

Extreme Inflation and Time-Varying Disaster Risk

Ilya Dergunov*

Christoph Meinerding**

Christian Schlag*

This version: May 8, 2017

Abstract

Low consumption growth tends to occur together with either very high or very low inflation. The probability of low expected consumption growth estimated from a Markov chain for consumption growth and inflation is highly correlated with a measure for the likelihood of consumption disasters suggested by Wachter (2013). A simple asset pricing model with recursive utility and unobservable states reproduces the time variation in volatilities and correlations of stock and bond returns very well.

Keywords: Inflation, consumption disasters, recursive utility, filtering

JEL: E31, E44, G12

* Faculty of Economics and Business Administration and Research Center SAFE, Goethe University, 60323 Frankfurt am Main, Germany. E-mails: dergunov@safe.uni-frankfurt.de, schlag@finance.uni-frankfurt.de.

**Deutsche Bundesbank, Research Centre, 60431 Frankfurt am Main, Germany. E-mail: christoph.meinerding@bundesbank.de.

We would like to thank the seminar participants and colleagues at Goethe University Frankfurt and Deutsche Bundesbank as well as the participants of the 2017 MFA Annual Meeting and the 2017 Colloquium on Financial Markets in Cologne for their comments and suggestions. Special thanks go to Fernando Duarte, Philipp Illeditsch, Dongho Song (discussant), Jessica Wachter, and last, but not least, to Franziska Schulte and Iliya Karaivanov for excellent research assistance. 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 or its staff.

1 Introduction

The risk of sudden large declines in consumption, so-called rare disasters, can contribute to the explanation of otherwise puzzling stylized facts of asset returns like the high unconditional equity premium or the low unconditional real risk-free rate in equilibrium models, as has been documented in the seminal papers by Rietz (1988) and Barro (2006, 2009). Moreover, there is mounting evidence that time variation in the risk of rare disasters can account for some other properties of asset returns, for instance excessive volatility, return predictability, or exchange rate volatility.¹

Yet, by definition rare disasters are hardly observed in samples of macroeconomic data, which makes it hard to come up with a proper estimate for the probability of their occurrence. While it is plausible to estimate the *unconditional* probability of a rare disaster by simply counting the number of respective events in historical samples (and this is the route taken, e.g., by Barro (2006)), it remains a challenge to come up with an empirical proxy for the time-varying *conditional* probability of consumption disasters that is solely based on fundamental data and not, in some way, extracted from financial data, i.e. from the quantities that the dynamics of the fundamentals are supposed to explain in the very end.² Consequently, disaster risk serves as a prominent example in the seminal paper on “dark matter” in asset pricing models by Chen, Dou, and Kogan (2015). These authors deem disaster risk models as fragile because the parameters governing the disaster risk process can only be estimated by taking asset price data into account, and not from fundamental macroeconomic data alone.

We make three major contributions. First, we show that inflation observations serve as a useful signal, which allows to better infer the time-varying probability of very low or even negative consumption growth. Empirically, such low growth tends to occur together with either very high or very low inflation (actually, even deflation). So observing extreme values for inflation increases the estimated probability of the economy being in a very bad state.³ Our estimation has the advantage that it does not require disasters to actually

¹See, e.g. the recent papers by Gabaix (2012), Wachter (2013), or Branger, Kraft, and Meinerding (2016). Tsai and Wachter (2015) provide an overview of the recent advances of this literature.

²In a recent paper, Marfe and Penasse (2017) apply consumption data from a panel of 42 countries to come up with an estimate for time-varying disaster risk in the US. Similar to the idea of Barro (2006) this approach is based on the underlying assumption that the latent disaster probability in the US can be deduced from realizations of disasters in foreign countries.

³We do not take a stand on structural explanations for this link or on potential transmission mechanisms. For further reading, we refer to survey articles such as Blanchard (2009), providing an overview of the different strands of macroeconomic research that have emerged over the past decades.

materialize in the sample, which circumvents the difficulties outlined above.

Second, the fact that inflation can be either extremely high or extremely low in periods of adverse real growth has important implications for the joint dynamics of stocks and treasury bonds, in particular for the second moments of their returns. A positive inflation shock can be a good or a bad signal for expected real growth, depending on the overall state of the economy. We propose a simple general equilibrium asset pricing model with recursive preferences and learning that features this time-varying signaling role of inflation. The model replicates empirical patterns of both stock return volatilities and stock-bond return correlations closely. Stated differently, the time-varying disaster risk paradigm which explains stock return volatilities as outlined by Wachter (2013) and Gabaix (2012) can be extended towards the time-varying nature of the stock-bond return correlation if one takes the signaling role of inflation into account.

Third, we generate these findings with a very clean and parsimonious research design. Our estimation is based solely on two fundamental time series (consumption growth and inflation) and is not distorted by the use of asset price information, i.e., the asset pricing tests can be viewed as true out-of-sample tests for the macroeconomic estimation. None of the parameters of our model is calibrated to match asset price moments, the preference parameters are adopted from state-of-the-art asset pricing literature. Moreover, we do not impose any constraints in the estimation and just rely on the Bayes Information Criterion (BIC) to identify the number of states that best fits the macro data. Finally, we use the longest possible sample of quarterly US consumption data starting in 1947, but additional analyses with subsamples, with GDP growth instead of consumption growth, or with monthly consumption data largely confirm our findings. If anything, the results from our benchmark specification can be considered conservative in terms of the fit to macroeconomic and the asset pricing performance.

The first contribution is presented graphically in Figure 1. It shows two time series, the implied disaster intensity, which Wachter (2013) obtains from reverse engineering of her asset pricing model (blue line)⁴, and the (5-year rolling window) probability of being in a bad consumption state as estimated (i.e., filtered) from a Markov switching model for consumption growth and inflation (red line). A high value of the disaster intensity in the Wachter (2013) model of course indicates a bad state for the economy. The correlation between the two series is 0.88. One can further see that the blue and the red line share basically all key features concerning the evolution over time, like the pronounced downturn in disaster risk during the 1950s, the low values in the 1960s, the sharp increase during the 1970s up to around 1982, and then, basically following the same kind of cycle, the

⁴We thank Jessica Wachter for sharing her data with us.

sharp decline and low level during the Great Moderation, followed by the recent spike at the beginning of the Great Recession. We conclude from this that the probability for a bad state can be inferred from a simple bivariate Markov model based exclusively on fundamental data just as well as by reverse engineering a sophisticated asset pricing model.

To see how inflation indeed has an impact on the pricing of real assets, one has to look at the estimation results for the Markov model in more detail. We estimate the Markov model using quarterly data from the U.S. since 1947, and for this sample a specification with four states is favored by the Bayes Information Criterion. Table 1 shows the parameter estimates for this 4-state model. One can immediately see that states 3 and 4 have a low unconditional probability and are characterized by low (state 3) or negative (state 4) expected consumption growth, so that they represent rare, bad states for the economy. At the same time, the estimates for expected inflation in these states are rather extreme, with states 3 and 4 characterizing a high-inflation and a deflationary regime, respectively. Intuitively, this means that observing inflation helps to identify the bad consumption states, since they tend to occur together with very high or very low inflation.

A comparison of Figures 1 and 2 shows that, although the quantity of interest is a purely real variable, a Markov model based on consumption only does not do the job. The blue lines in Figures 1 and 2 are the same, but the red line in Figure 2 is computed based on a 2-state Markov model estimated from consumption data only, i.e., without inflation. The correlation is still around 0.33, but now there are substantial differences between the two series. Most strikingly, the model using consumption only does a poor job from about the early 1970s onwards. It misses the increase up to the globally highest level in the late 1970s and early 1980s. During the 1990s the consumption-implied probability of being in a bad state is also not really close to the value implied by asset prices.

The recent literature on disaster risk and asset pricing (see, e.g., Wachter (2013), Gabaix (2012)) has put an emphasis on linking time variation in disaster risk with time variation in second moments of returns. Motivated by the time series result outlined above, we thus embed the signaling mechanism into a state-of-the-art equilibrium asset pricing model, where the representative investor is equipped with recursive preferences, and where consumption growth and inflation are the fundamental sources of risk. Their drifts follow a continuous-time Markov chain. As indicated above, we assume that the representative investor cannot observe the state of the Markov chain and thus has to filter the respective probabilities from the data.

This brings about the second major contribution of our paper. Having inflation in the

model allows us to take a close look at the correlation of stock and bond returns, which has been the object of the investigation in a number of recent papers, e.g., Burkhardt and Hasseltoft (2012) and David and Veronesi (2013). Now the two inflation states coupled with low consumption growth and their filtered probabilities become relevant individually. Intuitively, when the high-inflation state becomes more likely, nominal bonds will exhibit negative returns, and the same will be true for stocks, since the representative investor will perceive the economy to enter a bad state. The returns on both assets will thus tend to be negative, implying a positive correlation. When the deflation state becomes more likely, the opposite is true. Bonds are likely to exhibit positive returns, while stocks will again not perform well due to the indication of low growth. Taken together, this will rather lead to a negative stock-bond correlation. While the response of bond prices to the perceived inflation state is straightforward since expected inflation is a key component of bond yields, the response of stock prices to inflation is one of the central asset pricing effects generated by our model.

To test the ability of our model to match stylized facts in the data we feed it with the empirical time series of observed consumption growth and inflation and then price stocks and bonds using our model-implied pricing kernel. Based on these model-generated prices we run regressions of stock return volatilities and stock-bond correlations on the filtered probabilities.

It turns out that the model does a pretty good job. In both the model-generated and the real data the coefficient of a regression of the stock-bond return correlation on the probability of being in the high-inflation state is positive and significant, while for the deflationary state the coefficient is negative and significant. In addition to reproducing the patterns of the stock-bond correlation, 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.

All of these features of the model are a direct consequence of the recursive utility specification coupled with unobservable drifts for consumption and inflation. We assume that the representative agent has a preference for early resolution of uncertainty. In particular, as is well known from the asset pricing literature with recursive preferences, in such a setup the intertemporal substitution effect dominates the income effect, which eventually leads to the desirable result that asset prices are monotonically increasing in expected consumption growth. In fact, in a robustness check with time-additive CRRA preferences, the regression coefficients for the stock-bond return correlation do not match

the empirical ones and sometimes even carry the wrong sign.

Other than specifying the drifts as unobservable and assuming recursive preferences for the representative investor, our model is standard and does not feature any special ingredients. Furthermore, the estimation of fundamental dynamics is performed without the use of any asset price information. All asset pricing results are actually generated endogenously via the equilibrium mechanism inside the model and not via the parametrization of exogenous variables. Our research design thus differs substantially from other frequently used estimation methods for dynamic asset pricing models like the Generalized Method of Moments, where typically both asset pricing moments and macroeconomic moments are used to identify the macroeconomic parameters of the model. Instead we rather set up the asset pricing exercise as an “out-of-sample” test for the Markov chain estimation based on consumption and inflation data.

Some other papers are related to our work in that they also deal with inflation in an asset pricing context. David and Veronesi (2013) propose a model similar to ours, but use time-additive CRRA preferences with money illusion in the spirit of Basak and Yan (2010) in the pricing kernel. They furthermore employ an estimation approach which relies on asset pricing data and thus delivers quite different dynamics for the fundamentals of the model compared to our parametrization. A similar framework is used in David and Veronesi (2014), where a central bank sets the interest rate according to a Taylor rule.

Burkhardt and Hasseltoft (2012) also propose a model with recursive utility, inflation and long-run risk. The model is able to produce time-varying stock-bond correlations as well, 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.

So in a sense our model may be viewed as a combination of David and Veronesi (2013), Burkhardt and Hasseltoft (2012), and Wachter (2013). We employ recursive utility so that state variables are priced, and we assume that the unobservable state variables (real expected consumption growth and expected inflation) follow a Markov chain with a finite number of states. The fact that the representative investor has to filter the state from consumption and inflation observations ensures that inflation shocks have an effect on the estimated state of the economy and thus also on the prices of real assets.

Piazzesi and Schneider (2006) discuss the role of inflation as a signal about future consumption growth, but they exclusively focus on the term structure of (nominal and real) interest rates. Their model and estimation misses out the very important role of deflation. Their data sample ranges from 1952 to 2005 and covers a period in which

consumption growth and inflation were mostly negatively correlated. In particular, they cannot (and do not) make any statements about the time-varying stock-bond correlation. Moreover, similar to the approach of Collin-Dufresne, Johannes, and Lochstoer (2016), their model includes learning about an unknown constant parameter, and the learning algorithm accounts for the possibility of structural breaks by not taking the whole history of observations into account. 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 patterns in bond and foreign exchange returns. Ehling, Gallmeyer, Heyerdahl-Larsen, and Illeditsch (2016) consider heterogeneous agents who disagree about inflation, and the authors show that this disagreement increases yields and yield volatilities at all maturities. 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. Fleckenstein, Longstaff, and Lustig (2013) study the pricing of deflation risk using market prices of inflation-linked derivatives.

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 (2014), Campbell, Sunderam, and Viceira (2013), Hasseltoft (2012), Ang and Ulrich (2012), and Marfe (2015), to name just a few. Koijen, Lustig, and van Nieuwerburgh (2015) and Boons, de Roon, Duarte, and Szymanowska (2017) provide evidence for inflation risk being priced in the cross-section of stock returns. The latter paper can be viewed as complementary to our paper, 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. Campbell, Pflueger, and Viceira (2015) analyze the stock-bond correlation in a New Keynesian production economy with habit formation preferences and monetary policy regimes, and Song (2014) studies a production economy model with recursive preferences, a regime-switching Taylor rule, and a time-varying inflation target. In both papers, however, asset price data is used to calibrate or estimate the model. Baele, Bekaert, and Inghelbrecht (2010) empirically analyze the determinants of the stock-bond return comovement.

2 Fundamental Dynamics

2.1 Consumption and inflation dynamics

The two fundamental sources of risk in our model are aggregate consumption and inflation. We assume that log aggregate consumption, $\ln C$, follows the process

$$d \ln C_t = \mu_t^C dt + \sigma^C \left(\sqrt{1 - \rho^2} dW_t^C + \rho dW_t^\pi \right), \quad (1)$$

while the dynamics of the rate of inflation π are given as

$$d\pi_t = \mu_t^\pi dt + \sigma^\pi dW_t^\pi. \quad (2)$$

W^C and W^π are the (independent) components of a standard bivariate Wiener process. The dynamics in (1) and (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 μ_t^C and μ_t^π are stochastic and follow a bivariate continuous-time Markov chain. In particular, there are n states (index by $i = 1, \dots, n$), with state-dependent drifts μ_i^C and μ_i^π . The Markov chain is represented by the $n \times n$ matrix $\Lambda = (\lambda_{ij})_{i,j=1,\dots,n}$ of transition intensities. By definition, $\lambda_{ii} = -\sum_{j \neq i} \lambda_{ij}$. In our benchmark empirical case, we will have $n = 4$.

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

$$\begin{pmatrix} d \ln C_t \\ d\pi_t \end{pmatrix} = \mu_t dt + \Sigma dW_t$$

$$\text{with } \mu_t = \begin{pmatrix} \mu_t^C \\ \mu_t^\pi \end{pmatrix}, \Sigma = \begin{pmatrix} \sigma^C \sqrt{1 - \rho^2} & \sigma^C \rho \\ 0 & \sigma^\pi \end{pmatrix}, \text{ and } dW_t = \begin{pmatrix} dW_t^C \\ dW_t^\pi \end{pmatrix}.$$

2.2 Filtering

We assume that the representative agent cannot observe μ_t^C and μ_t^π and has to filter her estimates from the data. Mathematically, there are two filtrations, \mathcal{F} and \mathcal{G} , where \mathcal{F} is generated by the processes $(C_t)_t$, $(\mu_t^C)_t$, $(\pi_t)_t$ and $(\mu_t^\pi)_t$, whereas $\mathcal{G} \subset \mathcal{F}$ is generated by the processes $(C_t)_t$ and $(\pi_t)_t$ only. The agent's decisions (and thus all equilibrium quantities) are based on the conditional expectations of the drifts given the investor's information, i.e., on $\hat{\mu}_t^C$ and $\hat{\mu}_t^\pi$ given as

$$\hat{\mu}_t^C = E[\mu_t^C | \mathcal{G}_t] = \sum_{i=1}^n \hat{p}_{i,t} \mu_i^C$$

and

$$\hat{\mu}_t^\pi = E[\mu_t^\pi | \mathcal{G}_t] = \sum_{i=1}^n \hat{p}_{i,t} \mu_i^\pi.$$

Here $\hat{p}_{i,t} \equiv E[p_i | \mathcal{G}_t]$ 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 $\hat{p}_{1,t}, \dots, \hat{p}_{n-1,t}$, 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 $\hat{p}_{i,t}$ follow from the so-called Wonham filter and are given by

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

A proof of the filtering equation based on Theorem 9.1 of Liptser and Shiryaev (2001) is provided in Appendix A.

The drift in (3) is a linear function of the transition intensities λ and the current estimates of the probabilities \hat{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, $\hat{p}_{it} \lambda_{ii} = -\hat{p}_{it} \sum_{j \neq i} \lambda_{ij}$, involves the intensities for a switch from state i to some other state $j \neq i$. Suppose \hat{p}_{it} is currently large, i.e. the investor is relatively certain to be in state i . Then, loosely speaking, the more time passes, the more likely it becomes that an unobserved switch from state i to some other state j has occurred in the meantime. This effect induces a negative drift of \hat{p}_{it} . The second term, $\sum_{j \neq i} \hat{p}_{jt} \lambda_{ji}$, captures the probabilities of entering state i , given that the economy is currently in a different state j . Suppose one of the \hat{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 \hat{p}_{it} . The overall sign of the drift of \hat{p}_{it} thus depends on the current estimate of all \hat{p} . In particular, the drift terms ensure that the probabilities \hat{p} never fall below 0 and never exceed 1.

The diffusive volatility of \hat{p}_i is a *quadratic* function of all probabilities \hat{p}_j ($j = 1, \dots, n$). The probability update is the largest if the investor is rather uncertain about the current state of the economy, i.e., for intermediate values of \hat{p}_i . If the investor is almost sure in which state the economy currently is (i.e., if one of the \hat{p}_j is close to one and the others are close to zero), there is almost no noise in the estimates. When the respective estimate

\hat{p}_i is close to zero, the diffusion term in (3) is obviously also close to zero, since \hat{p}_i is one factor of the product in front of the Wiener innovations. When \hat{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 \hat{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 $\ln C$ or π 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 the observations for $\ln C$ and π have to be adapted to both \mathcal{F} and \mathcal{G} , which implies $\mu_t dt + \Sigma dW_t = \hat{\mu}_t dt + \Sigma d\widehat{W}_t$.

Finally, in the context of our analysis it is very important to note that the update in the estimated probability \hat{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 channel will be the major driving force behind our asset pricing results described below.

2.3 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.⁵ Our sample period ranges from 1947Q1 to 2014Q1 and represents the longest period for which quarterly data are available.⁶ The upper graph in Figure 3 shows time series plots of the data.

Based on these data for consumption and inflation we estimate a joint Markov chain for expected consumption growth and expected inflation using maximum likelihood (see, e.g., Hamilton (1990)). We assume a constant variance-covariance matrix and only allow for time-varying drifts. Standard errors for the parameter estimates are computed via a

⁵We do not take a stance on which of the various inflation time series measures realized inflation more precisely. For a detailed discussion of this issue, we refer the reader to Piazzesi and Schneider (2006). As will become clear below, the main feature we need is the existence of extremely high and low observations in the sample, which is robust to the inflation measure.

⁶We have done the estimation with other data samples as well, in order to compare our findings to those obtained by others. These results are discussed in Section 2.5.

standard block bootstrap with a block length of ten quarters⁷ with potentially overlapping blocks and 5,000 repetitions.

The results are presented in Table 1. The first important finding is that, based on the Bayes Information Criterion, the algorithm clearly 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 state 1 lasts for around 11 quarters on average, while the average time spent in state 2 is 38 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.

The lower graphs in Figure 3 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 3.

2.4 Extreme inflation as a signal about disaster risk

The filtered probabilities contain important information relevant for asset pricing. To see this take a look at Figure 1. The blue line is the implied disaster intensity shown in Figure 8 in the paper of Wachter (2013). This time series is reverse engineered from asset prices based on Wachter’s 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 the time series of historical S&P 500 price-earnings ratios. For the sake of comparison we take averages of this time series over each quarter. The red line is the sum of the filtered probabilities $\hat{p}_3 + \hat{p}_4$ 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 knowing the parameters of the chain would have assumed at each point in time. The plot

⁷Varying the block length does not affect the results.

shows 5-year moving averages of these probabilities.⁸

The two series have a correlation of 0.88 over our sample period. This is particularly remarkable, given that they are computed from very different data and with very different methodologies. Furthermore, they share all important trends, peaks, and troughs over our almost 70 year sample period. In particular, both the high inflation regime and the deflation regime contribute a significant amount to the sum. For instance, the peak in the time series in the early 1980s can be traced back to the high probability of the high inflation regime in that period, and the deflation regime prevails towards 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 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 only consumption data. The Markov chain estimation with consumption only gives rise to a Markov chain with two states only with values for expected consumption growth of 2.34 and -1.80 percentage points, respectively.

Figure 2 presents the (moving average) 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 by far not matched as well as before. The correlation between Wachter's and the consumption-only series goes down to roughly 0.33, but more importantly, the times series of estimated probabilities for low consumption is off substantially during most of the 1970's and 1980's and also towards the end of the sample period.

2.5 Different samples and constrained model specifications

The finding that extreme inflation provides information about low expected real consumption growth is based on a single Markov chain estimation. It may thus generally hinge on the data sample or on assumptions made for the estimation. In order to alleviate these concerns, we check our results for robustness along two dimensions. First, we repeat the estimation with alternative data samples. Second, we analyze a constrained version of our model in which the number of parameters is reduced.

⁸In Wachter's model, as in any asset pricing model, the price-earnings ratio today depends on the distribution of all future values of the state variables. Inverting the empirical time series of price-earnings ratios therefore implies that the resulting implied disaster intensity depends on the history of all prior price-earnings ratios. Therefore we consider a moving average of our filtered probabilities to be the proper analogue to the time series of implied disaster intensities.

We consider four alternative data samples. The first two samples 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 is available from 1959 onwards. Finally, we estimate our model with GDP growth rates instead of consumption growth rates, which are available on a quarterly basis starting in 1947.

Figure 4 depicts the proxies for the time-varying disaster intensity that we get if we apply exactly the same methodology as before to these 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 4 for all samples. As in the benchmark estimation, we label the states in which the conditional expected consumption growth rate μ_c^i is below the unconditional average consumption growth rate as “bad states”. This refers to 2 of the 4 states in every case. We then compute the filtered probabilities \hat{p}_i for these two bad states and add them up. In the upper pictures in Figure 4, the red dashed line shows the sum of these two probabilities \hat{p}_i . The blue solid line is the same as in Figure 1. Finally, we also repeat the estimation without inflation data, i.e. with consumption or GDP data only. In these cases we have only two states, and we treat the state with the lower expected consumption growth rate μ_c^i as the bad state. The filtered probability of this bad state is depicted in the lower pictures.

We can draw two conclusions from this exercise. First, recovering the time series of implied disaster probabilities from Wachter (2013) is independent of the specific sample that we use. In each case 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.

Besides analyzing alternative samples, one might also consider imposing more structure on the model. Our benchmark model has 27 free parameters to estimate. But 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 equate expected consumption growth rates or expected inflation 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 research setup. Nevertheless,

given our interpretation of states 3 and 4 as the two bad states which we couple together to obtain the disaster risk time series, 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. $\mu_c^1 = \mu_c^2$ and $\mu_c^3 = \mu_c^4$. 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 $\mu_c^1 = \mu_c^2$ and $\mu_c^3 = \mu_c^4$ cannot be rejected at the 10% level. Moreover, among various other constrained models that we have tried (results not reported here for brevity), this particular constrained model has the lowest BIC (797) and should thus be preferred.

Figure 5 depicts 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 not reported here for brevity, we also solved the asset pricing model, which will be explained in the next subsection, with the parameters from this constrained model. None of our asset pricing findings changes if we use this specification. Therefore we stick to the unconstrained model in the following.

To sum up, the finding that extreme inflation helps to recover the time-varying probability of consumption disasters is robust across samples and robust to the constraint of equal growth rates in the two bad states. If anything, the benchmark specification on which we rely in the rest of the paper yields the most conservative estimates among all the samples studied here. Nevertheless we stick to this specification because it uses the longest quarterly sample available, and it does not impose any constraints on the identified consumption and inflation states.

3 Asset Pricing Implications

Disaster risk models like the one by Wachter (2013) 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 in a state-of-the-art asset pricing model. Having inflation in the model allows us to analyze the returns of stocks and nominal bonds jointly. The numerical results in this section are based on the parameter estimates from the benchmark specification, i.e. the unconstrained model estimated on the quarterly sample from 1947 until 2014.

3.1 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

$$J_t = E_t \left[\int_t^\infty f(C_s, J_s) ds \right],$$

where the aggregator f is given by

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

γ , ψ , and β denote the degree of relative risk aversion, the elasticity of intertemporal substitution (EIS), and the subjective time preference rate. 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 $\beta = 0.02$. In particular, the agent has a preference for early resolution of uncertainty, i.e., that $\gamma > \psi^{-1}$.

3.2 Real Pricing Kernel and Wealth-Consumption Ratio

As in Duffie and Epstein (1992a), the real pricing kernel is given by

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

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

Given a solution for I , the pricing kernel has dynamics

$$\begin{aligned} \frac{d\xi_t}{\xi_t} = & -\delta\theta dt - (1-\theta)I^{-1} dt - \gamma dC_t + \frac{1}{2}\gamma^2\sigma_c^2 dt \\ & - (1-\theta) \sum_{i=1}^{n-1} \frac{I_{\hat{p}_i}}{I} d\hat{p}_{i,t} + \frac{1}{2} \sum_{i=1}^{n-1} \sum_{j=1}^{n-1} (\theta-1) \left[\frac{I_{\hat{p}_i} I_{\hat{p}_j}}{I} + (\theta-2) \left(\frac{I_{\hat{p}_i} I_{\hat{p}_j}}{I^2} \right) \right] \sigma_{\hat{p}_i} \sigma'_{\hat{p}_j} dt \\ & - \gamma(\theta-1) \sum_{i=1}^{n-1} \frac{I_{\hat{p}_i}}{I} \sigma_{c,\hat{p}_i} dt. \end{aligned}$$

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.⁹

3.3 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. We therefore assume

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

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 $g(\xi, D, \omega) = \xi D e^\omega$. This yields

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

Using Ito's Lemma, we can translate this equation into a PDE for $\omega(\widehat{p})$. This PDE, together with details about its derivation, is given in Appendix D. 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 $\int_t^T e^{-\pi_s} ds$. The price of a nominal bond at time t is thus equal to

$$B_t^{\$,T} = E_t \left[\xi_{t,T} \int_t^T e^{-\pi_s} ds \right].$$

One can define the nominal pricing kernel $\xi_{t,T}^{\$}$ as $\xi_{t,T} \int_t^T e^{-\pi_s} ds$ and rewrite the pricing formula as

$$B_t^{\$,T} = E_t [\xi_{t,T}^{\$}].$$

⁹To generate an impact of inflation on the pricing kernel, David and Veronesi (2013) rely on the behavioral concept of money illusion. 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 comprises the original real pricing kernel and an adjustment for inflation.

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

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

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 a third term which arises from $\frac{d[\xi, \pi]}{\xi}$. This third term is nonzero if inflation shocks affect the real pricing kernel, as they do in our model. It can be interpreted as an inflation risk premium which nominal bonds earn in equilibrium.

The Euler equation and the Feynman-Kac formula applied to $H(\xi_t^{\$}, b_t^{\$}) = \xi_t^{\$} e^{b_t^{\$}}$ yield a partial differential equation for $b_t^{\$} \equiv \ln B_t^{T, \$}$. Details on this partial differential equation and its solution are given in Appendix F.

3.4 Results

3.4.1 General approach

Based on our finding that certain state probabilities proxy disaster risk, and given that Wachter (2013) documents the important role of time-varying disaster intensities for the dynamics of second moments of returns we start 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 3, 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 five-year (20-quarter) 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)¹⁰ 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 \hat{p} . Note that

¹⁰The data are available for download at <http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html>.

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. These monthly data are then aggregated to quarterly and used in the regressions in the same way as described before. We repeat this exercise 5,000 times to obtain the 90% confidence intervals.¹¹

3.4.2 Conditional stock return volatilities

A look at Figure 6 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, $\hat{p}_3 + \hat{p}_4$, exhibit a positive covariation with stock return volatilities. Second, this relationship is nonlinear and concave, both in the model and in the data. Given our result above that the sum $\hat{p}_3 + \hat{p}_4$ is highly correlated with a measure for the conditional probability of a consumption disaster, one might expect a result like that, but it is nevertheless worth noting that the second result also confirms a prediction from the model of Wachter (2013), namely that stock market volatility is a concave function of time-varying disaster risk.¹² Figure 7 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 \hat{p} . Table 2 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, which indicates that the model-implied unconditional stock return volatility is somewhat on the low side.¹³ Overall, we conclude that the relation between the conditional probability of a consumption disaster (proxied by $\hat{p}_3 + \hat{p}_4$) and stock market volatility is indeed nonlinear, and our model reproduces this

¹¹Due 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 that they sum to 1.

¹²See, for instance, Figure 4 in her paper.

¹³We discuss this issue in detail in Section 3.4.4.

stylized fact. Finally, given that our estimation is based on two macro time series only, the goodness of fit in Figure 7 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, note that in our model higher levels of \hat{p}_3 and \hat{p}_4 induce larger fluctuations in \hat{p}_3 and \hat{p}_4 . This follows from the filtering equation (3) which reveals that \hat{p}_3 and \hat{p}_4 fluctuate a lot more when they are at intermediate levels as compared to when they are close to 0 or 1. Given that a level of 1 for \hat{p}_3 and \hat{p}_4 is almost never reached empirically, higher values for \hat{p}_3 and \hat{p}_4 practically thus go together with high fluctuations of \hat{p}_3 and \hat{p}_4 . States 3 and 4 are the least persistent in our estimation, so that when the agent currently has a high estimated probability of being in one of these two states, the probability is likely going to move down rather quickly. Overall, following a negative consumption shock in the economy, there is more movement in state variables. This in turn means that uncertainty about the likelihood of negative consumption shocks generates additional volatility of the price-dividend ratio, which is a function of these state variables. In sum, learning about consumption disasters induces higher 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 $\log(\hat{p}_1 + \hat{p}_2)$, for stock return volatilities. The coefficients in the data and in the model (not reported here for brevity) are both significantly negative since $\hat{p}_1 + \hat{p}_2 = 1 - \hat{p}_3 - \hat{p}_4$.

Finally, Table 3 reports 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 the extreme entropy of the state distribution (defined as $\hat{p}_3 \ln \hat{p}_3 + \hat{p}_4 \ln \hat{p}_4$), which we propose as another proxy for uncertainty about consumption disasters. We also show results for expected inflation (defined as $\sum_{i=1}^4 \mu_i^\pi \hat{p}_i$) 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 able 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 a component capturing the risk of high inflation (like \hat{p}_3) and a component capturing the risk of a deflationary regime (like \hat{p}_4). The reason is that, as outlined above, inflation is positively correlated with \hat{p}_3 , but negatively correlated with \hat{p}_4 . Depending on the amount of observations from the deflation regime along a given sample path, the signaling role of inflation for stock market volatility

may be ambiguous.

Note also that 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 (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, both in terms of its overall level (i.e., addressing the excess volatility puzzle) and in terms of its dynamics (i.e., capturing the predictive power of implied volatilities, variance risk premia, and other related quantities).

Finally, it is important to note that our whole approach to the explanation of stock return volatility does not rely on a Peso problem story, which is very popular in the literature on disaster risk and its role in the explanation of the equity premium (see, e.g., Barro (2006) and Wachter (2013)). In particular, the bad states in our model are not devastatingly bad, so that the mechanism in our model does not build on the notion that agents fear very bad realizations of consumption growth or inflation, which have never been observed in US data. Our estimates for the state-dependent expected growth rates and inflation are rather moderate.

3.4.3 Conditional stock-bond return correlations

The results concerning stock market volatility are related to the overall probability of being in a bad state for expected consumption growth. When we now look at the stock-bond return correlation, the distinction between the two bad consumption states with respect to expected inflation will become relevant. Table 4 and Table 5 contain the results of our regression analyses, Figure 8 depicts the corresponding scatter plots.¹⁴ Figure 9 shows the time series of correlation in the data and in the model.

The most important result here is that both in the model and in the data the estimated coefficient for $\log(\hat{p}_3)$ is positive and significant, while the coefficient for $\log(\hat{p}_4)$ is negative and significant. Given that we did not use any asset price information to estimate the fundamental dynamics in our model, the similarity between model and data results 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, given that our estimation is based on two macro time series only,

¹⁴Note that we regress 'raw' correlation on the state variables. Correlation has a limited range, so that transformations like $\tilde{p} = \ln(\frac{1+\rho}{1-\rho})$ might seem warranted to guarantee that the regressions are well specified. We reran all our analyses using this transformation. All our results remain qualitatively unchanged.

the goodness of fit in Figure 9 is remarkable. The correlation between the data and the model-implied time series is 0.54.

What is the mechanism inside the model that generates these patterns? First, an increase in \hat{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, overall the nominal yield curve shifts upwards in response to an increase in \hat{p}_3 . The stock return upon a positive shock to \hat{p}_3 depends on real quantities only. A high \hat{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 \hat{p}_3 go in the same direction.

State 4 is a low inflation state with low growth, so the response of bond prices to a high \hat{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. If deflation becomes more likely, the level of the nominal yield curve must decrease. Altogether, the influence of \hat{p}_4 on bond returns is large and positive. At the same time, a high \hat{p}_4 signals a high likelihood of low (even negative) expected consumption growth, which depresses equity prices. In sum, the reactions of bond and stock prices to an increase in \hat{p}_4 are larger than those to an increase in \hat{p}_3 , and they are of opposite signs.

The time-varying nature of the stock-bond correlation is again a result in our model, for which recursive preferences are essential. In a robustness check with time-additive CRRA preferences (results not shown), our model generates a small positive regression coefficient for \hat{p}_3 and a large positive coefficient for \hat{p}_4 . To get the intuition behind this result, look again at the three components of the holding period bond return as described above.

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 slight increase in \hat{p}_3 or \hat{p}_4 , i.e., a slight decrease in expected consumption growth, can lead to a massive decline in the overall level of the real yield curve in a CRRA economy. As is well known, a model

with CRRA preferences cannot explain the empirically observed smoothness of the real risk-free rate. Altogether, bond returns are thus positively related to both \hat{p}_3 and \hat{p}_4 (and the effect is stronger for \hat{p}_4 , the probability of being in a deflationary regime). Concerning stock returns, note that a high value for \hat{p}_3 or \hat{p}_4 signals a low expected consumption growth rate, while consumption volatility is not affected. With the usual popular CRRA parametrizations (most importantly $\gamma > 1$) a lower expected consumption growth rate implies a higher stock price. Therefore, stock returns are also positively related to both \hat{p}_3 and \hat{p}_4 in a CRRA economy (and the effect is stronger for \hat{p}_4 because state 4 has the lowest expected consumption growth rate). Altogether, we can conclude that with CRRA preferences the stock-bond return correlation reacts positively to an increase in both \hat{p}_3 and \hat{p}_4 .

So to obtain results similar to ours, but in a CRRA model, one has to include a feature like money illusion to make inflation and related states 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), who actually restrict expected consumption growth to be the same in all states. When we rely on recursive utility, we obtain model-implied results that are very well in line with the data without having 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.

The above findings concerning the role of \hat{p}_3 and \hat{p}_4 are very well in line with the literature. In a purely empirical paper, Baele, Bekaert, and Inghelbrecht (2010) try to fit the correlations of daily stock and bond returns with a multi-factor model. They find that macro variables (in particular the output gap and inflation) do not add much explanatory power. Our findings may be related to theirs. As our results show, the risk of low expected consumption growth, proxied by $\log(\hat{p}_3 + \hat{p}_4)$, does not predict correlation, both in the model and in the data. The coefficients on realized and expected inflation are positive and significant in the model and in the data, but the confidence bounds always include 0, which indicates that the results are not robust to different data samples obtained through the bootstrap approach. One reason why Baele, Bekaert, and Inghelbrecht (2010) find that macroeconomic variables do not explain correlation may thus be that they do not explicitly take time variation in the signaling role of inflation into account.

For both model and data, the regressions with extreme entropy do not work as well. 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 correlation. An aggregate measure of uncertainty cannot capture these two opposing effects adequately.

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 available from the Fed and analyzed by Gürkaynak, Sack, and Wright (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 long-term Treasury bonds' returns from the Ibbotson Stocks, Bonds, Bills and Inflation Yearbook. These bond returns are available starting in 1926, 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 10 depicts the resulting time series of rolling window correlations. As one can see from the figure and also from regressions identical to those presented above (results not reported here for brevity), the model-implied correlation tracks the correlation in the data closely. The two time series have a correlation of 0.54. Interestingly, in the data we now see another period with negative correlation and our model captures this negative correlation as well. We thus conclude that our approach is robust to longer time series of Treasury bond returns.

3.4.4 Unconditional moments

The unconditional asset pricing moments generated by our model are computed via the same Monte Carlo simulation as described above. The results are shown in Table 6. 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 state variables like long-run consumption risk or stochastic volatility. So it should not come as a surprise when the model does not perfectly match the data with respect to unconditional risk premia or volatilities.

Nevertheless with our baseline parametrization we already obtain an expected real return on equity of 3.6 percentage points, slightly less than half of what we see in the data. The main reason for the fact that the expected stock return is somewhat low in the model is that, based on the point estimates for the entries of the transition matrix (see Table 1), the bad states are not very persistent. In particular the average time spent in state 4 is

only $1/(1 - 0.666) \approx 3$ quarters. However, the estimate of element (4,4) in the transition matrix has a standard error of around 0.19. In fact, in robustness checks not reported here, we find that the expected equity return already increases by around 1.5 percentage points if we raise the transition probability (4,4) by one standard error. Moreover, note that the results reported here are based on the most conservative macro estimation, which uses quarterly consumption growth rates since 1947. With quarterly GDP growth rates, for instance, the bad states are characterized by lower expected growth rates and therefore the equity premium rises by another percentage point.

With respect to equity return volatility we find pretty much the same picture as for expected equity returns. The volatility generated by our model is somewhat low, but again a higher persistence for the bad states would lead to higher return volatility. 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.

4 Conclusion

Low consumption growth tends to occur together with either very high or very low inflation. Building on this key insight, our paper makes three contributions.

First, we document that the filtered probabilities from a simple four-state Markov chain for expected consumption growth and expected inflation contain important information about the real economy. In particular, our Markov chain model exhibits two bad states in which expected consumption growth is low or even negative. The sum of the probabilities for the two bad states is impressively close to the implied disaster intensity computed by Wachter (2013). Moreover, we document that inflation data indeed provides additional information about the real economy since the replication of the time series of disaster intensities fails with consumption data only.

With respect to inflation the two bad states are special, since they are the ones with the highest and the lowest (even negative) expected change in the price level. This brings about our second set of findings. We analyze the role of inflation in a standard asset pricing model with learning about unobservable states. The model does not contain any behavioral components like money illusion, but instead we equip the representative investor with recursive preferences.

In line with the intuition behind asset pricing models like the ones of Wachter (2013) or Gabaix (2012), we find in the data that stock return volatilities are high when the

investor perceives the sum of the probabilities of the two bad states to be high. Moreover, feeding the model with the observed time series of consumption growth and inflation produces patterns for stock return volatilities that are very similar to the data. An object of particular interest in the macro-based asset pricing literature is the time-varying nature of the return correlation between equity and nominal bonds. Our model is also able to match this stylized fact. 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 becomes relevant here. 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 securities.

As our third contribution, we wish to emphasize that our estimation is solely based on fundamentals, i.e., on the time series of consumption growth and inflation, so that no asset price data are used to calibrate the model. Furthermore, we do not impose any constraints in the estimation and just rely on the Bayes Information Criterion (BIC) to identify the number of states that best fits the data. Our research design thus differs substantially from other estimation methods for dynamic asset pricing models where both asset pricing moments and macroeconomic moments are used to identify the macroeconomic parameters of the model, which may confound the macroeconomic fit considerably. Moreover, we use the longest possible sample of quarterly US consumption data, starting in 1947, but our findings are robust to alternative data samples.

In summary, our paper shows that some fraction of the “dark matter” in disaster risk models as quantified by Chen, Dou, and Kogan (2015) may actually be attributable to inflation risk. Given this, the time-varying disaster risk paradigm, which has recently been documented to explain stock return volatilities, can then be extended towards the time-varying nature of the stock-bond return correlation if one takes the signaling role of inflation into account.

A Dynamics of the state variables

The consumption and inflation dynamics can be rewritten as

$$\begin{aligned} d \ln C_t &= \left(\sum_{i=1}^n \mu_i^c p_{it} \right) dt + \sigma_c \left(\sqrt{1 - \rho^2} dW_t^c + \rho dW_t^\pi \right) \\ d\pi_t &= \sum_{i=1}^n \mu_i^\pi p_{it} + \sigma_\pi dW_t^\pi, \end{aligned}$$

where $p_{it} = 1$ if the economy is in state i at time t and $p_{it} = 0$ otherwise ($i = 1, \dots, n$). In matrix form, this becomes

$$\begin{pmatrix} d \ln C_t \\ d\pi_t \end{pmatrix} = \begin{pmatrix} \sum_{i=1}^n \mu_i^c p_{it} \\ \sum_{i=1}^n \mu_i^\pi p_{it} \end{pmatrix} dt + \Sigma \times \begin{pmatrix} dW_t^c \\ dW_t^\pi \end{pmatrix},$$

where Σ is given by

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

and $d[W_c, W_\pi] = 0$.

The inverse of the volatility matrix 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, the processes p_{it} (and thus also μ_t^c and μ_t^π) are adapted to the filtration \mathcal{F} . But there is a subfiltration $\mathcal{G} \subset \mathcal{F}$ to which they are not adapted. This subfiltration is the filtration generated by the processes $(C_t)_t$ and $(\pi_t)_t$, whereas the large filtration \mathcal{F} is the filtration generated by the processes $(C_t)_t$, $(\pi_t)_t$ and $(p_{it})_t$. The equilibrium in the economy is thus based on the dynamics of $(C_t)_t$ and $(\pi_t)_t$ under the investor filtration \mathcal{G} , i.e. on the projections

$$\hat{\mu}_t^c = E[\mu_t^c | \mathcal{G}_t] = \sum_{i=1}^n \hat{p}_{i,t} \mu_i^c, \quad \hat{\mu}_t^\pi = E[\mu_t^\pi | \mathcal{G}_t] = \sum_{i=1}^n \hat{p}_{i,t} \mu_i^\pi$$

An application of Theorem 9.1 of Liptser and Shiryaev (2001) yields that the projected (henceforth also called “subjective”) probabilities have the following dynamics:

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

where

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

In particular, $d[\widehat{W}_c, \widehat{W}_\pi] = 0$. Under the investor filtration, log consumption has dynamics

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

For notational convenience, we abbreviate

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

Quadratic covariations are abbreviated by

$$\begin{aligned} \frac{d[C_t, \widehat{p}_{it}]}{dt} &= \begin{pmatrix} \rho \sigma_c, & \sqrt{1 - \rho^2} \sigma_c \end{pmatrix} \times \sigma'_{\widehat{p}_i} \equiv \sigma_{c, \widehat{p}_i} \\ \frac{d[\pi_t, \widehat{p}_{it}]}{dt} &= \begin{pmatrix} 0, & \sigma_\pi \end{pmatrix} \times \sigma'_{\widehat{p}_i} \equiv \sigma_{\pi, \widehat{p}_i}. \end{aligned}$$

B Wealth-consumption ratio

The indirect utility function of the investor is given by

$$J(t) = E_t \left[\int_t^\infty f(C_s, J(s)) ds \right].$$

$J \left(C_t + \int_0^t f(C_s, J(X_s, C_s)) ds, X_t \right)$ is a martingale, therefore we have the Bellman equation

$$E[dJ(C_t, X_t) + f(C_t, J(C_t, X_t))dt] = 0,$$

or equivalently

$$\frac{\mathcal{A}J(C_t, X_t)}{J(C_t, X_t)} + \frac{f(C_t, J(C_t, X_t))}{J(C_t, X_t)} = 0, \quad (4)$$

where the operator \mathcal{A} is the infinitesimal generator. We conjecture a functional form for J :

$$J(t) = \frac{C_t^{1-\gamma}}{1-\gamma} (\delta e^{v_t})^\theta$$

where in the end v_t will prove to be the log wealth-consumption ratio. This functional form implies

$$\frac{f(C, J)}{J} = \theta e^{-v_t} - \theta \delta.$$

With $I \equiv e^v$ the derivatives of J are:

$$\begin{aligned}
J_c &= C_t^{-\gamma} (\delta e^{v_t})^\theta \\
J_{cc} &= -\gamma C_t^{-\gamma-1} (\delta e^{v_t})^\theta \\
J_{\hat{p}_i} &= \frac{C_t^{1-\gamma}}{1-\gamma} \delta^\theta \theta (I(\hat{p}))^{\theta-1} I_{\hat{p}_i} \\
J_{\hat{p}_i \hat{p}_j} &= \frac{C_t^{1-\gamma}}{1-\gamma} \delta^\theta \theta [(\theta-1) I^{\theta-2} I_{\hat{p}_i}^2 + I_{\hat{p}_i \hat{p}_j} I^{\theta-1}] \\
J_{c \hat{p}_i} &= C_t^{-\gamma} \delta^\theta \theta I^{\theta-1} I_{\hat{p}_i}.
\end{aligned}$$

These derivatives result in the following quadratic variation and covariation terms:

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

Plugging everything into (4) results in the following partial differential equation for I :

$$\begin{aligned}
0 &= \left[(1-\gamma) \sum_{i=1}^n \mu_i^c \hat{p}_i + \frac{1}{2} (1-\gamma)^2 \sigma_c^2 - \delta \theta \right] + \theta I^{-1} \\
&+ \sum_{i=1}^{n-1} \theta \frac{I_{\hat{p}_i}}{I} \left(\hat{p}_{it} \lambda_{ii} + \sum_{\substack{j=1 \\ j \neq i}}^n \hat{p}_{jt} \lambda_{ji} \right) + \sum_{i=1}^{n-1} \theta (1-\gamma) \frac{I_{\hat{p}_i}}{I} \sigma_{c, \hat{p}_i} \\
&+ \frac{1}{2} \sum_{i=1}^{n-1} \sum_{j=1}^{n-1} \theta \left[(\theta-1) \left(\frac{I_{\hat{p}_i} I_{\hat{p}_j}}{I} \right) + \frac{I_{\hat{p}_i \hat{p}_j}}{I} \right] \sigma_{\hat{p}_i} \sigma'_{\hat{p}_j},
\end{aligned} \tag{5}$$

Note that there are only $n-1$ state variables due to the restriction $\sum_{i=1}^n \hat{p}_i = 1$. We solve the PDE with a Chebyshev approximation similar to Benzoni, Collin-Dufresne, and Goldstein (2011). We guess the following functional form for I as a function of the vector \hat{p} :

$$\begin{aligned}
I(\hat{p}) &= \exp(B(\hat{p})) \\
B(\hat{p}) &= \sum_{j=0}^d \alpha_j T_j(\hat{p}),
\end{aligned}$$

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

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

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 this guess are

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

Plugging the guess into (5) gives

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

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 .

C Pricing kernel

The pricing kernel is given by:

$$\xi_{0,t} = \exp \left(\int_0^t -\delta \theta - (1 - \theta) I^{-1}(\hat{p}_s) ds \right) C_t^{-\gamma} (I(\hat{p}_t))^{\theta-1}$$

and has dynamics

$$\begin{aligned}
\frac{d\xi_{0,t}}{\xi_{0,t}} &= -\delta\theta dt - (1-\theta)I^{-1}dt - \gamma dc + \frac{1}{2}\gamma^2\sigma_c^2 dt \\
&- (1-\theta)\sum_{i=1}^{n-1}\frac{I_{\hat{p}_i}}{I}d\hat{p}_i + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}(\theta-1)\left[\frac{I_{\hat{p}_i\hat{p}_j}}{I} + (\theta-2)\left(\frac{I_{\hat{p}_i}I_{\hat{p}_j}}{I^2}\right)\right]\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}dt \\
&- \gamma(\theta-1)\sum_{i=1}^{n-1}\frac{I_{\hat{p}_i}}{I}\sigma_{c,\hat{p}_i}dt,
\end{aligned}$$

where $\hat{p} = (\hat{p}_1, \hat{p}_2, \dots, \hat{p}_n)$. For later use, we abbreviate the drift terms

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

D Price-dividend ratio

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

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

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

$$\frac{Dg(\xi, D, \omega)}{g(\xi, D, \omega)} + e^{-\omega} = 0. \quad (6)$$

Itô's Lemma gives

$$\frac{Dg}{g} = \mu_\xi + \mu_D + \mu_\omega + \frac{1}{2}\frac{d[\omega]}{dt} + \frac{d[\xi, D]}{\xi D dt} + \frac{d[\omega, D]}{D dt} + \frac{d[\omega, \xi]}{\xi dt}.$$

Another application of Itô's Lemma leads to

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

where the subscripts \hat{p}_i and $\hat{p}_i\hat{p}_j$ denote first and second derivatives with respect to the respective state variables \hat{p}_i . We can formulate the drift μ_ω as a function of the derivatives $\omega_{\hat{p}_i}$ and $\omega_{\hat{p}_i\hat{p}_j}$ and \hat{p}_i :

$$\mu_\omega = \sum_{i=1}^{n-1}\omega_{\hat{p}_i}\left(\hat{p}_{it}\lambda_{ii} + \sum_{\substack{j=1 \\ j \neq i}}^n\hat{p}_{jt}\lambda_{ji}\right) + \frac{1}{2}\sum_{i=1}^{n-1}\sum_{j=1}^{n-1}\omega_{\hat{p}_i\hat{p}_j}\sigma_{\hat{p}_i}\sigma'_{\hat{p}_j}.$$

The quadratic variation terms are:

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

Plugging everything into (6) gives the following PDE for ω :

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

Similar to the wealth-consumption ratio, we approximate the price-dividend ratio $U(\hat{p}) = e^\omega$ with a multivariate Chebyshev polynomial expansion:

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

and solve the PDE numerically.

E Pricing of real bonds

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

$$\xi_{t,T} = \delta^\theta \left(\frac{C_T}{C_t} \right)^{-\gamma} \exp \left\{ -\delta\theta(T-t) + (\theta-1) \left(\int_t^T \exp(-v(\hat{p})) ds + v(\hat{p}) \right) \right\}$$

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

$$0 = \mathcal{A}H = \mu_\xi + \mu_b + \frac{1}{2} \frac{d[b_t]}{dt} + \frac{d[\xi, b_t]}{\xi dt}. \quad (7)$$

The dynamics of b_t are

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

Plugging everything into (7) gives the following PDE for b_t :

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

Note that the PDE for the bond price involves a time derivative. We approximate the bond price B_t^T at each time point t by multivariate Chebyshev polynomials:

$$B_t^T = \exp \left\{ \sum_{j=0}^n \alpha_{j,t,T} T_j(\hat{p}) \right\}.$$

We use an explicit Euler discretization for the time derivative and solve the PDE recursively backwards in time, starting from the boundary condition $e^{b_T} = 1$, i.e. $\alpha_{j,T,T} = 0$.

F Pricing of nominal bonds

Let the price of the nominal bond expiring at time T be

$$B_t^{T,\$} = E_t[\xi_{t,T}^\$] = E_t \left[\xi_{t,T} \times \exp \left(\int_t^T -\pi_\tau d\tau \right) \right]$$

and $b_t^\$ = \ln B_t^{T,\$}$. Then the Feynman-Kac formula applied to $H(\xi_t^\$, b_t^\$) = \xi_t^\$ e^{b_t^\$}$ yields the partial differential equation

$$0 = \mathcal{A}H = \mu_{\xi^\$} + \mu_{b^\$} + \frac{1}{2} \frac{d[b_t^\$]}{dt} + \frac{d[\xi^\$, b_t^\$]}{\xi dt} \quad (8)$$

Notice that

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

The dynamics of $b_t^\$$ are

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

Plugging everything into (8) gives the following PDE for $b_t^\$$:

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

Again, this PDE involves a time derivative. As for the prices of real bonds, we approximate $B_t^{T,\$}$ at each time point t by multivariate Chebyshev polynomials, use an explicit Euler discretization for the time derivative and solve the PDE recursively backwards in time, starting from the boundary condition $e^{b_T^\$} = 1$.

References

- ANG, A., AND M. ULRICH (2012): “Nominal Bonds, Real Bonds and Equity,” Working Paper.
- 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 (2014): “Survey Forecasts and the Time-Varying Second Moments of Stock and Bond Returns,” Working Paper.
- 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): “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.
- (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.
- 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.
- BLANCHARD, O. (2009): “The State of Macro,” Annual Review of Economics, 1, 209–228.
- 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 (2015): “Monetary Policy Drivers of Bond and Equity Risks,” Working Paper.
- CAMPBELL, J., A. SUNDERAM, AND L. VICEIRA (2013): “Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds,” Working Paper.
- CHEN, H., W. DOU, AND L. KOGAN (2015): “Measuring the “Dark Matter” in Asset Pricing Models,” Working Paper.
- COLLIN-DUFRESNE, P., M. JOHANNES, AND L. A. LOCHSTOER (2016): “Parameter learning in general equilibrium: The asset pricing implications,” American Economic Review, 106(3), 664–698.
- DAVID, A., AND P. VERONESI (2013): “What Ties Return Volatilities to Price Valuations and Fundamentals,” Journal of Political Economy, 121(4), 682–746.
- (2014): “Investors’ and Central Bank’s Uncertainty Measures Embedded in Index Options,” The Review of Financial Studies, 27(6), 1661–1716.
- 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.
- (1992b): “Stochastic Differential Utility,” Econometrica, 60(2), 353–394.
- EHLING, P., M. GALLMEYER, C. HEYERDAHL-LARSEN, AND P. ILLEDITSCH (2016): “Disagreement about Inflation and the Yield Curve,” Working Paper.
- ERAKER, B., I. SHALIASTOVICH, AND W. WANG (2016): “Durable Goods, Inflation Risk and Equilibrium Asset Prices,” Review of Financial Studies, 29(1), 193–231.
- FLECKENSTEIN, M., F. LONGSTAFF, AND H. LUSTIG (2013): “Deflation Risk,” Working Paper.
- 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): “The U.S. Treasury yield curve: 1961 to the present,” Journal of Monetary Economics, 54, 2291–2304.

- HAMILTON, J. (1990): “Analysis of time series subject to changes in regime,” Journal of Econometrics, 45(1), 39–70.
- HASSELTOFT, H. (2012): “Stocks, Bonds, and Long-Run Consumption Risks,” Journal of Financial and Quantitative Analysis, 47(2), 309–332.
- KOLJEN, R., H. LUSTIG, AND S. VAN NIEUWERBURGH (2015): “The Cross-Section and Time Series of Stock and Bond Returns,” Working Paper.
- 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.
- MARFE, R., AND J. PENASSE (2017): “The Time-Varying Risk of Macroeconomic Disasters,” Working Paper.
- PIAZZESI, M., AND M. SCHNEIDER (2006): “Equilibrium Yield Curves,” NBER Macroeconomics Annual, pp. 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.
- SONG, D. (2014): “Bond Market Exposures to Macroeconomic and Monetary Policy Risks,” Working Paper.
- TSAI, J., AND J. WACHTER (2015): “Disaster risk and its implications for asset pricing,” Annual Review of Financial Economics, 7, 219–252.
- WACHTER, J. (2013): “Can time-varying risk of rare disasters explain aggregate stock market volatility?,” Journal of Finance, 68(3), 987–1035.




Figure 1: Time-varying disaster probabilities (estimated from consumption growth and inflation)

The blue solid line depicts the time-varying disaster intensity which Wachter (2013) extracts from asset price data via reverse engineering. Our special thanks go to Jessica Wachter for making this time series available to us. The red dashed line shows 5-year moving averages of the estimated $\hat{p}_3 + \hat{p}_4$ from our model for the period from 1947 to 2014. The correlation between the two time series is 0.88.




Figure 2: Time-varying disaster probabilities (estimated from consumption growth only)

This figure depicts essentially the same as Figure 1. But now the estimation is based on consumption data only. In that case, the information criterion favors a two-state Markov chain with expected consumption growth rates of 2.34 or -1.80 percentage points annually. The correlation between the two time series is now 0.33.

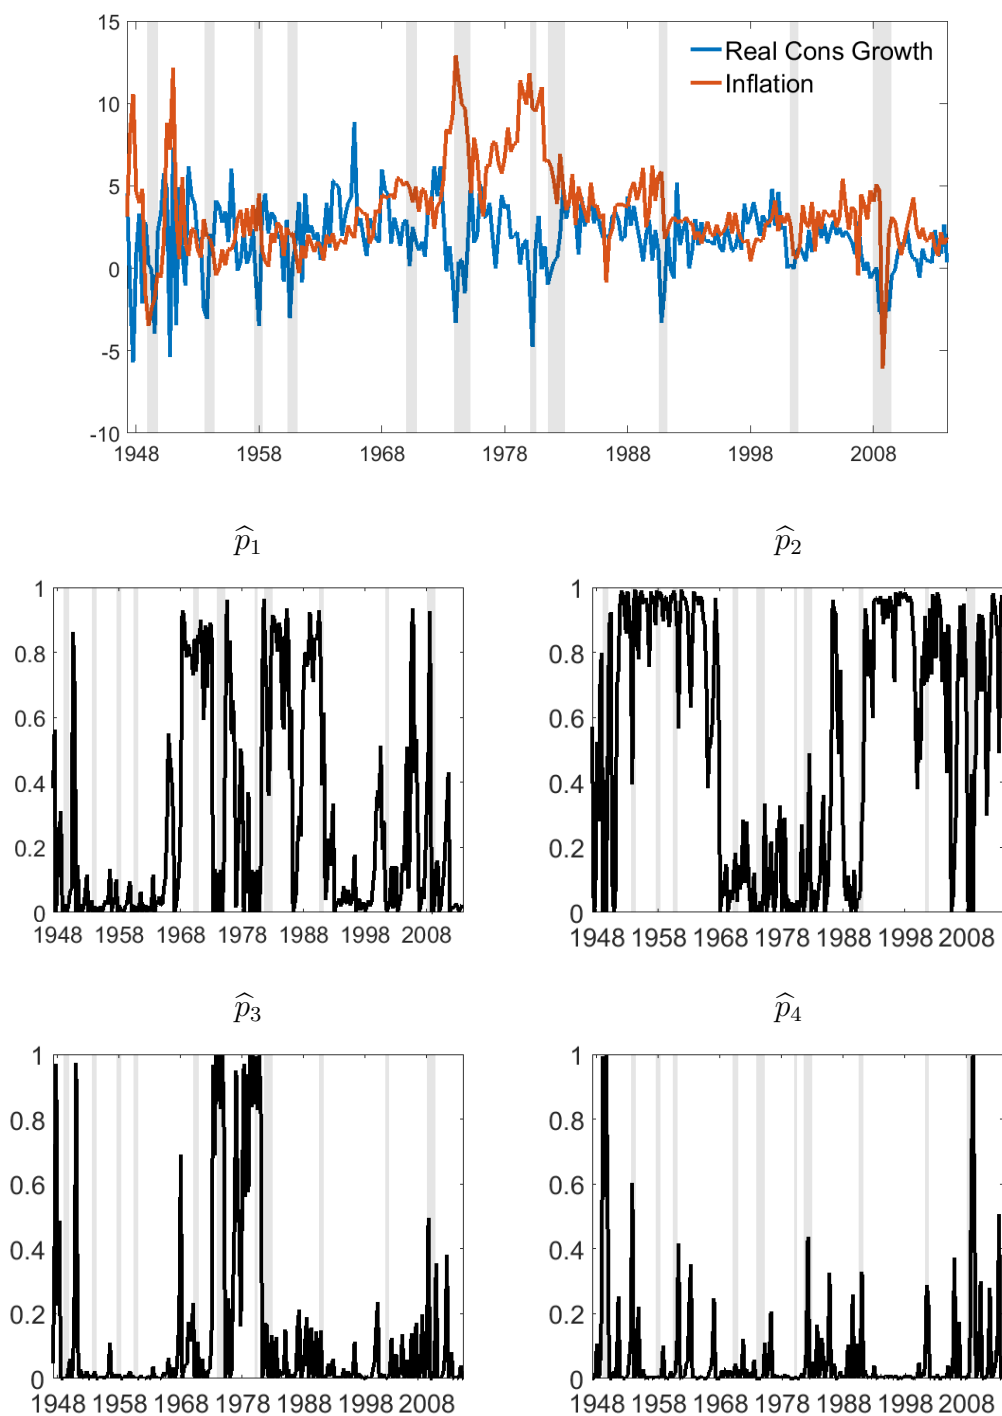



Figure 3: Fundamental data and filtered probabilities


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. Shaded areas indicate NBER recessions.

Estimation with inflation data


quarterly gdp growth
(1947-2014)




monthly cons growth
(1959-2014)



quarterly cons growth
(1962-2014)




quarterly cons growth
(1965-2014)




Estimation without inflation data


quarterly gdp growth
(1947-2014)



monthly cons growth
(1959-2014)



quarterly cons growth
(1962-2014)



quarterly cons growth
(1965-2014)




Figure 4: Time-varying disaster probabilities (robustness for different samples)

The figure depicts the results from applying the same estimation methodology as in Section 2.3 to different samples. The upper graphs show the time series 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 use monthly consumption growth rates instead of quarterly consumption growth rates, which are available from 1959 onwards. The four pictures on the right hand side are obtained from the subsamples of quarterly consumption growth rates starting in 1962 and 1965.




Figure 5: Time-varying disaster probabilities (estimated from consumption growth and inflation)

The blue solid line depicts the time-varying disaster intensity which Wachter (2013) extracts from asset price data via reverse engineering. The red dashed line shows 5-year moving averages of the estimated $\hat{p}_3 + \hat{p}_4$ from the constrained model for the period from 1947 to 2014. The correlation between the two time series is 0.88.




Figure 6: Scatter plots of stock return volatilities

The figure depicts scatter plots for the regressions presented in Table 2. 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 $\hat{p}_3 + \hat{p}_4$ over the same 20 quarters periods. The left figure labeled “Data” is based on the estimated time series of the \hat{p}_i depicted in Figure 3 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve. The right figure labeled “Model” is based on the same time series of the \hat{p}_i , but uses 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.




Figure 7: Conditional stock return volatilities

The figure depicts conditional 20-quarter rolling window stock return volatilities. The red dashed line is based on the estimated time series of the \hat{p}_i depicted in Figure 3 as well as the CRSP value-weighted index. The blue solid line is based on the same time series of the \hat{p}_i , but uses the returns which our model would have implied given this path 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.

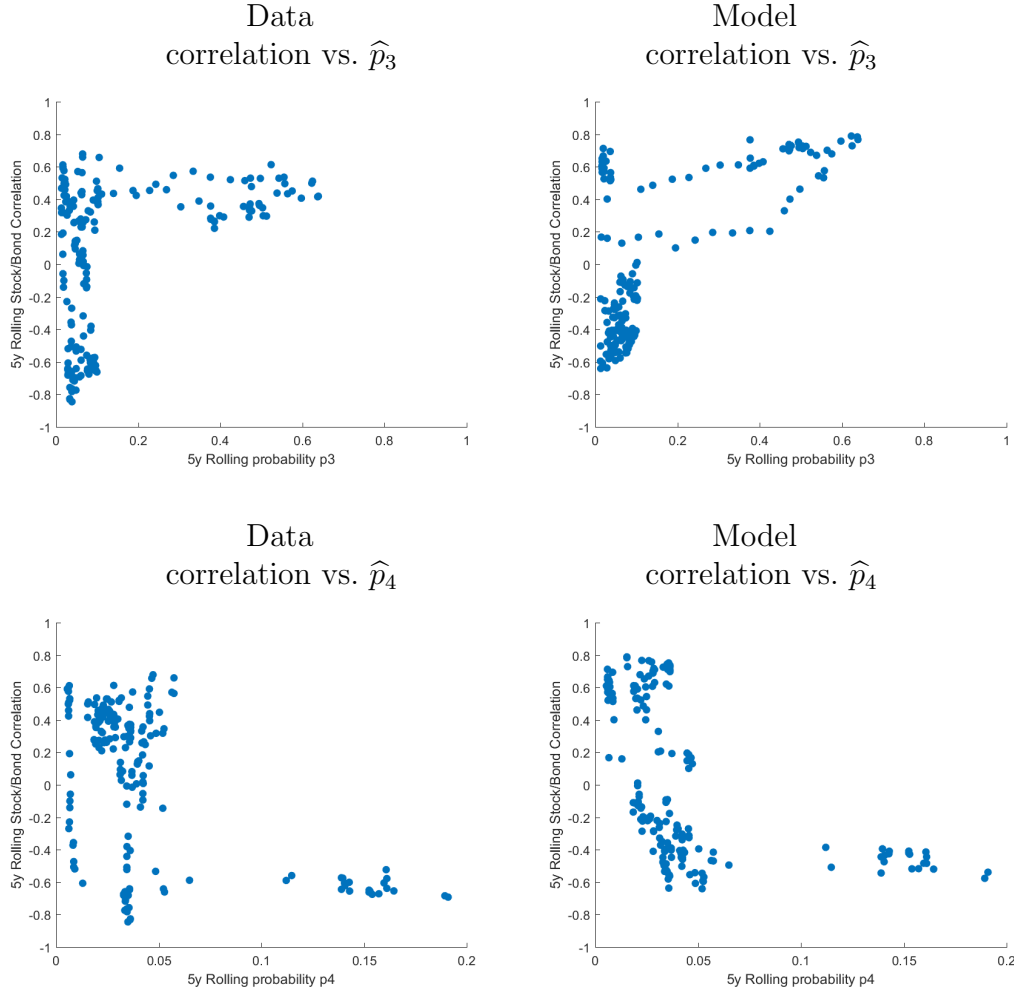


Figure 8: Scatter plots of stock-bond return correlations

The figure depicts scatter plots for the regressions presented in Table 4. 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 \hat{p}_3 and \hat{p}_4 over the same 20 quarters periods. The left figures labeled “Data” are based on the estimated time series of the \hat{p}_i depicted in Figure 3 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve. The right figures labeled “Model” are based on the same time series of the \hat{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.




Figure 9: 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 estimated time series of the \hat{p}_i depicted in Figure 3 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve. The blue solid line is based on the same time series of the \hat{p}_i , but uses the returns which our model would have implied given this path 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 0.54.




Figure 10: 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 estimated time series of the \hat{p}_i depicted in Figure 3 as well as the CRSP value-weighted index and data on long-term Treasury bonds' returns from the Ibbotson Stocks, Bonds, Bills and Inflation Yearbook from 1947 till 2014. The blue solid line is based on the same time series of the \hat{p}_i , but uses the returns which our model would have implied given this path 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 0.54.

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 (0.487)	1.898 (0.155)	0.444 (0.315)	-0.997 (0.631)	1.016 (0.099)	-0.136 (0.030)	-0.218 (0.103)
Inflation	4.704 (0.975)	2.161 (0.846)	9.514 (1.668)	-2.917 (2.438)	0.382 (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 (0.202)	0.027 (0.087)	0.039 (0.160)	0.025 (0.049)
from state 2	0.022 (0.029)	0.970 (0.056)	0.008 (0.027)	0 (0.024)
from state 3	0.135 (0.064)	0.037 (0.077)	0.828 (0.088)	0 (0.028)
from state 4	0 (0.065)	0.337 (0.253)	0 (0.152)	0.663 (0.203)

Panel C: Optimal number of states				
	3 states	4 states	5 states	6 states
Bayes Information Criterion	835	805	854	896

Table 1: Markov chain estimation

This table reports the results from our baseline Markov chain estimation. Growth rates are given in percentage points and annualized. Data are from 1947 to 2014 at quarterly frequency. The numbers in parantheses give standard errors for all the estimated parameters. They have been obtained from a standard block bootstrap with block length of 10 quarters.

Panel A: Model				
const.	$\log(\hat{p}_3)$	$\log(\hat{p}_4)$	$\log(\hat{p}_3 + \hat{p}_4)$	Adj. R^2
0.114 (12.627) [0.085, 0.174]	0.009 (10.833) [0.004, 0.020]	0.012 (4.724) [0.000, 0.023]		0.733 [0.220, 0.785]
0.079 (12.143) [0.071, 0.129]			0.015 (5.997) [0.007, 0.029]	0.615 [0.104, 0.763]
Panel B: Data				
const.	$\log(\hat{p}_3)$	$\log(\hat{p}_4)$	$\log(\hat{p}_3 + \hat{p}_4)$	Adj. R^2
0.278 (12.679)	0.012 (1.416)	0.023 (4.475)		0.304
0.214 (9.056)			0.022 (1.976)	0.214

Table 2: 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 \hat{p}_i over the same 20 quarters periods. The regressions labeled “Data” are based on the estimated time series of the \hat{p}_i depicted in Figure 3 as well as the CRSP value-weighted index and the interpolated yield curve data from the Federal Reserve. The regressions labeled “Model” are based on the same time series of the \hat{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 denote Newey-West-adjusted t-statistics (20 lags). The numbers in brackets denote 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).

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

Table 3: Regressions of stock return volatilities 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 ($\hat{p}_3 \ln \hat{p}_3 + \hat{p}_4 \ln \hat{p}_4$), 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 2.

Panel A: Model				
const.	$\log(\hat{p}_3)$	$\log(\hat{p}_4)$	$\log(\hat{p}_3 + \hat{p}_4)$	Adj. R^2
-0.837 (-4.825) [-1.833, 0.556]	0.290 (6.776) [0.065, 0.453]	-0.449 (-5.108) [-0.702, -0.089]		0.788 [0.172, 0.751]
0.325 (0.881) [-0.867, 0.962]			0.161 (0.903) [-0.275, 0.473]	0.087 [-0.009, 0.481]
Panel B: Data				
const.	$\log(\hat{p}_3)$	$\log(\hat{p}_4)$	$\log(\hat{p}_3 + \hat{p}_4)$	Adj. R^2
-0.512 (-1.61)	0.182 (2.731)	-0.294 (-3.608)		0.363
0.239 (0.923)			0.091 (0.774)	0.041

Table 4: 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 \hat{p}_i over the same 20 quarters periods. “Data” and “Model” have the same meaning as in Table 2. The numbers in parentheses denote Newey-West-adjusted t-statistics (20 lags). The numbers in brackets denote 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).

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

Table 5: Regressions of stock-bond return correlations 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 ($\hat{p}_3 \ln \hat{p}_3 + \hat{p}_4 \ln \hat{p}_4$), 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 2.

	Data	Model
Average real equity return	0.082	0.037 (0.005)
Volatility of real equity returns	0.164	0.073 (0.005)
Average nominal 3m rate	0.043	0.060 (0.005)
Volatility of nominal rate	0.050	0.006 (0.001)
Average yield spread (5y - 3m)	0.010	-0.005 (0.003)
Stock-bond correlation	0.114	-0.027 (0.158)

Table 6: 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, Sack, and Wright (2007) for bonds. The model-implied values are computed via Monte Carlo simulation where the model is parametrized according to Table 1. 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.