^d_{t+1}} = \mathbb{E}_t[\tilde{d}_{t+1} + p_{t+1} \mid \mathcal{I}_t],$$

and

$$\frac{\partial \tilde{c}_t}{\partial s^d_{t+1}} = -p_t.$$

Plugging the last 5 equations into equation (B.1) yields after re-arranging

$$0 = \mathbb{E}_t\left[m_{t+1,t}\frac{\tilde{d}_{t+1} + p_{t+1}}{p_t} - 1 \,\middle|\, \mathcal{I}_t\right],$$

where

$$m_{t+1,t} = \beta\left(\frac{F_{t+1}^{1-\gamma}}{\mathbb{E}_t\left[F_{t+1}^{1-\gamma} \mid \mathcal{I}_t\right]}\right)^{1-\frac{1}{\theta}}\left(\frac{u_{t+1}}{u_t}\right)^{\frac{1-\gamma}{\theta}-1}\frac{\frac{\partial u_{t+1}}{\partial \tilde{c}_{t+1}}}{\frac{\partial u_t}{\partial \tilde{c}_t}}.$$

These are the Lucas equation (16) and the stochastic discount factor (17).

B.1.2 Optimal labour supply

The first order condition with respect to l_t yields

$$\frac{\partial W_t}{\partial u_t}\left[\frac{\partial u_t}{\partial \tilde{c}_t}w_t + \frac{\partial u_t}{\partial l^s_t}\right] = 0. \tag{B.2}$$

We can then recover

$$\frac{\partial u_t}{\partial \tilde{c}_t} = \kappa\tilde{c}_t^{\kappa-1}(1-l^s_t)^{1-\kappa}$$

and

$$\frac{\partial u_t}{\partial l_t^s} = -\tilde{c}_t^{\kappa}(1-\kappa)(1-l_t^s)^{-\kappa}.$$

Plugging the last 2 equations into equation (B.2) yields after re-arranging

$$\tilde{c}_t = \frac{\kappa}{(1-\kappa)}(1-l_t^s)w_t,$$

which is the labor supply function (15).

B.2 Derivation of the return on equity

Firm's expected cash flow at the beginning of the period is defined as

$$\tilde{f}_t = \tilde{d}_t s_t - p_t(s_{t+1} - s_t) = \tilde{y}_t - w_t l_t - i_t.$$

From equation (11), it holds that $w_t l_t = (1-\alpha)\tilde{y}_t$. Hence we can simplify the above equation for expected cash flow to become

$$\tilde{f}_t = \alpha\tilde{y}_t - i_t.$$

Using equations (10) and (7), and due to the specific capital adjustment costs we apply, we can write

$$\begin{aligned}
q_t k_{t+1} =& \mathbb{E}_t \left\{ m_{t,t+1} \left[\tilde{A}_{t+1} l_{t+1}^{1-\alpha} \alpha k_{t+1}^{\alpha} - i_{t+1} + q_{t+1} \left(1 - \delta + \Phi\left(\frac{i_{t+1}}{k_{t+1}} \right) \right) k_{t+1} \right] \middle| \mathcal{I}_t \right\} \\
\Leftrightarrow q_t k_{t+1} =& \mathbb{E}_t \left[m_{t,t+1}(\alpha y_{t+1} - i_{t+1} + q_{t+1} k_{t+2}) \mid \mathcal{I}_t \right] \\
\Leftrightarrow q_t k_{t+1} =& \mathbb{E}_t \left[m_{t,t+1}(f_{t+1} + q_{t+1} k_{t+2}) \mid \mathcal{I}_t \right].
\end{aligned}$$

Iterating forward, we obtain

$$q_t k_{t+1} = \mathbb{E}_t \left[\sum_{i=1}^{+\infty} m_{t,t+i} f_{t+i} \middle| \mathcal{I}_t \right], \tag{B.3}$$

assuming that $\lim_{i\to+\infty} \mathbb{E}_t[m_{t,t+i}q_{t+i}k_{t+i+1} \mid \mathcal{I}_t] = 0$. Following Altug and Labadie (2008), the value of a firm on the stock market is equal its present value of future discounted cash flow. This allows us to rewrite equation (B.3) as

$$q_t k_{t+1} = p_t s_{t+1}.$$

Finally, using the above expressions, we can derive a formulation for the return on equity which depends on variables that have been pinned down uniquely in firm's and household's maximization problems

$$\begin{aligned}
\mathbb{E}_t\left[\frac{\tilde{d}_{t+1}+p_{t+1}}{p_t}\,\middle|\,\mathcal{I}_t\right] =& \mathbb{E}_t\left[\frac{s_{t+1}\tilde{d}_{t+1}-p_{t+1}(s_{t+2}-s_{t+1})+s_{t+2}p_{t+1}}{s_{t+1}p_t}\,\middle|\,\mathcal{I}_t\right] \\
=& \mathbb{E}_t\left[\frac{f_{t+1}+k_{t+2}q_{t+1}}{k_{t+1}q_t}\,\middle|\,\mathcal{I}_t\right] \\
=& \mathbb{E}_t\left[\frac{q_{t+1}k_{t+2}+y_{t+1}-w_{t+1}l^d_{t+1}-i_{t+1}}{q_t k_{t+1}}\,\middle|\,\mathcal{I}_t\right],
\end{aligned}$$

which is equation (20) in the main body.

B.3 Derivation of functional forms for parameters a_1 and a_2

The parameters a_1 and a_2 are calibrated to ensure that adjustment costs are zero in steady state, so that steady state investment and Tobin's q are $i = \delta k$ and $q = 1$. From equations (7) and (9), we can see that the latter is satisfied if

$$\Phi(\delta) = \delta \ \text{ and } \ \Phi^{'}(\delta) = 1.$$

Given the functional form of Φ, this implies

$$\frac{a_1}{1-\chi}\delta^{1-\chi} + a_2 = \delta \ \text{ and } \ a_1\delta^{-\chi} = 1,$$

from where we deduce

$$a_1 = \delta^{\chi} \ \text{ and } \ a_2 = -\frac{\delta\chi}{1-\chi}.$$

C Additional model-based results

Table 8 corresponds to Table 5 in the main body and reports statistics for second moments that are computed based on the cyclical components extracted using a one-sided HP(1600) filter instead of the Hamilton (2018) filter with the improvement suggested by Quast and Wolters (2020). The discussion about the adequacy and performance of these filters is ongoing. Based on simulated data, Hodrick (2020) points out that for basic time series the Hamilton (2018) filter dominates the HP-filter, but in more complex models the reverse is true. It is reassuring that our reported results are not substantially affected by the choice of filtering methodology. We find statistics across second moments are extremely similar across the two tables.

Table 8: Key moments of the risk premium and macroeconomic aggregates

	Relative std deviation	Correlation with output	1st order auto-cor.	Skewness
Panel A: U.S. data				
Risk premium	2.177	-0.486	0.772	0.122
Output	1.000	1.000	0.836	-0.523
Investment	4.373	0.898	0.817	-0.697
Hours	1.226	0.854	0.909	-0.965
Consumption	0.795	0.872	0.862	-0.672
Panel B: Baseline model (with learning and varying noise-to-signal)				
Risk premium	2.972 (0.028)	-0.511 (0.006)	0.640 (0.007)	0.156 (0.028)
Output	1.000 (0.000)	1.000 (0.000)	0.932 (0.003)	-0.591 (0.058)
Investment	2.038 (0.007)	0.735 (0.005)	0.890 (0.004)	-0.461 (0.063)
Hours	0.209 (0.001)	0.696 (0.006)	0.859 (0.004)	-0.453 (0.067)
Consumption	0.944 (0.004)	0.911 (0.001)	0.825 (0.007)	-0.137 (0.040)
Panel C: Model with learning but constant noise-to-signal				
Risk premium	2.854 (0.081)	-0.419 (0.025)	0.206 (0.024)	0.008 (0.015)
Output	1.000 (0.000)	1.000 (0.000)	0.822 (0.024)	-0.041 (0.046)
Investment	2.101 (0.023)	0.627 (0.022)	0.812 (0.021)	-0.066 (0.054)
Hours	0.188 (0.018)	0.661 (0.028)	0.791 (0.013)	-0.083 (0.076)
Consumption	0.895 (0.024)	0.876 (0.015)	0.783 (0.019)	0.027 (0.063)
Panel D: Model without learning				
Risk premium	3.676 (0.057)	-0.456 (0.006)	0.185 (0.007)	-0.054 (0.018)
Output	1.000 (0.000)	1.000 (0.000)	0.921 (0.004)	-0.029 (0.036)
Investment	1.992 (0.005)	0.983 (0.001)	0.908 (0.004)	-0.059 (0.049)
Hours	0.205 (0.001)	0.933 (0.003)	0.883 (0.000)	-0.071 (0.054)
Consumption	0.725 (0.002)	0.989 (0.000)	0.923 (0.003)	0.042 (0.046)

Values reported in parentheses are standard errors. The sample in panel A is 1957Q3 - 2019Q2. Statistics shown for the risk premium in Panel A are based on the historical average measure with one quarter investment horizon. The models in panels B, C and D are simulated 500 times over 298 periods after which the first 50 periods are discarded. Second moments are calculated based on percentage deviations from (one-sided) HP(1600) filter trend. Skewness is calculated from log first-differenced series.