DEUTSCHE BUNDESBANK
EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 16/2019

Extreme inflation and time-varying consumption growth

Ilya Dergunov
(Goethe University Frankfurt)

Christoph Meinerding
(Deutsche Bundesbank)

Christian Schlag
(Goethe University Frankfurt)

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet http://www.bundesbank.de

ISBN 978–3–95729–583–5 (Printversion)
ISBN 978–3–95729–584–2 (Internetversion)

Non-technical summary

Research Question

Whether and how changes in expected inflation affect expectations concerning the real economy has always been an intensely debated question. In this paper, we provide a new perspective on this issue by embedding inflation in one of the workhorse models of real equilibrium asset pricing, namely the long-run risk framework. This class of models has been widely applied in the literature, but has rarely been extended towards the pricing of nominal assets.

Results

We find that, when the long-run risk model is augmented by inflation, time variation in expected consumption growth can explain time variation in the stock-bond return correlation. Key to this finding is the empirical observation that extreme expected inflation – both very high and very low – is linked to low average consumption growth. Depending on the state of the economy, an increase in expected inflation can be either a good or a bad signal for real expected consumption growth, and this signalling channel drives changes in the return correlation between real and nominal assets. As a validation of our channel, we also find that the long-run risk variable that we extract exclusively from macro data via the estimation of our model is highly correlated with long-run risk proxies suggested in the literature, with the important difference that those are usually extracted from asset price data.

Contribution

We show that one of the workhorse real equilibrium asset pricing models can explain the stock-bond return correlation when we properly condition on inflation as a signal for expected consumption growth. The probability of being in a low expected consumption growth state is closely linked to implicit long-run risk variables in the literature that have been extracted from asset prices. All the aforementioned results are derived from macro data only. This is key for our analysis, since asset price data have the potential to severely confound a macro estimation by imposing the strong parametric structure of an asset pricing model.

Nichttechnische Zusammenfassung

Fragestellung

Ob und wie sich Änderungen der erwarteten Inflation auf realwirtschaftliche Erwartungen auswirken, ist nach wie vor eine umstrittene Frage. In diesem Papier beleuchten wir dieses Thema aus einer neuen Perspektive, indem wir Inflation in ein Asset-Pricing-Gleichgewichtsmodell mit langfristigen Konsumrisiken einbetten. Diese Klasse von Asset-Pricing-Modellen gehört inzwischen zum Standard in der Literatur, ist aber bislang kaum auf die Bewertung von nominalen Wertpapieren angewendet worden.

Ergebnisse

Wir zeigen, dass die zeitliche Variation im erwarteten Konsumwachstum die zeitliche Variation in der Korrelation zwischen Aktien- und Anleihenrenditen erklären kann, sofern das reale Standard-Modell um Inflation erweitert wird. Der entscheidende Schritt ist hierbei die empirische Beobachtung, dass extreme erwartete Inflation - sowohl sehr hohe als auch sehr niedrige - mit einem niedrigen durchschnittlichen Konsumwachstum einhergeht. Abhängig vom Zustand der Volkswirtschaft kann ein Anstieg der erwarteten Inflation damit entweder als positives oder als negatives Signal für das reale erwartete Konsumwachstum interpretiert werden. Dies führt zu Zeitvariation in der Korrelation der Renditen von realen und nominalen Wertpapieren. Um diesen Mechanismus zu validieren, zeigen wir, dass die Variable für langfristiges Konsumrisiko, die wir ausschließlich aus Makrodaten durch die Schätzung unseres Modells extrahieren, in hohem Maße mit ähnlichen in der Literatur vorgeschlagenen Variablen korreliert, mit dem wichtigen Unterschied, dass jene normalerweise aus Wertpapierpreisdaten extrahiert werden.

Beitrag

Wir zeigen, dass eines der wichtigsten realen Asset-Pricing-Gleichgewichtsmodelle zeitliche Variation in der Korrelation von Aktien- und Anleiherenditen erklären kann, wenn wir Inflation als Signal für erwartetes Konsumwachstum angemessen berücksichtigen. Die Wahrscheinlichkeit eines niedrigen erwarteten Konsumwachstums ist eng mit impliziten Variablen für langfristiges Konsumrisiko verknüpft, die in der Literatur bisher zumeist aus Wertpapierpreisen extrahiert wurden. Unsere Schätzung basiert ausschließlich auf makroökonomischen Daten. Dies ist für unsere Analyse von zentraler Bedeutung, da Asset-Preisdaten das Potenzial haben, die makroökonomischen Implikationen eines Asset-Pricing-Modells stark zu beeinträchtigen, wenn zu strenge parametrische Annahmen getroffen werden.

Deutsche Bundesbank Discussion Paper No 16/2019

Extreme Inflation and Time-Varying Consumption Growth*

Ilya Dergunov
Goethe University Frankfurt

Christoph Meinerding
Deutsche Bundesbank

Christian Schlag
Goethe University Frankfurt

Abstract

In a parsimonious regime switching model, expected consumption growth varies over time. Adding inflation as a conditioning variable, we uncover two states in which expected consumption growth is low, one with high and one with negative expected inflation. Embedded in a general equilibrium asset pricing model with learning, these dynamics replicate the observed time variation in stock return volatilities and stock-bond return correlations. Furthermore, they provide an alternative way to come up with a measure of time-varying disaster risk in the spirit of Wachter (2013). Our findings imply that both the disaster and the long-run risk paradigm can be extended towards explaining movements in the stock-bond return correlation.

Keywords: Long-run risk, inflation, recursive utility, filtering, disaster risk

JEL classification: E31, E44, G12.

*Contact address: Deutsche Bundesbank, Research Centre, Wilhelm-Epstein-Str. 14, 60431 Frankfurt am Main, Germany. E-mail: christoph.meinerding@bundesbank.de. We would like to thank the seminar participants and colleagues at Goethe University, Deutsche Bundesbank, Wharton, BI Oslo, IWH Halle, SNB and OeNB as well as the participants of the WFA 2018 Coronado, SGF 2018, German Economists Abroad Conference 2017, Paris December Finance Meeting 2017, DGF 2017, ESSFM Gerzensee 2017, MFA 2017 , AFFI 2017, and the 2017 Colloquium on Financial Markets in Cologne for their comments and suggestions. Special thanks go to Angela Abbate (discussant), Geert Bekaert, Fernando Duarte, Stefan Eichler (discussant), Philipp Illeditsch, Jan Kragt (discussant), Mattia Landoni, Sydney Ludvigson, Emanuel Mönch, Olesya Grishchenko (discussant), Philippe Mueller, Athanasios Orphanides, Marcel Priebsch, Dongho Song (discussant), Marta Szymanowska, Jules van Binsbergen, Jessica Wachter (discussant), Todd Walker, Rüdiger Weber (discussant), Amir Yaron, and to Franziska Schulte and Iliya Karaivanov for excellent research assistance. A previous version of this paper was circulated under the title "Extreme inflation and time-varying disaster risk". We gratefully acknowledge research and financial support from the Research Center SAFE, funded by the State of Hessen initiative for research LOEWE. This work represents the authors' personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank.

1 Introduction

Whether and how changes in expected inflation affect expectations concerning the real economy has always been an intensely debated question. In this paper, we provide a new perspective on this issue by embedding inflation in one of the workhorse models of equilibrium asset pricing, namely the long-run risk framework. We find that, when a regime-switching model featuring long-run risk is augmented by inflation, time variation in expected consumption growth can explain time variation in the stock-bond return correlation and in aggregate stock market volatility.

Key to this finding is the empirical observation that extreme expected inflation – both very high and very low – is linked to low average consumption growth. Including inflation in a parsimonious regime switching model thus significantly affects the estimated dynamics of the conditional mean of consumption growth relative to a model based on consumption only. We embed the estimated regime-switching dynamics in an otherwise standard equilibrium asset pricing model with learning, thereby extending the long-run risk framework towards explaining the joint dynamics of real and nominal assets. This very stylized model allows us to match, among other things, the empirically observed time-varying nature of the return correlation between stocks and nominal bonds. As a further validation of our channel, we also find that the long-run risk variable that we extract (exclusively) from macro data via the estimation of the regime-switching model is highly correlated with long-run risk proxies suggested in the literature, with the important difference that those are usually reverse-engineered from asset price data.

Long-run risk asset pricing models rest on the key assumption that the conditional distribution of consumption growth, in particular its mean, is time-varying. Following the seminal publication of Bansal and Yaron (2004), a lot of papers dealt with the issue of detecting such time variation.[1] In this paper, we take a step back and document the existence of long-run risk by fitting a parsimonious regime-switching model for consumption growth to standard quarterly NIPA aggregate consumption data. We show that the hypothesis of constant expected consumption growth can be rejected at any conventional significance level.

Based on this, our paper then makes the following four key contributions. First, augmenting the time series model by inflation as a second variable, we detect two states in which expected consumption growth is low, one with very high expected inflation (so-called "stagflation") and one with negative expected inflation ("deflation"). The fact that there are *two* very different states with low expected consumption growth turns out to be important for asset pricing.

Second, by embedding the estimated dynamics in a standard general equilibrium asset pricing model with recursive preferences and learning, we show that imperfect information about expected consumption growth drives time variation in aggregate stock market volatility and in the stock-bond correlation when we properly condition on inflation as a signal.

Third, the probability of being in a low expected consumption growth state that we

[1] A nice synopsis of arguments against or in favor of long-run risk is given by the two competing papers of Beeler and Campbell (2012) and Bansal, Kiku, and Yaron (2012), respectively. Recent advances were made by Schorfheide, Song, and Yaron (2018), who document and analyze the persistent component in expected consumption growth employing sophisticated Bayesian mixed frequency techniques.

obtain from our macro estimation with consumption and inflation data tracks the historical price-earnings ratio of US equity well. It is therefore closely linked to implicit long-run risk variables that have been obtained through reverse engineering from asset prices in the literature. As an example, we provide an alternative derivation and interpretation of the "time-varying disaster risk" presented in Wachter (2013), which is one prime example of such a variable.

Finally, we wish to emphasize that we produce all the aforementioned results within an rather simple setup, where no asset price data are used in the estimation. This is key for our analysis, since asset price data have the potential to severely confound a macro estimation, e.g., via certain moment conditions to represent the parametric structure of an asset pricing model included in an application of GMM.

We now give some more details on our contributions. The Markov regime switching model is a simple and convenient representation of the case of time-varying expected consumption growth. The estimation gives two states for the conditional mean, and a Wald test rejects the hypothesis of the conditional means being equal at any conventional significance level. We then add inflation as a second variable, and in this extended model, expected inflation and expected consumption growth are related in a nonlinear way: both extremely high and extremely low expected inflation are coupled to low expected consumption growth, and in particular the possibility of negative inflation is an important driver of many of our results.

The recent literature on long-run risk in asset pricing, e.g. Wachter (2013), puts an emphasis on the link between time variation in the conditional distribution of consumption growth and time variation in second moments of returns. Motivated by the time series results above, we embed the fundamental dynamics into a state-of-the-art equilibrium asset pricing model, where the representative investor is equipped with recursive preferences. Given the usual values for risk aversion and the elasticity of intertemporal substitution, the substitution effect dominates, so that prices are decreasing in the amount of aggregate risk.

In line with papers like Detemple (1986), or Croce, Lettau, and Ludvigson (2012), we make the assumption that the representative investor knows the structural parameters of the model, but cannot observe the true state and thus can only filter the respective probabilities from the data. Given our estimation, extreme (high or low) inflation observations serve as a useful signal, which allows to better infer the time-varying probability of very low or even negative consumption growth. The fact that low real growth can be linked to high or low expected inflation is key to producing a time-varying stock-bond return correlation. When very high expected inflation occurs together with low expected real growth, both bonds and stocks will tend to have negative returns, resulting in a positive correlation. On the other hand, low growth expectations in periods with low expected inflation will lead to negative equity returns, but positive bond returns, implying a negative correlation. Stated differently, the long-run risk paradigm can be extended towards the time-varying nature of the stock-bond return correlation when the signaling role of inflation is taken into account properly.

We test the asset pricing implications of our model by feeding it with the empirical time series of observed consumption growth and inflation and then pricing stocks and bonds using our model-implied pricing kernel. We then compare the dependence of stock return volatilities and stock-bond correlations on the filtered probabilities in the model and the data. In both model and real data the stock-bond return correlation is positively and significantly linked to the probability of being in the high-inflation state, while for the deflationary state we find exactly the opposite. The model also qualitatively matches two key results in Wachter (2013). First, stock market volatility is increasing in the (filtered) probability of being in a bad consumption growth state. Second, the relation between stock market volatility and this probability is nonlinear both in the model and in the data.

Finally, our approach using only macroeconomic data also offers a nice alternative derivation of the state variable "time-varying disaster probability" that Wachter (2013) has produced in her paper. The correlation between her variable and our filtered probability of being in a bad consumption growth state is 0.88, and this co-movement is mirrored in the close co-movement between the dividend-price ratio in the data and in our model. Although there is no explicit role for disaster risk in our model, we choose this time series for comparison because its interpretation as a "probability" is much closer to the basic state variables in our model than, for instance, estimated time series of mean consumption growth. We draw the tentative conclusion that the "implied disaster probability" in Wachter (2013), which is obtained through a transformation of the historical dividend-price ratio of the S&P 500, is closely linked to time variation in expected consumption growth, which we estimate from consumption and inflation data alone. This interpretation also reinforces the findings of Branger, Kraft, and Meinerding (2016) who argue that disaster risk and long-run risk are intertwined in historical data. The episodes of high marginal utility which Wachter (2013) labels as episodes of elevated disaster risk are characterized by low expected consumption growth according to our Markov chain estimation.

We close the paper with a number of robustness checks concerning both the estimation and the asset pricing model. First, we show that a Markov switching model for consumption growth only (i.e., without inflation) cannot replicate the time series of the dividend-price ratio. Second, our results do not depend on expected consumption growth being particularly low (high) in one of the two bad (good) states. A model where expected growth is constrained to be the same in the two good and in the two bad states, respectively, delivers qualitatively the same results. Third, additional analyses with subsamples, with GDP growth instead of consumption growth, or with monthly consumption data also largely confirm our findings.[2] Fourth, on the asset pricing side, we document that our two major assumptions – recursive preferences and learning from inflation observations – are both key to generating our main results. In a nested model with CRRA preferences, the regression of correlations on state variables delivers coefficients whose signs are opposite to the data. A nested model in which the state of the economy is perfectly observable has problems to generate the recently observed negative correlation in the first place.

[2] We also estimate the Markov chain model on a long sample of annual data and identify a deflationary state there as well. As is clear from the discussion above, the identification of a deflationary state is crucial to match the asset price data. In our benchmark estimation, this identification largely rests on the deflation observations during the Financial Crisis in 2008. However, there are also deflationary episodes in pre-war data, which are not contained in our benchmark sample of quarterly data starting in 1947.

2 Related Literature

Our paper contributes to and links two major strands of literature, namely the asset pricing literature about inflation as a priced risk factor and the asset pricing literature featuring long-run risk and its empirical estimation. We do not aim at giving a full review of the numerous papers in the latter field. Contributions there have been made by Bansal and Yaron (2004), Bansal, Kiku, and Yaron (2016) Beeler and Campbell (2012), Bansal et al. (2012), Constantinides and Ghosh (2011), Ortu, Tamoni, and Tebaldi (2013), and Schorfheide et al. (2018), among many others.

Concerning equilibrium asset pricing with inflation risk, we mostly build on the following papers. Bansal and Shaliastovich (2013) propose a long-run risk model with expected inflation and expected growth as risk factors and use it to explain the empirically observed predictability patters in bond and foreign exchange returns. Eraker (2008) proposes an affine jump-diffusion model with jumps in consumption volatility. Eraker, Shaliastovich, and Wang (2016) discuss a long-run risk model with inflation as a risk factor, but their focus is on differences between durable and non-durable consumption and their implications for equity and bond prices in these sectors. Ehling, Gallmeyer, Heyerdahl-Larsen, and Illeditsch (2018) consider heterogeneous agents who disagree about inflation, and the authors show that this disagreement increases yields and yield volatilities at all maturities. Burkhardt and Hasseltoft (2012) propose a model with recursive utility, inflation and long-run risk similar to ours, but, given the way in which the authors introduce inflation risk premia, the asset pricing results seem to a certain degree hardwired into the model. We consider our approach less restrictive in terms of the specification of inflation and consumption growth. Piazzesi and Schneider (2006) discuss the role of inflation as a signal about future consumption growth, but they focus on the term structure of (nominal and real) interest rates and do not address time variation in the stock-bond correlation.

Song (2017) studies an endowment economy model with recursive preferences, a regime-switching Taylor rule, and a time-varying inflation target. Campbell, Pflueger, and Viceira (2018) analyze the stock-bond correlation in a New Keynesian production economy with habit formation preferences and monetary policy regimes. Complementary to our paper, Constantinides and Ghosh (2017) assess the ability of several macroeconomic predictor variables to improve the performance of consumption-based equilibrium asset pricing models, and they find that inflation data helps to generate the (non-)predictability of price-dividend ratios by cash flow growth rates. In all these papers, however, asset price data is used to calibrate or estimate the model. Ermolov (2018) estimates an external habit model with macroeconomic data only, but his explanation of time variation in the stock-bond correlation is very different from ours. He assumes that consumption and inflation can be hit by two different shocks labeled as supply and demand shocks, whereas we rely on different regimes for expected growth rates.

David and Veronesi (2013) propose Markov switching dynamics for fundamentals, but they assume a model featuring a representative agent with time-additive CRRA preferences who suffers from bounded rationality in the form of money illusion in the spirit of Basak and Yan (2010). Since their GMM estimation relies on asset price data, it delivers quite different dynamics for the fundamentals compared to our estimation, most importantly state-independent expected consumption growth. This is because the estimation targets empirical dividend-price ratios and thus has to shut down the intertemporal substitution channel, i.e. there is no role for long-run risk in their estimated model. We show in our empirical results that the hypothesis of constant expected consumption growth can be rejected at any conventional significance level if the estimation is based on macro data only.

Boons, de Roon, Duarte, and Szymanowska (2017) provide empirical evidence for inflation risk being priced in the cross-section of stock returns. Their paper can be viewed as complementary to ours, since the market price of inflation risk estimated from the cross-section of stock returns switches sign and is linked to the stock-bond correlation in the data. It can be considered a stylized fact that the correlation between inflation and other variables can change the sign of the stock-bond correlation, and we provide a model-theoretic explanation for this result. Other papers in this area include Schmeling and Schrimpf (2011), Balduzzi and Lan (2016), Campbell, Sunderam, and Viceira (2017), Hasseltoft (2012), Ang and Ulrich (2012), and Marfe (2015), to name just a few. Baele, Bekaert, and Inghelbrecht (2010) empirically analyze the determinants of the stock-bond return comovement. Fleckenstein, Longstaff, and Lustig (2016) study the pricing of deflation risk using market prices of inflation-linked derivatives.

Through the re-interpretation of our state variables as measuring time-varying disaster risk, our paper is also related to this area of research. Rietz (1988) and Barro (2006, 2009) rationalize a high equity premium in the disaster risk framework. Extensions of their basic model have been studied by Chen, Joslin, and Tran (2012) and Julliard and Ghosh (2012), among others. Constantinides (2008) criticizes that historically consumption disasters rather unfold over several years instead of just one point in time. Similarly to the critique of Constantinides (2008), the assumption of extreme jumps is also questioned by Backus, Chernov, and Martin (2011). As a response, Branger et al. (2016) combine disaster risk and long-run risk and show that the equity premium puzzle can still be solved with multi-period disasters. Similarly, Gabaix (2012), Wachter (2013), and Tsai and Wachter (2015) analyze models with time-varying jump intensities and recursive preferences. Our results imply that the disaster risk paradigm may be extended towards an explanation of the time-varying stock-bond return correlation, when the effect of inflation on real asset prices is captured properly. Finally, in this regard, our paper may also contribute to the discussion about "dark matter" in asset prices started by Chen, Dou, and Kogan (2017) in the sense that a large fraction of this dark matter may be attributed to uncertainty about extreme inflation.

3 Fundamental Dynamics

3.1 Consumption and inflation

The two fundamental sources of risk in our model are aggregate consumption and inflation. In the baseline version without inflation, we assume that log aggregate real consumption, $\ln C$, follows the process

$$d \ln C_t = \mu^C(S_t)dt + \sigma^C dW_t^C. \tag{1}$$

W^C is a standard Wiener process, the volatility σ^C is constant. The conditional drift rate $\mu^C(S_t)$ is stochastic and follows a continuous-time Markov chain whose current state is denoted by S_t. There are n states (indexed by $i = 1, \ldots, n$), with state-dependent drifts

μ_i^C. The Markov chain transitions are governed by counting processes whose intensities are collected in the $(n \times n)$-matrix $\Lambda = (\lambda_{ij})_{i,j=1,\dots,n}$. Following the usual convention, we define the diagonal elements $\lambda_{ii} := -\sum_{j\neq i} \lambda_{ij}$ so that the rows of Λ sum to 0. In our benchmark empirical case, we will have $n = 2$.

In the full model, the joint dynamics of log aggregate real consumption and of the log price level π are given as

$$\begin{aligned} d\ln C_t &= \mu^C(S_t)dt + \sigma^C \left(\sqrt{1-\rho^2}dW_t^C + \rho dW_t^\pi\right) \\ d\pi_t &= \mu^\pi(S_t)dt + \sigma^\pi dW_t^\pi. \end{aligned} \tag{2}$$

Here W^C and W^π are the (independent) components of a standard bivariate Wiener process. The dynamics in (2) imply that the increments to $\ln C$ and to π are correlated with correlation parameter ρ. The volatilities σ^C and σ^π are assumed constant. The conditional drift rates $\mu^C(S_t)$ and $\mu^\pi(S_t)$ now follow a bivariate continuous-time Markov chain whose current state is again denoted by S_t. Keeping the rest of the notation as above, the number of states in the full model will later turn out to be $n = 4$.

We will often use the vector representation of the above dynamics, which can be written as

$$\begin{pmatrix} d\ln C_t \\ d\pi_t \end{pmatrix} = \mu(S_t)\,dt + \Sigma\,dW_t$$

with

$$\mu(S_t) = \begin{pmatrix} \mu^C(S_t) \\ \mu^\pi(S_t) \end{pmatrix}, \quad \Sigma = \begin{pmatrix} \sigma^C\sqrt{1-\rho^2} & \sigma^C\rho \\ 0 & \sigma^\pi \end{pmatrix}, \quad dW_t = \begin{pmatrix} dW_t^C \\ dW_t^\pi \end{pmatrix}.$$

3.2 Markov chain estimation

To estimate the dynamics of the fundamentals we use quarterly real consumption growth rates from NIPA and quarterly inflation rates constructed according to the Piazzesi and Schneider (2006) mechanism.[3] Our sample period ranges from 1947Q1 to 2014Q1 and represents the longest period for which quarterly data are available.[4] The upper graph in Figure 1 shows time series plots of the data.

Based on these data for consumption and inflation we estimate the two models (1) and (2) using maximum likelihood.[5] We assume a constant variance-covariance matrix and only allow for time-varying drifts. Instead of the transition intensities Λ, the estimation gives us an $(n \times n)$-matrix $Q = (q_{ij})_{i,j=1\dots,n}$ of transition probabilities, which are linked to the intensities via $\lambda_{ij} = -\log(1 - q_{ij})$ for $j \neq i$. The diagonal elements q_{ii} of the transition probability matrix are set such that the rows sum to 1. Standard errors for the

[3]The choice of this inflation time series is in line with the literature on consumption-based asset pricing with a focus on inflation risk, e.g. Song (2017), David and Veronesi (2013), Burkhardt and Hasseltoft (2012). For a detailed discussion of this issue, we refer the reader to Piazzesi and Schneider (2006).

[4]We have performed the estimation also with alternative samples to compare our findings to those stated in other papers. These results are discussed in Section 6.4. Besides, we have also estimated various constraint versions of the models in which the number of parameters is reduced. This does not change any of our results qualitatively. Details on these constraint models are presented in Section 6.5.

[5]For details about the estimation procedure, see Online Appendix A.

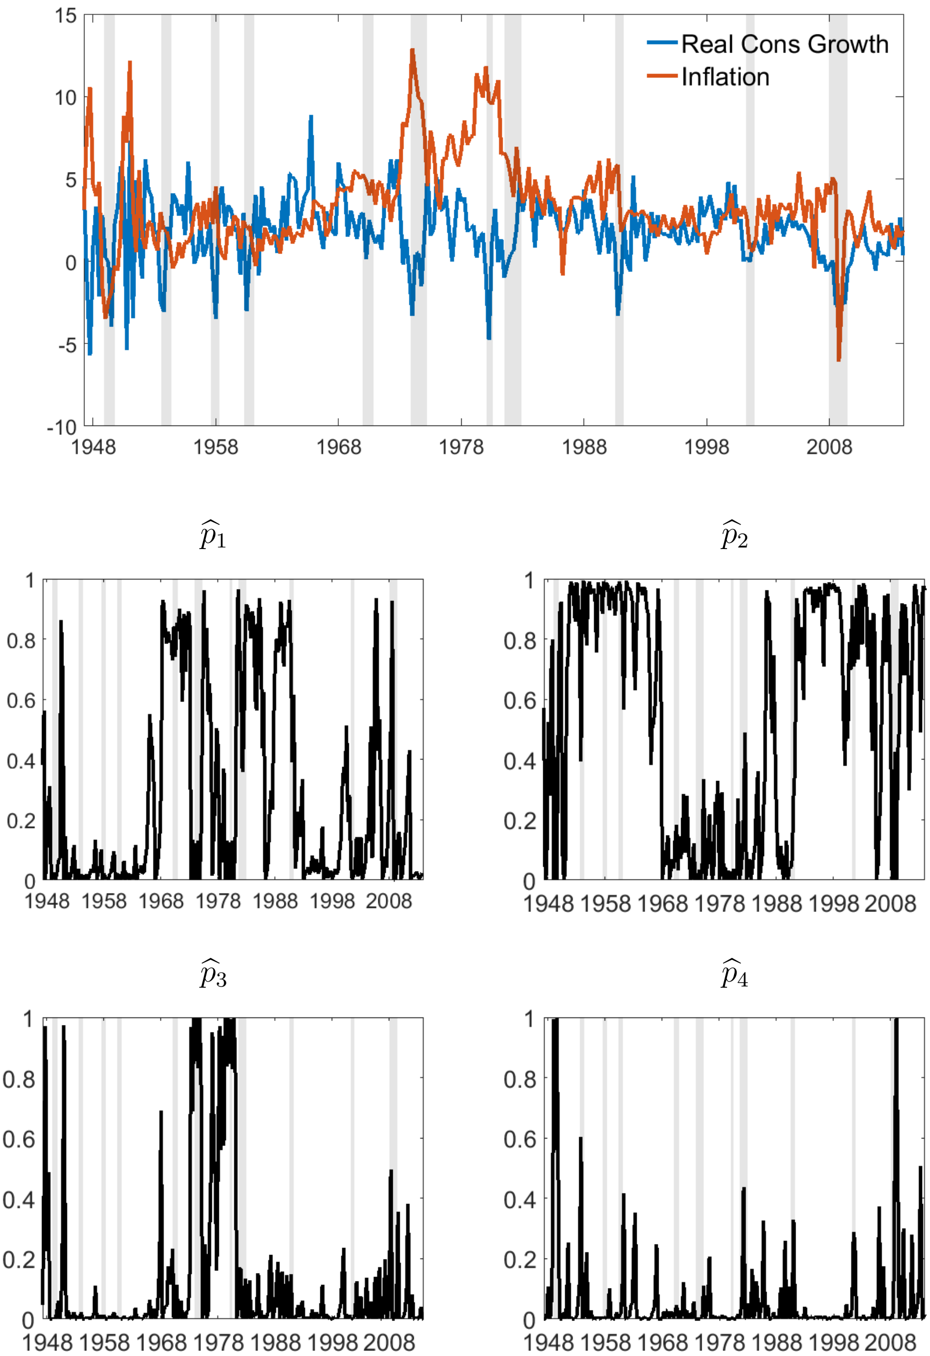

Figure 1: Fundamental data and filtered probabilities

The upper graph shows time series plots of the data for consumption growth and inflation over our sample period from 1947 to 2014. The lower graphs present the real-time filtered probabilities for each of the four states estimated for our Markov switching models with parameters shown in Table 2. Shaded areas indicate NBER recessions.

Panel A: Consumption and inflation parameters

	μ_1^C	μ_2^C	$(\sigma^C)^2$
Consumption growth	2.340	-1.800	0.708
	(0.188)	(0.615)	(0.085)

Panel B: Markov chain transition probabilities

	to state 1	to state 2
from state 1	0.939	0.061
	(0.092)	(0.092)
from state 2	0.467	0.533
	(0.217)	(0.217)

Panel C: Optimal number of states

	2 states	3 states
Log likelihood	-195.44	-184.14
Penalty term	39.14	72.68
Bayes Information Criterion	**430.02**	440.96
$(= -2 \cdot \log L + \text{penalty term})$		

Table 1: Markov chain estimation

This table reports the results from our univariate Markov chain estimation for consumption growth only. Growth rates are given in percentage points and annualized. The data span the period from 1947 to 2014 at the quarterly frequency. Numbers in parantheses are standard errors obtained from a standard block bootstrap with block length of 10 quarters.

parameter estimates are computed via a standard block bootstrap with a block length of ten quarters[6] with potentially overlapping blocks and 5,000 repetitions.

The results are presented in Tables 1 and 2. The first important finding is that in the univariate case, based on the Bayes Information Criterion (BIC), the algorithm clearly identifies two regimes with values for expected consumption growth of 2.34 and -1.80 percentage points, respectively. Moreover, a Wald test based on the bootstrapped standard errors clearly rejects the hypothesis that these two values are equal (p-value 0.002). We regard this test as clear evidence that the conditional mean of consumption growth is varying over time, which is the first key contribution of our paper.

Adding inflation to the model, the estimation results change significantly. For the bivariate model the algorithm identifies four regimes: high growth–medium inflation (state 1), medium growth–low inflation (state 2), low growth–high inflation (state 3), and negative growth–negative inflation (state 4). The estimated transition probabilities imply that

[6] Varying the block length does not affect the results.

Panel A: Consumption and inflation parameters

	μ_1^i	μ_2^i	μ_3^i	μ_4^i	$(\sigma^i)^2$	$\rho\sigma^C\sigma^\pi$	ρ
Consumption growth	2.365	1.898	0.444	-0.997	1.016	-0.136	-0.218
	(0.487)	(0.155)	(0.315)	(0.631)	(0.099)	(0.030)	(0.103)
Inflation	4.704	2.161	9.514	-2.917	0.382		
	(0.975)	(0.846)	(1.668)	(2.438)	(0.035)		

Panel B: Markov chain transition probabilities

	to state 1	to state 2	to state 3	to state 4
from state 1	0.909	0.027	0.039	0.025
	(0.202)	(0.087)	(0.160)	(0.049)
from state 2	0.022	0.970	0.008	0
	(0.029)	(0.056)	(0.027)	(0.024)
from state 3	0.135	0.037	0.828	0
	(0.064)	(0.077)	(0.088)	(0.028)
from state 4	0	0.337	0	0.663
	(0.065)	(0.253)	(0.152)	(0.203)

Panel C: Optimal number of states

	3 states	4 states	5 states	6 states
Log likelihood	-357.80	-317.65	-307.60	-287.75
Penalty term	119.40	169.70	238.80	320.50
Bayes Information Criterion	835.00	**805.00**	854.00	896.00
($= -2 \cdot \log L +$ penalty term)				

Table 2: Markov chain estimation

This table reports the results from our bivariate Markov chain estimation for consumption growth and inflation. Growth rates are given in percentage points and annualized. The data span the period from 1947 to 2014 at the quarterly frequency. Numbers in parantheses are standard errors obtained from a standard block bootstrap with block length of 10 quarters.

	Data		Model	
	Cons. growth	Inflation	Cons. growth	Inflation
Mean	1.826	3.471	1.825	3.283
			[1.640, 2.036]	[2.822, 3.526]
Volatility	1.064	1.363	1.023	1.264
			[1.006, 1.095]	[1.125, 1.339]
Correlation	-0.138		-0.113	
			[-0.282, 0.078]	
AC(1)	0.261	0.756	0.0637	0.637
			[-0.049, 0.169]	[0.430, 0.779]
AC(2)	0.237	0.631	0.044	0.533
			[-0.063, 0.147]	[0.297, 0.700]
AC(3)	0.208	0.578	0.029	0.450
			[-0.082, 0.131]	[0.221, 0.631]

Table 3: Unconditional moments of consumption growth and inflation

The table reports unconditional moments of consumption growth and inflation. The columns labeled "Data" are based on quarterly data from 1947 to 2014. The columns labeled "Model" have been obtained by Monte Carlo simulation using the parameters given in Table 2 (5,000 paths of 68 years each). The numbers in parentheses give 90% confidence bounds around the point estimates.

state 1 lasts for around 11 quarters on average, while the average time spent in state 2 is 33 quarters. The other two states are not very persistent with an average occupation time of around 6 and 3 quarters, respectively. So most of the time, the economy is in state 1 or 2, but it is the rare states 3 and 4 which are very important in the context of asset pricing, since they feature low (or even negative) expected consumption growth. State 3 is a high-inflation state with low growth (sometimes labeled 'stagflation'), whereas in state 4 the expected change in the price level is negative, i.e., there is deflation on average.

Table 3 reports unconditional moments of consumption and inflation in the data and in the estimated time series model (2). Note that the maximum likelihood estimation does not explicitly target these moments. Still, the general fit of the time series model is good. Only the autocorrelations of consumption growth in the model are a little too low, due to the rather simple Markov chain structure.

Since the focus of our paper is on asset pricing, we do not want to go too much into detail about macroeconomic interpretations of our time series model. However, our estimation results may also challenge our understanding of basic macroeconomics. In traditional New-Keynesian models, for instance, inflation and consumption growth are typically positively correlated, whereas finance researchers like Wachter (2006) argue that a negative correlation is necessary to match asset prices. Admittedly, our model does not involve a time-varying correlation of shocks. But our estimation with time-varying expected growth rates may be interpreted in such a way that there are times in which inflation and consumption growth comove positively (e.g. around the deflation episodes) and other times in which they comove negatively (e.g. around the stagflation episodes). Moreover, our estimation is also in line with policy debates which document that negative inflation (i.e., deflation) has re-entered the mindset of policymakers during the recent zero lower bound episode.

4 Asset Pricing Model

Long-run risk models in which expected consumption growth is time-varying are supposed to explain the time series behavior of asset prices and returns. In the following we will therefore embed the dynamics estimated in the previous section into a state-of-the-art asset pricing model with recursive preferences and learning. Having inflation in the model allows us to analyze the returns of stocks and nominal bonds jointly. The quantitative results presented in Section 5 are based on the parameter estimates from the benchmark specification in Table 2.[7]

4.1 Preferences

The economy is populated by an infinitely-lived representative investor with stochastic differential utility as introduced by Duffie and Epstein (1992b). The investor has the indirect utility function

$$J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}) = E_t \left[\int_t^{\infty} f(C_s, J(C_s, \widehat{p}_{1s}, \ldots, \widehat{p}_{ns}))ds \right],$$

where the aggregator f is given by

$$f(C, J) = \frac{\beta C^{1-\frac{1}{\psi}}}{\left(1 - \frac{1}{\psi}\right) \left[(1-\gamma)J\right]^{\frac{1}{\theta}-1}} - \beta\theta J \qquad \text{with } \gamma \neq 1 \text{ and } \psi \neq 1$$

γ, ψ, and β denote the degree of relative risk aversion, the elasticity of intertemporal substitution (EIS), and the subjective time preference rate, respectively. We define $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$. The special case of time-separable CRRA preferences is represented by $\theta = 1$, i.e., by $\gamma = \psi^{-1}$. Throughout the paper, we assume $\gamma = 10$, $\psi = 1.7$, and $\beta = 0.02$.[8] With this parameter choice, the agent has a preference for early resolution of uncertainty, since $\gamma > \psi^{-1}$.

4.2 Filtering

We assume that the representative agent cannot observe S_t (and thus $\mu^C(S_t)$ and $\mu^\pi(S_t)$) and has to filter her estimates from the data.[9] We add learning first and foremost because a full information economy does not generate reasonable time variation in price-dividend

[7] Alternative parameterizations are discussed in Section 6.

[8] A nested version of our model with CRRA preferences ($\theta = 1$) is discussed in Section 6.2.

[9] As pointed out in the introduction and as it is standard in the literature, we assume that the representative agent knows the structural parameters of the model, but does not know the current state of the economy.

ratios. Incomplete information generates an additional layer of uncertainty that is key for our results. Besides the risk to *switch* to a bad state next period, which would also be present in a full information model, we add the uncertainty about the current regime and thus about the *probability* of switching to a bad regime.[10]

Mathematically, there are two filtrations, $\mathcal{F}$ and $\mathcal{G}$, where $\mathcal{F}$ is generated by the processes $(C_t)_t$, $(\pi_t)_t$ and $(S_t)_t$, whereas $\mathcal{G} \subset \mathcal{F}$ is generated by the processes $(C_t)_t$ and $(\pi_t)_t$ only. The conditional expectations of the drifts given the investor's information, $\widehat{\mu}_t^C$ and $\widehat{\mu}_t^\pi$, are given as

$$\widehat{\mu}_t^C = \mathbb{E}\left[\mu^C(S_t)|\mathcal{G}_t\right] = \sum_{i=1}^n \widehat{p}_{it}\mu_i^C$$

and

$$\widehat{\mu}_t^\pi = \mathbb{E}\left[\mu^\pi(S_t)|\mathcal{G}_t\right] = \sum_{i=1}^n \widehat{p}_{it}\mu_i^\pi.$$

Here $\widehat{p}_{it} = \mathbb{E}\left[\mathbf{1}_{\{S_t=i\}}|\mathcal{G}_t\right]$ denotes the subjective conditional probability of being in state i at time t, and these conditional probabilities will serve as state variables in our economy. Since probabilities always sum up to 1, we will have $n-1$ state variables $\widehat{p}_1, \ldots, \widehat{p}_{n-1}$, whose support is the standard simplex in $\mathbb{R}^{n-1}$.

Consumption growth and inflation realizations are observable and serve as a signal for the aggregate state. The dynamics of $\widehat{p}_{it}$ follow from the so-called Wonham filter and are given by

$$d\widehat{p}_{it} = \left(\lambda_{ii}\widehat{p}_{it} + \sum_{j\neq i}\lambda_{ji}\widehat{p}_{jt}\right)dt + \widehat{p}_{it}\left[\begin{pmatrix}\mu_i^C\\ \mu_i^\pi\end{pmatrix} - \sum_{j=1}^n \widehat{p}_{jt}\begin{pmatrix}\mu_j^C\\ \mu_j^\pi\end{pmatrix}\right]'(\Sigma')^{-1}\begin{pmatrix}d\widehat{W}_t^C\\ d\widehat{W}_t^\pi\end{pmatrix}. \tag{3}$$

with the "subjective" Brownian motions

$$\begin{pmatrix}d\widehat{W}_t^C\\ d\widehat{W}_t^\pi\end{pmatrix} = \Sigma^{-1}\left[\begin{pmatrix}\mu_i^C\\ \mu_i^\pi\end{pmatrix} - \sum_{j=1}^n \widehat{p}_{jt}\begin{pmatrix}\mu_j^C\\ \mu_j^\pi\end{pmatrix}\right]dt + \begin{pmatrix}dW_t^C\\ dW_t^\pi\end{pmatrix}.$$

A proof of the filtering equation based on Theorem 9.1 of Liptser and Shiryaev (2001) and a discussion of its properties are provided in Online Appendix B.

In the context of our analysis it is essential to note that the update in the estimated probability $\widehat{p}_i$ depends on both signals, i.e., on both realized consumption growth and realized inflation. Inflation observations have an impact on the perceived probability of being in state i and thus on the conditional expected consumption growth rate. This will be the key driver for our asset pricing results described below.

The lower graphs in Figure 1 show the filtered estimates for the probabilities of the four states, i.e., the estimates the investor would have computed based on information up to and including time t. These estimates are the key quantities analyzed in the following subsection. They will also serve as the explanatory variables in our regression analyses in Section 5. First of all, there is considerable variation in each of the four time series, i.e., the probability of being in state i changes substantially over time. State 1 with the highest

[10] In Section 6.3 we compare our results to those from a nested model with full information, in which the agent can observe the economic state at any point in time.

expected real growth rate, but also above-average inflation is considered most likely by the investor during the 1960s and much of the 1970s. The investor furthermore perceives a high probability to be in the regime 2 with low inflation and stable growth for extended periods during the 1950s and much of the 1990s, but this probability is very low during the 1970s. Not surprisingly, there is a very high probability for the high inflation state 3 during the latter period. The deflation state 4 is seen as very likely in the beginning of the sample right after the war and as well towards the end during the Great Recession.

4.3 Real Pricing Kernel and Wealth-Consumption Ratio

As shown in Duffie and Epstein (1992a), the real pricing kernel depends on the log wealth-consumption ratio v and is given by

$$\xi_t = C_t^{-\gamma} e^{-\beta\theta t + (\theta-1)\left(\int_0^t e^{-v_u} du + v_t\right)}.$$

The wealth-consumption ratio $I \equiv e^v$ depends on the estimated expected consumption growth $\widehat{\mu}^C$, and therefore in particular on the estimated probabilities $\widehat{p}_i$. It solves a nonlinear partial differential equation given in Online Appendix C.1. A proof and details concerning the numerical solution using a Chebyshev polynomial approximation are also presented in Online Appendix C.1.

Given a solution for I, the pricing kernel has dynamics

$$\begin{aligned}\frac{d\xi_t}{\xi_t} &= -\beta\theta dt - (1-\theta)I^{-1}\,dt - \gamma\frac{dC_t}{C_t} + \frac{1}{2}\gamma^2(\sigma^C)^2 dt - (1-\theta)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\,d\widehat{p}_{it} \\ &\quad + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}(\theta-1)\left[\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + (\theta-2)\left(\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I^2}\right)\right]\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}dt - \gamma(\theta-1)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i}\,dt.\end{aligned}$$

with the dynamics of $d\widehat{p}_{it}$ given in Equation (3). Importantly, shocks to the state variables $\widehat{p}_i$ affect the pricing kernel. Since these shocks are themselves driven by both consumption and inflation observations, realized inflation indirectly enters the pricing kernel through the learning mechanism.

4.4 Pricing the Assets in the Economy

We are mainly interested in two types of assets, equity and nominal bonds. Equity is defined as a claim to real dividends. When defining dividends, one has to be careful not to alter the informational setup of the model. Dividends are observable, and if they provided a non-redundant signal about the state of the economy, this would affect the initial filtering problem. Technically, this requires the two systems of equations

$$\begin{pmatrix} d\widehat{W}_t^C \\ d\widehat{W}_t^\pi \end{pmatrix} = \Sigma^{-1}\left[\begin{pmatrix}\mu_i^C \\ \mu_i^\pi\end{pmatrix} - \sum_{j=1}^{n}\widehat{p}_{jt}\begin{pmatrix}\mu_j^C \\ \mu_j^\pi\end{pmatrix}\right]dt + \begin{pmatrix} dW_t^C \\ dW_t^\pi\end{pmatrix}$$

and

$$\begin{pmatrix} d\widehat{W}_t^C \\ d\widehat{W}_t^\pi \\ d\widehat{W}_t^D \end{pmatrix} = \Sigma_*^{-1} \left[\begin{pmatrix} \mu_i^C \\ \mu_i^\pi \\ \mu_i^D \end{pmatrix} - \sum_{j=1}^n \widehat{p}_{jt} \begin{pmatrix} \mu_j^C \\ \mu_j^\pi \\ \mu_j^D \end{pmatrix} \right] dt + \begin{pmatrix} dW_t^C \\ dW_t^\pi \\ dW_t^D \end{pmatrix}$$

to yield the same solution for $d\widehat{W}_t^C$ and $d\widehat{W}_t^\pi$. Here the superscript D denotes terms related to dividend dynamics and $\Sigma_*\Sigma_*'$ is the covariance matrix of innovations to $\ln C$, π and $\ln D$.

The above condition for the redundance of dividends is satisfied by assuming

$$d\ln D_t = \bar{\mu}dt + \phi \left(\sum_{i=1}^n (\mu_i^C - \bar{\mu})\widehat{p}_{it} \right) dt + \phi\sigma^C \left(\sqrt{1-\rho^2} d\widehat{W}_t^C + \rho d\widehat{W}_t^\pi \right).$$

Similar to Bansal and Yaron (2004), the deviation of the drift from its long-term average $\bar{\mu}$ is levered by a factor of ϕ, and like Bansal and Yaron (2004) we assume $\phi = 3$.

Let ω denote the log price-dividend ratio. Starting from the Euler equation for the price of the dividend claim, we can apply the Feynman-Kac formula to $g(\xi, D, \omega) \equiv \xi D e^\omega$. This yields

$$\frac{\mathcal{A}g(\xi, D, \omega)}{g(\xi, D, \omega)} + e^{-\omega} = 0,$$

where $\mathcal{A}$ denotes the infinitesimal generator. Using Ito's Lemma, we can translate this equation into a PDE for $\omega(\widehat{p})$. This PDE, together with details regarding its derivation, is given in Online Appendix C.3. We solve this PDE again numerically using a Chebyshev approximation.

A nominal bond pays off one unit of money at maturity T, which, in real terms, is equal to $\exp\left(-\int_t^T d\pi_s ds\right) = \exp(\pi_t - \pi_T)$. The price of a nominal bond at time t is thus equal to

$$B_t^{\$,T} = E_t \left[\frac{\xi_T}{\xi_t} \exp(\pi_t - \pi_T) \right].$$

Equivalently, one can define the nominal pricing kernel as $\frac{\xi_T^\$}{\xi_t{}^\$} \equiv \frac{\xi_T}{\xi_t} \exp(\pi_t - \pi_T)$ and rewrite the pricing formula as

$$B_t^{\$,T} = E_t \left[\frac{\xi_T^\$}{\xi_t{}^\$} \right].$$

The dynamics of the nominal pricing kernel then follow from Ito's lemma:

$$\frac{d\xi_t^\$}{\xi_t^\$} = \frac{d\xi_t}{\xi_t} - d\pi_t + \frac{1}{2} d[\pi]_t - \frac{d[\xi, \pi]_t}{\xi_t}.$$

Importantly, the nominal risk-free short rate, i.e., the negative of the drift of $\xi^\$$, is not just the sum of the real short rate and expected inflation, but involves the covariation between the real pricing kernel and inflation, $d[\xi, \pi]$ (and the quadratic variation of π). The covariation is nonzero if inflation shocks affect the real pricing kernel, as they do in our model, and can be interpreted as an equilibrium inflation risk premium.

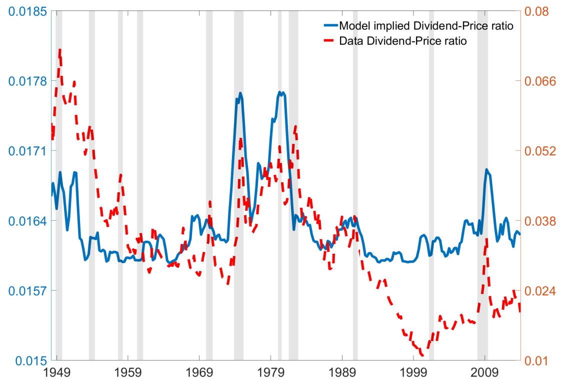

Figure 2: Dividend-price ratio

The figure depicts quarterly time series of dividend-price ratios. The blue solid line is the dividend-price ratio of the S&P 500 index. The red dashed line is obtained by plugging the historical paths of consumption, inflation, and our state variables $\widehat{p}_i$ into the numerical solution of the model. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is 0.43.

The Euler equation and the Feynman-Kac formula applied to $H(\xi_t^{\$}, b_t^{T,\$}) = \xi_t^{\$} e^{b_t^{T,\$}}$ yield a partial differential equation for the (log) price $b_t^{T,\$} \equiv \ln B_t^{T,\$}$ of a nominal zero coupon bond. Details on this partial differential equation and its solution are given in Online Appendix C.5.

5 Results

5.1 Dividend-price ratios

In order to see how inflation risk influences real asset prices through the long-run risk channel that we have established via the Markov chain estimation, it is instructive to start by comparing the dividend-price ratios generated by our model with those observed in the data. To this end, we plug the historical quarterly consumption and inflation time series and our estimated state variables $\widehat{p}_i$ into the numerical solution of the model.

Figure 2 shows the model-implied dividend-price ratio together with the historical dividend-price ratio of the S&P 500 index. The correlation between the two time series is 0.43 and the two time series share all major upward and downward trends. Given that we have not used any asset price data in the estimation, this finding is remarkable. We take this as a first piece of evidence that our proposed state variables $\widehat{p}_i$, embedded into an otherwise standard asset pricing model featuring recursive Epstein-Zin preferences, indeed capture the time variation in valuation ratios well.

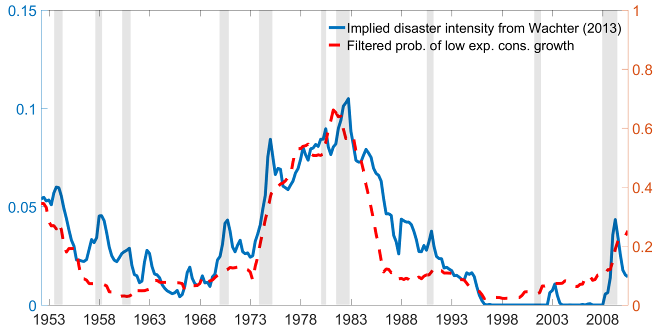

Figure 3: Time-varying disaster probabilities

The solid (blue) line depicts the time-varying disaster intensity, which Wachter (2013) extracts from asset price data. The dashed (red) line shows 20-quarter moving averages of the sum of estimated probabilities $\widehat{p}_3 + \widehat{p}_4$ from our Markov switching model using consumption and inflation data for the period from 1947 to 2014. To obtain $\widehat{p}_3$ and $\widehat{p}_4$ we plug realized consumption growth and inflation data into our filtering equations and compute the probabilities, which a Bayesian learner would have assumed at each point in time. The correlation between the two time series is 0.88.

5.2 Extreme inflation as a signal about disaster risk

Based on the similarity between model-implied and empirical dividend-price ratios, we first turn towards a deeper analysis of the time series pattern of the state variables in our model. The main result of this section is depicted in Figure 3. The solid blue line is the implied disaster intensity shown in Figure 8 (p. 1017) in the paper of Wachter (2013).[11] She reverse-engineers this quantity from asset prices based on her model, where the intensity of rare consumption disasters follows a mean-reverting process and serves as a state variable. Given the parameters of her model, she recovers monthly implied values for this state variable from a smoothed time series of historical S&P 500 price-earnings ratios. Although our model does not explicitly feature "disaster risk", we choose this particular time series for our comparison because its interpretation as a "probability" (of very low consumption growth) is more in line with the state variables $\widehat{p}_i$ in our model than, for instance, time-varying conditional means of consumption growth, which are typically presented in long-run risk papers. To make her series comparable to our estimates we take averages of the monthly implied disaster intensities over each quarter.

[11]We thank Jessica Wachter for sharing her data with us.

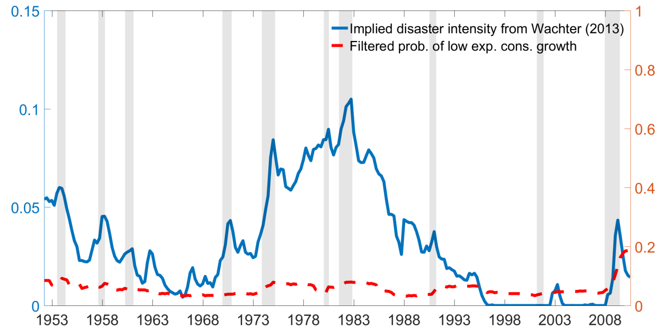

Figure 4: Disaster probabilities estimated from consumption growth only

The solid (blue) line depicts the time-varying disaster intensity, which Wachter (2013) extracts from asset price data. The dashed (red) line shows 20-quarter moving averages of the estimated probability for the state with low expected consumption growth from a Markov switching model using only consumption data for the period from 1947 to 2014. To obtain this probability we plug realized consumption growth into our filtering equations and compute the probabilities, which a Bayesian learner would have assumed at each point in time. The correlation between the two time series is 0.33.

The dashed red line is the sum of the filtered probabilities $\widehat{p}_3 + \widehat{p}_4$ from our model. To obtain these estimates we plug realized consumption growth and inflation data into our filtering equations and compute the probabilities, which a Bayesian learner would have assumed at each point in time. The plot shows 5-year moving averages of these probabilities.[12]

The two series have a correlation of 0.88 over our sample period covering almost 70 years. This is particularly remarkable, given that they are computed from very different data and using rather different approaches. Furthermore, they share all important trends, peaks, and troughs. There is a pronounced downturn during the 1950s, followed by rather low values in the 1960s, a sharp increase during the 1970s up to around 1982, and then, basically following the same kind of cycle, we observe the sharp decline and low level during the Great Moderation, followed by the recent spike at the beginning of the Great Recession. In particular, both the high inflation regime and the deflation regime and their respective probabilities are relevant. For instance, a look at the time series plots of the filtered probabilities in Figure 1 shows that the peak of the two series in the early 1980's can be traced back to the high probability of the high inflation regime (state 3) over that period, and the deflation regime prevails towards both the beginning and the end of the sample period. This result strongly supports the notion that inflation can serve as a signal for expected real consumption growth in that it allows to quantify the probability of large negative future consumption shocks.

To check whether it is indeed inflation that is important here, and not just certain special characteristics of the consumption time series, we redo the analysis based on

[12] We use moving averages in order to account for the smoothing in Wachter (2013). More precisely, for her reverse engineering exercise, she uses the ratio of prices to the previous 10 years of earnings. The two time series depicted in Figure 3 both have an autocorrelation of 0.99.

only consumption data, i.e. the univariate baseline model presented in the beginning. Figure 4 presents the time series of the estimated probability for the state with low expected growth. Already from a first rough inspection it becomes clear that the disaster intensity is matched much less precisely than before. The correlation between the series based on Wachter (2013) and the filtered probability of being in a bad state derived from the consumption-only series goes down to roughly 0.33, but more importantly, the times series of estimated probabilities for low consumption growth is substantially off during basically all periods when the risk of the economy being in a bad state is actually high, e.g., during most of the 1970's and 1980's and also to a certain degree towards the end of the sample period. These findings are a clear indication that information about inflation in necessary to obtain reliable estimates for the probability of low real growth.

5.3 Conditional stock return volatilities

Given that Wachter (2013) documents the important role of time-varying disaster intensities for the dynamics of second moments of returns, we continue the discussion of our asset pricing results by analyzing second moments.

We proceed in the following way. We take the time series of filtered probabilities as shown in Figure 1, plug them into our model solution and compute model-implied real prices for equity and for nominal bonds with five years to maturity. From these time series of real prices we compute model-implied quarterly real log returns for these two assets. More precisely, with S_t and $B_t(20)$ denoting the price of the equity claim and the 20-quarter (five-year) nominal zero coupon bond in quarter t, the returns from quarter t to quarter $t+1$ are computed as $\ln(S_{t+1}+D_{t+1})-\ln S_t$ and $\ln B_{t+1}(19)-\ln B_t(20)$. We then add log realized inflation to the real returns to obtain nominal returns. The corresponding quantities in the data are quarterly returns of the CRSP value-weighted index and log bond returns computed from the US Treasury yield curve data provided by Gürkaynak, Sack, and Wright (2007)[13] from 1962 on. As the final input to our analyses we compute 20-quarter rolling window return volatilities and correlations and regress them on (the logarithm of) 20-quarter moving averages of the relevant state probabilities $\widehat{p}$. Note that these right-hand variables are the same for model and data in all the regressions reported below.

For the regressions in the model and in the data we state Newey-West adjusted t-statistics with 20 lags, but in addition we also provide confidence intervals derived from a Monte Carlo simulation of the model (shown in square brackets below the respective coefficient). Here we first simulate the model given the dynamics for the fundamentals and the filtered probabilities in Equations (1) to (3) with monthly time increments over a time span of 68 years, corresponding to the length of our sample period for the macroeconomic variables. These monthly data are then aggregated to quarterly and used in the regressions in the same way as described before, i.e., we only use the later 50 years of each sample path, corresponding to the period over which financial market data are available. We repeat this exercise 5,000 times to obtain the 90% confidence intervals.[14]

[13] The data are available for download at `http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html`.

[14] Due to the discretization error in the simulation it sometimes happens that the sum of the filtered probabilities exceeds 1 by a very small amount. In that case we rescale the filtered probabilities such

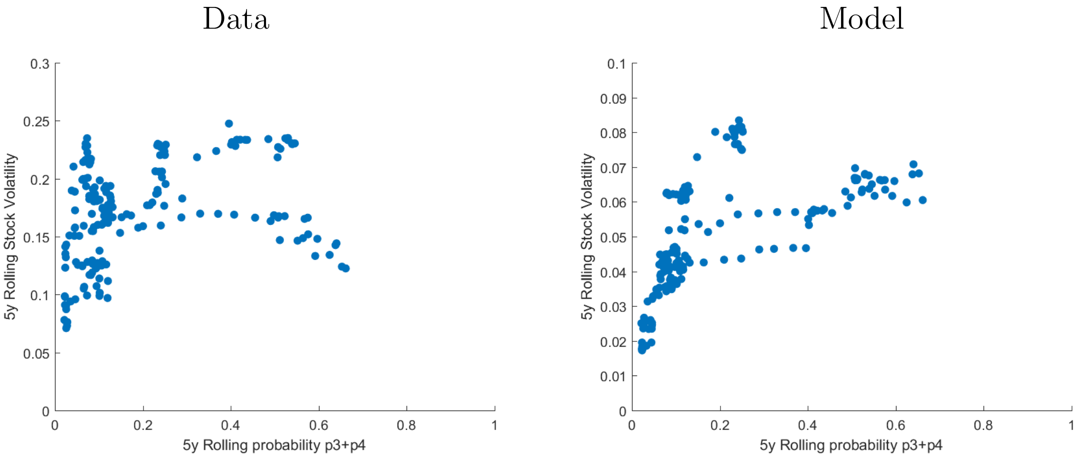

Figure 5: Stock return volatilities and probability of low expected growth

The figure depicts scatter plots for the regressions presented in Table 4. The dependent variables in each regression are volatilities of quarterly stock returns computed over rolling windows of 20 quarters. The independent variable is the logarithm of the averages of $\widehat{p}_3 + \widehat{p}_4$ over the same 20 quarters periods. The left figure labeled "Data" is based on the estimated time series of the $\widehat{p}_i$ depicted in Figure 1 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve (see Gürkaynak et al. (2007)). The right figure labeled "Model" is based on the same time series of the $\widehat{p}_i$, but uses the returns which our model would have implied given this path of consumption, inflation, and the state variables. The financial data for these regressions starts in 1965. The model parameters are estimated using macroeconomic data since 1947.

A look at Figure 5 shows that our model nicely reproduces the patterns of state-dependent stock return volatilities in the data along two important dimensions. First, the estimated probabilities for the states with low consumption growth, $\widehat{p}_3 + \widehat{p}_4$, exhibit a positive covariation with stock return volatilities. Second, this relationship is nonlinear and concave, both in the model and in the data. It is worth noting that the second result confirms another prediction from the model of Wachter (2013), namely that stock market volatility is a concave function of time-varying disaster risk.[15] Figure 6 shows the time series of volatilities in the data and in the model.

Motivated by these scatter plots, we regress stock return volatilities on the logarithm of the moving averages of the relevant state probabilities $\widehat{p}$. Table 4 reports the results. The regression coefficients are positive and significant in both model and data, the R^2 is high both in the model and in the data, and almost all of the regression coefficients from the data are within the simulated confidence bounds for the model. The only exception with respect to this last point is the low constant in our model regressions, which indicates that the model-implied unconditional stock return volatility is somewhat on the low side.[16] Overall, we conclude that the relation between the conditional probability of

that they exactly sum to 1.

[15] See, for instance, Figure 4 (p. 1002) in Wachter (2013).

[16] We discuss this issue in detail in Section 6.6.

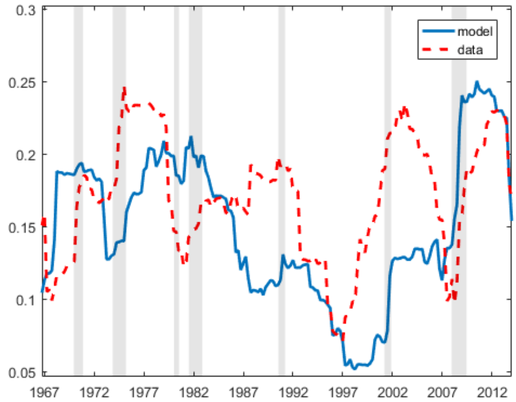

Figure 6: Conditional stock return volatilities

The figure depicts conditional 20-quarter rolling window stock return volatilities. The red dashed line is based on the CRSP value-weighted index. The blue solid line is based on the returns which our model would have implied given the historical paths of consumption, inflation, and state variables (this series is multiplied by 3). The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is 0.41.

low consumption growth (given by the sum $\widehat{p}_3 + \widehat{p}_4$) and stock market volatility is indeed nonlinear, and our model reproduces this stylized fact. Finally, given that our estimation is based on two macro time series only, the goodness of fit in Figure 6 is also remarkable. The correlation between the data and the model-implied time series is 0.41.

To see how our model generates these results, a look at the filtering equation (3) is instructive. $\widehat{p}_3$ and $\widehat{p}_4$ fluctuate a lot more when they are at intermediate levels as compared to when they are close to 0 or 1. Moreover, the unconditional averages of both $\widehat{p}_3$ and $\widehat{p}_4$ are rather low, so that higher values for $\widehat{p}_3$ and $\widehat{p}_4$ more or less automatically go together with high fluctuations in these quantities. States 3 and 4 are the least persistent in our estimation, so that when the agent currently has a high estimate for the probability of being in one of these two states, this probability is likely to decrease again quickly.

Altogether, there is thus more movement in state variables when the likelihood of low consumption growth is large. This in turn means that a high likelihood of low consumption growth coincides with a high volatility of the price-dividend ratio, which is a function of these state variables. In sum, the higher the likelihood of low consumption growth, the higher the equity return volatility. Note that our model of course also reproduces the relevance of the probability of being in a good state, i.e., of $\log(\widehat{p}_1 + \widehat{p}_2)$, for stock return volatilities. The coefficients in the data and in the model (not reported here for brevity) are both significantly negative which follows from the simple fact that $\widehat{p}_1 + \widehat{p}_2 = 1 - \widehat{p}_3 - \widehat{p}_4$.

Finally, Table 5 presents the results from additional regressions to investigate the notion of a signaling role of inflation for second moments of stock returns. Here we regress the same left-hand side variable as before on rolling averages of what we call the extreme entropy of the state distribution (defined as $\widehat{p}_3 \ln \widehat{p}_3 + \widehat{p}_4 \ln \widehat{p}_4$). We propose this quantity as another proxy for uncertainty about low consumption growth. We also show

Panel A: Model				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
0.114	0.009	0.012		0.733
(12.627)	(10.833)	(4.724)		
[0.085, 0.174]	[0.004, 0.020]	[0.000, 0.023]		[0.220, 0.785]
0.079			0.015	0.615
(12.143)			(5.997)	
[0.071, 0.129]			[0.007, 0.029]	[0.104, 0.763]

Panel B: Data				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
0.278	0.012	0.023		0.304
(12.679)	(1.416)	(4.475)		
0.214			0.022	0.214
(9.056)			(1.976)	

Table 4: Regressions of stock return volatilities on state variables

The table reports results from time series regressions. The dependent variables in each regression are volatilities of quarterly stock returns computed over rolling windows of 20 quarters. The independent variables are logarithms of the averages of the $\widehat{p}_i$ over the same 20 quarters periods. The regressions labeled "Data" are based on the estimated time series of the $\widehat{p}_i$ depicted in Figure 1 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve (see Gürkaynak et al. (2007)). The regressions labeled "Model" are based on the same time series of the $\widehat{p}_i$, but use the returns which our model would have implied given this path of consumption, inflation, and state variables. The financial data for these regressions starts in 1965. The model parameters are estimated using macroeconomic data since 1947. The numbers in parentheses are t-statistics adjusted following Newey and West (1987) (20 lags). The numbers in brackets denote 90% confidence bounds around the regression coefficients and are obtained from a Monte Carlo simulation of the model (5,000 paths of 68 years each, with the last 50 years of each path used in the regressions).

Panel A: Model				
const.	extreme entropy	expected inflation	realized inflation	Adj. R^2
0.001	0.269			0.538
(0.106)	(5.859)			
$[-0.038, 0.053]$	[0.041, 0.618]			[0.002, 0.609]
0.032		0.004		0.159
(2.504)		(1.930)		
$[0.009, 0.064]$		[-0.001, 0.015]		[-0.007, 0.584]
0.036			0.003	0.147
(3.130)			(1.801)	
$[0.018, 0.063]$			[-0.001, 0.012]	[-0.007, 0.578]
Panel B: Data				
const.	extreme entropy	expected inflation	realized inflation	Adj. R^2
0.097	0.399			0.187
(2.607)	(2.320)			
0.147		0.006		0.036
(4.814)		(0.855)		
0.153			0.004	0.040
(6.010)			(0.781)	

Table 5: Regressions on alternative explanatory variables

The table reports results from time series regressions. The dependent variables in each regression are volatilities of quarterly stock returns computed over rolling windows of 20 quarters. The independent variables are the average of the extreme entropy ($\widehat{p}_3 \ln \widehat{p}_3 + \widehat{p}_4 \ln \widehat{p}_4$), the average expected inflation and the average realized inflation, always taken over the same 20 quarter periods. "Data" and "Model" have the same meaning as in Table 4.

results for expected inflation (defined as $\sum_{i=1}^{4} \mu_i^{\pi} \widehat{p}_i$) and realized inflation as explanatory variables.

The results for extreme entropy can be interpreted in the way that uncertainty about extreme inflation and low consumption growth is a main driver of stock return volatilities, and our model provides an economic equilibrium mechanism to explain this stylized fact. On the other hand, expected and realized inflation as the right-hand side variables do not have explanatory power in the data and are only marginally significant in the model as well. This provides additional support for our model. Inflation itself cannot explain stock return volatilities, unless it is decomposed into components capturing the risk of very high inflation (represented by $\widehat{p}_3$) and of a deflationary regime (represented by $\widehat{p}_4$), respectively. Our Markov chain estimation implies that inflation is positively correlated with $\widehat{p}_3$, but negatively correlated with $\widehat{p}_4$. Depending on the amount of observations from the deflation regime along a given sample path, the coefficient from a regression of equity return volatility on inflation in model-generated data can thus be positive or negative.

Recursive preferences are a key ingredient of our model. With respect to this feature, our paper is closely related to recent studies like Benzoni, Collin-Dufresne, and Goldstein

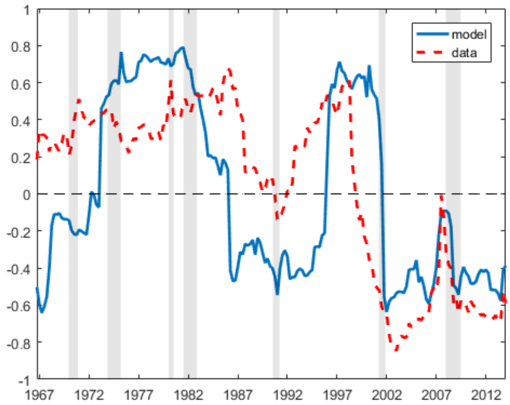

Figure 7: Conditional stock-bond return correlation

The figure depicts conditional 20-quarter rolling window correlations of stock returns and returns of 5-year nominal bonds. The red dashed line is based on the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve. The blue solid line is based on the returns which our model would have implied given the historical paths of consumption, inflation, and state variables. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is 0.54.

(2011) and Drechsler (2013), where it has been shown that models featuring recursive preferences, coupled with learning about fundamentals, are very well able to match stylized facts about stock return volatility, in particular in terms of its dynamics (i.e., capturing the predictive power of implied volatilities, variance risk premia, and other related conditional quantities). The impact of the preferences will also be analyzed in more detail in Section 6.2.

5.4 Conditional stock-bond return correlations

The results concerning stock market volatility are related to the total estimated probability of being in a bad state for expected consumption growth, given by the sum $\widehat{p}_3 + \widehat{p}_4$. When we now look at the stock-bond return correlation, the distinction between the two bad consumption states with respect to expected inflation becomes relevant. Figure 7 shows the time series of correlation in the data and in the model. Tables 6 and 7 contain the results of our regression analyses, Figure 8 presents the corresponding scatter plots.[17]

The most important result here is that both in the model and in the data the estimated coefficient for $\log(\widehat{p}_3)$ is positive (and highly significant), while the coefficient for $\log(\widehat{p}_4)$ is negative (and also highly significant). Given that we do not use any asset price information to estimate the fundamental dynamics in our model, the similarity between model

[17]Note that we regress 'raw' correlation ρ on the state variables. Since correlation is bounded between -1 and 1, transformations like $\tilde{\rho} = \ln(\frac{1+\rho}{1-\rho})$ might seem warranted to guarantee that the regressions are well specified. Rerunning all our regressions using this transformation leaves the results qualitatively unchanged.

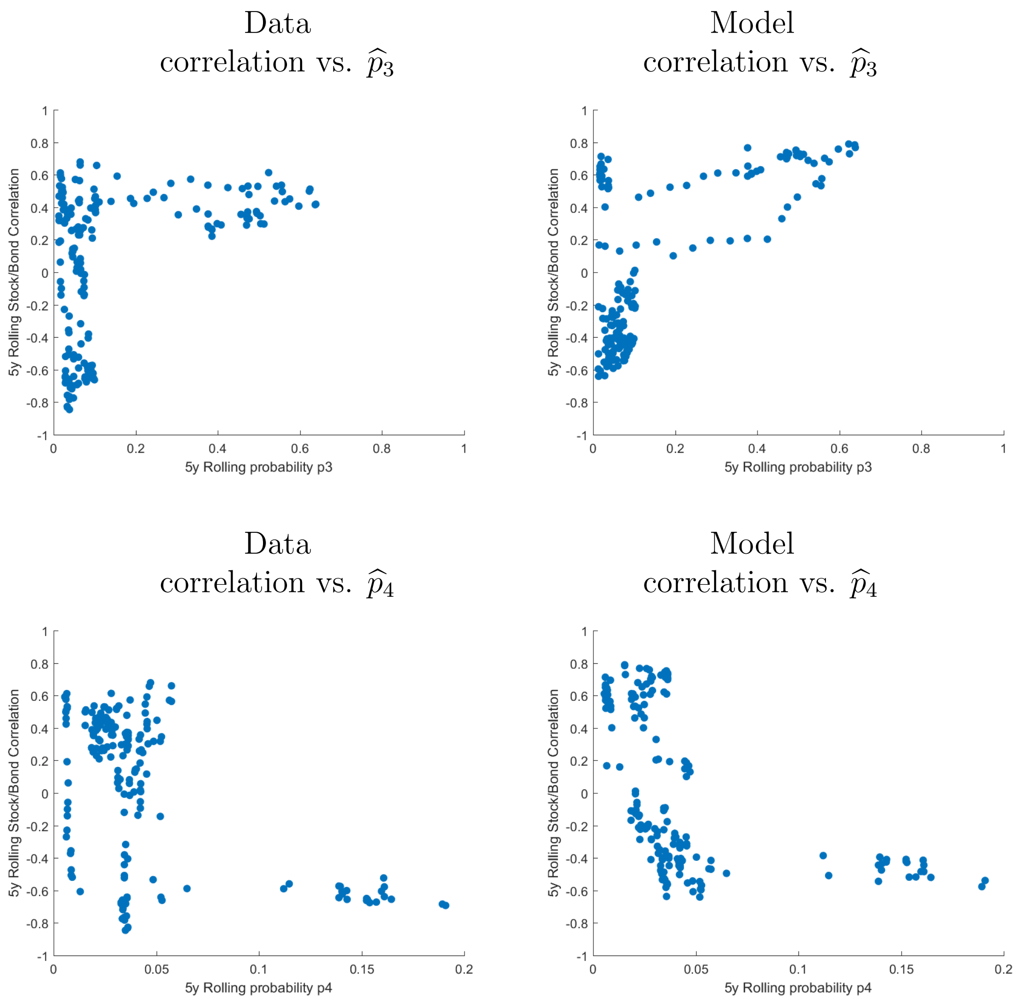

Figure 8: Stock-bond correlations and probabilities for high or low inflation

The figure depicts scatter plots for the regressions presented in Table 6. The dependent variables in each regression are correlations of quarterly holding-period returns of stocks and 5-year nominal bonds computed over rolling windows of 20 quarters. The independent variables are the logarithms of the averages of $\widehat{p}_3$ and $\widehat{p}_4$, respectively, over the same 20-quarter periods. The left figures labeled "Data" are based on the estimated time series of the $\widehat{p}_i$ depicted in Figure 1 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve (see Gürkaynak et al. (2007)). The right figures labeled "Model" are based on the same time series of the $\widehat{p}_i$, but use the returns which our model would have implied given this path of consumption, inflation, and the state variables. The financial data for these regressions starts in 1965. The model parameters are estimated using macroeconomic data since 1947.

Panel A: Model				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3+\widehat{p}_4)$	Adj. R^2
-0.837	0.290	-0.449		0.788
(-4.825)	(6.776)	(-5.108)		
$[-1.833, 0.556]$	[0.065, 0.453]	[-0.702, -0.089]		[0.172, 0.751]
0.325			0.161	0.087
(0.881)			(0.903)	
$[-0.867, 0.962]$			[-0.275, 0.473]	[-0.009, 0.481]

Panel B: Data				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3+\widehat{p}_4)$	Adj. R^2
-0.512	0.182	-0.294		0.363
(-1.61)	(2.731)	(-3.608)		
0.239			0.091	0.041
(0.923)			(0.774)	

Table 6: Regressions of stock-bond return correlations on state variables

The table reports results from time series regressions. The dependent variables in each regression are correlations of quarterly holding-period returns of stocks and 5-year nominal bonds computed over rolling windows of 20 quarters. The independent variables are logarithms of the averages of the $\widehat{p}_i$ over the same 20 quarters periods. "Data" and "Model" have the same meaning as in Table 4. Numbers in parentheses are t-statistics adjusted following Newey and West (1987) (20 lags), numbers in brackets are 90% confidence bounds around the regression coefficients and have been obtained from a Monte Carlo simulation of the model (5,000 paths of 68 years each, with the last 50 years of each path used in the regressions).

and data appears remarkable. Moreover, the scatter plots again reveal a pronounced nonlinearity in the relation between correlation and the filtered probabilities, both in the model and in the data. Finally, the goodness of fit in Figure 7 is again remarkable. The correlation between the data and the model-implied time series is 0.54.

What is the mechanism inside the model responsible for these patterns? First, an increase in $\widehat{p}_3$ makes it subjectively more likely for the investor that the economy is in the high inflation state. In this case the bond return over the next quarter is composed of a positive 'carry' component (which is, if nothing changes, equal to the yield of the bond) and a negative component due to an upward shift in the nominal yield curve. In general the second effect dominates the first, so that bond prices tend to go down. Note that the upward shift in the nominal yield curve itself is the composite of two effects: an increase in expected inflation and a slight decrease in the level of the real yield curve. The second of these two effects is typically negligible with recursive preferences, and therefore, the

Panel A: Model				
const.	extreme entropy	expected inflation	realized inflation	Adj. R^2
0.294	-1.613			0.020
(0.720)	(-0.696)			
$[-1.301, 1.390]$	[-8.620, 7.868]			[-0.010, 0.248]
-0.776		0.195		0.40
(-2.617)		(3.523)		
$[-1.186, 0.167]$		[-0.066, 0.322]		[-0.007, 0.533]
-0.645			0.160	0.410
(-2.630)			(3.750)	
$[-1.015, 0.104]$			[-0.036, 0.276]	[-0.007, 0.547]
Panel B: Data				
const.	extreme entropy	expected inflation	realized inflation	Adj. R^2
-0.047	0.560			0.012
(-0.096)	(0.220)			
-0.613		0.167		0.334
(-2.028)		(3.172)		
-0.488			0.135	0.328
(-1.904)			(3.324)	

Table 7: Regressions on alternative explanatory variables

The table reports results from time series regressions. The dependent variables in each regression are correlations of quarterly holding-period returns of stocks and 5-year nominal bonds computed over rolling windows of 20 quarters. The independent variables are the average of the extreme entropy ($\widehat{p}_3 \ln \widehat{p}_3 + \widehat{p}_4 \ln \widehat{p}_4$), the average expected inflation and the average realized inflation, always taken over the same 20 quarter periods. "Data" and "Model" have the same meaning as in Table 4. Numbers in parentheses are t-statistics adjusted following Newey and West (1987) (20 lags), numbers in brackets are 90% confidence bounds around the regression coefficients and have been obtained from a Monte Carlo simulation of the model (5,000 paths of 68 years each, with the last 50 years of each path used in the regressions).

nominal yield curve shifts upwards in response to an increase in $\widehat{p}_3$. The stock return upon a positive shock to $\widehat{p}_3$ depends on real quantities only. A high $\widehat{p}_3$ implies that the economy is more likely to be in a low consumption growth regime, and stock prices tend to be low in such an environment. Taken together, the reactions of bond and stock prices to an increase in $\widehat{p}_3$ go in the same direction, implying a positive correlation.

State 4 is a low inflation state with low growth, so the response of bond prices to a high $\widehat{p}_4$ is different. Again, there is the positive carry return. But now there is also an additional positive return because the nominal yield curve shifts downwards in response to a higher probability for deflation. When deflation becomes more likely, the level of the nominal yield curve must decrease. Altogether, the impact of a high probability $\widehat{p}_4$ on bond returns is large and positive. At the same time, such a high $\widehat{p}_4$ signals a high likelihood of low (even negative) expected consumption growth, which depresses equity prices. In sum, the impact of a high likelihood for the deflationary regime is strong on both stock and bond prices, but it is of opposite signs, implying a negative correlation between the two types of assets.

The above findings concerning the role of $\widehat{p}_3$ and $\widehat{p}_4$ for the conditional stock-bond correlation are well in line with the literature. In a purely empirical paper, Baele et al. (2010) try to fit the correlations of daily stock and bond returns with a multi-factor model. They find that macro factors (in particular output gap and inflation) do not add much explanatory power when the loadings of stock and bond returns are assumed to be constant over time. But the performance of the macro factors improves in a regime switching estimation when the loadings are allowed to switch sign. Our findings are potentially related to theirs in the sense that we show that the overall risk of low expected consumption growth, proxied by the sum $\widehat{p}_3 + \widehat{p}_4$, does not predict correlation, neither in the model nor in the data. The coefficients on realized and expected inflation are positive, but the simulated confidence bounds always include zero.

For both model and data, the regressions with extreme entropy generate the expected result with insignificant coefficient estimates. The reason is again that this measure captures general uncertainty about bad consumption growth states. Uncertainty about being in the deflation state, however, decreases correlation, whereas uncertainty about the high inflation state increases it. An aggregate measure of uncertainty cannot capture these two opposing effects adequately.

6 Robustness

6.1 Overview

In the previous sections, we have shown that long-run risk and time-varying disaster risk can be linked to time variation in the stock-bond return correlation if one accounts for the signaling role of inflation. Since these results have been obtained within the framework of a particular asset pricing model, we are going to discuss why the major modeling assumptions are necessary to obtain this result and how our findings would change in simplified versions of our model. Finally, we will also present evidence from additional data samples to verify that our findings are not due to a specific choice of these samples.

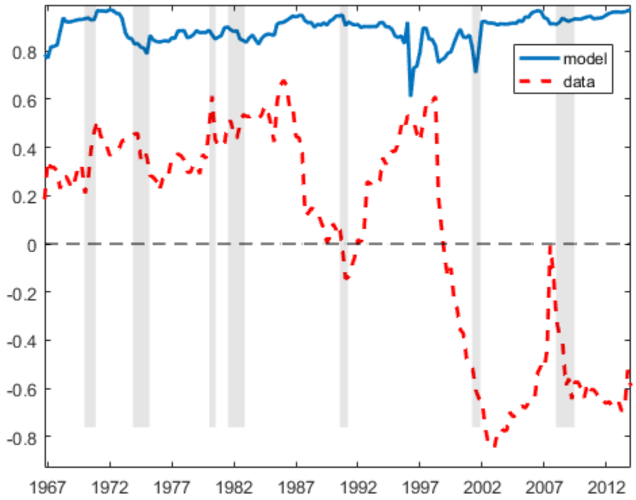

Figure 9: Conditional stock-bond return correlation with CRRA preferences

The figure depicts conditional 20-quarter rolling window correlations of stock returns and returns of 5-year nominal bonds. The red dashed line is based on the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve. The blue solid line is based on the returns which a constrained version of our model with CRRA preferences (i.e. $\gamma = \frac{1}{\psi}$ would have implied given the historical paths of consumption, inflation, and state variables. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is -0.34.

6.2 Recursive Preferences, Money Illusion, CRRA preferences

Recursive preferences are essential to match the time-varying stock-bond correlation in our model. We have also solved a version of our model with constant relative risk aversion preferences, where we set the elasticity of intertemporal substitution to $\psi = \frac{1}{10}$ so that $\theta = 1$, but leave the parameters that we obtain from the Markov chain estimation unchanged.

Figure 9 depicts the time series of model-implied return correlations. The constrained model does not match the historical time series at all, the correlation between the historical time series and the model-implied time series is even negative at -0.34. Table 8 presents results from regressions analogous to those presented in Table 6. The constrained model generates an insignificant (and slightly negative) regression coefficient for $\widehat{p}_3$ and a large and significantly positive coefficient for $\widehat{p}_4$. To get the intuition behind this result, look again at the three components of the holding period bond return as described in Section 5.

Both the carry component and the change in expected inflation are independent of the representative agent's preferences, but the change in the real yield curve is clearly not, since real bond prices are determined in equilibrium. A slightly higher current value of $\widehat{p}_3$ or $\widehat{p}_4$ results in a reduced estimate of conditionally expected consumption growth, which leads to a massive decline in the overall level of the real yield curve in a CRRA economy. It is well known that CRRA models have a hard time matching the empirically observed smoothness of the real risk-free rate. Altogether, in a CRRA economy bond returns are

Panel A: Model with CRRA preferences				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
1.034	-0.004	0.043		0.341
(38.33)	(-0.58)	(5.42)		
[0.922, 2.070]	[-0.011, 0.169]	[0.008, 0.278]		[0.237, 0.803]
0.912			0.009	0.016
(25.874)			(0.54)	
[0.874, 1.849]			[0.002, 0.170]	[0.004, 0.509]

Panel B: Data				
const.	$\log(\widehat{p}_3)$	$\log(\widehat{p}_4)$	$\log(\widehat{p}_3 + \widehat{p}_4)$	Adj. R^2
-0.512	0.182	-0.294		0.363
(-1.61)	(2.731)	(-3.608)		
0.239			0.091	0.041
(0.923)			(0.774)	

Table 8: Regressions of return correlations with CRRA preferences

The table reports results from time series regressions. The dependent variables in each regression are correlations of quarterly holding-period returns of stocks and 5-year nominal bonds computed over rolling windows of 20 quarters. The independent variables are logarithms of the averages of the $\widehat{p}_i$ over the same 20 quarters periods. "Data" and "Model" have the same meaning as in Table 4. Numbers in parentheses are t-statistics adjusted following Newey and West (1987) (20 lags). The numbers in brackets denote 90% confidence bounds around the regression coefficients and are obtained from a Monte Carlo simulation of the model (5,000 paths of 68 years each, with the last 50 years of each path used in the regressions).

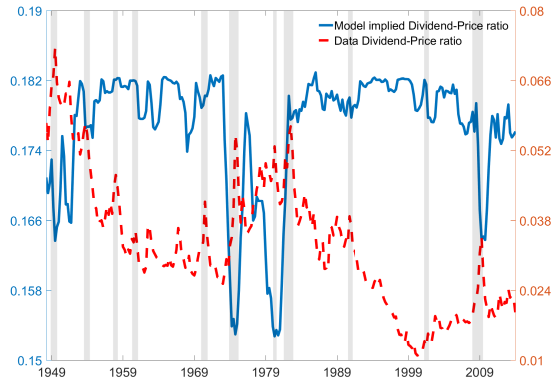

Figure 10: Dividend-price ratio with CRRA preferences

The figure depicts quarterly time series of dividend-price ratios. The blue solid line is the dividend-price ratio of the S&P 500 index. The red dashed line is obtained by plugging the historical paths of consumption, inflation, and our state variables $\widehat{p}_i$ into the numerical solution of the constrained model with CRRA preferences. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is -0.45.

thus slightly negatively related to $\widehat{p}_3$, the probability of being in a stagflation regime, and strongly positively related to $\widehat{p}_4$, the probability of being in a deflationary regime.

Concerning stock returns, high values for $\widehat{p}_3$ or $\widehat{p}_4$ signal low expected consumption growth, while consumption volatility is not affected. With the usual popular CRRA parameterization of $\gamma > 1$ and, consequently, $\psi = \gamma^{-1} < 1$, the income effect dominates the substitution effect, and a lower expected consumption growth rate actually implies a *higher* stock price. This means that stock returns will be positively related to both $\widehat{p}_3$ and $\widehat{p}_4$ in a CRRA economy (and the effect is stronger for $\widehat{p}_4$, since state 4 has the lowest expected consumption growth rate). This can also be seen from Figure 10, which plots the model-implied dividend-price ratio in the constrained CRRA model together with the historical dividend-price ratio in the data. The two time series are actually negatively correlated. Altogether, we can conclude that with CRRA preferences the stock-bond return correlation reacts slightly negatively to an increase in $\widehat{p}_3$ and strongly positively to an increase in $\widehat{p}_4$, which is at odds with the empirical data.

So to obtain results similar to ours, but in a CRRA model, one has to include a feature like money illusion to make inflation enter the pricing kernel, and one has to keep the variation in the expected consumption growth rate small enough to mitigate the consequences of the typical counterintuitive CRRA result that prices are lower in higher growth states. This is exactly the path taken by David and Veronesi (2013). They assume that the agent (in their setup irrationally) bases real decisions partly on nominal variables. Basak and Yan (2010) show that, with CRRA utility, this assumption results in a pricing kernel which is composed of the original real pricing kernel and an adjustment for inflation. Since the estimation in David and Veronesi (2013) relies on both asset pricing and macroeconomic data, they find that expected consumption growth is hardly varying across states, so they constrain it to be equal across states in their following numerical evaluation (p. 703). This clearly contradicts our estimation results presented above. Relying on recursive utility, we do not have to restrict the fundamental dynamics in such a way, and we also do not need to assume any sort of bounded rationality on the part of the representative investor.

6.3 Full Information

In our model we make the assumption that the current state of the economy is unobservable and has to be filtered from macro data. To see why this assumption is necessary, we have also solved a version of our model with full information. The theoretical solution is presented in Appendix D, the proof is a slight modification of the proof in Appendix A of Branger et al. (2016). In an economy with four states, we obtain four possible values for the wealth-consumption ratio, for the price-dividend ratio and for the price of a nominal bond with a given maturity.

To evaluate the model with full information, we perform the following exercise, which is similar to the exercise in Section 5. Our Markov chain estimation above gives time series of ex-post probabilities that the economy was in one of the four states at a particular historical point in time. We round these probabilities (which are already close to 0 or 1) to exactly 0 or 1, respectively, so that we get a "clean" historical time series of economic states. Combining this time series with the model solution gives us historical time series of wealth-consumption ratios, price-dividend ratios and bond prices under the assumption of full information. From these time series and historical consumption and inflation shocks, we then again compute time series of model-implied returns, and from these model-implied returns we compute rolling-window correlations and volatilities as before.

For brevity, we discuss the stock-bond return correlation results only because the failure of a full information model becomes most evident here. Figure 11 shows the respective time series from model-generated and empirical data. The two time series differ substantially throughout the sample. Most importantly, the model with full information has problems to generate the negative correlation between stock and bond returns that we see in the data in the most recent years.

The reason for this failure can be traced back to the estimated Markov chain transition probabilities (Table 2). According to the estimation, there is a zero probability to enter the deflationary state 4 from state 2 or 3. In an economy with full information, this implies that the threat of entering the worst possible economic state in the next quarter is shut down completely for about 70% of the time. In an economy with partial information, on the other hand, there is always a nonzero subjective probability assigned to being in a state from which the economy can slide into state 4. Consequently, the average level of disaster risk is much lower in the full information economy, and moreover the pricing effect of the fear of deflation, as described in Section 5, is reduced substantially. For instance, in the full information economy the price-dividend ratio of the stock is the largest in states 2 and 3 (155.5 and 148.1, respectively) and the smallest in states 1 and 4 (145.6 and 143.0). Since this fear of deflation is the key mechanism that generates negative comovement between stock and bond returns, the full information model cannot reproduce the negative stock-bond return correlation.

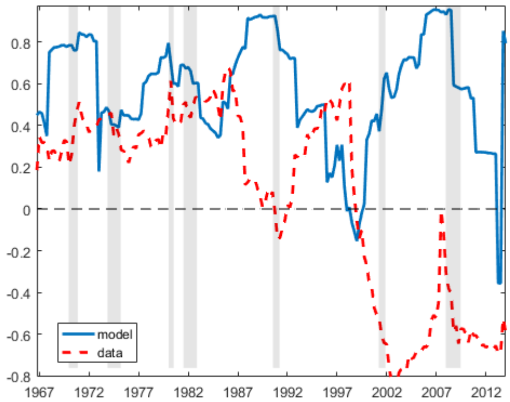

Figure 11: Conditional stock-bond return correlation with full information

The figure depicts conditional 20-quarter rolling window correlations of stock returns and returns of 5-year nominal bonds. The red dashed line is based on the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve. The blue solid line is based on the returns which a version of our model with full information would have implied given the historical paths of consumption, inflation, and state variables. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is -0.05.

6.4 Different samples

The finding that extreme inflation provides information about low expected real consumption growth is, up to now, based on a single Markov chain estimation. It may thus hinge on the data sample. In order to alleviate this concern, we repeat the estimation with four alternative data samples. The first two are subsamples of our quarterly consumption and inflation data starting in 1962 and 1965, respectively, which have been used by David and Veronesi (2013) and Burkhardt and Hasseltoft (2012). Moreover, we analyze monthly US consumption and inflation data, which are available from 1959 onwards. Finally, we re-estimate our model with GDP growth rates instead of consumption growth rates, which are available on a quarterly basis starting in 1947.

Figure 12 shows the proxies for the time-varying disaster intensity obtained by applying exactly the same methodology as before to the alternative samples. More precisely, we proceed as follows. For every sample, we estimate the time series model as defined in Equations (1) and (2). The number of states identified by the Bayes Information Criterion is four for all samples. As in the benchmark estimation, we label the states in which the conditional expected consumption growth rate μ_i^C is below the unconditional average consumption growth rate as "bad states". This refers to two of the four states in every case. We then compute the filtered probabilities $\widehat{p}_i$ for these two bad states and add them. In the upper pictures in Figure 12, the red dashed line shows the sum of these two probabilities $\widehat{p}_i$. The blue solid line is the same as in Figure 3. Finally, we also repeat the estimation without inflation data, i.e., with only consumption or GDP data. In these cases we have only two states, and we treat the state with the lower expected consumption

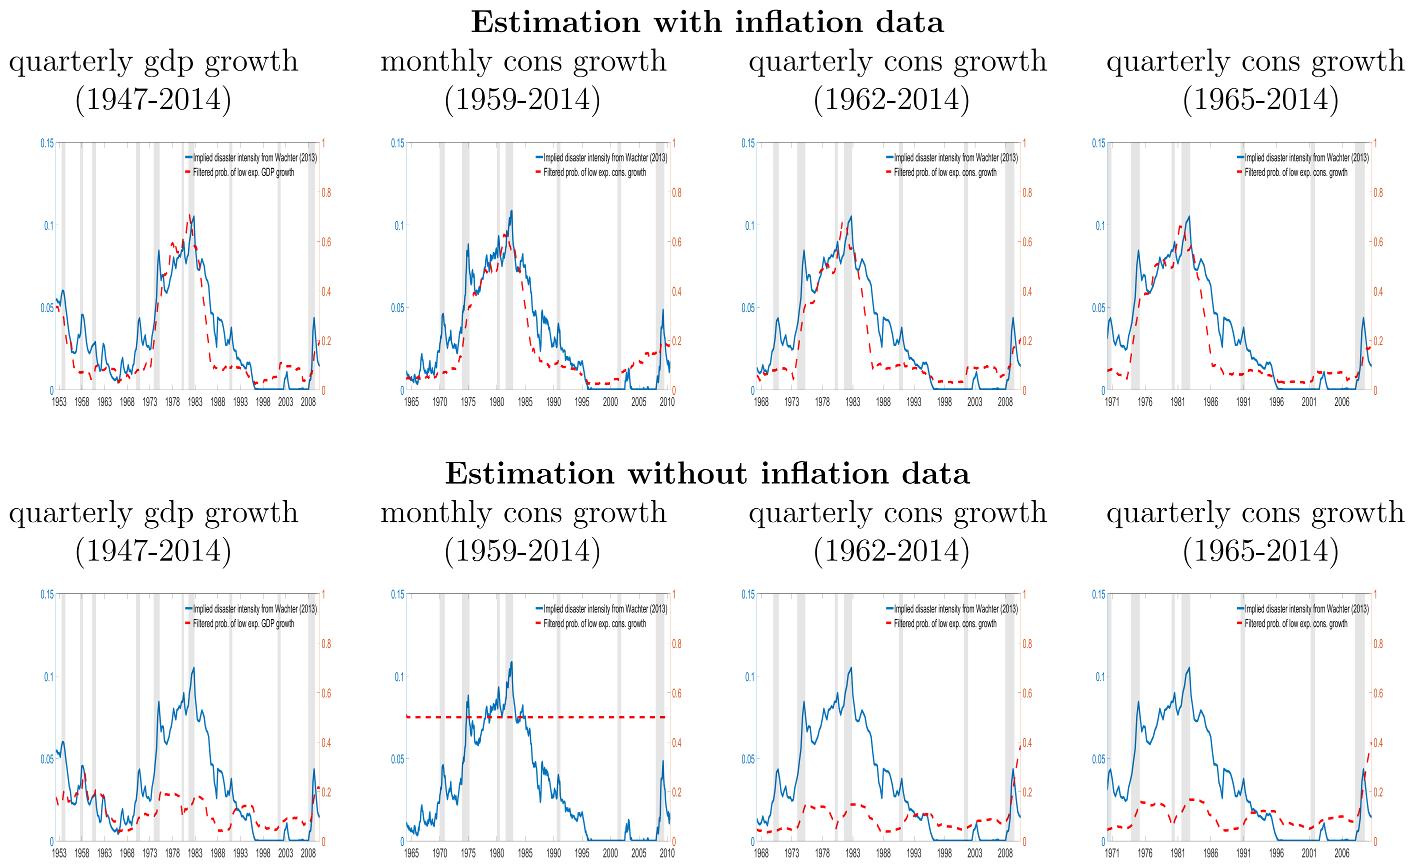

Figure 12: Time-varying disaster probabilities (robustness for different samples)

The figure depicts the results from applying the estimation methodology described in Section 3.2 to different samples. The upper graphs show the time series for the filtered probabilities for states with low expected consumption growth obtained using both consumption (or GDP) and inflation data. For the lower graphs we use consumption (or GDP) data only. The two pictures on the left hand side are based on quarterly GDP growth rates (1947Q1 to 2014Q1) instead of consumption growth rates. The next two pictures are based on an estimation using monthly consumption growth rates instead of quarterly consumption growth rates, which are available from 1959 onwards. The final four pictures are obtained from the subsamples of quarterly consumption growth rates starting in 1962 and 1965, respectively.

growth rate as the bad state. The filtered probability of this state is depicted in the lower row of pictures.

We can draw two conclusions from this exercise. First, recovering the time series of implied disaster probabilities from Wachter (2013) is to a very large degree independent of the specific sample that we use. In each of the graphs in Panel A, the red dashed line tracks the blue line very closely, the correlations between the two time series are in fact even higher than for the benchmark sample (0.88, 0.90, 0.89, and 0.89, respectively). Second, the result that the replication fails with consumption (or GDP) data only is also confirmed. The best fit is obtained in the case with GDP instead of consumption data, but the correlation between the two time series is 0.51 only. The monthly consumption data is too noisy to replicate the time-varying disaster probability.

6.5 Constrained model specifications

Besides analyzing alternative samples, one might also consider imposing more structure on the Markov chain model. The benchmark specification has 27 free parameters to estimate, so there may be constrained versions of the model in which the number of parameters can be reduced without losing too much explanatory power. One may even allow for more states, but, e.g., restrict expected consumption growth rates or expected inflation to be the same across some of these states. Generally, the number of possible constrained models is infinitely large, and we think that an unconstrained estimation provides the cleanest setup. Nevertheless, given our interpretation of states 3 and 4 as the two bad states, an obvious candidate constrained model is one in which expected consumption growth is equal across the two good states and across the two bad states, i.e. $\mu_1^C = \mu_2^C$ and $\mu_3^C = \mu_4^C$. This constraint is also justified by the fact that our estimates of the expected consumption growth rates in the two good states and in the two bad states are relatively close together. For instance, using the bootstrapped sample paths, an F-test yields that the joint hypothesis $\mu_1^C = \mu_2^C$ and $\mu_3^C = \mu_4^C$ cannot be rejected at the 10%-level. Moreover, among the many different constrained models we have estimated, this particular one exhibits the lowest BIC value and should thus be preferred.

Figure 13 presents the proxy for time-varying disaster risk that we obtain in the constrained model. This time series is very similar to those that were generated using the benchmark unconstrained model. Moreover, in additional tests, we have also solved the asset pricing model with the parameters from this constrained model. None of our asset pricing findings change when we use this specification.

To sum up, the finding that extreme inflation helps to recover the time-varying probability of consumption disasters is robust across different samples and robust to the constraint of equal expected real growth rates in the two bad states. If anything, the benchmark specification, on which we rely throughout the paper, yields conservative estimates with respect to the correlation between our and Wachter's measure for the probability of being in a bad real growth state. Nevertheless we stick to this specification, since it uses the longest quarterly sample available, and it does not impose any constraints on the identified consumption and inflation states.

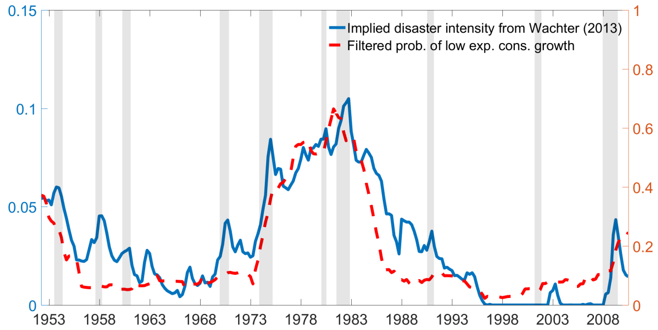

Figure 13: Time-varying disaster probabilities for constrained model

The blue solid line depicts the time-varying disaster intensity which Wachter (2013) extracts from asset price data. The red dashed line shows the 20-quarter moving average of the estimated $\widehat{p}_3 + \widehat{p}_4$ from the constrained model with equal expected consumption growth in states 1 and 2 and in states 3 and 4, respectively, for the period from 1947 to 2014. The correlation between the two time series is 0.88.

6.6 Unconditional asset pricing moments

The unconditional asset pricing moments generated by our model are computed via the same Monte Carlo simulation as described in Section 5. The results are shown in Table 9. When interpreting the numbers, one has to keep in mind that our model is estimated only on the basis of fundamental data for consumption and inflation, i.e., it is not calibrated to match unconditional return moments, and that there are no additional risk factors like disaster risk or stochastic volatility. So it should not come as a surprise that the model does not match the data perfectly with respect to unconditional risk premia or volatilities. The equity premium generated by our model, for instance, is roughly 1 percentage point and the equity return volatility is 6.7 percentage points, which reflects the fact that there is relatively little variation in the market prices of risk. The average spread between bonds with a maturity of 5 years and those with 3 months is small on average in the data and in the model (where it is basically equal to zero). Finally, the unconditional stock-bond correlation is matched pretty well by the model.

We can improve the unconditional asset pricing moments considerably with a few very slight modifications of the parametrization. Besides the unconditional moments from our benchmark parametrization, Table 9 presents unconditional moments from three such modifications. Parametrization 2 is the same as the benchmark parametrization except that we lower the expected consumption growth in states 3 and 4 by one standard error, i.e. we set $\mu_3^C = 0.129$ and $\mu_4^C = -1.628$. Such a parametrization reflects results from an estimation with GDP growth instead of consumption growth data (details not shown here for brevity). Parametrization 3 is the same as the benchmark parametrization except that we increase the probability of the Markov chain to remain in state 3 or 4 by one respective standard error compared to the benchmark estimate. The new third and fourth row of the transition matrix are then given as $(0.057, 0.027, 0.916, 0.000)$ and $(0.000, 0.137, 0.000, 0.863)$, respectively. This parametrization reflects the fact that the

	Data	Benchmark	Param. 2	Param. 3	Param. 4
Avg. nominal equity return	0.116	0.070	0.104	0.087	0.080
		(0.005)	(0.006)	(0.006)	(0.009)
Vol. of nominal equity returns	0.148	0.067	0.174	0.099	0.149
		(0.005)	(0.013)	(0.007)	(0.012)
Avg. nominal 3m rate	0.043	0.060	0.076	0.071	0.060
		(0.005)	(0.004)	(0.007)	(0.005)
Vol. of nominal rate	0.050	0.006	0.007	0.011	0.006
		(0.001)	(0.001)	(0.002)	(0.001)
Avg. yield spread (5y - 3m)	0.010	-0.005	-0.005	-0.025	-0.005
		(0.003)	(0.002)	(0.019)	(0.003)
Stock-bond correlation	0.114	-0.027	0.43	-0.017	-0.027
		(0.158)	(0.114)	(0.172)	(0.158)

Table 9: Unconditional asset pricing moments

The table shows unconditional asset pricing moments. "Data" refers to the CRSP value-weighted index for stocks and to the data set provided by Gürkaynak et al. (2007) for bonds. The model-implied values are computed via Monte Carlo simulation. All numbers are computed based on monthly observations and then annualized. In the data the average yield spread (5y - 3m) is available from 1952 onwards. The correlation between nominal stock and 5y-bond returns is based on five-year rolling window estimates with data from 1962 onwards. The column labeled Benchmark presents results for the model with the estimated parameters from Table 2. Parametrization 2 is the same as the benchmark parametrization except that we lower the expected consumption growth in states 3 and 4 by one standard error, i.e. we set $\mu_3^C = 0.129$ and $\mu_4^C = -1.628$. Parametrization 3 is the same as the benchmark parametrization except that we increase the transition probabilities from states 3 and 4 to the other states by one standard error, i.e. we set the third and fourth row of the transition matrix to (0.057 0.027 0.916 0.000) and (0.000 0.137 0.000 0.863), respectively. Parametrization 4 is the same as the benchmark parametrization except for a leverage parameter $\phi = 7$.

bad states (in particular state 4) are relatively rare and their persistence is thus estimated relatively imprecisely (see Table 2). Parametrization 4 is the same as the benchmark parametrization except for a leverage parameter $\phi = 7$.

In all three modifications, the equity premium is about 1-2 percentage points higher than in the benchmark cases. The equity return volatility increases considerably and matches the volatility in the data. The largest overall effect on unconditional moments can be seen in Parametrization 2. This is in line with existing research on asset pricing models with disaster risk, in which parameters governing the severeness of disasters (i.e. the size of potential losses in consumption growth upon a disaster) have the largest effect on the equity premium (see, e.g., the discussion in Barro (2006)). Calibrating such a model to post-war consumption data only, one may considerably underestimate the true macroeconomic volatility.

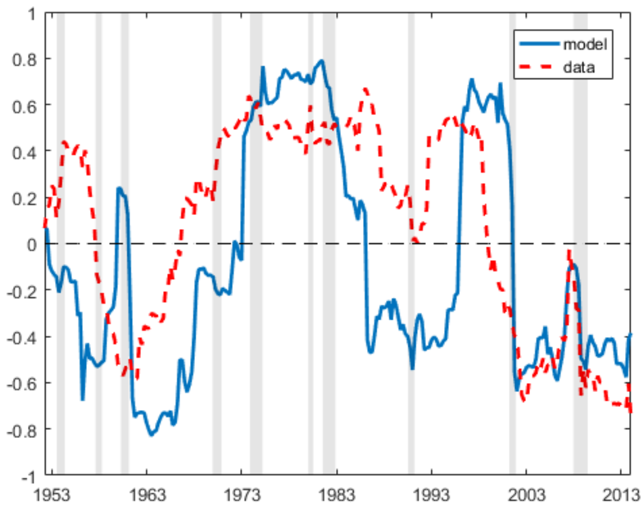

Figure 14: Conditional stock-bond return correlation

The figure depicts conditional 20-quarter rolling window correlations of stock returns and returns of long-term Treasury bonds. The red dashed line is based on the CRSP value-weighted index and the Ibbotson U.S. Long-Term Government Bond Index from 1947 till 2014. The blue solid line is based on the returns which our model would have implied given the historical paths of consumption, inflation, and state variables. The model parameters are estimated using macroeconomic data since 1947. The correlation between the data and the model-implied time series is 0.54.

6.7 Bond return data

Finally, we provide another robustness check with regard to the bond return data. Our benchmark data is the sample of interpolated term structures of U.S. Treasury bonds analyzed by Gürkaynak et al. (2007). This sample is shorter than our macroeconomic data sample. Given that our macro sample starts in 1947, we can in principle compute model-implied asset returns for the whole time period from 1947 to 2014. As a robustness check, we therefore determine the stock-bond correlation in the data from an alternative sample, namely the holding period return of the Ibbotson U.S. Long-Term Government Bond Index which measures the performance of 20-year maturity U.S. Treasury bonds. These index returns are available from 1926 onwards, but since our macro estimation uses quarterly data after 1947, we constrain ourselves to the sample from 1947 until 2014.

We repeat the whole asset pricing analysis from above using this bond index, and Figure 14 shows the resulting time series of rolling window correlations. As one can see from the figure, the model-implied correlation tracks the correlation in the data closely. The two time series have a correlation of 0.54, i.e., they exhibit the same correlation over this longer sample as over our shorter benchmark sample. Interestingly, in the data we now see another period with negative correlation in the very early part of the sample, and our model captures this negative correlation as well. We thus conclude that our results are not specific to the time series of bond prices chosen for the analysis.

7 Conclusion

Long-run risk models for asset pricing rest on the key assumption that the conditional distribution of consumption growth is time-varying. We provide evidence in favor of this assumption using a very stylized Markov regime switching model for expected consumption growth. While already interesting by itself, this reduced-form approach turns out to be very fruitful when it comes to explaining the joint dynamics of real and nominal asset prices. Augmenting the time series model by inflation as a second macro variable significantly alters the estimated regimes. In particular, we find two states in which expected consumption growth is low, one with high expected inflation and one with negative expected inflation. Embedding the estimated dynamics in a standard general equilibrium asset pricing model with recursive preferences and learning allows us to match time series of aggregate stock return volatility and the stock-bond return correlation.

The basic intuition underlying our results is that low consumption growth tends to occur together with either very high or very low inflation. In contrast to the volatility of stock returns, where mainly the *overall* probability of the two bad states for expected consumption growth matters, it is the distinction between the two with respect to expected inflation which is relevant for the stock-bond return correlation. In the high expected inflation state, stocks and bonds will both tend to have negative returns, so that their correlation will be positive, while in the deflationary state, stocks will still do poorly, but nominal bonds will exhibit positive returns, resulting in a negative correlation between the two types of assets.

Our research design differs substantially from other approaches to calibrate dynamic asset pricing models, where often both asset pricing moments and macroeconomic moments are used to identify the deep parameters of the model. This often leads to a parametrization where macro dynamics are not matched very well. Given the parsimony of the research design, in particular the fact that the model is estimated from two macro time series only, we regard our results as evidence that the long-run risk or disaster risk paradigm in asset pricing can be extended towards an explanation of the time-varying stock-bond return correlation when the signaling role of inflation is properly accounted for. In particular, we also document that our filtered probability of being in a bad consumption growth regime closely tracks the evolution of state variables like the "time-varying disaster risk" which Wachter (2013) obtains through reverse-engineering from asset prices.

In summary, our paper shows that a large part of the variation in asset prices may actually be attributable to inflation risk, and that the long-run risk model class provides a promising framework to study the link between inflation, consumption growth and both real and nominal asset returns.

APPENDIX

A. Maximum Likelihood Estimation

The maximum likelihood estimation is based on Hamilton (1994). The parameters of the model and the transitions probabilities p_{ij} are collected in the vector Θ. This vector Θ is estimated based on the data Y_T that is observed until time T. This data forms a $T \times 2$ matrix. Let $P(s_t = j|\Theta, Y_t)$ define the probability of state s_t conditional on all data until t and conditional on the knowledge of the parameters Θ. The econometrician assigns $P(s_t = j|\Theta, Y_t)$ to the possibility that the observation at time t is generated by state j. These conditional probabilities are collected in the $n \times 1$ vector $\hat{\xi}_{t|t}$, where n denotes the number of states, which is fixed throughout the entire estimation. We have

$$\begin{aligned} \hat{\xi}_{t|t} &= \frac{\hat{\xi}_{t|t-1} \odot \eta_t}{\mathbf{1'} \cdot (\hat{\xi}_{\mathbf{t|t-1}} \odot \eta_{\mathbf{t}})} \\ \hat{\xi}_{t+1|t} &= \mathbf{Q} \cdot \hat{\xi}_{\mathbf{t|t}}. \end{aligned}$$

Here η_t is the conditional density, whose jth element is

$$f(y_t|s_t = j; \Theta) = \frac{1}{\sqrt{2\pi|\Omega|}} \exp\left\{0.5 \cdot (y_t - \mu_j)'\Omega^{-1}(y_t - \mu_j)\right\},$$

where y_t is an $n \times 1$ vector, $\mu_j = (\mu_j^C, \mu_j^\pi)'$ and $\Omega = \Sigma\Sigma'$ where

$$\Sigma = \begin{pmatrix} \sigma^C\sqrt{1-\rho^2} & \sigma^C\rho \\ 0 & \sigma^\pi \end{pmatrix}$$

$\mathbf{Q}$ is the transition probability matrix

$$\mathbf{Q} = \begin{pmatrix} q_{11} & \cdots & q_{1n} \\ \vdots & \ddots & \vdots \\ q_{n1} & \cdots & q_{nn} \end{pmatrix},$$

$\odot$ denotes element-by-element multiplication, and $\mathbf{1}$ denotes a vector of ones. The jth element of the product $\hat{\xi}_{t|t-1} \odot \eta_t$ is interpreted as the conditional joint density of $\mathbf{y_t}$ and s_t

$$P(s_t = j|\Theta) \times f(y_t|s_t = j; \Theta) = p(\mathbf{y_t}, \mathbf{s_t} = \mathbf{j}|\mathbf{\Theta}).$$

The density of the observed vector $\mathbf{y_t}$ is the sum

$$f(y_t|\Theta) = \mathbf{1'} \cdot (\hat{\xi}_{\mathbf{t|t-1}} \odot \eta_{\mathbf{t}})$$

The objective is to find a maximum of the log likelihood function

$$\mathbb{L}(\Theta) = \sum_{t=1}^{T} \log f(y_t|\Theta).$$

The starting value for the maximization, $\hat{\xi}_{1|0}$, is set to $\frac{1}{n} \cdot \mathbf{1}$.

B. Filtering

B.1. Dynamics of the state variables

The dynamics of consumption and the log prive level can be rewritten as

$$\begin{aligned} d\ln C_t &= \left(\sum_{i=1}^{n} \mu_i^C \mathbb{1}_{\{S_t=i\}}\right) dt + \sigma^C \left(\sqrt{1-\rho^2} dW_t^C + \rho dW_t^\pi\right) \\ d\pi_t &= \sum_{i=1}^{n} \mu_i^\pi \mathbb{1}_{\{S_t=i\}} + \sigma^\pi dW_t^\pi, \end{aligned}$$

where $\mathbb{1}_{\{S_t=i\}}$ is the indicator function equal to one, if the economy is in state i at time t, $(i = 1, \ldots, n)$, and equal to zero otherwise. In matrix form, this becomes

$$\begin{pmatrix} d\ln C_t \\ d\pi_t \end{pmatrix} = \begin{pmatrix} \sum_{i=1}^{n} \mu_i^C \mathbb{1}_{\{S_t=i\}} \\ \sum_{i=1}^{n} \mu_i^\pi \mathbb{1}_{\{S_t=i\}} \end{pmatrix} dt + \Sigma \begin{pmatrix} dW_t^C \\ dW_t^\pi \end{pmatrix},$$

with

$$\Sigma = \begin{pmatrix} \sigma^C \sqrt{1-\rho^2} & \sigma^C \rho \\ 0 & \sigma^\pi \end{pmatrix}$$

and $d[W^C, W^\pi] = 0$.
The inverse of the of the volatility matrix Σ is

$$\Sigma^{-1} = \frac{1}{\sigma^C \sigma^\pi \sqrt{1-\rho^2}} \begin{pmatrix} \sigma^\pi & -\sigma^C \rho \\ 0 & \sigma^C \sqrt{1-\rho^2} \end{pmatrix}.$$

We assume that the drift rates are unobservable and need to be filtered by the investor. Mathematically, there are two filtrations, $\mathcal{F}$ and $\mathcal{G}$, where $\mathcal{F}$ is generated by the processes $(C_t)_t$, $(\pi_t)_t$ and $(S_t)_t$, whereas $\mathcal{G} \subset \mathcal{F}$ is generated by the processes $(C_t)_t$ and $(\pi_t)_t$ only. The equilibrium in the economy is based on the dynamics of $(C_t)_t$ and $(\pi_t)_t$ under the investor filtration $\mathcal{G}$, i.e. on the projections $\widehat{\mu}_t^C = \mathbb{E}\left[\mu^C(S_t)|\mathcal{G}_t\right] = \sum_{i=1}^{n} \widehat{p}_{it} \mu_i^C$ and $\widehat{\mu}_t^\pi = \mathbb{E}\left[\mu^\pi(S_t)|\mathcal{G}_t\right] = \sum_{i=1}^{n} \widehat{p}_{it} \mu_i^\pi$, where $\widehat{p}_{it} = \mathbb{E}\left[\mathbb{1}_{\{S_t=i\}}|\mathcal{G}_t\right]$.

An application of Theorem 9.1 on p. 355 of Liptser and Shiryaev (2001) yields dynamics for the projected (henceforth also called "subjective") probabilities:

$$d\widehat{p}_{it} = \left(\widehat{p}_{it}\lambda_{ii} + \sum_{j \neq i} \widehat{p}_{jt}\lambda_{ji}\right) dt + \widehat{p}_{it} \left[\begin{pmatrix} \mu_i^C \\ \mu_i^\pi \end{pmatrix} - \sum_{j=1}^{n} \widehat{p}_{jt} \begin{pmatrix} \mu_j^C \\ \mu_j^\pi \end{pmatrix}\right]' \times \Sigma'^{-1} \begin{pmatrix} d\widehat{W}_t^C \\ d\widehat{W}_t^\pi \end{pmatrix},$$

where

$$\begin{pmatrix} d\widehat{W}_t^C \\ d\widehat{W}_t^\pi \end{pmatrix} = \Sigma^{-1} \left[\begin{pmatrix} \mu_i^C \\ \mu_i^\pi \end{pmatrix} - \sum_{j=1}^{n} \widehat{p}_{jt} \begin{pmatrix} \mu_j^C \\ \mu_j^\pi \end{pmatrix}\right] dt + \begin{pmatrix} dW_t^C \\ dW_t^\pi \end{pmatrix}.$$

In particular, $d[\widehat{W}^C, \widehat{W}^\pi] = 0$. Under the investor's filtration, log consumption dynamics are given as

$$d\ln C_t = \sum_{i=1}^{n} \mu_i^C \widehat{p}_{it} dt + \sigma^C \left(\sqrt{1-\rho^2} d\widehat{W}^C + \rho d\widehat{W}^\pi \right)$$

For notational convenience, we define

$$\sigma_{\hat{p}_i} d\widehat{W}_t \equiv \widehat{p}_{it} \left[\begin{pmatrix} \mu_i^C \\ \mu_i^\pi \end{pmatrix} - \sum_{j=1}^{n} \widehat{p}_{jt} \begin{pmatrix} \mu_j^C \\ \mu_j^\pi \end{pmatrix} \right]' \times \Sigma'^{-1} \times \begin{pmatrix} d\widehat{W}_t^C \\ d\widehat{W}_t^\pi \end{pmatrix}$$

and

$$\begin{aligned} \sigma_{c,\hat{p}_i} &\equiv \left(\rho\sigma^C, \quad \sqrt{1-\rho^2}\sigma^C \right) \times \sigma'_{\hat{p}_i} \\ \sigma_{\pi,\hat{p}_i} &\equiv (0, \quad \sigma^\pi) \times \sigma'_{\hat{p}_i}. \end{aligned}$$

B.2. Discussion

In order to understand the dynamics of $\widehat{p}_{it}$, it is instructive to analyze the drift and the diffusion component separately. The drift in (3) is a linear function of the transition intensities λ and the current estimates of the probabilities $\widehat{p}$. Since the states (and consequently also switches between states) are unobservable, the subjective probability of being in state i changes deterministically over time, depending on the conditional probabilities to enter or exit state i. The drift therefore comprises two terms. The first term, $\widehat{p}_{it}\lambda_{ii} = -\widehat{p}_{it} \sum_{j \neq i} \lambda_{ij}$, involves the intensities for a switch from state i to some other state $j \neq i$. Loosely speaking, the more time goes by, the higher the chance that an unobserved switch from state i to some other state j has occurred in the meantime. This effect induces a negative drift for $\widehat{p}_{it}$, i.e., the estimated probability of still being in state i decreases in expectation. The second term, $\sum_{j \neq i} \widehat{p}_{jt}\lambda_{ji}$, captures the probabilities of entering state i, given that the economy is currently in a state j different from i. Suppose one of the $\widehat{p}_{jt}$ $(j \neq i)$ is currently large. Then, as time passes and if no other conflicting signals arrive, it becomes more and more likely that an unobserved switch to state i has occurred in the meantime. This effect induces a positive drift in $\widehat{p}_{it}$. The overall sign of the drift of $\widehat{p}_{it}$ thus depends on the current estimate of all state probabilities. In particular, the drift terms ensure that the probabilities $\widehat{p}$ will always be between zero and one.

The volatility of the change in $\widehat{p}_i$ is a *quadratic* function of all probabilities $\widehat{p}_j$ $(j = 1, \ldots, n)$. The probability update is largest when the investor is rather uncertain about the current state of the economy, i.e., for intermediate values of $\widehat{p}_i$. When the investor is almost sure about the current state of the economy (i.e., when one of the $\widehat{p}_j$ is close to one and the others are close to zero), the estimated probability will change in a basically deterministic fashion. To see this, note that when the respective estimate $\widehat{p}_i$ is close to zero, the diffusion term in (3) is obviously also close to zero, since $\widehat{p}_i$ is one factor of the product in front of the Wiener innovations. When $\widehat{p}_i$ is in turn close to one, the term in square brackets in (3) will be very close to zero, since in the sum only the term involving $\widehat{p}_i$ will remain, whereas all the others will vanish.

The volatility of the innovation in the filtered probability also depends on the precision of the signals. When the signals are very imprecise, i.e., when the volatilities σ^C and σ^π are large, an observed innovation in (log) consumption growth or inflation delivers less information about the true state, and the investor will put less weight on them when computing the new estimate for p_i.

The sign of the diffusion term depends on the sign of the 'observed' Brownian shocks $d\widehat{W}$. These are defined via the restriction that $\ln C$ and π are observable, so technically speaking they have to be adapted to both $\mathcal{F}$ and $\mathcal{G}$, which implies $\mu(S_t)dt + \Sigma dW_t = \widehat{\mu}_t dt + \Sigma d\widehat{W}_t$.

C. Solution of the Model with Partial Information

C.1. Wealth-consumption ratio

The indirect utility function of the investor is given by

$$J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}) = E_t \left[\int_t^\infty f(C_s, J(C_s, \widehat{p}_{1s}, \ldots, \widehat{p}_{ns}))\, ds \right].$$

$J\left(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}\right) + \int_0^t f\left(C_s, J(C_s, \widehat{p}_{1s}, \ldots, \widehat{p}_{ns})\right)\, ds$ is a martingale, therefore we have the Bellman equation

$$E[dJ(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}) + f(C_t, J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}))\, dt] = 0,$$

or, equivalently,

$$\frac{\mathcal{A}J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt})}{J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt})} + \frac{f(C_t, J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}))}{J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt})} = 0, \tag{C.1}$$

where $\mathcal{A}$ is the infinitesimal generator. The aggregator function is given by

$$f(C, J) = \frac{\beta C^{1-\frac{1}{\psi}}}{\left(1 - \frac{1}{\psi}\right) \left[(1-\gamma)J\right]^{\frac{1}{\theta}-1}} - \beta\theta J.$$

We conjecture a functional form for J:

$$J(C_t, \widehat{p}_{1t}, \ldots, \widehat{p}_{nt}) \quad = \quad \frac{C_t^{1-\gamma}}{1-\gamma} \left(\beta e^{v(\widehat{p}_{1t}, \ldots, \widehat{p}_{nt})}\right)^\theta$$

where in the end v will turn out to be the log wealth-consumption ratio. This functional form together with the definition of $f(C, J)$ implies

$$\frac{f(C, J)}{J} \quad = \quad \theta e^{-v} - \theta\beta.$$

With $I \equiv e^v$ the partial derivatives of J (denoted by subscripts) are

$$\begin{aligned}
J_c &= C^{-\gamma} \left(\beta e^v\right)^\theta \\
J_{cc} &= -\gamma C^{-\gamma-1} \left(\beta e^v\right)^\theta \\
J_{\widehat{p}_i} &= \frac{C^{1-\gamma}}{1-\gamma} \beta^\theta \theta I^{\theta-1} I_{\widehat{p}_i} \\
J_{\widehat{p}_i\widehat{p}_j} &= \frac{C^{1-\gamma}}{1-\gamma} \beta^\theta \theta \left[(\theta-1) I^{\theta-2} I_{\widehat{p}_i}^2 + I_{\widehat{p}_i\widehat{p}_j} I^{\theta-1}\right] \\
J_{c\widehat{p}_i} &= C^{-\gamma} \beta^\theta \theta I^{\theta-1} I_{\widehat{p}_i},
\end{aligned}$$

resulting in

$$\begin{aligned}
\frac{J_c}{J}dC_t &= (1-\gamma)\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it}dt + \frac{1}{2}(\sigma^C)^2(1-\gamma)dt + (1-\gamma)\sigma^C\left(\sqrt{1-\rho^2}d\widehat{W}_t^C + \rho d\widehat{W}_t^\pi\right) \\
\frac{J_{cc}d[C,C]_t}{J} &= (\sigma^C)^2(-\gamma)(1-\gamma)dt \\
\frac{J_{c\widehat{p}_i}d[C,\widehat{p}_i]_t}{J} &= \theta(1-\gamma)\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i}dt \\
\frac{J_{\widehat{p}_i}d\widehat{p}_{it}}{J} &= \theta\frac{I_{\widehat{p}_i}}{I}d\widehat{p}_{it} \\
\frac{J_{\widehat{p}_i\widehat{p}_j}d[\widehat{p}_i,\widehat{p}_j]_t}{J} &= \frac{1}{2}\theta d[\widehat{p}_i,\widehat{p}_j]_t\left[(\theta-1)\left(\frac{I_{\widehat{p}_i}}{I}\right)^2 + \frac{I_{\widehat{p}_i\widehat{p}_j}}{I}\right].
\end{aligned}$$

Plugging everything into (C.1) results in the following partial differential equation for I:

$$\begin{aligned}
0 = &\left[(1-\gamma)\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it} + \frac{1}{2}(1-\gamma)^2(\sigma^C)^2 - \beta\theta\right] + \theta I^{-1} \\
&+\sum_{i=1}^{n-1}\theta\frac{I_{\widehat{p}_{it}}}{I}\left(\widehat{p}_i\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) + \sum_{i=1}^{n-1}\theta(1-\gamma)\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i} \\
&+\frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\theta\left[(\theta-1)\left(\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I}\right) + \frac{I_{\widehat{p}_i\widehat{p}_j}}{I}\right]\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j},
\end{aligned} \tag{C.2}$$

There are $n-1$ state variables due to the restriction $\sum_{i=1}^{n}\widehat{p}_{it} = 1$. We solve the PDE with a Chebyshev approximation similar to Benzoni et al. (2011). We guess the following functional form for I as a function of the vector $\widehat{p} = (\widehat{p}_1, \widehat{p}_2, ..., \widehat{p}_{n-1})$ with:

$$\begin{aligned}
I(\widehat{p}) &= \exp(B(\widehat{p})) \\
B(\widehat{p}) &= \sum_{j=0}^{d}\alpha_j T_j(\widehat{p}),
\end{aligned} \tag{C.3}$$

where the $T_j(\widehat{p})$ are multivariate Chebyshev polynomials. For the interval $[-1,1]$, the univariate Chebyshev polynomials are defined recursively via

$$\begin{aligned}
T_0(x) &= 1 \\
T_1(x) &= x \\
T_{d+1}(x) &= 2xT_d(x) - T_{d-1}(x).
\end{aligned}$$

Univariate Chebyshev polynomials for the general interval $[a,b]$ are given by transformations

$$T_d\left(\frac{2x-b-a}{b-a}\right).$$

Multivariate versions of the Chebyshev polynomials are defined as sums of products of the univariate ones. The derivatives of the guess in (C.3) are

$$
\begin{aligned}
I_{\widehat{p}_i} &= e^{B(\widehat{p})} B_{\widehat{p}_i} = e^{B(\widehat{p})} \sum_{j=1}^{d} \alpha_j \frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p}) \\
I_{\widehat{p}_i\widehat{p}_j} &= e^{B(\widehat{p})} \left[(B_{\widehat{p}_i})^2 + B_{\widehat{p}_i\widehat{p}_j}\right] = e^{B(\widehat{p})} \left[\left(\sum_{j=1}^{d} \alpha_j \frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p}) \right)^2 + \left(\sum_{j=2}^{d} \alpha_j \frac{\partial^2 T_j}{\partial \widehat{p}_i \partial \widehat{p}_j}(\widehat{p}) \right) \right].
\end{aligned}
$$

Plugging the partial derivatives into (C.2) gives

$$
\begin{aligned}
0 &= \left[(1-\gamma)\sum_{i=1}^{n} \mu_i^C \widehat{p}_i + \frac{1}{2}(1-\gamma)^2(\sigma^C)^2 - \beta\theta\right] \quad + \quad e^{-\sum_{j=0}^{d}\alpha_j T_j(\widehat{p})}\theta \\
&+ \sum_{i=1}^{n-1} \theta \left(\sum_{j=1}^{d} \alpha_j \frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p}) \right) \left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n} \widehat{p}_{jt}\lambda_{ji} \right) \quad + \quad \sum_{i=1}^{n-1} \theta(1-\gamma) \left(\sum_{j=1}^{d} \alpha_j \frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p}) \right) \sigma_{c,\hat{p}_i} \\
&+ \frac{1}{2}\sum_{i=1}^{n-1}\sum_{k=1}^{n-1} \theta \left[(\theta-1)\sum_{j=1}^{d} \alpha_j \frac{\partial T_j}{\partial \widehat{p}_i}(\widehat{p}) \sum_{j=1}^{d} \alpha_j \frac{\partial T_j}{\partial \widehat{p}_k}(\widehat{p}) + \sum_{j=2}^{d} \alpha_j \frac{\partial^2 T_j}{\partial \widehat{p}_i \partial \widehat{p}_k}(\widehat{p}) \right] \sigma_{\hat{p}_i}\sigma'_{\hat{p}_k}.
\end{aligned}
$$

This equation is defined on the simplex Δ^{n-1}. We partition this simplex by choosing grid points according to the Chebyshev methodology. Evaluating the equation on every grid point leaves us with a number of algebraic equations, whose solution gives the Chebyshev coefficients α_j. For the multivariate Chebyshev polynomials we choose the order of $d = 4$. We have also tried higher values for d but the solution remained unchanged.

C.2. Pricing kernel

The pricing kernel is given by

$$
\xi_t = \exp\left(\int_0^t -\beta\theta - (1-\theta) I^{-1}(\widehat{p}_s) ds \right) C_t^{-\gamma} (I(\widehat{p}_t))^{\theta-1}
$$

with dynamics

$$
\begin{aligned}
\frac{d\xi_t}{\xi_t} &= -\beta\theta dt - (1-\theta)I^{-1}\, dt - \gamma \frac{dC_t}{C_t} + \frac{1}{2}\gamma^2(\sigma^C)^2 dt \\
&- (1-\theta)\sum_{i=1}^{n-1} \frac{I_{\widehat{p}_i}}{I} d\widehat{p}_{it} + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}(\theta-1)\left[\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + (\theta-2)\left(\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I^2}\right)\right]\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j} dt \\
&- \gamma(\theta-1)\sum_{i=1}^{n-1} \frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i} dt,
\end{aligned}
$$

where $\widehat{p} = (\widehat{p}_1, \widehat{p}_2, ..., \widehat{p}_n)$. For later use, we abbreviate the drift term as

$$\mu_{\xi,t} \equiv -\beta\theta \quad - (1-\theta)I^{-1}(\widehat{p}_t) \quad - \gamma \sum_{i=1}^{n} \mu_i^C \widehat{p}_{it} \quad + \frac{1}{2}\gamma^2(\sigma^C)^2 \quad - \sum_{i=1}^{n-1}(1-\theta)\frac{I_{\widehat{p}_i}}{I}\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n} \widehat{p}_{jt}\lambda_{ji}\right)$$
$$+ \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}(\theta-1)\Big[\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + (\theta-2)\left(\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I^2}\right)\Big]\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j} \quad - \gamma(\theta-1)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i}.$$

C.3. Price-dividend ratio

We want to price a claim on levered consumption. Under the investor's filtration, dividends follow the process

$$d\ln D_t = \bar{\mu}dt + \phi\left(\sum_{i=1}^{n}(\mu_i^C - \bar{\mu})\widehat{p}_{it}\right)dt + \phi\sigma^C\left(\sqrt{1-\rho^2}d\widehat{W}_t^C + \rho d\widehat{W}_t^\pi\right).$$

Let ω denote the log price-dividend ratio. For $g(\xi, D, \omega) \equiv \xi D e^{\omega}$, the Feynman-Kac formula yields

$$\frac{\mathcal{A}g(\xi, D, \omega)}{g(\xi, D, \omega)} + e^{-\omega} = 0. \tag{C.4}$$

Itô's Lemma gives

$$\frac{\mathcal{A}g_t}{g_t} = \mu_{\xi,t} + \mu_{D,t} + \mu_{\omega,t} + \frac{1}{2}\frac{d[\omega]_t}{dt} + \frac{d\,[\xi, D]_t}{\xi D dt} + \frac{d\,[\omega, D]_t}{D dt} + \frac{d\,[\omega, \xi]_t}{\xi dt}.$$

Another application of Itô's Lemma leads to

$$d\omega_t = \sum_{i=1}^{n-1}\omega_{\widehat{p}_i}d\hat{p}_{it} + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\omega_{\widehat{p}_i\widehat{p}_j}\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}dt$$

where the subscripts denote partial derivatives with respect to the state variables $\widehat{p}_i$. We can rewrite the drift of ω as a function of the derivatives $\omega_{\widehat{p}_i}$ and $\omega_{\widehat{p}_i,\widehat{p}_j}$ and the state variables $\widehat{p}_i$:

$$\mu_{\omega,t} = \sum_{i=1}^{n-1}\omega_{\widehat{p}_i}\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) \quad + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\omega_{\widehat{p}_i\widehat{p}_j}\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}.$$

The quadratic variation terms are:

$$
\begin{aligned}
d[\omega]_t &= \sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\omega_{\widehat{p}_i}\omega_{\widehat{p}_j}\sigma_{\hat{p}_i}\sigma_{\hat{p}_j}' dt \\
\frac{d[\xi,\omega]_t}{\xi_t} &= -(1-\theta)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\omega_{\widehat{p}_i}\sigma_{\hat{p}_i}\sigma_{\hat{p}_i}' dt + \sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\omega_{\widehat{p}_j}\sigma_{\hat{p}_i}\sigma_{\hat{p}_j}' dt - \gamma\sum_{i=1}^{n-1}\omega_{\widehat{p}_i}\sigma_{c,\hat{p}_i} dt \\
\frac{d[\xi,D]_t}{\xi_t D_t} &= -\gamma\phi(\sigma^C)^2 - (1-\theta)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\phi\sigma_{c,\hat{p}_i} dt \\
\frac{d[\omega,D]_t}{D_t} &= \sum_{i=1}^{n-1}\omega_{\widehat{p}_i}\phi\sigma_{c,\hat{p}_i} dt.
\end{aligned}
$$

Plugging everything into (C.4) gives the following PDE for ω:

$$
\begin{aligned}
0 \;=\;& -\beta\theta - (1-\theta)I^{-1} + e^{-\omega} - \gamma\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it} + \bar{\mu} + \phi\left(\sum_{i=1}^{n}(\mu_i^C - \bar{\mu})\widehat{p}_{it}\right) + \frac{1}{2}(\phi-\gamma)^2(\sigma^C)^2 \\
&+ \sum_{i=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i}}{I} + \omega_{\widehat{p}_i}\right)\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1 \\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) + \sum_{i=1}^{n-1}\left((\phi-\gamma)(\theta-1)\frac{I_{\widehat{p}_i}}{I} + (\phi-\gamma)\omega_{\widehat{p}_i}\right)\sigma_{c,\hat{p}_i} \\
&+ \sum_{i=1}^{n-1}\left(\frac{1}{2}(\theta-1)(\theta-2)\left(\frac{I_{\widehat{p}_i}}{I}\right)^2 + (\theta-1)\frac{I_{\widehat{p}_i}}{I}\omega_{\widehat{p}_i} + \frac{1}{2}\omega_{\widehat{p}_i}^2\right)\sigma_{\hat{p}_i}\sigma_{\hat{p}_i}' \\
&+ \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left((\theta-1)(\theta-2)\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I^2} + \omega_{\widehat{p}_i}\omega_{\widehat{p}_j}\right)\sigma_{\hat{p}_i}\sigma_{\hat{p}_j}' \\
&+ \frac{1}{2}\sum_{i=1}^{n-1}\left((\theta-1)\left(\frac{I_{\widehat{p}_i\widehat{p}_i}}{I}\right) + \omega_{\widehat{p}_i\widehat{p}_i} + \frac{1}{2}\omega_{\widehat{p}_i}^2\right)\sigma_{\hat{p}_i}\sigma_{\hat{p}_i}' + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + \omega_{\widehat{p}_i\widehat{p}_j}\right)\sigma_{\hat{p}_i}\sigma_{\hat{p}_j}'.
\end{aligned}
$$

Similar to the wealth-consumption ratio, we approximate the price-dividend ratio $U \equiv e^{\omega}$ via a multivariate Chebyshev polynomial expansion, i.e., we approximate the function $U(\widehat{p})$ as

$$
U(\widehat{p}) \;=\; \exp\left\{\sum_{j=0}^{d}\beta_j T_j(\widehat{p})\right\},
$$

and solve the PDE numerically.

C.4. Pricing real bonds

Let the price of a real bond expiring at time T be denoted by $B_t^T = E_t[\frac{\xi_T}{\xi_t}]$ with the real pricing kernel

$$
\frac{\xi_T}{\xi_t} = \beta^{\theta}\left(\frac{C_T}{C_t}\right)^{-\gamma}\exp\left\{-\beta\theta(T-t) + (\theta-1)\left(\int_t^T I^{-1}(\hat{p}_s)ds\right)\right\}(I(\hat{p}_t))^{\theta-1}
$$

Let $b_t \equiv \ln B_t^T$ be the log real bond price (we will omit the superscript T in the following). The

Feynman-Kac formula applied to $H(\xi_t, b_t) = \xi_t e^{b_t}$ yields the partial differential equation:

$$0 \;=\; \mathcal{A}H \quad = \mu_{\xi,t} + \mu_{b,t} + \frac{1}{2}\frac{d[b]_t}{dt} + \frac{d[\xi,b]_t}{\xi_t dt}, \tag{C.5}$$

where μ_ξ and μ_b are the drifts of the processes ξ_t and b_t respectively. The dynamics of b_t are

$$db_t = \frac{\partial b_t}{\partial t}dt + \sum_{i=1}^{n-1} b_{\widehat{p}_i} d\widehat{p}_{it} + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1} b_{\widehat{p}_i\widehat{p}_j}\sigma_{\widehat{p}_i}\sigma'_{\widehat{p}_j}dt.$$

Plugging everything into (C.5) gives the following PDE for b_t:

$$\begin{aligned}
0 \;=\; & -\beta\theta - (1-\theta)I^{-1} - \gamma\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it} + \frac{1}{2}\gamma^2(\sigma^C)^2 + \frac{\partial b_t}{\partial t} \\
+ & \sum_{i=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i}}{I} + b_{\widehat{p}_i}\right)\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) \quad - \gamma(\theta-1)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i} - \gamma\sum_{i=1}^{n-1}\omega_{\widehat{p}_i}\sigma_{c,\hat{p}_i} \\
+ & \sum_{i=1}^{n-1}\left(\frac{1}{2}(\theta-1)(\theta-2)\left(\frac{I_{\widehat{p}_i}}{I}\right)^2 + (\theta-1)\frac{I_{\widehat{p}_i}}{I}b_{\widehat{p}_i} + \frac{1}{2}b^2_{\widehat{p}_i}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_i} \\
+ & \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left((\theta-1)(\theta-2)\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I} + b_{\widehat{p}_i}b_{\widehat{p}_j}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j} \\
+ & \sum_{i=1}^{n-1}\frac{1}{2}\left((\theta-1)\left(\frac{I_{\widehat{p}_i\widehat{p}_i}}{I}\right) + b_{\widehat{p}_i\widehat{p}_i} + \frac{1}{2}b^2_{\widehat{p}_i}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_i} \quad + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + b_{\widehat{p}_i\widehat{p}_j}\right)\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}.
\end{aligned}$$

We approximate the bond price e^{b_t} at each time point t by a multivariate Chebyshev polynomial, i.e.,

$$b_t \;=\; \sum_{j=0}^{n}\alpha_{j,t}^{(T)}T_j(\hat{p}).$$

Note that the PDE for the bond price involves a time derivative. We use an explicit Euler discretization for this time derivative and solve the PDE recursively backwards in time, starting from the boundary condition $b_T = 0$, i.e. $\alpha_{j,T}^{(T)} = 0$ for $j = 0, \ldots, n$.

C.5. Pricing nominal bonds

Let the price of the nominal bond maturing at time T be denoted by $B_t^{T,\$}$ with

$$\begin{aligned}
B_t^{T,\$} \;&=\; E_t\left[\frac{\xi_T^\$}{\xi_t^\$}\right] \\
&=\; E_t\left[\frac{\xi_T}{\xi_t}\frac{e^{\pi_t}}{e^{\pi_T}}\right],
\end{aligned}$$

where $\xi_t^\$ \equiv \xi_t e^{-\pi_t}$ is the nominal pricing kernel. Define $b_t^\$ \equiv \ln B_t^{T,\$}$ (we again omit the superscript T in the following). Then the Feynman-Kac formula applied to $H(\xi_t^\$, b_t^\$) = \xi_t^\$ e^{b_t^\$}$

yields the partial differential equation

$$0 \quad = \quad \mathcal{A}H \quad = \mu_{\xi^\$,t} + \mu_{b^\$,t} + \frac{1}{2}\frac{d[b_t^\$]}{dt} + \frac{d[\xi^\$, b^\$]_t}{\xi_t^\$ dt}, \tag{C.6}$$

where $\mu_{\xi^\$}$ and $\mu_{b^\$}$ are drifts of the processes $\xi_t^\$$ and $b_t^\$$ respectively. Notice that

$$\frac{d\xi_t^\$}{\xi_t^\$} = \frac{d\xi_t}{\xi_t} - d\pi_t + \frac{1}{2}d[\pi]_t - \frac{d[\xi,\pi]_t}{\xi_t}.$$

The dynamics of $b_t^\$$ are

$$db_t^\$ = \frac{\partial b_t^\$}{\partial t}dt + \sum_{i=1}^{n-1} b_{\widehat{p}_i}^\$ d\hat{p}_{it} + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1} b_{\widehat{p}_i\widehat{p}_j}^\$ \sigma_{\hat{p}_i}\sigma_{\hat{p}_j}' dt.$$

Plugging everything into (C.6) yields the following PDE for $b_t^\$$:

$$\begin{aligned}
0 \quad = \quad & -\beta\theta - (1-\theta)I^{-1} - \gamma\sum_{i=1}^{n}\mu_i^C\widehat{p}_{it} + \frac{1}{2}\gamma^2(\sigma^C)^2 + \frac{\partial b_t^\$}{\partial t} - \sum_{i=1}^{n}\mu_i^\pi\widehat{p}_{it} + \frac{1}{2}(\sigma^\pi)^2 + \gamma\rho\sigma^C\sigma^\pi \\
& + \sum_{i=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i}}{I} + b_{\widehat{p}_i^\$}\right)\left(\widehat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1\\ j\neq i}}^{n}\widehat{p}_{jt}\lambda_{ji}\right) \\
& - \gamma(\theta-1)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\sigma_{c,\hat{p}_i} - \gamma\sum_{i=1}^{n-1}b_{\widehat{p}_i}^\$\sigma_{c,\hat{p}_i} - (\theta-1)\sum_{i=1}^{n-1}\frac{I_{\widehat{p}_i}}{I}\sigma_{\pi,\hat{p}_i} - \sum_{i=1}^{n-1}b_{\widehat{p}_i}^\$\sigma_{\pi,\hat{p}_i} \\
& + \sum_{i=1}^{n-1}\left(\frac{1}{2}(\theta-1)(\theta-2)\left(\frac{I_{\widehat{p}_i}}{I}\right)^2 + (\theta-1)\frac{I_{\widehat{p}_i}}{I}b_{\widehat{p}_i}^\$ + \frac{1}{2}(b_{\widehat{p}_i}^\$)^2\right)\sigma_{\hat{p}_i}\sigma_{\hat{p}_i}' \\
& + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left((\theta-1)(\theta-2)\frac{I_{\widehat{p}_i}I_{\widehat{p}_j}}{I} + b_{\widehat{p}_i}^\$ b_{\widehat{p}_j}^\$\right)\sigma_{\hat{p}_i}\sigma_{\hat{p}_j}' \\
& + \frac{1}{2}\sum_{i=1}^{n-1}\left((\theta-1)\left(\frac{I_{\widehat{p}_i\widehat{p}_i}}{I}\right) + b_{\widehat{p}_i\widehat{p}_i}^\$ + \frac{1}{2}(b_{\widehat{p}_i}^\$)^2\right)\sigma_{\hat{p}_i}\sigma_{\hat{p}_i}' \\
& + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\left((\theta-1)\frac{I_{\widehat{p}_i\widehat{p}_j}}{I} + b_{\widehat{p}_i\widehat{p}_j}^\$\right)\sigma_{\hat{p}_i}\sigma_{\hat{p}_j}'.
\end{aligned}$$

Again, this PDE involves a time derivative. As for the prices of real bonds, we approximate $e^{b_t^\$}$ at each time point t by multivariate Chebyshev polynomials, use an explicit Euler discretization for this time derivative and solve the PDE recursively backwards in time, starting from the boundary condition $b_T^\$ = 0$.

D. Solution of the Model with Full Information

The dynamics of consumption and the log prive level are

$$\begin{aligned} d\ln C_t &= \left(\sum_{i=1}^{n} \mu_i^C \mathbb{1}_{\{S_t=i\}}\right) dt + \sigma^C \left(\sqrt{1-\rho^2} dW_t^C + \rho dW_t^\pi\right) \\ d\pi_t &= \sum_{i=1}^{n} \mu_i^\pi \mathbb{1}_{\{S_t=i\}} + \sigma^\pi dW_t^\pi, \end{aligned}$$

For ease of notation, we relabel the indicator variable as $p_{i,t} = \mathbb{1}_{\{S_t=i\}}$. These are the state variables in a model in which the representative agent perfectly knows the state of the economy at every point in time. The dynamics of these state variables are

$$dp_{i,t} = -p_{i,t} dN_{i,i} + (1-p_{i,t}) \sum_{j\neq i} dN_{j,i}$$

where the counting process $N_{j,i}$ counts transitions from state j to state i and has the intensity $\lambda_{j,i}$. The counting process $N_{i,i}$ counts the transitions from state i to any other state $j \neq i$ and has the intensity $\sum_{j\neq i} \lambda_{i,j}$. These dynamics imply that the state variables almost surely only take the values 0 and 1.
The indirect utility function of the investor is given by

$$J(C_t, p_{1t}, \ldots, p_{nt}) = E_t \left[\int_t^\infty f(C_s, J(C_s, p_{1s}, \ldots, p_{ns}))\, ds\right].$$

With the same reasoning as in the incomplete information economy (i.e. defining the indirect utility function, setting up the Bellman equation, conjecturing the same functional form for the indirect utility function) and after some simplifications, we arrive at the following algebraic equations (one equation for each state i) for the four unknowns v_i (the log wealth-consumption ratios in each state i):

$$0 = \left[(1-\gamma)\mu_i^C + \frac{1}{2}(1-\gamma)^2(\sigma^C)^2 - \beta\theta\right] + \theta e^{-v_i} + \sum_{j\neq i} \lambda_{ij} \cdot \left(e^{\theta(v_j - v_i)} - 1\right).$$

The proof is only a very slight modification of the proof in Appendix A of Branger et al. (2016). Similarly, for the log price-dividend ratio ω_i of the equity claim, we get the equations (one for each state i):

$$\begin{aligned} 0 = & -\beta\theta - (1-\theta)e^{-v_i} + e^{-\omega_i} - \gamma\mu_i^C + \bar{\mu} + \phi\left(\mu_i^C - \bar{\mu}\right) + \frac{1}{2}(\phi-\gamma)^2(\sigma^C)^2 \\ & + \sum_{j\neq i} \lambda_{ij} \cdot \left(e^{(\theta-1)(v_j-v_i)} \cdot e^{\omega_j - \omega_i} - 1\right). \end{aligned}$$

For the prices of nominal bonds $b_i^\$$ we get the four ordinary differential equations

$$\begin{aligned} 0 = & -\beta\theta - (1-\theta)e^{-v_i} - \gamma\mu_i^C + \frac{1}{2}\gamma^2(\sigma^C)^2 + \frac{\partial b_i^\$}{\partial t} - \mu_i^\pi + \frac{1}{2}(\sigma^\pi)^2 + \gamma\rho\sigma^C\sigma^\pi \\ & + \sum_{j\neq i} \lambda_{ij} \cdot \left(e^{(\theta-1)(v_j-v_i)} \cdot e^{b_j^\$ - b_i^\$} - 1\right). \end{aligned}$$

REFERENCES

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

Backus, D., M. Chernov, and I. Martin (2011). Disasters implied by equity index options. *Journal of Finance 66*(6), 1969–2012.

Baele, L., G. Bekaert, and K. Inghelbrecht (2010). The determinants of stock and bond return comovements. *Review of Financial Studies 23*(6), 2374–2428.

Balduzzi, P. and C. Lan (2016). Survey forecasts and the time-varying second moments of stock and bond returns. *Working Paper*.

Bansal, R., D. Kiku, and A. Yaron (2012). An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review 1*(1), 183–221.

Bansal, R., D. Kiku, and A. Yaron (2016). Risks for the Long Run: Estimation with Time Aggregation. *Journal of Monetary Economics 82*(1), 52–69.

Bansal, R. and I. Shaliastovich (2013). A long-run risks explanation of predictability puzzles in bond and currency markets. *Review of Financial Studies 26*(1), 1–33.

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

Barro, R. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics 121*(3), 823–866.

Barro, R. (2009). Rare disasters, asset prices, and welfare costs. *American Economic Review 99*(1), 243–64.

Basak, S. and H. Yan (2010). Equilibrium asset prices and investor behavior in the presence of money illusion. *Review of Economic Studies 77*(3), 914–936.

Beeler, J. and J. Campbell (2012). The long-run risks model and aggregate asset prices: An empirical assessment. *Critical Finance Review 1*(1), 142–182.

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein (2011). Explaining asset pricing puzzles associated with the 1987 market crash. *Journal of Financial Economics 101*(3), 552–573.

Boons, M., F. de Roon, F. Duarte, and M. Szymanowska (2017). Time-varying inflation risk and the cross-section of stock returns. *Working Paper*.

Branger, N., H. Kraft, and C. Meinerding (2016). The dynamics of crises and the equity premium. *Review of Financial Studies 29*(1), 232–270.

Burkhardt, D. and H. Hasseltoft (2012). Understanding asset correlations. *Working Paper*.

Campbell, J., C. Pflueger, and L. Viceira (2018). Monetary policy drivers of bond and equity risks. *Working Paper*.

Campbell, J., A. Sunderam, and L. Viceira (2017). Inflation bets or deflation hedges? the changing risks of nominal bonds. *Critical Finance Review 6*(2), 263–301.

Chen, H., W. Dou, and L. Kogan (2017). Measuring the "dark matter" in asset pricing models. *Working Paper*.

Chen, H., S. Joslin, and N.-K. Tran (2012). Rare disasters and risk sharing. *Review of Financial Studies 25*, 2189–224.

Constantinides, G. (2008). Comment on "macroeconomic crises since 1870" by barro and ursua. *Brookings Papers on Economic Activity*, 341–50.

Constantinides, G. and A. Ghosh (2011). Asset pricing tests with long-run risk in consumption growth. *Review of Asset Pricing Studies 1*(1), 96–136.

Constantinides, G. and A. Ghosh (2017). What information drives asset prices? *Working Paper*.

Croce, M., M. Lettau, and S. Ludvigson (2012). Investor information, long-run risk and the term structure of equity. *Review of Financial Studies 28*(3), 706–742.

David, A. and P. Veronesi (2013). What ties return volatilities to price valuations and fundamentals. *Journal of Political Economy 121*(4), 682–746.

Detemple, J. B. (1986). Asset pricing in a production economy with incomplete information. *Journal of Finance 41*, 383–391.

Drechsler, I. (2013). Uncertainty, time-varying fear and asset prices. *Journal of Finance 68*(5), 1843–1889.

Duffie, D. and L. G. Epstein (1992a). Asset pricing with stochastic differential utility. *Review of Financial Studies 5*(3), 411–436.

Duffie, D. and L. G. Epstein (1992b, March). Stochastic differential utility. *Econometrica 60*(2), 353–394.

Ehling, P., M. Gallmeyer, C. Heyerdahl-Larsen, and P. Illeditsch (2018). Disagreement about inflation and the yield curve. *Journal of Financial Economics 127*(3), 459–484.

Eraker, B. (2008). Affine general equilibrium models. *Management Science 54*(12), 2068–2080.

Eraker, B., I. Shaliastovich, and W. Wang (2016). Durable goods, inflation risk and equilibrium asset prices. *Review of Financial Studies 29*(1), 193–231.

Ermolov, A. (2018). Time-varying risks of nominal bonds: How important are macroeconomic shocks. *Working Paper*.

Fleckenstein, M., F. Longstaff, and H. Lustig (2016). Deflation risk. *forthcoming: Review of Financial Studies*.

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

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007, November). The u.s. treasury yield curve: 1961 to the present. *Journal of Monetary Economics 54*, 2291–2304.

Hamilton, J. D. (1994). *Time series analysis*. Princeton University Press.

Hasseltoft, H. (2012). Stocks, bonds, and long-run consumption risks. *Journal of Financial and Quantitative Analysis 47*(2), 309–332.

Julliard, C. and A. Ghosh (2012). Can rare events explain the equity premium puzzle? *Review of Financial Studies 25*, 3037–3076.

Liptser, R. and A. Shiryaev (2001). *Statistics of Random Processes I: General Theory*. Springer-Verlag, Berlin.

Marfe, R. (2015). Labor rigidity, inflation risk and bond returns. *Working Paper*.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–708.

Ortu, F., A. Tamoni, and C. Tebaldi (2013). Long-run risk and the persistence of consumption shocks. *Review of Financial Studies 26*(11), 2876–2915.

Piazzesi, M. and M. Schneider (2006). Equilibrium yield curves. *NBER Macroeconomics Annual*, 389–442.

Rietz, T. (1988). The equity premium: A solution. *Journal of Monetary Economics 22*(1), 117–31.

Schmeling, M. and A. Schrimpf (2011). Expected inflation, expected stock returns, and money illusion: What can we learn from survey expectations? *European Economic Review 55*(5), 702–719.

Schorfheide, F., D. Song, and A. Yaron (2018). Identifying long-run risks: A bayesian mixed-frequency approach. *Econometrica 86*(2), 617–654.

Song, D. (2017). Bond market exposures to macroeconomic and monetary policy risks. *Review of Financial Studies 30*(8), 2761–2817.

Tsai, J. and J. Wachter (2015). Disaster risk and its implications for asset pricing. *Annual Review of Financial Economics 7*, 219–252.

Wachter, J. (2006). A consumption-based model of the term structure of interest rates. *Journal of Financial Economics 79*(2), 365–399.

Wachter, J. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance 68*(3), 987–1035.