ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 86, No. 2 (March, 2018), 617–654

IDENTIFYING LONG-RUN RISKS: A BAYESIAN MIXED-FREQUENCY APPROACH

FRANK SCHORFHEIDE
Department of Economics, University of Pennsylvania

DONGHO SONG
Department of Economics, Boston College

AMIR YARON
The Wharton School, University of Pennsylvania

Econometrica, Vol. 86, No. 2 (March, 2018), 617–654

IDENTIFYING LONG-RUN RISKS: A BAYESIAN MIXED-FREQUENCY APPROACH

FRANK SCHORFHEIDE
Department of Economics, University of Pennsylvania

DONGHO SONG
Department of Economics, Boston College

AMIR YARON
The Wharton School, University of Pennsylvania

We document that consumption growth rates are far from i.i.d. and have a highly persistent component. First, we estimate univariate and multivariate models of cash-flow (consumption, output, dividends) growth that feature measurement errors, time-varying volatilities, and mixed-frequency observations. Monthly consumption data are important for identifying the stochastic volatility process; yet the data are contaminated, which makes the inclusion of measurement errors essential for identifying the predictable component. Second, we develop a novel state-space model for cash flows and asset prices that imposes the pricing restrictions of a representative-agent endowment economy with recursive preferences. To estimate this model, we use a particle MCMC approach that exploits the conditional linear structure of the approximate equilibrium. Once asset return data are included in the estimation, we find even stronger evidence for the persistent component and are able to identify three volatility processes: the one for the predictable cash-flow component is crucial for asset pricing, whereas the other two are important for tracking the data. Our model generates asset prices that are largely consistent with the data in terms of sample moments and predictability features. The state-space approach allows us to track over time the evolution of the predictable component, the volatility processes, the decomposition of the equity premium into risk factors, and the variance decomposition of asset prices.

KEYWORDS: Asset pricing, Bayesian inference, consumption dynamics, long-run risks, measurement errors, mixed frequency observations, nonlinear state-space model, particle MCMC, stochastic volatility.

1. INTRODUCTION

THE DYNAMICS OF AGGREGATE CONSUMPTION play a key role in business cycle models, tests of the permanent income hypothesis, and asset pricing. Perhaps surprisingly, there is a significant disagreement about the basic time series properties of consumption. First, while part of the profession holds a long-standing view that aggregate consumption follows a random walk (e.g., Hall (1978) and Campbell and Cochrane (1999)), the recent literature on long-run risks (LRR) (e.g., Bansal and Yaron (2004) and Hansen, Heaton, and Li (2008)) emphasizes the presence of a small persistent component in consumption

Frank Schorfheide: schorf@ssc.upenn.edu
Dongho Song: dongho.song@bc.edu
Amir Yaron: yaron@wharton.upenn.edu

We thank a co-editor, four anonymous referees, Bent J. Christensen, Ian Dew-Becker, Frank Diebold, Emily Fox, Roberto Gomez Cram, Lars Hansen, Arthur Lewbel, Lars Lochstoer, Ivan Shaliastovich, Neil Shephard, Minchul Shin, and participants at many seminars and conferences for helpful comments and discussions. Schorfheide gratefully acknowledges financial support from the National Science Foundation under Grant SES 1424843. Yaron thanks the Rodney White and Jacob Levy centers for financial support.

 https://doi.org/10.3982/ECTA14308

growth.[1] Second, while time-varying volatility was a feature that until recently was mainly associated with financial time series, there is now a rapidly growing literature stressing the importance of stochastic volatility in macroeconomic aggregates (e.g., Bansal and Yaron (2004), Bloom (2009), and Fernández-Villaverde and Rubio-Ramírez (2011)) and the occurrence of rare disasters (e.g., Barro (2009) and Gourio (2012)).

Studying consumption growth dynamics leads to the following challenge. On the one hand, it is difficult to identify time variation in volatility based on time-aggregated data (e.g., Drost and Nijman (1993)), which favors the use of high-frequency monthly data. On the other hand, monthly consumption growth data are contaminated by measurement errors (e.g., Slesnick (1998) and Wilcox (1992)), which mask the dynamics of the underlying process. We address this challenge by developing and estimating a novel Bayesian state-space model with an elaborate measurement error component that is consistent with the view that annual and quarterly consumption data are more accurate than monthly data. The model is tailored toward monthly data, but a mixed-frequency approach enables us to accommodate the longest span of annual consumption growth data starting from the Great Depression era.

In the first part of our empirical analysis, we provide strong evidence for a persistent component of consumption growth as well as its time-varying volatility, which contradicts the commonly held view that consumption follows a random walk. The combination of measurement errors and the local-level component in "true" consumption growth in our empirical model allows us to generate the strong second-order moving average (MA(2)) component in observed consumption growth. Our basic empirical finding is robust across a wide range of model specifications that include univariate models for consumption growth as well as bivariate models with either output or dividend growth as second observable. The bivariate models feature a common persistent factor in cash-flow growth rates. An important conclusion from our analysis is that plausible models of observed monthly consumption growth need to contain a MA(2) component, while macroeconomic models that confront monthly data should filter out the high-frequency movements attributable to measurement errors.

In the second part of our empirical analysis, we embed the cash-flow process into a representative-agent endowment economy as in Bansal and Yaron (2004). This model is referred to as long-run risks (LRR) model. When asset returns are added to the set of observables and the LRR model is jointly estimated with the dynamics of consumption and dividend growth, the credible intervals for a common persistent component in cash-flow growth rates are further sharpened and three separate volatility components are identified: one governing dynamics of the persistent cash-flow growth component, and the other two controlling temporally independent shocks to consumption and dividend growth. The stochastic volatility process for the persistent component is important for asset prices, while the other two volatility processes are important for tracking the data. We show that the estimated LRR model generates asset prices that are largely consistent with the data. Moreover, we demonstrate that if we replace the parameters of the cash-flow processes from the joint estimation by those obtained from the cash-flow-only estimation, the LRR model still has by-and-large realistic asset pricing implications.

In addition to the empirical results, our paper also contains an important technical innovation. To incorporate market returns and the risk-free rate into the state-space model that is used for the second part of the empirical analysis, we have to solve for the asset

[1]The literature on robustness (e.g., Hansen and Sargent (2007)) highlights that merely contemplating low-frequency shifts in consumption growth can be important for macroeconomic outcomes and asset prices.

pricing implications of the LRR model to obtain measurement equations for these two series.[2] Unlike in the cash-flow-only specifications, the model with asset prices has the feature that the volatility processes also affect the conditional means of the asset prices. This considerably complicates the evaluation of the likelihood function with a nonlinear filter as well as the implementation of Bayesian inference. In fact, due to the high-dimensional state space that arises from our measurement error model and the mixed-frequency setting, this nonlinear filtering is a seemingly daunting task. We show how to exploit the partially linear structure of the state-space model to derive a very efficient sequential Monte Carlo (particle) filter.

Unlike the generalized method of moments (GMM) approach that is common in the LRR literature, our sophisticated state-space approach lets us track the predictable component x_t as well as the stochastic volatilities over time. In turn, this allows us to construct period-by-period decompositions of risk premia and asset price variances. Our Bayesian approach allows us to account for three types of statistical uncertainties in a unified manner: parameter uncertainty, uncertainty about the hidden states, and uncertainty about future (or hypothetical shocks). These three types of uncertainty feature prominently in our empirical results. Depending on the question at hand, we present in some instances credible bands for our results reflecting multiple sources of uncertainty, for example, when we provide bands for the predictable component of cash flows; and in other instances, to facilitate clear comparisons across parameterizations, we focus on the dominating source of uncertainty, for example, shock uncertainty when we examine the model-implied sample moments of asset prices or the sampling distribution of R^2's from predictability regressions.

Our empirical analysis starts with the estimation of a state-space model according to which consumption growth is the sum of an i.i.d. and an AR(1) component, focusing on the persistence ρ of the AR(1) component. We show that once we include monthly measurement errors that average out at the annual frequency, the fit of the model improves significantly, and we obtain an estimate of ρ around 0.92.[3] Using a battery of model specifications, we show that our measurement error model in which measurement errors account for half of the variation in monthly consumption is the preferred one. We further show that the estimation of the monthly model with measurement errors leads to a more accurate estimate of ρ than the estimation with time-aggregated data. Importantly, adding stochastic volatility leads to a further improvement in model fit, a reduction in the posterior uncertainty about ρ, and an increase in the point estimate of ρ to 0.95. Because consumption is generally viewed as being influenced by output fluctuations, we use our framework to show that a similar persistent component is also important for characterizing quarterly GDP dynamics. When we estimate a common persistent component in consumption and output growth (imposing cointegration), inference regarding ρ is essentially the same as in the consumption-only specifications. When we augment the state-space model to include a measurement equation for dividend growth as a precursor to ultimately pricing equity, the joint estimation based on consumption and dividend growth based on post-1959 data leads the estimate of ρ to rise to 0.97.

[2] In order to solve the model, we approximate the exponential Gaussian volatility processes by linear Gaussian processes such that the standard analytical solution techniques that have been widely used in the LRR literature can be applied. The approximation of the exponential volatility process is used only to derive the coefficients in the law of motion of the asset prices.

[3] Without accounting for measurement errors, the estimate of ρ using monthly consumption growth data is insignificantly different from 0, which can partly account for some view that consumption growth is an i.i.d. process.

The LRR model used in the second part of the empirical analysis distinguishes itself from the existing LRR literature in two important dimensions: our model for the cash flows includes measurement errors and three volatility processes to improve the fit. Moreover, we specify an additional process for variation in the time rate of preference as in Albuquerque, Eichenbaum, Luo, and Rebelo (2016), which generates risk-free rate variation that is independent of cash flows and leads to an improved fit for the risk-free rate. The estimation of the LRR model delivers several important empirical findings. First, the estimate of ρ, that is, the autocorrelation of the persistent cash-flow component, is 0.987, which is higher than what we obtained based on the cash-flow-only estimation.[4] Importantly, we show that the time path of the persistent component looks very similar with and without asset price data.

Second, the volatility processes partly capture heteroscedasticity of innovations, and in part they break some of the tight links that the model imposes on the conditional mean dynamics of asset prices and cash flows. This feature significantly improves the model implications for consumption and return predictability. An important feature of our estimation is that the likelihood focuses on conditional correlations between the risk-free rate and consumption—a dimension often not directly targeted in the literature. We show that because consumption growth and its volatility determine the risk-free rate dynamics, one requires another independent process to account for the weak correlation between consumption growth and the risk-free rate. The independent time-rate-of-preference shocks mute the model-implied correlation further and improve the model fit in regard to the risk-free rate dynamics.

Third, it is worth noting that the median posterior estimate for risk aversion is around 9, while it is around 2 for the intertemporal elasticity of substitution (IES). These estimates are broadly consistent with the parameter values highlighted in the LRR literature (see Bansal, Kiku, and Yaron (2012), and Bansal, Kiku, and Yaron (2014)). Fourth, at the estimated preference parameters and those characterizing the consumption and dividend dynamics, the model is able to successfully generate many key asset pricing moments, and improve model performance relative to previous LRR models along several dimensions. In particular, the posterior median of the equity premium is 8.2%, while the model's posterior predictive distribution is consistent with the observed large volatility of the price-dividend ratio at 0.45, and the R^2's from predicting returns and consumption growth by the price-dividend ratio.

Our paper is connected to several strands of the literature. In terms of the LRR literature, Bansal, Kiku, and Yaron (2014) utilized data that are time-aggregated to annual frequency to estimate the LRR model by GMM, and Bansal, Gallant, and Tauchen (2007) pursued an approach based on the efficient method of moments (EMM). Both papers use cash-flow *and* asset price data jointly for the estimation of the parameters of the cash-flow process. Our likelihood-based approach provides evidence which is broadly consistent with the results highlighted in those papers and other calibrated LRR models, for example, Bansal, Kiku, and Yaron (2012). Our likelihood function implicitly utilizes a broader set of moments than earlier GMM or EMM estimation approaches. These moments include the entire sequence of autocovariances as well as higher-order moments of the time series used in the estimation, and let us measure the time path of the predictable component of cash flows as well as the time path of the innovation volatilities. Rather than asking the model to fit a few selected moments, we are raising the bar and force the

[4]The corresponding half-lives of the cash-flow-only (0.97) and asset pricing based (0.987) estimates for ρ are 1.9 and 4.4 years, respectively.

model to track cash-flow and asset return time series. Finally, it is worth noting that our paper distinguishes itself from previous LRR literature in showing that even by just using monthly consumption growth data with an appropriate measurement error structure, we are able to estimate the highly persistent predictable component. In complementary research, Nakamura, Sergeyev, and Steinsson (2015) showed that an estimation based on a long cross-country panel of annual consumption data also yields large estimates of the autocorrelation of the persistent component.

To implement Bayesian inference, we embed a particle-filter-based likelihood approximation into a Metropolis–Hastings algorithm as in Fernández-Villaverde and Rubio-Ramírez (2007) and Andrieu, Doucet, and Holenstein (2010). This algorithm belongs to the class of particle Markov chain Monte Carlo (MCMC) algorithms. Because our state-space system is linear conditional on the volatility states, we can use Kalman-filter updating to integrate out a subset of the state variables. The genesis of this idea appears in the mixture Kalman filter of Chen and Liu (2000). Particle filter methods are also utilized in Johannes, Lochstoer, and Mou (2016), who estimated an asset pricing model in which agents have to learn about the parameters of the cash-flow process from consumption growth data. While Johannes, Lochstoer, and Mou (2016) examined the role of parameter uncertainty for asset prices, which is ignored in our analysis, they used a more restrictive version of the cash-flow process and did not utilize mixed-frequency observations.[5]

Our state-space setup makes it relatively straightforward to utilize data that are available at different frequencies. The use of state-space systems to account for missing monthly observations dates back to at least Harvey (1989) and has more recently been used in the context of dynamic factor models (see, e.g., Mariano and Murasawa (2003) and Aruoba, Diebold, and Scotti (2009)) and VARs (see, e.g., Schorfheide and Song (2015)). Finally, there is a growing and voluminous literature in macro and finance that highlights the importance of volatility for understanding the macroeconomy and financial markets (see, e.g., Bansal, Khatacharian, and Yaron (2005), Bloom (2009), Fernández-Villaverde and Rubio-Ramírez (2011), Bansal, Kiku, and Yaron (2012), and Bansal, Kiku, Shaliastovich, and Yaron (2014)). Our volatility specification that accommodates three processes further contributes to identifying the different uncertainty shocks in the economy.

The remainder of the paper is organized as follows. Section 2 introduces the state-space model for consumption growth and presents the empirical findings based on consumption growth data. In Section 3, we consider multivariate cash-flow models and examine the evidence for a predictable growth rate component in specifications that include GDP growth and dividend growth. Section 4 introduces the LRR asset pricing model, and describes the model solution and the particle MCMC approach used to implement Bayesian inference. Section 5 discusses the empirical findings obtained from the estimation of the LRR model, and Section 6 provides concluding remarks. A description of our data sources, analytical derivations, a detailed description of the state-space representations and posterior inference, and additional empirical results are relegated to the Supplemental Material (Schorfheide, Song, and Yaron (2018)).

[5]Building on our approach, Creal and Wu (2015) used gamma processes to model time-varying volatilities and estimated a yield curve model using particle MCMC. Doh and Wu (2015) estimated a nonlinear asset pricing model in which all the asset prices and the consumption process are quadratic rather than linear function of the states.

2. MODELING CONSUMPTION GROWTH

The first step of our analysis is to develop an empirical state-space model for consumption growth. We take the length of the period to be one month. The use of monthly data is important for identifying stochastic volatility processes. Unfortunately, consumption data are less accurate at monthly frequency than at the more widely used quarterly or annual frequencies. In this regard, the main contribution in this section is a novel specification of a measurement error model for consumption growth, which has the feature that monthly measurement errors average out under temporal aggregation. Moreover, because monthly consumption data have only been published since 1959, we use annual consumption growth rates prior to 1959 and adapt the measurement equation to the data availability.[6] We develop our measurement error model in Section 2.1 and present the empirical results in Section 2.2.

2.1. *A Measurement Equation for Consumption*

We proceed in two steps. First, we derive a measurement equation for consumption growth at the annual frequency, which is used for pre-1959 data. Second, we specify a measurement equation for consumption growth at the monthly frequency, which is used for post-1959 data. We use C_t^o and C_t to denote the observed and the "true" level of consumption, respectively. Moreover, we represent the monthly time subscript t as $t = 12(j-1)+m$, where $m = 1, \ldots, 12$. Here j indexes the year and m the month within the year.

Measurement of Annual Consumption Growth. We define annual consumption as the sum of monthly consumption over the span of one year, that is, $C^a_{(j)} = \sum_{m=1}^{12} C_{12(j-1)+m}$. Log-linearizing this relationship around a monthly value C_* and defining lowercase c as percentage deviations from the log-linearization point, that is, $c = \log C/C_*$, we obtain $c^a_{(j)} = \frac{1}{12}\sum_{m=1}^{12} c_{12(j-1)+m}$. Defining monthly consumption growth as the log difference

$$g_{c,t} = c_t - c_{t-1},$$

we can deduce that annual growth rates are given by

$$g^a_{c,(j)} = c^a_{(j)} - c^a_{(j-1)} = \sum_{\tau=1}^{23}\left(\frac{12-|\tau-12|}{12}\right) g_{c,12j-\tau+1}. \tag{1}$$

We assume a multiplicative i.i.d. measurement error model for the level of annual consumption, which implies that, after taking log differences, observed annual consumption growth (o superscript)

$$g^{a,o}_{c,(j)} = g^a_{c,(j)} + \sigma^a_\epsilon\left(\epsilon^a_{(j)} - \epsilon^a_{(j-1)}\right). \tag{2}$$

Measurement of Monthly Consumption Growth. Consistent with the practice of the Bureau of Economic Analysis (BEA), we assume that the levels of monthly consumption are constructed by distributing annual consumption over the 12 months of a year. We approximate the BEA's data construction by assuming that this distribution is based on an observed monthly proxy series z_t, where z_t is a noisy measure of monthly consumption.

[6] In principle, we could utilize the quarterly consumption growth data from 1947 to 1959, but we do not in this version of the paper.

The monthly levels of consumption are determined such that the growth rates of monthly consumption are proportional to the growth rates of the proxy series and monthly consumption adds up to annual consumption. A measurement error model that is consistent with this assumption is the following (a detailed derivation is provided in the Supplemental Material):

$$\begin{aligned} g^o_{c,12(j-1)+1} &= g_{c,12(j-1)+1} + \sigma_\epsilon(\epsilon_{12(j-1)+1} - \epsilon_{12(j-2)+12}) \\ &\quad - \frac{1}{12}\sum_{m=1}^{12} \sigma_\epsilon(\epsilon_{12(j-1)+m} - \epsilon_{12(j-2)+m}) + \sigma^a_\epsilon\left(\epsilon^a_{(j)} - \epsilon^a_{(j-1)}\right), \qquad (3) \\ g^o_{c,12(j-1)+m} &= g_{c,12(j-1)+m} + \sigma_\epsilon(\epsilon_{12(j-1)+m} - \epsilon_{12(j-1)+m-1}), \quad m = 2, \ldots, 12. \end{aligned}$$

The term $\epsilon_{12(j-1)+m}$ can be interpreted as the error made by measuring the level of monthly consumption through the monthly proxy variable, that is, in log deviations $c_{12(j-1)+m} = z_{12(j-1)+m} + \epsilon_{12(j-1)+m}$. The summation of monthly measurement errors in the second line of (3) ensures that monthly consumption sums up to annual consumption. It can be verified that converting the monthly consumption growth rates into annual consumption growth rates according to (1) averages out the measurement errors and yields (2).

2.2. *Empirical Analysis*

We use the per capita series of real consumption expenditure on nondurables and services from the NIPA tables available from the Bureau of Economic Analysis. Annual observations are available from 1929 to 2014, quarterly from 1947:Q1 to 2014:Q4, and monthly from 1959:M1 to 2014:M12. Growth rates of consumption are constructed by taking the first difference of the corresponding log series.

Autocorrelation of Consumption Growth. Figure 1 displays the sample autocorrelation of consumption growth for monthly, quarterly, and annual data, respectively. The figure clearly demonstrates that at the annual frequency, consumption growth is strongly positively autocorrelated, while at the monthly frequency, consumption growth has a negative first autocorrelation. These autocorrelation plots provide prima facie evidence for a negative moving average component at the monthly frequency, which is consistent with the

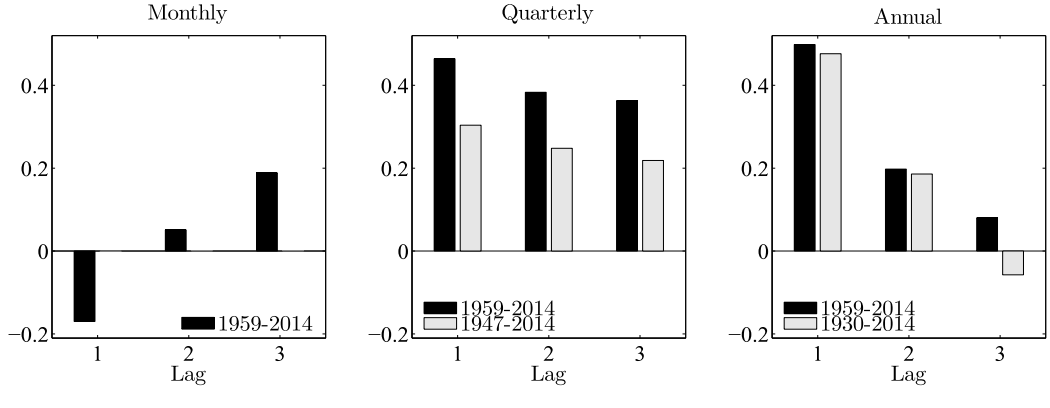

FIGURE 1.—Sample autocorrelation. *Notes*: Monthly data available from 1959:M2 to 2014:M12, quarterly from 1947:Q2 to 2014:Q4, annual from 1930 to 2014.

measurement error model described in Section 2.1. Our measurement error model can reconcile the monthly negative autocorrelation with a strongly positive autocorrelation for time-aggregated annual consumption. The right panel in Figure 1 also shows that the strong positive autocorrelation in annual consumption growth is robust to using the long pre-war sample as well as the post-war data. Given these features of the data, we focus our analysis of measurement errors in consumption using the post-1959 monthly series.

A State-Space Model for Consumption Growth. In our subsequent analysis, we will consider several different laws of motion for "true" consumption growth. The benchmark specification takes the following form:

$$\begin{aligned} g_{c,t+1} &= \mu_c + x_t + \sigma_{c,t}\eta_{c,t+1}, \quad \eta_{c,t+1} \sim N(0,1), \\ x_{t+1} &= \rho x_t + \sqrt{1-\rho^2}\varphi_x \sigma_{c,t}\eta_{x,t+1}, \quad \eta_{x,t+1} \sim N(0,1), \\ \sigma_{c,t} &= \sigma \exp(h_{c,t}), \qquad h_{c,t+1} = \rho_{h_c} h_{c,t} + \sigma_{h_c} w_{c,t+1}, \quad w_{c,t+1} \sim N(0,1). \end{aligned} \tag{4}$$

This specification is based on Bansal and Yaron (2004) and decomposes consumption growth $g_{c,t+1}$ into a persistent component, x_t, and a transitory component, $\sigma_{c,t}\eta_{c,t+1}$. The state-transition equation is augmented by the measurement equations (2) and (3) to form a state-space model.

The combination of measurement and state-transition equations leads to a high-dimensional state-space model; see the Supplemental Material for details. The data that we are using for the estimation have the property that monthly consumption is consistent with annual consumption. While the statistical agency may have access to the monthly proxy series z_t in real time, it can only release the monthly consumption series that is consistent with the corresponding annual consumption observation at the end of each year. Thus, we specify the state-space model such that, every 12 months, the econometrician observes 12 consumption growth rates. This implies that in addition to tracking the monthly measurement errors $\epsilon_{12(j-1)+m}$ for $m = 1, \ldots, 12$, the state-space model also has to track 12 lags of x_t.

Throughout this paper, we use Bayesian inference for the model parameters and the hidden states. In addition to the latent monthly consumption growth rates and measurement errors, the state space also comprises the latent volatility process $h_{c,t}$. Define the parameter vectors

$$\Theta_{cf} = [\mu_c, \rho, \varphi_x, \sigma]', \qquad \Theta_h = [\rho_{h_c}, \sigma_{h_c}],$$

and the sequence of latent volatilities $H_{0:T-1}$. To sample from the posterior distribution of $(\Theta_{cf}, \Theta_h, H_{0:T-1})$, we use a Metropolis-within-Gibbs algorithm that iterates over three conditional distributions: First, a Metropolis–Hastings step is used to draw from the posterior of Θ_{cf} conditional on $(Y, \Theta_h^{(s-1)}, H_{0:T-1}^{(s-1)})$. Here the likelihood $p(Y|\Theta_{cf}, H_{0:T-1}^{(s-1)})$ is evaluated with the Kalman filter. Second, we draw the sequence of stochastic volatilities $H_{0:T-1}$ conditional on $(Y, \Theta_{cf}^{(s)}, \Theta_h^{(s-1)})$ using the algorithm developed by Kim, Shephard, and Chib (1998). This step involves the use of a simulation smoother (e.g., Carter and Kohn (1994)) for a linear state-space model to obtain draws from the conditional posterior distributions of the "residuals" $g_{c,t+1} - \mu_c - x_t$ and $x_{t+1} - \rho x_t$. Conditional on these residuals, it is possible to draw from the posterior of $H_{0:T-1}$. Finally, we draw from the posterior of the coefficients of the stochastic volatility processes, Θ_h, conditional on $(Y, H_{0:T-1}^{(s)}, \Theta_{cf}^{(s)})$.

The Likelihood Function. We simplify the law of motion of consumption growth in (4) by assuming that the innovations are homoscedastic, that is, $\sigma_{h_c} = 0$ and $h_{c,t} = 0$ for all t.

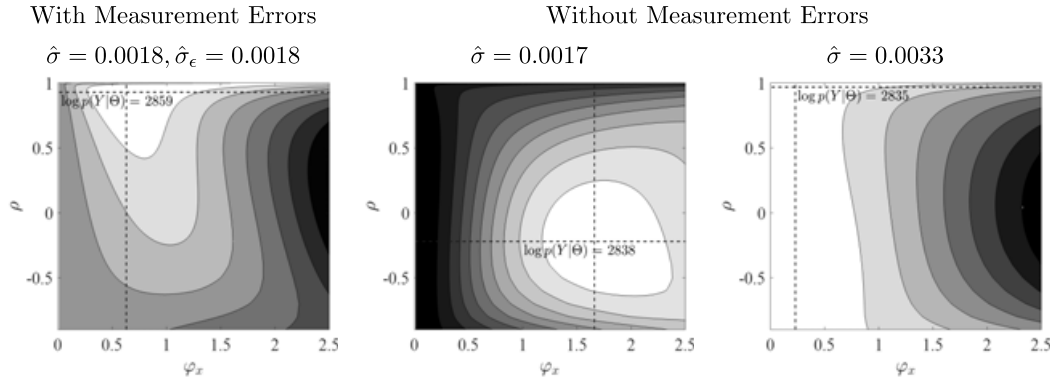

FIGURE 2.—Log-likelihood contour. *Notes*: We use maximum likelihood estimation to estimate the homoscedastic version ($\sigma_{h_c} = 0$, $h_{c,t} = 0$) of model (4) with and without allowing for measurement errors. We then fix $\sigma = \hat{\sigma}$ and $\sigma_\epsilon = \hat{\sigma}_\epsilon$ at their point estimates and vary ρ and φ_x to plot the log-likelihood function contour. Without measurement errors, we find that the log-likelihood function is bimodal at positive and negative values of ρ. Therefore, we obtain two sets of $\hat{\sigma}$.

In Figure 2, we plot likelihood function contours based on a sample of monthly consumption growth rates that ranges from 1959:M2 to 2014:M12. We consider two specifications: with and without measurement errors. To isolate the role of measurement errors for inference about ρ, we set μ_c to the sample mean and fix σ and σ_ϵ to their respective maximum likelihood estimates, while varying the two parameters, ρ and φ_x, that govern the dynamics of x_t. In the absence of measurement errors, the log-likelihood function is bimodal. The first mode is located at $\rho = -0.23$, which matches the negative monthly sample autocorrelation (see Figure 1). The location of the second mode is at $\rho = 0.96$, but the log-likelihood function is flat across a large set of values of ρ between -1 and 1. Importantly, when we allow for monthly measurement errors according to (3), setting $\sigma_\epsilon^a = 0$, the log-likelihood function has a very sharp peak, displaying a very persistent expected consumption growth process with $\rho = 0.92$. Measurement errors at the monthly frequency help identify a large persistent component in consumption by allowing the model to simultaneously match the negative first-order autocorrelation observed at the monthly frequency and the large positive autocorrelation at the annual frequency.

Bayesian Estimation of Homoscedastic Models. We now proceed with the Bayesian estimation of variants of (4) using the monthly sample ranging from 1959:M2 to 2014:M12. Table I reports quantiles of the prior distribution[7] as well as posterior median estimates of the model parameters. Estimates for the benchmark specification with monthly and annual measurement errors are reported in column (1). We briefly comment on some important aspects of the prior distribution. The prior for ρ (persistence of x_t) is uniform over the interval $(-1, 1)$ and encompasses values that imply near i.i.d. consumption growth as well as values for which x_t is almost a unit root process. The parameter φ_x can be interpreted as the square root of a "signal-to-noise ratio," meaning the ratio of the variance of x_t over the variance of the i.i.d. component $\sigma_c \eta_{c,t+1}$. We use a uniform prior for φ_x that allows for "signal-to-noise ratios" between 0 and 1. At an annualized rate, our a priori 90% credible interval for σ and σ_ϵ ranges from 0.3% to 2.1% and the prior

[7]In general, our priors attempt to restrict parameter values to economically plausible magnitudes. The judgment of what is *economically plausible* is, of course, informed by some empirical observations, in the same way the choice of the model specification is informed by empirical observations.

TABLE I

POSTERIOR MEDIAN ESTIMATES OF CONSUMPTION GROWTH PROCESSES[a]

					Posterior Estimates					
					State-Space Model					
		Prior Distribution			M&A	No ME AR(2)	M $\rho_\epsilon \neq 0$	M $\rho_\eta \neq 0$	IID	ARMA(1, 2)
	Distr.	5%	50%	95%	(1)	(2)	(3)	(4)	(5)	(6)
μ_c	N	−0.007	0.0016	0.0100	0.0016	0.0016	0.0016	0.0016	0.0016	0.0016
ρ	U	−0.90	0	0.90	0.918	−0.684	0.918	0.919	-	0.913
ρ_2	U	−0.90	0	0.90	-	−0.353	-	-	-	-
φ_x	U	0.05	0.5	0.95	0.681		0.704	0.644	-	-
	U	0.1	1.0	1.9	-	0.482	-	-	-	-
σ	IG	0.0008	0.0019	0.0061	0.0018	0.0027	0.0017	0.0019	0.0033	0.0032
σ_ϵ	IG	0.0008	0.0019	0.0061	0.0018		0.0019	0.0018	-	-
σ_ϵ^a	IG	0.0007	0.0029	0.0386	0.0011	-	-	-	-	-
ρ_ϵ	U	−0.90	0	0.90	-	-	0.060	-	-	-
ρ_η	U	−0.90	0	0.90	-	-	-	−0.046	-	-
ζ_1	N	−8.2	0	8.2	-	-	-	-	-	−1.14
ζ_2	N	−8.2	0	8.2	-	-	-	-	-	0.302
$\ln p(Y)$					2887.1	2870.3	2883.9	2885.8	2863.2	2884.0

[a]The estimation sample is from 1959:M2 to 2014:M12. We denote the persistence of the growth component x_t by ρ (and ρ_2 if follows an AR(2) process), the persistence of the measurement errors by ρ_ϵ, and the persistence of $\eta_{c,t}$ by ρ_η. We report posterior median estimates for the following measurement error specifications of the state-space model: (1) monthly and annual measurement errors (M&A); (2) no measurement errors with AR(2) process for x_t (no ME AR(2)); (3) serially correlated monthly measurement errors (M, $\rho_\epsilon \neq 0$); (4) serially correlated consumption shocks $\eta_{c,t}$ (M, $\rho_\eta \neq 0$, $\rho > \rho_\eta$). In addition, we report results for the following models: (5) consumption growth is i.i.d.; (6) consumption growth is ARMA(1, 2).

for the σ_ϵ^a covers the interval 0.07% to 3.9%. For comparison, the sample standard deviations of annualized monthly consumption growth and annual consumption growth are approximately 1.1% and 2%.

The estimate of ρ obtained from our benchmark specification is approximately 0.92, pointing toward a fairly persistent predictable component in consumption growth. The estimate of φ_x is 0.68 and implies that the variance of x_t is roughly 50% smaller than the variance of the i.i.d. component of consumption growth. At first glance, the large estimate of ρ in the benchmark model may appear inconsistent with the negative sample autocorrelation of monthly consumption growth reported in Figure 1. However, the sample moment confounds the persistence of the "true" consumption growth process and the dynamics of the measurement errors. Our state-space framework is able to disentangle the various components of the observed monthly consumption growth, thereby detecting a highly persistent predictable component x_t that is hidden under a layer of measurement errors.

To assess the robustness of this finding, we now modify the benchmark specification in several dimensions. If we shut down the measurement errors and generalize x_t to an AR(2) process (see column (2)), then the estimates of the autoregressive coefficients turn negative, thereby confirming the graphical pattern in Figure 2. Reverting back to an AR(1) process for x_t and allowing for serially correlated measurement errors (see column (3)) does not change the estimates of the benchmark model. In fact, the estimated

autocorrelation for the measurement error is close to zero. Likewise, if we allow for some serial correlation in the transitory component of "true" consumption growth (see column (4)), the estimate of ρ stays around 0.92. Finally, in the last two columns of Table I, we report estimates for an i.i.d. model of consumption growth and an ARMA(1, 2) model. We compute log marginal data densities for each specification. Differentials of $\ln p(Y)$ can be interpreted as log posterior odds (under the assumption that the prior odds are 1). The last row of Table I shows that the benchmark specification dominates all of the alternatives. In particular, the i.i.d. specification and the state-space model without measurement errors are strongly dominated with log odds of 23.9 and 16.8 in favor of the benchmark.[8]

In order to examine the degree to which measurement errors contribute to the variation in the observed consumption growth, we conduct variance decomposition of monthly and annual consumption growth using measurement error specification of column (1) in Table I. We find that more than half of the observed monthly consumption growth variation is due to measurement errors.[9] For annual consumption growth data, this fraction drops below 1%. On the other hand, the opposite pattern holds true for the persistent growth component. While the variation in the persistent growth component only accounts for 13% of the monthly consumption growth variation, this fraction increases to 87% for annual consumption growth data.

Informational Gain Through Temporal Disaggregation. The observation that monthly consumption growth data are strongly contaminated by measurement errors which to a large extent average out at quarterly or annual frequency, suggests that one might be able to estimate ρ equally well based on time-aggregated data. We examine this issue in Table II. The first row reproduces the ρ estimate from Specification (1) of Table I. However, we now also report the 5% and 95% quantile of the posterior distribution. Keeping the length of a period equal to a month in the state-space model, we change the measurement equation to link it with quarterly and annual consumption growth data. As the data frequency drops from monthly to annual, the posterior median estimate of ρ falls from 0.92 to 0.89. Moreover, the width of the equal-tail probability 90% credible interval increases from 0.12 to 0.36, highlighting that the use of high-frequency data sharpens inference about ρ.

Hansen, Heaton, and Li (2008) estimated a cointegration model for log consumption and log earnings to extract a persistent component in consumption. The length of a time period in their reduced-rank vector autoregression (VAR) is a quarter and the model is estimated based on quarterly data. The authors found that the ratio of long-run to short-run response of log consumption to a persistent growth shock, $\eta_{x,t}$ in our notation, is about 2, which would translate into an estimate of ρ of approximately 0.5 for a quarterly model. As a robustness check, we estimate two quarterly versions of the homoscedastic state-space model: without quarterly measurement errors and with quarterly measurement errors. The posterior median estimates of ρ are 0.649 and 0.676, respectively. These results are by and large consistent with the low value reported in Hansen,

[8] In a preliminary data analysis, we estimated a variety of ARMA(p, q) models using maximum likelihood estimation. We used the Schwarz Information Criterion (BIC) to estimate p and q, which led us to the ARMA(1, 2) specification. In the Supplemental Material, we report results for other variants of the benchmark state-space model. Among them, the specification in which the monthly measurement errors are not restricted to average out over the year is at par with the benchmark specification in terms of fit. For reasons explained above, we find the benchmark specification more appealing.

[9] Wilcox (1992) found that more than a quarter of the variation in the retail sales series from Detroit and Philadelphia is due to measurement error. For New York, this figure is around 50%, and for LA, around 67%. Our estimates are of a similar order of magnitude.

TABLE II

INFORMATIONAL GAIN THROUGH HIGH-FREQUENCY OBSERVATIONS[a]

Data	Posterior of ρ			
Frequency	5%	50%	95%	90% Intv. Width
		Without Stochastic Volatility		
Monthly	0.847	0.918	0.963	0.116
Quarterly	0.843	0.917	0.962	0.119
Annual	0.621	0.893	0.983	0.362
		With Stochastic Volatility		
Monthly	0.904	0.951	0.980	0.076
Quarterly	0.872	0.931	0.971	0.099

[a]The estimation sample ranges from 1959:M2 to 2014:M12. The model frequency is monthly. For monthly data, we use both monthly and annual measurement errors (specification (1) in Table I). For quarterly (annual) data, we use quarterly (annual) measurement errors only. The model specification is provided in (4).

Heaton, and Li (2008) as well as the estimate in Hansen (2007) under the "loose" prior. Using a crude cube-root transformation, the quarterly ρ estimates translate into 0.866 and 0.878 at monthly frequency and thereby somewhat lower than the estimates obtained by estimating a monthly model on quarterly data.

Accounting for Stochastic Volatility. We now re-estimate the benchmark model (4) allowing for stochastic volatility. Our prior interval for the persistence of the volatility processes ranges from 0.27 to 0.999. The prior for the standard deviation of the consumption volatility process implies that the volatility may fluctuate either relatively little, over the range of 0.7 to 1.2 times the average volatility, or substantially, over the range of 0.4 to 2.4 times the average volatility.

According to Table II, the width of the 90% credible interval for ρ shrinks from 0.116 to 0.076 for monthly data and from 0.119 to 0.099 for quarterly data.[10] At the same time, the posterior median of ρ increases from 0.918 to 0.951 for monthly data and from 0.917 to 0.931 for quarterly data. Without stochastic volatility, sharp movements in consumption growth must be accounted for by large temporary shocks reducing the estimate of ρ; however, the presence of stochastic volatility allows the model to account for these sharp movements by fluctuations in the conditional variance of the shocks enabling ρ to be large. We conclude that allowing for heteroscedasticity reduces the posterior uncertainty about ρ and raises the point estimate.

As a by-product, we also obtain an estimate for the persistence, ρ_{h_c}, of the stochastic volatility process in (4). The degree of serial correlation of the volatility also has important implications for asset pricing. Starting from a truncated normal distribution that implies a 90% prior credible set ranging from 0.27 to 0.99, based on monthly observations the posterior credible set ranges from 0.955 to 0.999, indicating that the data favor a highly persistent volatility process $h_{c,t}$. Once the observation frequency is reduced from monthly to quarterly, the sample contains less information about the high-frequency volatility pro-

[10]We found that the state-space model with stochastic volatility is poorly identified if the observation frequency is annual, which is why we do not report this case in Table I.

cess and there is less updating of the prior distribution. Now the 90% credible interval ranges from 0.73 to 0.97.[11]

Estimation Based on Mixed-Frequency Data. To measure the small persistent component in consumption growth, one would arguably want to use the longest span of data possible. Adopting a mixed-frequency approach, we now add annual consumption growth data from 1930 to 1959 to our estimation sample. It is well known from Romer (1986) and Romer (1989) that pre-war data on consumption are known to be measured with significantly greater error that exaggerates the size of cyclical fluctuations. To cope with the criticism, we allow for annual measurement errors during 1930–1948 but restrict them to be zero afterwards. This break in measurement errors is also motivated by Amir-Ahmadi, Matthes, and Wang (2016) who provided empirical evidence for larger measurement errors in the early sample before the end of World War II. Importantly, we always account for monthly measurement errors whenever we use monthly data.

Prior credible intervals and posterior estimates are presented in Table III. Note that the ρ estimate reported under the 1959:M2–2014:M12 posterior is the same as the estimate reported in Table II based on monthly data and the model with stochastic volatility. Extending the sample period reduces the posterior median estimate of ρ slightly, from 0.95 to 0.94. We attribute this change to the large fluctuations around the time of the Great Depression. The width of the credible interval stays approximately the same. Note that at this stage, we are adding 30 annual observations to a sample of 671 monthly observations (and we are losing 11 monthly observations from 1959). The standard deviation

TABLE III

POSTERIOR ESTIMATES: CONSUMPTION ONLY[a]

		Prior			Posterior 1959:M2–2014:M12			Posterior 1930–1959 1960:M1–2014:M12		
	Distr.	5%	50%	95%	5%	50%	95%	5%	50%	95%
					Consumption Growth Process					
μ_c	N	−0.007	0.0016	0.0100	0.0009	0.0016	0.0019	0.0010	0.0016	0.0018
ρ	U	−0.9	0	0.9	0.904	0.951	0.980	0.891	0.940	0.971
φ_x	U	0.05	0.50	0.95	0.357	0.509	0.778	0.369	0.535	0.759
σ	IG	0.0008	0.0019	0.0061	0.0017	0.0021	0.0025	0.0017	0.0022	0.0028
ρ_{h_c}	N^T	0.27	0.80	0.999	0.955	0.988	0.999	0.949	0.984	0.996
$\sigma^2_{h_c}$	IG	0.0011	0.0043	0.0283	0.0007	0.0014	0.0030	0.0022	0.0054	0.0242
					Consumption Measurement Error					
σ_ϵ	IG	0.0008	0.0019	0.0061	0.0010	0.0013	0.0016	0.0010	0.0013	0.0016
σ^a_ϵ	IG	0.0007	0.0029	0.0386	0.0010	0.0015	0.0020	0.0010	0.0198	0.0372

[a]We report estimates of model (4). We adopt the measurement error model of Section 2.1. N, N^T, G, IG, and U denote normal, truncated (outside of the interval $(-1, 1)$) normal, gamma, inverse gamma, and uniform distributions, respectively. We allow for annual consumption measurement errors ϵ^a_t during the periods from 1930 to 1948. We impose monthly measurement errors ϵ_t when we switch from annual to monthly consumption data from 1960:M1 to 2014:M12.

[11]We conducted a small Monte Carlo experiment in which we repeatedly simulated data from a consumption growth model with stochastic volatility and then estimated models without and with stochastic volatility. For both specifications, the estimate of ρ is downward biased, and for the misspecified version without stochastic volatility, the downward bias is slightly larger.

of the monthly measurement error σ_ϵ is estimated to be about half of σ and is robust to different estimation samples because it is solely identified from monthly consumption growth data. The standard deviation of the annual measurement error is larger than that of monthly measurement error by a factor of 4 (recall that to compare σ_ϵ and σ_ϵ^a, one needs to scale the latter by $\sqrt{12}$). This finding is consistent with Amir-Ahmadi, Matthes, and Wang (2016) who found significant presence of measurement errors in output growth during 1930 and 1948.

3. INFORMATION FROM OTHER CASH-FLOW SERIES

Because aggregate consumption is typically thought of as an endogenous variable that responds to fluctuations in aggregate income, we examine in Section 3.1 whether our evidence for a predictable component in consumption growth can be traced back to GDP and whether estimating a joint model for consumption and GDP has important effects on our inference. In Section 3.2, we include dividend growth data in the estimation of the cash-flow model to set the stage for the subsequent asset pricing analysis. Finally, we provide a brief summary of the cash-flow estimation results in Section 3.3. Posterior inference for the models considered in this section is implemented with a Metropolis-within-Gibbs sampler that is similar to the one described in Section 2.2.

3.1. *Real GDP*

We begin the analysis with a monthly model for GDP growth $g_{y,t+1}$ that is identical to the benchmark consumption growth model in (4). Because GDP is only available at quarterly frequency, the measurement equation is

$$g^o_{y,t+1} = \sum_{j=1}^{5} \left(\frac{3 - |j-3|}{3} \right) g_{y,t+2-j}, \quad t = 1, 4, 7, \ldots. \tag{5}$$

We estimate this model without measurement errors (this was the preferred specification based on marginal data density comparisons) using observations on per capita GDP growth from 1959:Q1 to 2014:Q4.[12] The estimation results are provided in Table IV. The posterior median is 0.874 and the equal-tail 90% credible interval ranges from 0.698 to 0.966. These estimates can be compared to those obtained from quarterly consumption growth reported in Table II where the posterior median estimate of ρ is 0.921 (with stochastic volatility) and the upper bound of the credible interval is 0.963. Thus, while the median of ρ for GDP is smaller than for consumption, the 95% quantiles are in fact very similar.

So far, we have considered univariate models of consumption and income growth. Next, we examine the joint dynamics of these two series. In most models, consumption and income are cointegrated. We impose this cointegration relationship in the empirical analysis below. Specifically, the consumption dynamics are given by (4), while the log income-consumption ratio $yc_t \equiv y_t - c_t$ takes the form

$$yc_{t+1} = \mu_{yc} + \phi_{yc} x_{t+1} + s_{t+1}, \qquad s_{t+1} = \rho_s s_t + \sqrt{1-\rho_s^2}\sigma_{s,t}\eta_{s,t+1}, \quad \eta_{s,t+1} \sim N(0,1). \tag{6}$$

[12] We take log differences of the real GDP per capita series provided by FRED.

TABLE IV

POSTERIOR ESTIMATES: GDP GROWTH ONLY[a]

		Prior			Posterior		
	Distr.	5%	50%	95%	5%	50%	95%
μ_y	N	−0.007	0.0016	0.0100	0.0011	0.0017	0.0022
ρ	U	0.05	0.50	0.95	0.698	0.874	0.966
φ_x	U	0.05	0.50	0.95	0.117	0.259	0.418
σ	IG	0.0008	0.0019	0.0061	0.0040	0.0045	0.0051
ρ_{h_y}	N^T	0.27	0.80	0.999	0.928	0.970	0.992
$\sigma^2_{h_y}$	IG	0.0013	0.0043	0.0283	0.0026	0.0086	0.0228

[a]We report estimates of model (4) for GDP growth. N, N^T, IG, and U denote normal, truncated (outside of the interval $(-1, 1)$) normal, inverse gamma, and uniform distributions, respectively. The estimation sample ranges from 1959:Q1 to 2014:Q4.

We assume that the log of stochastic volatility $\sigma_{s,t}$ follows an AR(1) process and adopt the measurement error model of Section 2.1 for consumption growth. For GDP, the measurement equation time-aggregates monthly growth rates $g_{y,t} = g_{c,t} + \Delta yc_t$ to average quarterly growth rates as in (5).

The estimated parameters for the cointegration model based on monthly consumption growth and quarterly GDP growth data are reported in Table V. The posterior median estimate of ρ is 0.948 and the equal-tail probability 90% credible interval ranges from 0.913 to 0.970. Here, the strong evidence in monthly consumption in favor of a predictable

TABLE V

POSTERIOR ESTIMATES: CONSUMPTION AND OUTPUT[a]

		Prior			Posterior		
	Distr.	5%	50%	95%	5%	50%	95%
				Consumption			
μ	N	−0.007	0.0016	0.0100	0.0012	0.0016	0.0020
ρ	U	−0.9	0	0.9	0.913	0.948	0.970
φ_x	U	0.05	0.50	0.95	0.419	0.593	0.796
σ	IG	0.0008	0.0019	0.0061	0.0015	0.0018	0.0022
σ_ϵ	IG	0.0008	0.0019	0.0061	0.0013	0.0015	0.0018
ρ_{h_c}	N^T	0.27	0.80	0.999	0.949	0.979	0.996
$\sigma^2_{h_c}$	IG	0.0013	0.0043	0.0283	0.0024	0.0091	0.0235
				Output			
ρ_s	U	−0.9	0	0.9	0.943	0.965	0.983
φ_s	U	5	50	95	6.65	8.82	13.25
ϕ_{yc}	U	−90	0	90	−1.83	−1.57	−1.18
ρ_{h_s}	N^T	0.27	0.80	0.999	0.943	0.982	0.996
$\sigma^2_{h_s}$	IG	0.0013	0.0043	0.0283	0.0015	0.0039	0.0187

[a]The estimation sample ranges from 1959 to 2014. We report estimates of model (6). N, N^T, IG, and U denote normal, truncated (outside of the interval $(-1, 1)$) normal, inverse gamma, and uniform distributions, respectively.

component x_t seems to dominate the estimation result. There is no information in GDP growth that contradicts this information. The log-GDP consumption ratio itself is fairly persistent with median estimate of ρ_s of 0.965. Thus, deviations from the steady-state ratio are relatively long-lived.

How does our evidence relate to common views on GDP dynamics? U.S. GDP growth is well described as an AR(1) model with an autocorrelation coefficient of about 0.3. In our cointegration model, the implied posterior predictive quantiles (0.05%, 0.50%, and 0.95%) for the autocorrelation of output growth at the quarterly frequency are 0.156, 0.273, and 0.389, which is consistent with this conventional wisdom. Thus, on balance, we view the dynamics of output and consumption to be consistent with our LRR specification with both series containing a small persistent component, and with models that imply a transmission from income to consumption.[13]

3.2. *Dividends*

As our subsequent asset pricing analysis focuses on the U.S. aggregate equity market, we now include dividend growth data in the estimation of the cash-flow model. We use monthly observations of dividends of the CRSP value-weighted portfolio of all stocks traded on the NYSE, AMEX, and NASDAQ. Dividend series are constructed on the per share basis as in Campbell and Shiller (1988b) and Hodrick (1992). Following Robert Shiller, we smooth out dividend series by aggregating three months' values of the raw nominal dividend series.[14] We then compute real dividend growth as log difference of the adjusted nominal dividend series and subtract CPI inflation. Further details are provided in the Supplemental Material.

Measurement Equation for Dividend Growth. Dividend data are available at monthly frequency for the estimation period from 1930 to 2014. There is a consensus in the finance literature that aggregate dividend series for a broad cross-section of stocks exhibit a strong seasonality. This seasonality is generated by payout patterns which are not uniform over the calendar year. Much of this seasonality, in particular its deterministic component, can be removed by averaging observed dividend growth over the span of a year. To do so, we utilize the same "tent" function as for consumption growth in (1) and define

$$g^{a,o}_{d,t+1} = \sum_{j=1}^{23} \left(\frac{12 - |j-12|}{12} \right) g^{o}_{d,t-j+2}, \qquad g^{a}_{d,t+1} = \sum_{j=1}^{23} \left(\frac{12 - |j-12|}{12} \right) g_{d,t-j+2}. \tag{7}$$

Our measurement equation then takes the form

$$g^{a,o}_{d,t+1} = g^{a}_{d,t+1} + \sigma^{a}_{d,\epsilon} \epsilon^{a}_{d,t+1}, \quad \epsilon^{a}_{d,t+1} \sim N(0,1). \tag{8}$$

For computational reasons that arise in the estimation of the asset pricing model in Section 4, we allow for some additional measurement errors, which we assume to be i.i.d.

[13] It is well known that, in production models, consumption is not a Martingale sequence and the predictable component in consumption growth can be generated by a predictable component in productivity growth; see Croce (2014). In a frictionless environment, both labor and capital income are determined by their respective marginal products, which in turn depend on the exogenous productivity process.

[14] We follow Shiller's approach despite the use of the annualization in (8) because we found that the annualization did not remove all the anomalies in the data.

across periods. We fix these measurement errors at 1% of the sample variance of dividend growth rates. Note that (8) does not imply $g^o_{d,t+1} = g_{d,t+1}$, even for $\sigma^a_{d,\epsilon} = 0$. For instance, there could be a deterministic seasonal pattern in the observed monthly dividend growth data $g^o_{d,t+1}$ that is not part of the model-implied process $g_{d,t+1}$. The tent-shaped transformation would remove the seasonal component from observed data such that we are effectively equating the non-seasonal component of the observed data to the model-implied data.

State-Transition Equation. We will model consumption and dividend growth as a joint process. The law of motion for consumption growth is identical to (4), except for the fact that now we will have separate volatility processes for the persistent and transitory components. Dividend streams have levered exposures to both x_t and $\eta_{c,t+1}$, which is captured by the parameters ϕ and π, respectively. We allow $\sigma_{d,t}\eta_{d,t+1}$ to capture idiosyncratic movements in dividend streams. Overall, the cash-flow dynamics follow:

$$\begin{aligned} g_{c,t+1} &= \mu_c + x_t + \sigma_{c,t}\eta_{c,t+1}, \\ x_{t+1} &= \rho x_t + \sqrt{1-\rho^2}\sigma_{x,t}\eta_{x,t+1}, \\ g_{d,t+1} &= \mu_d + \phi x_t + \pi\sigma_{c,t}\eta_{c,t+1} + \sigma_{d,t}\eta_{d,t+1}, \\ \sigma_{i,t} &= \varphi_i\sigma\exp(h_{i,t}), \qquad h_{i,t+1} = \rho_{h_i}h_{i,t} + \sigma_{h_i}w_{i,t+1}, \quad i = \{c, x, d\}, \end{aligned} \tag{9}$$

where the shocks are assumed to be $\eta_{i,t+1}, w_{i,t+1} \sim N(0,1)$, $i = \{c, x, d\}$ and we impose the normalization $\varphi_c = 1$. For now, we will also restrict $h_{x,t} = h_{c,t}$ and only report estimates for $\rho_{h,c}$ and $\sigma^2_{h,c}$.

Estimation Results. Table VI provides percentiles of the prior distribution and the posterior distribution for the post-1959 estimation sample and for the mixed frequency sample starting in 1930. The priors for ϕ and π, parameters that determine the comovement of dividend and consumption growth, are uniform distributions on the interval $[-10, 10]$. The parameter φ_d determines the standard deviation of the i.i.d. component of dividend growth relative to consumption growth. Here we use a prior that is uniform on the interval [0, 10], thereby allowing for dividends to be much more volatile than consumption. The prior for the standard deviation of the dividend volatility process implies that the volatility may fluctuate either relatively little, over the range of 0.5 to 2.1 times the average volatility, or substantially, over the range of 0.1 to 13 times the average volatility. Finally, we fix the measurement error variance $(\sigma^a_{d,\epsilon})^2$ at 1% of the sample variance of dividend growth.

The most important finding is that the posterior median ρ increases as we add dividend growth data in the estimation. In addition, we find significant reduction in our uncertainty about ρ captured by the distance between 95% and 5% posterior quantiles. The posterior median of ρ is around 0.97 for the post-1959 sample and is 0.95 for the longer sample, both of which are higher than those in Table III. The 5%–95% distance dropped from 0.076 to 0.054 as we include dividend growth in the estimation (compare with Table III). The posterior median of the standard deviation of the unconditional volatility of the persistent component φ_x is around 0.43, slightly lower than before.

The dividend leverage ratio on expected consumption growth ϕ is estimated to be around 2.8–3.5 and the standard deviation of the idiosyncratic dividend shocks φ_d is around 4.8. The estimation results also provide strong evidence for stochastic volatility. According to the posteriors reported in Table VI, both $\sigma_{c,t}$ and $\sigma_{d,t}$ exhibit significant time variation. The posterior medians of ρ_{h_c} and ρ_{h_d} range from 0.98 to 0.99.

Cointegration of Dividends and Consumption. In our analysis, aggregate consumption is measured per capita and dividends are computed per share. Thus, there is no theo-

TABLE VI

POSTERIOR ESTIMATES: CONSUMPTION AND DIVIDEND GROWTH[a]

		Prior			Posterior 1959:M2–2014:M12			Posterior 1930–1959 1960:M1–2014:M12		
	Distr.	5%	50%	95%	5%	50%	95%	5%	50%	95%
					Consumption Growth Process					
ρ	U	−0.90	0	0.90	0.937	0.967	0.991	0.923	0.952	0.978
φ_x	U	0.05	0.50	0.95	0.285	0.430	0.834	0.291	0.430	0.684
σ	IG	0.0008	0.0019	0.0061	0.0019	0.0022	0.0025	0.0021	0.0029	0.0036
ρ_{h_c}	N^T	0.27	0.80	0.999	0.952	0.985	0.997	0.976	0.992	0.998
$\sigma^2_{h_c}$	IG	0.0013	0.0043	0.0283	0.0015	0.0053	0.0185	0.0013	0.0034	0.0132
					Dividend Growth Process					
ϕ	U	−9.0	0.0	9.0	2.13	2.85	3.55	3.31	3.52	3.64
π	U	−9.0	0.0	9.0	0.136	0.358	0.751	0.642	0.819	0.932
φ_d	U	0.50	5.0	9.5	3.51	4.69	6.16	3.31	4.82	7.66
ρ_{h_d}	N^T	0.27	0.80	0.999	0.939	0.977	0.994	0.951	0.977	0.992
$\sigma^2_{h_d}$	IG	0.015	0.0445	0.208	0.0166	0.0418	0.1076	0.0146	0.0357	0.0835
					Consumption Measurement Error					
σ_ϵ	IG	0.0008	0.0019	0.0062	0.0009	0.0011	0.0014	0.0009	0.0012	0.0015
σ^a_ϵ	IG	0.0007	0.0029	0.0389	-	-	-	0.0006	0.0067	0.0134

[a]We utilize the mixed-frequency approach in the estimation: For consumption we use annual data from 1930 to 1959 and monthly data from 1960:M1 to 2014:M12; we use monthly dividend annual growth data from 1930:M1 to 2014:M12. For consumption we adopt the measurement error model of Section 2.1. We allow for annual consumption measurement errors ϵ^a_t during the periods from 1930 to 1948. We impose monthly measurement errors ϵ_t when we switch from annual to monthly consumption data from 1960:M1 to 2014:M12. We fix $\mu_c = 0.0016$ and $\mu_d = 0.0010$ at their sample averages. Moreover, we fix the measurement error variance $(\sigma^a_{d,\epsilon})^2$ at 1% of the sample variance of dividend growth. N, N^T, G, IG, and U denote normal, truncated (outside of the interval $(-1, 1)$) normal, gamma, inverse gamma, and uniform distributions, respectively.

retical reason for the two series to be cointegrated. Nonetheless, we examined the presence of a cointegration relationship between the observed series c^o_t and d^o_t. First, we conducted two frequentist cointegration tests based on post-1959 monthly data. The first test is an augmented Dickey–Fuller test that imposes the cointegration vector $[1, -1]$ and the second test is an Engle–Granger test based on an estimated cointegration vector of $[1, -0.55]$. None of these tests can reject the null hypothesis of no cointegration. Second, we estimate a modified state-space model with a hardwired cointegration restriction. This model retains the consumption growth dynamics of (9), but the law of motion of dividends is modified as follows:

$$dc_{t+1} = \mu_{dc} + \phi_{dc} x_{t+1} + s_{t+1}, \qquad s_{t+1} = \rho_s s_t + \sqrt{1-\rho_s^2}\sigma_{s,t}\eta_{s,t+1}, \quad \eta_{s,t+1} \sim N(0,1), \tag{10}$$

where $dc_{t+1} \equiv d_{t+1} - c_{t+1}$ is the error correction representation for dividends and consumption. dc_{t+1} loads on x_{t+1}, and a stationary AR(1) process s_{t+1}, which has its own stochastic volatility process $\sigma_{s,t}$. Under this structure, it can be easily verified that dividend growth can be written as $g_{d,t+1} = \Delta dc_{t+1} + g_{c,t+1}$. The measurement equation for dividends then follows equation (8). A full set of estimates of the cointegration specification is reported in the Supplemental Material. The estimate of ρ, as well as the other

estimates of the consumption parameters, are essentially unaffected by the cointegration specification. The marginal likelihood for the cointegration specification is 6041.1, whereas the marginal likelihood for the original specification is 6101.2. Thus, there is no evidence in the data in favor of the cointegration restriction.

3.3. *Summary of Cash-Flow-Only Analysis*

Aggregate consumption is a key macroeconomic variable, and it is therefore important for macroeconomists to understand its dynamic properties. There are several important implications that are robust across our analyses of consumption, consumption and output, and consumption and dividends. At the monthly frequency, consumption growth has a very strong MA(2) component. Ignoring this MA(2) component distorts inference. There is clear evidence against the hypothesis that consumption is a random walk at monthly frequency. Our interpretation of the MA(2) component is that it is generated by MA(1) measurement errors and a highly persistent "local level" component.[15] Empirically, our measurement error specification is preferred to the ARMA(1, 2) specification. Thus, if the goal is to create a reduced-form model of consumption, it is important to capture the MA component. If the goal is to confront a macro model with monthly consumption data, it is important to apply a "filter" that removes the high-frequency movements in consumption growth that we attribute to measurement error, because a typical macro model is not equipped to capture these dynamics. Overall, the posterior interval for the parameter estimates essentially encompass those used in the LRR literature (e.g., Bansal, Kiku, and Yaron (2012)). Importantly, the various estimation results (univariate consumption, consumption and GDP, and consumption and dividends) provide supportive evidence for a small persistent component in both consumption growth rate and its stochastic volatility. Consistent with previous LRR work, this evidence is distinctly different from a commonly held view in which consumption growth is an i.i.d. process.

4. THE LONG-RUN RISKS MODEL

We now embed the cash-flow process (9) into an endowment economy, which allows us to price financial assets. The preferences of the representative household are described in Section 4.1. Section 4.2 describes the model solution. Section 4.3 presents the state-space representation of the asset pricing model and its Bayesian estimation.

4.1. *Representative Agent's Optimization*

We consider an endowment economy with a representative agent that has Epstein and Zin (1989) recursive preferences and maximizes her lifetime utility,

$$V_t = \max_{C_t} \left[(1-\delta)\lambda_t C_t^{\frac{1-\gamma}{\theta}} + \delta \left(\mathbb{E}_t \left[V_{t+1}^{1-\gamma} \right] \right)^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}},$$

subject to budget constraint

$$W_{t+1} = (W_t - C_t) R_{c,t+1},$$

[15] This interpretation is consistent with studies that examine the quality of consumption data (e.g., Wilcox (1992)), but from a pure time series perspective, we cannot rule out that the MA(2) component is partly due to transient dynamics in "true" consumption growth.

where W_t is the wealth of the agent, $R_{c,t+1}$ is the return on all invested wealth, γ is risk aversion, $\theta = \frac{1-\gamma}{1-1/\psi}$, and ψ is intertemporal elasticity of substitution. As highlighted in Albuquerque et al. (2016), we also allow for a preference shock, λ_t, to the time rate of preference. The endowment stream is given by the law of motion for consumption and dividend growth in (9), and the growth rate of the preference shock, denoted by $x_{\lambda,t}$, follows an AR(1) process with shocks that are independent of the shocks to cash flows:

$$x_{\lambda,t+1} = \rho_\lambda x_{\lambda,t} + \sigma_\lambda \eta_{\lambda,t+1}, \quad \eta_{\lambda,t+1} \sim N(0,1). \tag{11}$$

The Euler equation for any asset $r_{i,t+1}$ takes the form

$$\mathbb{E}_t\big[\exp(m_{t+1} + r_{i,t+1})\big] = 1, \tag{12}$$

where $m_{t+1} = \theta \log \delta + \theta x_{\lambda,t+1} - \frac{\theta}{\psi} g_{c,t+1} + (\theta - 1) r_{c,t+1}$ is the log of the real stochastic discount factor (SDF), and $r_{c,t+1}$ is the log return on the consumption claim. We reserve $r_{m,t+1}$ for the log market return—the return on a claim to the market dividend cash flows.[16]

4.2. *Solution*

Our goal is to devise a solution method that strikes a balance between accuracy and computational time. The solution described subsequently meets this requirement: it can be computed quickly because it relies on analytical approximations; it leads to a conditionally linear state-space representation for which the likelihood function can be efficiently evaluated with a particle filter (see below); and it is accurate for the empirically relevant parameter values.

Conditional on the cash-flow dynamics in (9) and the Euler equation (12), we have to derive the asset prices for the model economy. In order to fit the cash-flow specification to consumption and dividend growth data, we assumed that the volatilities follow log Gaussian processes: $\sigma_{i,t} = \varphi_i \sigma \exp(h_{i,t})$, where $h_{i,t}$ is a linear autoregressive process with normally-distributed innovations. This specification has been empirically successful in capturing conditional heteroscedasticity in a broad set of financial and macroeconomic time series.

The advantage of the exponential transformation is that it ensures that volatilities are nonnegative. The disadvantage is that, under this specification, the expected value of the level of consumption and dividends is infinite, which creates problems with the existence of continuation values in the endowment economy. This issue has been pointed out, for instance, in Chernov, Gallant, Ghysels, and Tauchen (2003) and Andreasen (2010), who proposed to splice the exponential transformation of $h_{i,t}$ together with a non-exponential function, for example, a square-root function, for $h_{i,t}$ exceeding some large threshold $\bar{h}_i$. To obtain a solution for the asset prices, we proceed slightly differently, by taking a

[16]Formally, markets are complete in the sense that all income and assets are tradable and add up to total wealth for which the return is $R_{c,t}$. In particular, let $R_{j,t+1} = (d_{j,t+1} + p_{j,t+1})/p_{j,t}$ be the return to a claim that pays the dividend stream $\{d_{j,\tau}\}_{\tau=t}^{\infty}$ and has the price $p_{j,t}$. Let $q_{j,t}$ be the number of shares. Then $W_t - C_t = \sum_j p_{j,t} q_{j,t}$. Wealth next period, W_{t+1}, equals $\sum_j p_{j,t} q_{j,t} R_{j,t+1}$, and it follows that $R_{c,t+1} = \frac{\sum_j p_{j,t} q_{j,t} R_{j,t+1}}{\sum_j p_{j,t} q_{j,t}}$. As in Lucas (1978), we normalize all shares $q_{j,t}$ to 1 and the risk-free asset to be in zero net supply such that in equilibrium $C_t = D_m + D_o$, where D_m are the dividends to all tradable financial assets and D_o are dividends on all other assets (e.g., labor, housing, etc.). R_m, the return we utilize in our empirical work, is the return on the claims that pay dividends D_m.

linear approximation of the exponential transformation $\sigma_{i,t}^2 = (\varphi_i \sigma)^2 \exp(2h_{i,t})$ around the steady state $h_{i,*} = 0$ and replacing the innovation variances in (9) with a process that follows Gaussian dynamics:

$$\sigma_{i,t+1}^2 = (\varphi_i \sigma)^2 (1 - \rho_{h_i}) + \rho_{h_i} \sigma_{i,t}^2 + \sigma_{w_i} w_{i,t+1}, \quad i = \{c, x, d\}. \tag{13}$$

Just as the above-mentioned approaches of modifying the exponential transformation, the linear approximation effectively generates thinner tails for the variance processes and facilitates valuations to be finite.[17] After this linearization, the volatility process is identical to the one used in Bansal and Yaron (2004) and the subsequent work that builds on their paper.

Asset prices can now be derived by using the approximate analytical solution described in Bansal, Kiku, and Yaron (2012) which utilizes the Campbell and Shiller (1988a) log-linear approximation for returns. This solution serves our purpose well, because it can be computed very quickly, which facilitates the Bayesian estimation below. The log price-consumption ratio takes the form

$$pc_t = A_0 + A_1 x_t + A_{1,\lambda} x_{\lambda,t} + A_{2,c} \sigma_{c,t}^2 + A_{2,x} \sigma_{x,t}^2. \tag{14}$$

As discussed in Bansal and Yaron (2004), $A_1 = \frac{1 - \frac{1}{\psi}}{1 - \kappa_1 \rho}$, the elasticity of prices with respect to growth prospects, will be positive whenever the IES, ψ, is greater than 1. $A_{1,\lambda} = \frac{\rho_\lambda}{1 - \kappa_1 \rho_\lambda}$, the elasticity of prices with respect to the growth rate of the preference shock, is always positive. Further, the elasticity of pc_t with respect to the two volatility processes $\sigma_{c,t}^2$ and $\sigma_{x,t}^2$ is $\frac{\theta}{2} \frac{(1 - \frac{1}{\psi})^2}{1 - \kappa_1 \nu_c}$ and $\frac{\theta}{2} \frac{(\kappa_1 A_1)^2}{1 - \kappa_1 \nu_x}$, respectively; both will be negative—namely, prices will decline with uncertainty—whenever θ is negative. A condition that guarantees a negative θ is that $\gamma > 1$, and $\psi > 1$—a configuration relevant for our parameter estimates and one in which agents exhibit preference for early resolution of uncertainty. The innovations to the log stochastic discount factor (SDF) are linear in the shocks to consumption growth η_c, η_x, the preference shocks η_λ, and the shocks to volatilities w_c and w_x. Denoting λ's as their respective market prices of risk, it is instructive to note that $\lambda_c = \gamma$, $\lambda_x = \frac{(\gamma - \frac{1}{\psi})\kappa_1}{1 - \kappa_1 \rho}$, $\lambda_\lambda = -\frac{\theta - \kappa_1 \rho_\lambda}{1 - \kappa_1 \rho_\lambda}$ (and λ_{w_c} and λ_{w_x}) are positive (negative) whenever $\gamma > 1$ and $\psi > 1$. Furthermore, when preferences are time separable, namely, when $\theta = 1$, λ_x, λ_{w_x}, and λ_{w_c} are all zero.

Risk premia are determined by the negative covariation between the innovations to returns and the innovations to the SDF. It can be shown that the risk premium for the market return, $r_{m,t+1}$, is

$$\begin{aligned} &\mathbb{E}_t(r_{m,t+1} - r_{f,t}) + \frac{1}{2} \operatorname{var}_t(r_{m,t+1}) \\ &= -\operatorname{cov}_t(m_{t+1}, r_{m,t+1}) \\ &= \underbrace{\beta_{m,c} \lambda_c \sigma_{c,t}^2}_{\text{short-run risk}} + \underbrace{\beta_{m,x} \lambda_x \sigma_{x,t}^2}_{\text{long-run growth risk}} + \underbrace{\beta_{m,\lambda} \lambda_\lambda \sigma_\lambda^2}_{\text{preference risk}} + \underbrace{\beta_{m,w_x} \lambda_{w_x} \sigma_{w_x}^2 + \beta_{m,w_c} \lambda_{w_c} \sigma_{w_c}^2}_{\text{volatility risks}}, \end{aligned} \tag{15}$$

where the β's reflect the exposures of the market return to the underlying consumption risks. Equation (15) highlights that the conditional equity premium can be attributed to (i)

[17] A quantitative comparison among the various approaches of thinning the tails of the $\sigma_{i,t}^2$ processes is beyond the scope of this paper.

short-run consumption growth, (ii) long-run growth, (iii) preference shock, (iv) short-run and long-run volatility risks.

A key variable for identifying the model parameters is the risk-free rate. Under the assumed dynamics in (9), the risk-free rate is affine in the state variables and follows

$$r_{f,t} = B_0 + B_1 x_t + B_{1,\lambda} x_{\lambda,t} + B_{2,c}\sigma_{c,t}^2 + B_{2,x}\sigma_{x,t}^2. \tag{16}$$

It can be shown that $B_1 = \frac{1}{\psi} > 0$ and the risk-free rate rises with good economic prospects, while $B_{1,\lambda} = -\rho_\lambda < 0$ and the risk-free rate falls with positive preference shock. Under $\psi > 1$ and $\gamma > 1$, $B_{2,c}$ and $B_{2,x}$ are negative so the risk-free rate declines with a rise in economic uncertainty. Further details of the solution are provided in the Supplemental Material (Schorfheide, Song, and Yaron (2018)).

The accuracy of the log-linearization depends on the parameterization of the LRR model. Taking the linear volatility process in (13) as given, Pohl, Schmedders, and Wilms (2016) compared the quantitative implications of our model solution to that of a nonlinear solution obtained by a projection method. They found that discrepancies for key asset pricing moments between the solutions are small (less than 6%) conditional on a parameterization that is similar to the posterior medians reported in Table VII. However, the discrepancies become larger if the persistence parameters ρ and ρ_h are increased and pushed toward the upper bound of credible sets derived from marginal posterior distributions. Thus, strictly speaking, the parameter estimates that we are reporting below should be interpreted as parameter estimates for the approximation.[18]

4.3. *State-Space Representation and Bayesian Inference*

While the state-space models for the cash-flow dynamics analyzed in Sections 2 and 3 can be analyzed with a fairly straightforward Metropolis-within-Gibbs sampler, posterior computations for the model with asset returns are considerably more complicated because the stochastic volatility process $h_t = [h_{c,t}, h_{x,t}, h_{d,t}]'$ affects the conditional mean of the asset prices.

The measurement equation can be expressed as

$$y_{t+1} = A_{t+1}\big(D + Zs_{t+1} + Z^v s_{t+1}^v(h_{t+1}, h_t) + \Sigma^u u_{t+1}\big), \quad u_{t+1} \sim N(0, I). \tag{17}$$

The vector of observables y_t comprises consumption and dividend growth, the observed market return $r_{m,t}^o$, and the risk-free rate $r_{f,t}^o$. u_{t+1} is a vector of measurement errors and A_{t+1} is a selection matrix that accounts for deterministic changes in the data availability. $s_{t+1}^v(\cdot)$ is a vector of conditional variances that depend on the log volatilities of the cash flows, h_{t+1} and h_t. The remaining "linear" state variables are collected in the vector s_{t+1}, which essentially consists of the persistent cash-flow component x_t (see (9)) and the preference shock $x_{\lambda,t}$. However, in order to express the observables y_{t+1} as a linear function of s_{t+1}, to capture the elaborate measurement error model of consumption, and to account

[18] A similar issue arises in the literature on dynamic stochastic general equilibrium (DSGE) models: the vast majority of DSGE models are estimated based on log-linear approximations, which facilitate a speedy evaluation of the likelihood function with the Kalman filter. The caveat that, strictly speaking, the resulting parameter estimates are estimates for the log-linear approximation has been widely accepted in this literature. An immediate consequence is that one should not plug parameter estimates obtained from a log-linear approximation into the nonlinear version of the model, because it will lead to a mismatch between model-implied moments and the moments in the data; see An and Schorfheide (2007) for more details.

TABLE VII

PRIOR AND POSTERIOR ESTIMATES[a]

		Prior			Posterior		
	Distr.	5%	50%	95%	5%	50%	95%
			Household Preferences				
ψ	G	0.30	1.30	3.45	1.25	1.97	3.22
γ	G	2.75	7.34	15.46	5.44	8.89	14.44
			Preference Risk				
ρ_λ	U	−0.90	0	0.90	0.933	0.959	0.974
σ_λ^2	IG	0.0003	0.0005	0.0015	0.0003	0.0004	0.0005
			Consumption Growth Process				
ρ	U	−0.90	0	0.90	0.949	0.987	0.997
φ_x	U	0.05	0.50	0.95	0.120	0.215	0.382
σ	IG	0.0008	0.0019	0.0061	0.0027	0.0035	0.0042
ρ_{h_c}	N^T	0.27	0.80	0.999	0.977	0.991	0.998
$\sigma_{h_c}^2$	IG	0.0011	0.0043	0.0283	0.0075	0.0096	0.0109
ρ_{h_x}	N^T	0.27	0.80	0.999	0.982	0.992	0.998
$\sigma_{h_x}^2$	IG	0.0011	0.0043	0.0283	0.0022	0.0039	0.0044
			Dividend Growth Process				
ϕ	N	−9.0	0.0	9.0	2.14	3.65	6.43
π	N	−9.0	0.0	9.0	0.75	1.47	2.37
φ_d	U	0.50	5.0	9.5	3.19	4.54	6.55
ρ_{h_d}	N^T	0.28	0.80	0.999	0.943	0.969	0.974
$\sigma_{h_d}^2$	IG	0.015	0.0445	0.208	0.0404	0.0447	0.0565
			Consumption Measurement Error				
σ_ϵ	IG	0.0008	0.0019	0.0062	0.0009	0.0014	0.0020
σ_ϵ^a	IG	0.0007	0.0029	0.0389	0.0038	0.0141	0.0213

[a]The estimation results are based on annual consumption growth data from 1930 to 1960 and monthly consumption growth data from 1960:M1 to 2014:M12. We allow for annual consumption measurement errors ϵ_t^a during the periods from 1930 to 1948. We impose monthly measurement errors ϵ_t when we switch from annual to monthly consumption data from 1960:M1 to 2014:M12. For the other three series, we use monthly data from 1930:M1 to 2014:M12. We fix $\delta = 0.999$. We fix $\mu_c = 0.0016$ and $\mu_d = 0.0010$ at their sample averages. We also fix the measurement error variances $(\sigma_{d,\epsilon}^a)^2$ and $(\sigma_{f,\epsilon})^2$ at 1% of the sample variance of dividend growth and the risk-free rate, respectively. B, N, N^T, G, and IG are beta, normal, truncated (outside of the interval $(-1, 1)$) normal, gamma, and inverse gamma distributions, respectively.

for potentially missing observations, it is necessary to augment s_{t+1} by lags of x_t and $x_{\lambda,t}$ as well as the innovations for the cash-flow process and measurement errors. This leads to a high-dimensional state vector s_t (see the Supplemental Material for a precise definition).

The solution of the LRR model sketched in Section 4.2 provides the link between the state variables and the observables y_{t+1}. The state variables themselves follow vector autoregressive processes of the form

$$s_{t+1} = \Phi s_t + v_{t+1}(h_t), \qquad h_{t+1} = \Psi h_t + \Sigma_h w_{t+1}, \quad w_{t+1} \sim N(0, I), \tag{18}$$

where $v_{t+1}(\cdot)$ is an innovation process with a variance that is a function of the log volatility process h_t and w_{t+1} is the innovation of the stochastic volatility process.

The key difference between the nonlinear state-space model given by (17) and (18) and the state-space models estimated in Sections 2 and 3 is that the volatility states (h_{t+1}, h_t) enter the conditional mean of y_{t+1} through the model-implied asset returns. This means that the Metropolis-within-Gibbs sampler that we used previously is not valid for the model with asset prices. Instead, we will use a particle filter to approximate the likelihood function of the state-space model and then embed the likelihood approximation into a Metropolis–Hastings algorithm.

Our particle filter exploits the particular structure of the state-space model. Conditional on the volatility states (h_{t+1}, h_t), the model is linear. Building on ideas in Chen and Liu (2000), we use Kalman-filtering steps to track the Gaussian distribution of $s_t|(h_t^j, Y_{1:t})$, where $\{h_t^j, W_t^j\}_{j=1}^M$ is a set of particle values and weights for the volatility states. Because conditional on the three-dimensional volatility vector h_t^j one can integrate over the high-dimensional vector s_t analytically (Rao–Blackwellization), the particle filter approximation $\hat{p}(Y|\Theta)$ of the likelihood function tends to be sufficiently accurate so that it can be embedded into a random-walk Metropolis–Hastings algorithm. Here Θ comprises the parameters of the cash-flow process, the volatility parameters, and the preference parameters of the representative household. The resulting sampler belongs to the class of particle MCMC samplers. Andrieu, Doucet, and Holenstein (2010) have shown that the use of $\hat{p}(Y|\Theta)$ in MCMC algorithms can still deliver draws from the actual posterior $p(\Theta|Y)$ because these approximation errors essentially average out as the Markov chain progresses. Further details of the posterior sampler are provided in the Supplemental Material.

5. EMPIRICAL RESULTS BASED ON THE LONG-RUN RISKS MODEL

We now turn to the empirical analysis based on the LRR model. Section 5.1 describes the asset price data that are used in addition to the cash-flow data. We discuss the estimation results in Section 5.2 and present the asset pricing implications of the estimated model in Section 5.3.

5.1. *Data*

In addition to the consumption and dividend data used in Sections 2 and 3, we now also use financial market data from 1930:M1 to 2014:M12. This includes monthly observations of returns and prices of the CRSP value-weighted portfolio of all stocks traded on the NYSE, AMEX, and NASDAQ. Prices are also constructed on the per share basis as in Campbell and Shiller (1988b) and Hodrick (1992). The stock market data are converted to real using the consumer price index (CPI) from the Bureau of Labor Statistics. Finally, the ex ante real risk-free rate is constructed as a fitted value from a projection of the ex post real rate on the current nominal yield and inflation over the previous year. To run the predictive regression, we use monthly observations on the three-month nominal yield from the CRSP Fama Risk Free Rate tapes and CPI series. Data sources and summary statistics are available in the Supplemental Material.

5.2. *Model Estimation*

Parameter Estimates. The prior distribution for the parameters associated with the exogenous cash-flow process is the same as the ones used in Section 3.2. Thus, we focus

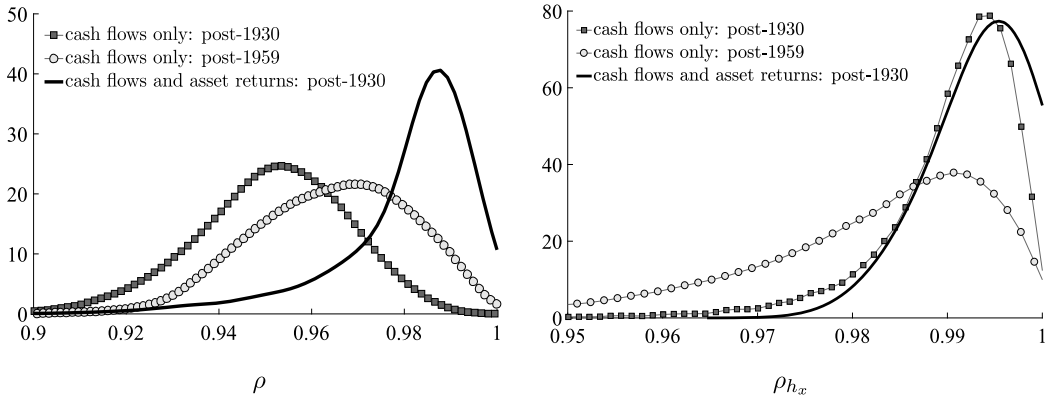

FIGURE 3.—Posterior distribution of ρ and ρ_{h_x}. *Notes*: We plot posterior densities of ρ from the estimation with cash-flow data only from post-1930 (squared-line) and from post-1959 samples (circled-line), respectively, and from the estimation with cash-flow and asset return data from post-1930 sample (solid-line).

on the preference parameters that affect the asset pricing implications of the model. Percentiles for the prior are reported in the left-side columns of Table VII. The prior for the discount rate δ reflects beliefs about the magnitude of the risk-free rate. For the asset pricing implications of our model, it is important whether the IES is below or above 1. Thus, we choose a prior that covers the range from 0.3 to 3.5. The 90% prior credible interval for the risk-aversion parameter γ ranges from 3 to 15, encompassing the values that are regarded reasonable in the asset pricing literature. The prior for the persistence and the innovation standard deviation of the preference shock is identical to the prior for the cash-flow parameters ρ and σ. Finally, we fix the variance $\sigma^2_{f,\epsilon}$ of the measurement error of the risk-free rate at 1% of the risk-free rate's sample variance.

The remaining columns of Table VII summarize the percentiles of the posterior distribution for the model parameters. While the estimated cash-flow parameters are, by and large, similar to those reported in Table VI when asset prices are not utilized, a few noteworthy differences emerge. First, the estimate of ρ, the persistence of the predictable cash-flow component, increases from 0.952 to 0.987 to better capture the equity premium and persistence of the price-dividend ratio. The left panel of Figure 3 overlays the posterior densities of ρ obtained with (post-1930 sample) and without asset prices (post-1930 and post-1959 samples, respectively).[19] Interestingly, the figure shows that although the mode of the posterior increases and shifts to the right when asset prices are used in estimation, the 90% credible interval ranging from 0.949 to 0.997 contains the posterior medians of ρ from the cash-flow-only estimations.[20]

Second, the right panel of Figure 3 shows the posterior distribution of ρ_{h_x}, the persistence of the stochastic volatility process for x_t. The modes of the three posteriors are

[19]Results from the post-1959 sample with asset prices are virtually identical to the results from the post-1930 sample. For this reason, they are not plotted separately in Figure 3.

[20]In the Supplemental Material, we present additional misspecification tests for the consumption dynamics. To assess the extent to which the increase in ρ leads to a decrease in fit of the consumption growth process, we re-estimate model (4) conditional on various choices of ρ between 0.90 and 0.99 and recompute the marginal data density for consumption growth. The key finding is that the drop in the marginal data density by changing ρ from $\hat{\rho}$ to 0.99 is small, indicating that there essentially is no tension between the parameter estimates obtained with and without asset prices.

quite similar, with the cash-flow-only posteriors having a longer left tail. Again, the posterior becomes more concentrated as asset returns are added to the estimation. Third, in the cash-flow-only estimation, we imposed a common stochastic log volatility process for the transitory and persistent component of consumption growth, that is, $h_{x,t} = h_{c,t}$, which led to an estimate $\hat{\sigma}^2_{h_c} = 0.0034$. Once we add the returns to the set of observables and remove the restriction, we obtain $\hat{\sigma}^2_{h_x} = 0.0039$ and $\hat{\sigma}^2_{h_c} = 0.0096$, reflecting asset price information about the volatility of volatilities. Fourth, the estimate of φ_x drops from 0.430 to 0.215, which reduces the model-implied predictability of consumption growth by the price-dividend ratio and brings it more in line with the data. Finally, the estimate of σ increases somewhat from 0.0029 to 0.0035 to explain the highly volatile asset price data.

Overall, the information from the market returns and risk-free rate reduces the posterior uncertainty about the cash-flow parameters and strengthens the evidence in favor of a time-varying conditional mean of cash-flow growth rates as well as time variation in the volatility components. Table VII also provides the estimated preference parameters. Importantly, the IES is estimated above 1 with a credible interval ranging from 1.3 to 3.2, while the posterior median estimate of the risk-aversion parameter γ is 8.9 and its interval estimate is 5.4 to 14.4.

Smoothed Mean and Volatility States. Figure 4 depicts smoothed estimates of the predictable growth component x_t. Because the estimate of x_t is, to a large extent, deter-

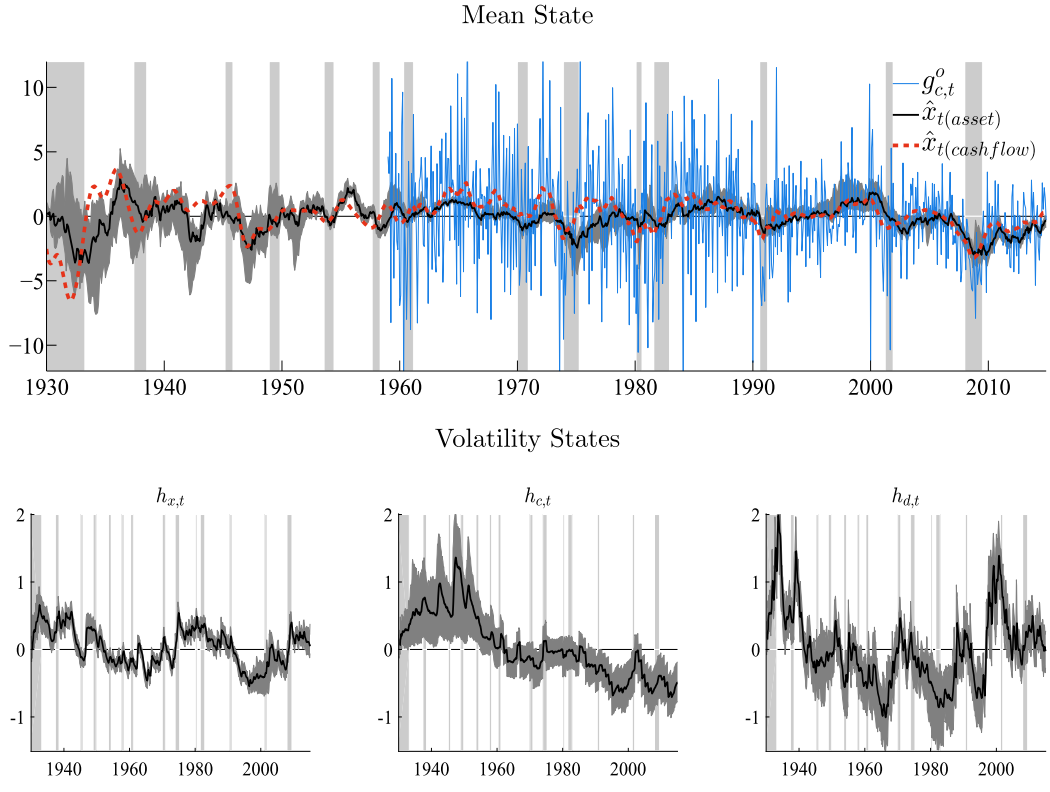

FIGURE 4.—Smoothed states. *Notes*: Black lines represent posterior medians of smoothed states and gray-shaded areas correspond to 90% credible intervals. Shaded bars indicate NBER recession dates. In the top panel, we overlay the smoothed state x_t obtained from the estimation without asset prices (red dashed line) and monthly consumption growth data (blue solid line).

mined by the time path of consumption, the 90% credible bands (reflecting uncertainty about parameters and the latent states) are much wider prior to 1960, when only annual consumption growth data were used in the estimation. Post 1959, x_t tends to fall in recessions (indicated by the shaded bars in Figure 4), but periods of falling x_t also occur during expansions. We overlay the smoothed estimate of x_t obtained from the estimation without asset price data. It is very important to note that the two estimates are similar, which highlights that x_t is, in fact, detectable based on cash-flow data only. We also depict the monthly consumption growth data post 1959, which confirms that x_t indeed captures low-frequency movements in consumption growth.

The smoothed volatility processes are plotted in Figure 4. Recall that our model has three independent volatility processes, $h_{c,t}$, $h_{d,t}$, and $h_{x,t}$, associated with the innovations to consumption growth, dividend growth, and the predictable component, respectively. The most notable feature of $h_{c,t}$ is that it captures a drop in consumption growth volatility that occurred between 1940 and 1960. In magnitude, this drop in volatility is much larger than a subsequent decrease around 1984, the year typically associated with the Great Moderation. The stochastic volatility process for dividend growth $h_{d,t}$ seems to exhibit more medium- and high-frequency movements than $h_{c,t}$. Finally, the volatility of the persistent component, $h_{x,t}$, exhibits substantial fluctuations over our sample period, and it tends to peak during NBER recessions.

5.3. *Asset Pricing Implications*

Risk-Free Rate Estimate and Preference Shock. Figure 5 overlays the actual risk-free rate, which is assumed to be subject to measurement errors, and the smoothed "true" or model-implied risk-free rate. We find that the measurement errors are fairly small. To highlight the importance of the preference shock, we also plot a counterfactual risk-free rate that would prevail in the absence of $x_{\lambda,t}$. It turns out that ex post much of the risk-free rate fluctuations are explained by the preference shock. In the absence of the preference shock, the process for the expected stochastic discount factor implied by the predictable component of cash-flow growth and the stochastic volatilities is too smooth relative to the observed risk-free rate. The preference shock can generate additional fluctuations in the expected discount factor without having a significant impact on asset returns (as we will see below).

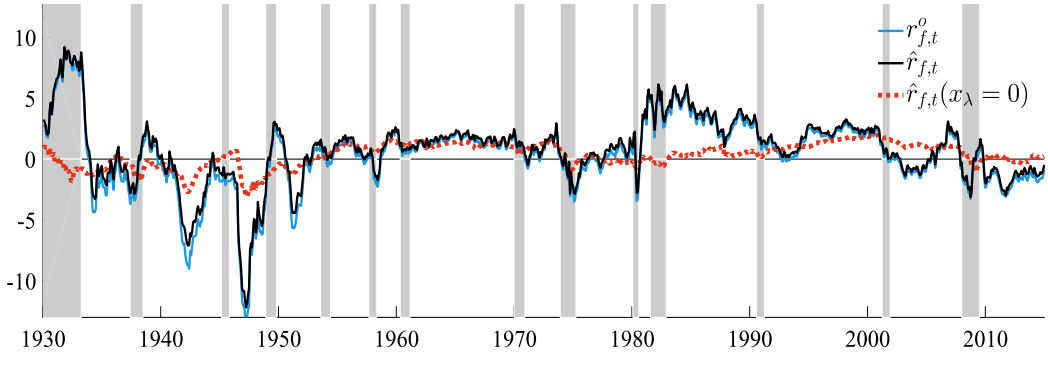

FIGURE 5.—Model-implied risk-free rate. *Notes*: Blue line depicts the actual risk-free rate, and black line depicts the smoothed, model-implied risk-free rate without measurement errors. Red dashed line depicts the model-implied risk-free rate with $x_{\lambda,t} = 0$. The parameters are fixed at their posterior median estimates.

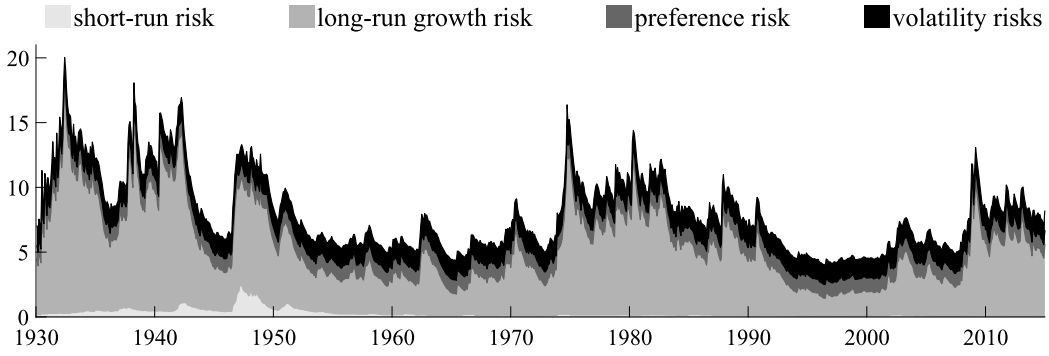

FIGURE 6.—Decomposition of the equity risk premium. *Notes*: We provide the decomposition of the risk premium (15). We compute β's and λ's based on the median posterior parameter estimates and multiply with the median volatility state estimates $\hat{\sigma}^2_{c,t}$ and $\hat{\sigma}^2_{x,t}$ to construct the model-implied risk premium. On average, the risk premium is accounted for by the short-run risk (0.3%), long-run growth risk (5.0%), preference risk (1.1%), and volatility risk (1.8%), respectively. The total in-sample market risk premium (annualized) is around 8.2%.

We assumed that the preference shock is independent of cash flows. In a production economy, this assumption will typically not be satisfied. Stochastic fluctuations in the discount factor generate fluctuations in consumption and investment, which in turn affect cash flows. We assess the independence as follows. First, we compute the ex post correlation between the smoothed preference shock innovations $\eta_{\lambda,t}$ and the cash-flow innovations $\eta_{c,t}$ and $\eta_{x,t}$. We can do so for every parameter draw Θ^s from the posterior distribution. The 90% posterior predictive intervals range from -0.09 to 0.03 for the correlation between $\eta_{\lambda,t}$ and $\eta_{c,t}$ and from 0 to 0.2 for the correlation between $\eta_{\lambda,t}$ and $\eta_{x,t}$. Second, we re-estimate our model under the assumption that $\eta_{\lambda,t}$ and $\eta_{x,t}$ are negatively correlated. The resulting parameter estimates as well as the asset pricing moments are essentially unaltered. According to a marginal data density comparison, the more parsimonious specification in which preference shocks and cash flows are independent is preferred. Based on these results, we conclude that there is no evidence that contradicts the independence assumption.

Determinants of the Equity Risk Premium. Figure 6 depicts the contribution of short-run risk, $\sigma^2_{c,t}$, the long-run growth risk, $\sigma^2_{x,t}$, the preference risk, σ^2_{λ}, and the volatility risks, $\sigma^2_{w_c}$ and $\sigma^2_{w_x}$, to the risk premium at the posterior median parameter estimates; see (15). We compute β's and λ's based on the median posterior parameter estimates and multiply them by the median volatility state estimates to construct the risk premium. The total (annualized) equity risk premium is around 8.2%.[21] The two major sources of the risk premium are the long-run growth risk and the volatility risks, and when combined they account for 83% of the risk premium. More specifically, the 8.2% equity premium can be decomposed as follows. On average, the long-run growth risk generates a premium of 5.0%, the volatility risks account for 1.8%, the preference shock generates 1.1%, and the short-run volatility risk contributes 0.3%.

[21]The gross equity premium is $\mathbb{E}[r_{m,t+1} - r_{f,t}] + 1/2\sigma^2_{r_m} \approx 0.0615 + 0.5 * 0.226^2 - 0.0047 = 8.2\%$. If we were to attribute the moving-average fluctuations in observed monthly consumption growth to "true" consumption growth instead of measurement errors, the asset pricing implications of the model would essentially remain unchanged. The equity premimum would rise by approximately 0.06%.

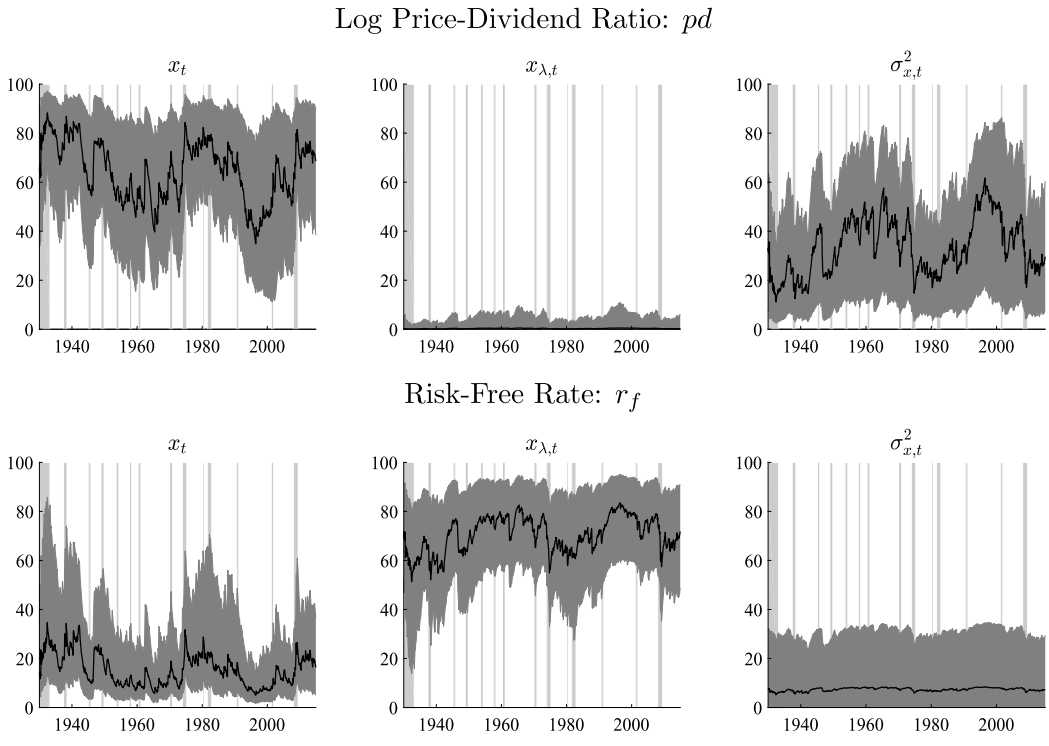

FIGURE 7.—Variance decomposition for market returns and risk-free rate. *Notes*: Fraction of volatility fluctuations (in percent) in the price-dividend ratio and the risk-free rate that is due to x_t, $x_{\lambda,t}$, and $\sigma^2_{x,t}$, respectively. We do not present the graphs for $\sigma^2_{c,t}$, $\sigma^2_{d,t}$ since their time-varying shares are less than 1% on average. See the main text for computational details.

Determinants of Asset Price Volatility. Figure 7 depicts the time-varying contribution of the fluctuations in growth prospects, x_t, the preference shock, $x_{\lambda,t}$, and the conditional variability of growth prospects, $\sigma_{x,t}$, to the volatility of the price-dividend ratio and the risk-free rate.[22] We generate counterfactual volatilities by shutting down the estimated x_t, $x_{\lambda,t}$, and $\sigma_{x,t}$ processes, respectively. The ratios of the counterfactual and the actual asset price volatilities measure the contribution of the non-omitted risk factors. We subtract this ratio from 1 to obtain the relative contribution of the omitted risk factor shown in Figure 7. The credible bands reflect parameter uncertainty and uncertainty about the latent states. While the preference shocks are important for the risk-free rate, they contribute very little to the variance of the price-dividend ratio. Most of the variability of the price-dividend ratio is, in equal parts, due to the variation in x_t and $\sigma_{x,t}$. The remaining risk factors $\sigma^2_{c,t}$ and $\sigma^2_{d,t}$ have negligible effects (less than 1% on average) on the asset price volatilities, but are important for tracking the consumption and dividend growth data.

Matching Asset Price Moments. While asset pricing moments implicitly enter the likelihood function of our state-space model, it is instructive to examine the extent to which sample moments implied by the estimated state-space model mimic the sample moments computed from our actual data set. To do so, we report percentiles of the posterior pre-

[22]The decomposition of market return volatility (not shown in Figure 7) is qualitatively similar to that of the price-dividend ratio.

TABLE VIII
ASSET RETURN MOMENTS[a]

		Parameter Estimates are Based On					
		Cash Flows & Asset Returns			Cash Flows Only		
	Data	5%	50%	95%	5%	50%	95%
Mean (r_m)	6.06	2.56	6.15	10.29	1.99	4.61	7.55
StdDev (r_m)	19.8	14.9	22.6	46.1	11.0	16.6	28.3
AC1 (r_m)	−0.01	−0.30	−0.05	0.18	−0.29	−0.02	0.19
Corr (Δc, r_m)	0.11	−0.10	0.10	0.29	−0.10	0.12	0.32
Mean (pd)	3.40	2.63	3.14	3.41	3.26	3.42	3.51
StdDev (pd)	0.45	0.17	0.32	0.76	0.11	0.18	0.39
AC1 (pd)	0.87	0.49	0.75	0.89	0.32	0.62	0.82
Mean (r_f)	0.37	−0.56	0.47	1.32	0.22	0.97	1.64
StdDev (r_f)	2.85	1.54	2.09	2.90	1.61	1.93	2.41
AC1 (r_f)	0.64	0.38	0.57	0.73	0.35	0.52	0.67

[a]We present descriptive statistics for log returns of the aggregate stock market (r_m), its correlation with consumption growth (Δc), the log risk-free rate (r_f), and the log price-dividend ratio (pd). We report means (Mean), standard deviations (StdDev), first-order sample autocorrelations (AC1), and correlations (Corr). Market returns, the risk-free rate, and the price-dividend ratio refer to 12-month averages (in percent). Computing asset pricing implications for the cash-flow-only estimates requires calibration of the preference parameters and the preference shock $x_{\lambda,t}$. We set δ, ψ, γ, ρ_λ, σ_λ^2 to the median posterior estimates from Table VII.

dictive distribution for various sample moments based on simulations from the posterior distribution of the same length as the data.[23] Typically, the posterior predictive distribution is computed to reflect both parameter and shock uncertainty. In our application, the effect of the parameter uncertainty is an order of magnitude smaller than the effect of the shock (or sampling) uncertainty. Thus, we decided to fix the parameters at their posterior median values as this facilitates a clear comparison between the two types of model parameterization.

Results are summarized in Table VIII. Means and standard deviations refer to annualized asset prices. We first focus on the results from estimating the full model based on cash-flow data and asset returns (full model estimation). All of the "actual" sample moments are within the 5th and the 95th percentile of the corresponding posterior predictive distribution.[24] In particular, the model generates a sizable mean log market return with median value of 6.2%, and a sizable equity risk premium with a median value of about 8.2%. Consistent with the data, the model's return variability is about 22%. The high volatility of the market returns translates into a large variability of the sample moments. As in the data, the model generates both a highly variable and persistent price-dividend ratio. The median and 95th percentile of the price-dividend volatility distribution are significantly larger than in other LRR calibrated models with Gaussian shocks. This feature

[23]This is called a posterior predictive check; see Geweke (2005) for a textbook treatment. Specifically, the percentiles are obtained using the following simulation: draw parameters Θ^s from the posterior distribution; for each Θ^s, simulate a trajectory Y^s (same number of observations as in the actual sample) and compute the sample statistics $\mathcal{S}(Y^s)$ of interest.

[24]Although not reported in the table, this is also the case for the mean, standard deviation, and first autocorrelation moments of consumption and dividend growth.

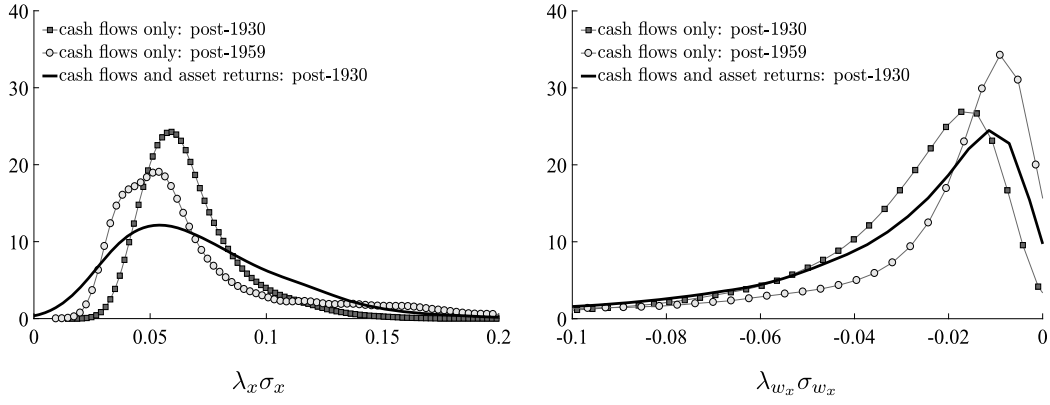

FIGURE 8.—Posterior distribution of market prices of risks. *Notes*: We plot posterior densities of $\lambda_x\sigma_x$ and $\lambda_{w_x}\sigma_{w_x}$ from the estimation with cash-flow data only from post-1930 (squared-line) and from post-1959 samples (circled-line), respectively, and from the estimation with cash-flow and asset return data from post-1930 sample (solid-line).

owes in part to the fact that the model contains an independent dividend volatility process. Finally, partly due to the preference shocks, the model is able to reproduce the observed sample moments of the risk-free rate.

In Section 5.2, we noted that the parameter estimates for the cash-flow processes change a bit once asset pricing data are included. To assess the economic implications of the parameter differentials, we combine the cash-flow process parameter estimates reported in Table VI (1930–2014 sample) with the posterior median estimates of the preference parameters and the preference shock $x_{\lambda,t}$ from the full estimation. Because the cash-flow-only model was estimated without the third volatility process $\sigma^2_{x,t}$, we set $h_{c,t} = h_{x,t}$ when recomputing the asset pricing implications of the LRR model. The last three columns of Table VIII show that, due to a lower persistence ρ, the cash-flow-only estimates generate a slightly lower mean and variance for the market return, and a slightly higher and less volatile price-dividend ratio and risk-free rate. The standard deviation and autocorrelation of the price-dividend ratio and the standard deviation of the risk-free rate lie just outside the posterior predictive bands, whereas all other sample moments continue to fall within the bands. This confirms that even the cash-flow-only estimates of the endowment process parameters can generate realistic asset price fluctuations.[25]

Figure 8 compares the posterior distributions for the appropriately-scaled market prices of risk, $\lambda_x\sigma_x$ and $\lambda_{w_x}\sigma_{w_x}$, based on the estimation with and without asset prices. The posterior densities are remarkably similar. In fact, the modes of the distributions are almost identical; the main difference lies in the dispersion of the densities. This indicates that the lower risk premium (see Table VIII) obtained under the cash-flow-only estimate is due to smaller return exposures to the shocks (β's in (15)).[26]

[25] Chen, Dou, and Kogan (2015) formalized this comparison by developing a measure of model fragility, roughly speaking based on the discrepancy between the posterior medians obtained under the cash-flow-only estimation and the estimation with asset returns.

[26] For completeness, using the SDF decomposition in Equation (A.18) in the Supplemental Material, we also report λ_i and $\lambda_i\sigma_i$ for $i = \{c, x, \lambda, w_w, w_c\}$ at our median parameter estimates. The resulting values are $\{8.9, 695.2, 406.3, -26{,}572{,}824, -3391.0\}$ and $\{0.03, 0.07, 0.16, -0.03, -0.01\}$, respectively. These figures are consistent with the variance decomposition of the risk-free rate and equity return presented earlier, whereby

Consumption Growth and Excess Return Predictability. Asset pricing models are often evaluated based on their implications for the predictability of future cash flows and returns. In the model, the price-dividend ratio is determined by multiple state variables. Consequently, a VAR-based predictive regression is a natural starting point. As in Bansal, Kiku, and Yaron (2012), we estimate a first-order VAR that includes consumption growth, the price-dividend ratio, the real risk-free rate, and the market return. Based on the estimated VAR coefficients, we compute R^2's for cumulative H-step-ahead consumption growth and excess returns:

$$\sum_{h=1}^{H} \Delta c_{t+h} \quad \text{and} \quad \sum_{h=1}^{H} (r_{m,t+h} - r_{f,t+h-1}).$$

While the VAR-based predictive checks are appealing from a theoretical perspective, much of the empirical literature focuses on R^2's from univariate predictive regressions using the price-dividend ratio as the only regressor. We subsequently consider both multivariate and univariate regressions.

Predictive checks are graphically summarized in Figure 9. We begin with a discussion of the results depicted in the four panels of the top row of the figure. The sample statistics considered are the R^2 values obtained from predictability regressions. The top and bottom ends of the boxes correspond to the 5th and 95th percentiles, respectively, of the posterior predictive distribution, and the horizontal bars signify the medians. The predictive intervals reflect the fact that we are repeatedly generating data from the model and computing a sample statistic for each of these simulated trajectories.[27] The small squares correspond to R^2 statistics computed from "actual" U.S. data.

The top left panel of Figure 9 depicts results for the VAR-based predictability regressions for consumption growth. Based on multiple variables, consumption growth is *highly* predictable in the data. At the one-year horizon, the R^2 is about 52% (see also Bansal et al. (2014)). While the predictability diminishes over time, it is still nontrivial with an R^2 of 12% at the 10-year horizon. The key finding is that the data R^2's lie within the 90% credible intervals constructed from the model-implied predictive distribution. At the one-year horizon, the median of the model-implied R^2 is somewhat lower than its data estimate, whereas over horizons of three years or more, the medians are slightly larger than the data estimates.

Panel 2 in the top row of Figure 9, labeled "Univariate," provides results for univariate consumption growth predictability regressions. As for the VAR-based predictability checks, we simulate the LRR model with all of its five state variables: x_t, $x_{\lambda,t}$, $\sigma^2_{x,t}$, $\sigma^2_{c,t}$, and $\sigma^2_{d,t}$. However, we only use the price-dividend ratio to predict future consumption growth. As is well known, when the price-dividend ratio is used as a single regressor, it produces low R^2's. They are less than 5% for horizons from one to eight years and reach almost 10% at the 10-year horizon.[28] The median R^2 values obtained from regressions on model-generated data are between 10% to 15%, slightly higher than in the actual data.

the former has a relatively large exposure (β) to the preference shock while the latter has a large exposure to the growth and volatility shock.

[27]To ease subsequent comparisons, we condition on the posterior median estimates of the LRR model. This is innocuous because the contribution of parameter uncertainty to the variability of the posterior predictive distributions is small.

[28]The univariate-based low R^2's for the first several years are consistent with the findings in Table 4 of Beeler and Campbell (2012)—the slight differences attributed to the longer sample available here.

However, the posterior predictive intervals range from 0 to 30% for the one-year horizon and from 0 to about 50% for horizons longer than three years, which means that there is no evidence in the data that contradicts our estimated asset pricing model.

Panels 3 and 4 in the top row of Figure 9 show the respective VAR and univariate predictive R^2's for future excess returns. It is noteworthy that the VAR median R^2's of the model-based estimates are almost perfectly aligned with the data-based estimates. The model also performs quite well in terms of the univariate excess return predictability regressions. Specifically, for all horizons, the median of the model-implied distribution of R^2's is quite close to actual data R^2's and the model-based credible intervals contain the R^2 obtained from the actual data. The good performance is obtained because, according to the model, the price-dividend ratio is the most important predictor of long-horizon excess returns among the observables.[29]

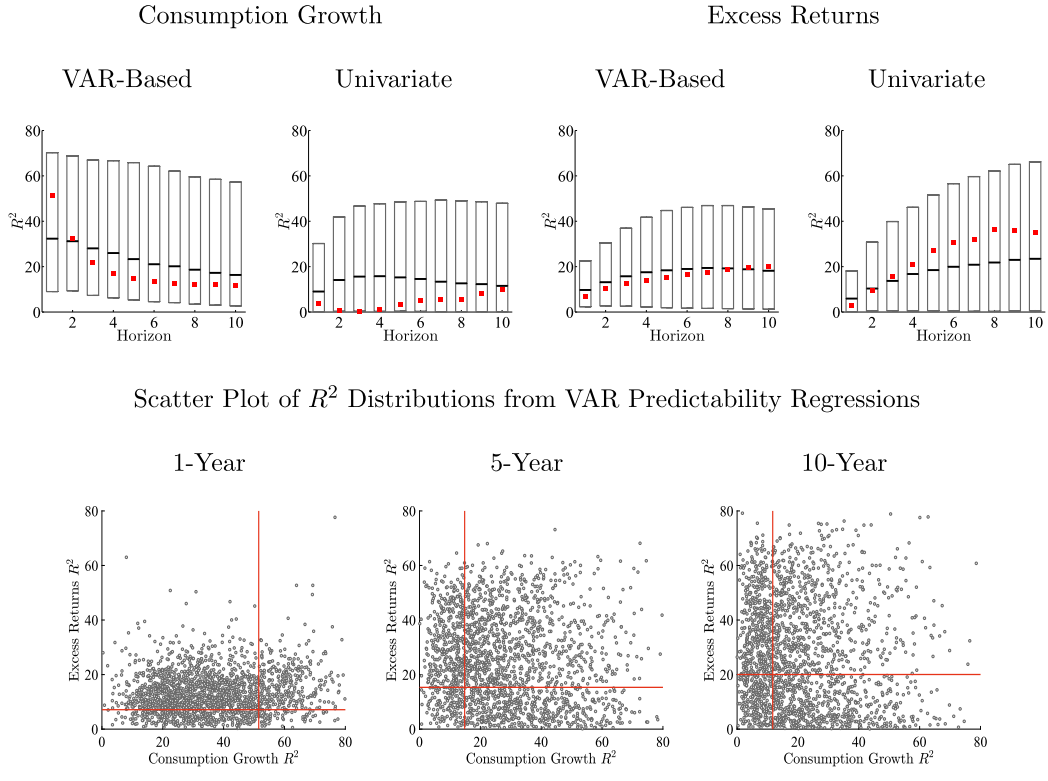

FIGURE 9.—Predictability checks. *Notes*: We fix the parameters at their posterior median estimates and simulate data sets. The red squares represent R^2 values obtained from the actual data. The boxes represent 90% posterior predictive intervals and the horizontal lines represent medians. The "Benchmark" case is based on simulations with all five state variables x_t, $x_{\lambda,t}$, $\sigma^2_{x,t}$, $\sigma^2_{c,t}$, and $\sigma^2_{d,t}$; the horizon is measured in years. The VAR-based R^2's are constructed as in Hodrick (1992). In the bottom panel, the intersection of the solid lines indicates the R^2 values obtained from the actual data.

[29]In the Supplemental Material, we explore the relative importance of "growth" and "volatility" risks by simulating model specifications that are only driven by (i) x_t and $\sigma^2_{x,t}$ or (ii) x_t. In Case (i), the posterior predictive distributions are quite close to the ones in the two "univariate" subplots of Figure 9 because x_t and $\sigma^2_{x,t}$ represent the key pricing state variables. In Case (ii), the credible intervals are often too small and do not encompass the data estimates. Driven by x_t only, the model generates too much consumption predictability,

While our model passes the predictive checks, the credible intervals depicted in Figure 9 are wide. The high variability of the sampling distribution of the R^2 measures under the LRR model implies that, despite their popularity, the predictability regressions have little power to detect model misspecifications. The diffuse and skewed sampling distributions of the R^2 statistics are caused by various non-standard features of predictive regressions. Due to overlapping time periods, residuals are typically serially correlated and lagged residuals may be correlated with the predictor. Moreover, the persistent component of the dependent variable (consumption growth or excess returns) is dominated by i.i.d. shocks and the right-hand-side regressor (price-dividend ratio) is highly persistent—a feature that can render the predictive regressions spurious (see Hodrick (1992) and Stambaugh (1999)).[30]

As a final check, the bottom row of Figure 9 illustrates the model-implied *joint* distribution of R^2's for predicting consumption growth and excess returns. Each dot in these scatter plots is obtained by computing the two R^2's based on a single model simulation. The intersection of the solid lines indicates the R^2 values computed from the actual data. The figure shows that the R^2 values at the 1-year and 5-year horizon are almost uniformly distributed over a rectangle. For every horizon, the observed R^2's do not lie in the far tails of the posterior predictive distribution, which means that the model is also able to jointly generate the observed consumption growth and excess return predictability.

Dividend-Growth Predictability. Cochrane (2011) argued that there is very little dividend-growth predictability at all horizons. This view is based on a univariate regression with the price-dividend ratio as a predictor of future dividend growth. The data feature modest predictability, with an R^2 in the range of 4% to 9%, depicted by the red squares in the left panel of Figure 10. However, dividend growth is found to be highly predictable both at short and long horizons, once additional predictors are included in a VAR-based predictive regression, with adjusted R^2's as large as 35% at the 10-year horizon (see Column 2 of Figure 10).[31] Importantly, in both the univariate and VAR-based predictive regressions, the model implications for dividend growth predictability line up with the data and cover the data R^2's.

The strong evidence for dividend growth predictability has important implications for the variability of the log dividend yield dp_t. Based on the Campbell and Shiller (1988a) approximate present value identity, it follows that

$$dp_t \approx \sum_{j=1}^{k} \varrho^{j-1} r_{t+j} - \sum_{j=1}^{k} \varrho^{j-1} \Delta d_{t+j} + \varrho^k dp_{t+k}, \tag{19}$$

where ϱ is an approximation constant based on the average dividend yield. Multiplying both sides of (19) by the log dividend yield and taking expectations implies that the variance of the current dividend yield can be attributed to its covariance with expected

thereby highlighting that volatility shocks play an important role in lowering the model-implied predictability to a more realistic level.

[30] Valkanov (2003) derived an asymptotic distribution of the R^2 under the assumption that the regressor follows a local-to-unity process. He showed that the goodness-of-fit measure converges to a random limit as the sample size increases. More recently, Bauer and Hamilton (2015) studied the sampling distribution of R^2 measures in predictive regressions for bond returns, which exhibit similar distortions.

[31] This evidence is consistent with Lettau and Ludvigson (2005), Koijen and van Binsbergen (2010), and Jagannathan and Liu (2016) who reported R^2 values from a VAR-based regression that range from 10% to 40%.

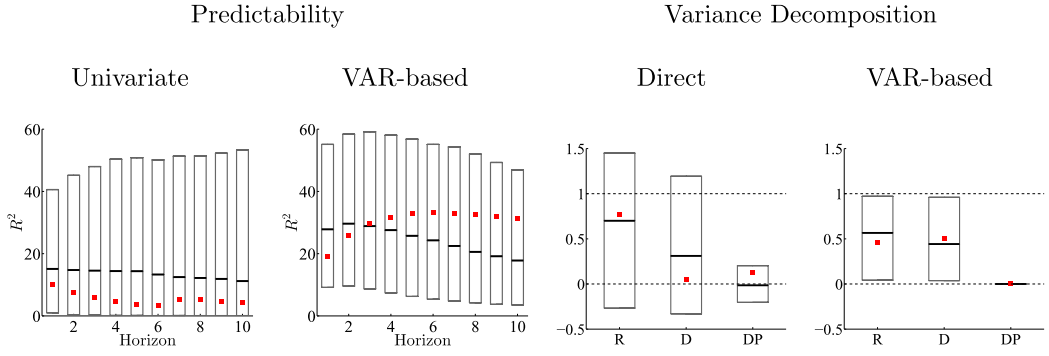

FIGURE 10.—Dividend growth predictability and dividend yield variance decomposition. *Notes*: (Predictability) We fix the parameters at their posterior median estimates and simulate data sets. The horizon is measured in years. We run a univariate regression with the price-dividend ratio as predictor of future dividend growth. For the multivariate regression, we consider a first-order VAR that includes consumption growth, dividend growth, the price-dividend ratio, and the real risk-free rate. Based on the estimated coefficients, we compute R^2's for cumulative H-step-ahead dividend growth. The red squares represent R^2 values obtained from the actual data. The boxes represent 90% posterior predictive intervals and the horizontal lines represent medians. The VAR-based R^2's are constructed as in Hodrick (1992). (Variance Decomposition, Direct) We regress 15-year ex post returns, dividend growth, and dividend yield, respectively, on a constant term and the dividend yield. (Variance Decomposition, VAR-based) We infer long-run coefficients ($k \to \infty$) from 1-year coefficients of the same VAR used for the predictability analysis. Using the Campbell–Shiller approximation, the fractions of dividend yield variation attributed to each source are provided as $1 \approx \frac{\text{Cov}(dp_t, \sum_{j=1}^{k} \varrho^{j-1} r_{t+j})}{\text{Var}(dp_t)} - \frac{\text{Cov}(dp_t, \sum_{j=1}^{k} \varrho^{j-1} \Delta d_{t+j})}{\text{Var}(dp_t)} + \frac{\varrho^k \text{Cov}(dp_t, dp_{t+k})}{\text{Var}(dp_t)}$. These components are marked as R, D, and DP, respectively.

future returns, dividend growth rates, and the expected future dividend yield, respectively, marked as "R," "D," and "DP" in Figure 10 (see figure notes for details). As k approaches infinity, the dividend yield variability is explained completely by covariation with expected returns and cash-flow growth. We compute the fraction of variability explained by the three covariances via "Direct" regression (setting k equal to 15 years and separately regressing the "R," "D," and "DP" components on the dividend yield) and "VAR-based" regression (inferring the $k = \infty$ decomposition from the coefficients of a VAR estimated based on annual data). The estimates based on the direct regressions attribute much of the variation in dividend yield to variation in discount rates (although not entirely), whereas the point estimates of the VAR attribute about half of the variation to discount rates and the other to dividend growth. Again, it is important to note that, in both cases, the model credible intervals contain the data point estimates. Moreover, in both cases, the credible intervals around the point estimates are consistent with a view in which a large portion (about half) of the dividend yield variability is driven by cash flows.[32]

6. CONCLUSION

We developed a nonlinear Bayesian state-space model that utilizes mixed-frequency data to study the time series dynamics of consumption and its implications for asset pricing. We show that after accounting for monthly measurement errors, there is strong evi-

[32] Albuquerque et al. (2016) also examined the 7- and 10-year correlations between cumulative return and cumulative consumption and dividend growth. For brevity, we defer this analysis to the Supplemental Material where we show the model's credible confidence bands contain the data estimates.

dence for both a small persistent predictable component as well as a stochastic volatility component in consumption growth. Importantly, this evidence emerges when the estimation uses information just from cash flows, namely, consumption, consumption and output, and consumption and dividends. It is further reinforced and sharpened when the estimation uses consumption, dividends, and asset return data jointly. The estimation identifies three volatility processes: one governing dynamics of the persistent cash-flow growth component, and the other two controlling temporally independent shocks to consumption and dividend growth. The model is able to successfully capture many asset pricing moments and improve upon key predictability moments of previous LRR models.

Our findings raise the broader question of whether DSGE models, more generally, should have a predictable component built into one or more of the exogenous processes that drive macroeconomic fluctuations. If the goal of the modeling endeavor is to capture business cycle fluctuations at the quarterly level, then the answer is no, because the signal in the data is not strong enough to render macroeconomic predictions from a model without this predictable component to be inaccurate. But if the goal is to rely on long-horizon implications of the model, for instance, with respect to asset prices, then the answer is affirmative.

REFERENCES

ALBUQUERQUE, R., M. EICHENBAUM, V. LUO, AND S. REBELO (2016): "Valuation Risk and Asset Pricing," *Journal of Finance*, 71 (6), 2861–2904. [620,636,651]

AMIR-AHMADI, P., C. MATTHES, AND M.-C. WANG (2016): "Drifts and Volatilities Under Measurement Error: Assessing Monetary Policy Shocks Over the Last Century," *Quantitative Economics*, 7 (2), 591–611. [629,630]

AN, S., AND F. SCHORFHEIDE (2007): "Bayesian Analysis of DSGE Models," *Econometric Reviews*, 26 (2–4), 113–172. [638]

ANDREASEN, M. (2010): "Stochastic Volatility and DSGE Models," *Economics Letters*, 108, 7–9. [636]

ANDRIEU, C., A. DOUCET, AND R. HOLENSTEIN (2010): "Particle Markov Chain Monte Carlo Methods (With Discussion)," *Journal of the Royal Statistical Society, Series B*, 72, 269–342. [621,640]

ARUOBA, S., F. DIEBOLD, AND C. SCOTTI (2009): "Real-Time Measurement of Business Conditions," *Journal of Business and Economic Statistics*, 27 (4), 417–427. [621]

BANSAL, R., AND A. YARON (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509. [617,618,624,637]

BANSAL, R., A. GALLANT, AND G. TAUCHEN (2007): "Rational Pessimism, Rational Exuberance, and Asset Pricing Models," *Review of Economic Studies*, 74, 1005–1033. [620]

BANSAL, R., V. KHATACHARIAN, AND A. YARON (2005): "Interpretable Asset Markets?" *European Economic Review*, 49, 531–560. [621]

BANSAL, R., D. KIKU, I. SHALIASTOVICH, AND A. YARON (2014): "Volatility, the Macroeconomy and Asset Prices," *Journal of Finance*, 69, 2471–2511. [621,648]

BANSAL, R., D. KIKU, AND A. YARON (2012): "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 1, 183–221. [620,621,635,637,648]

——— (2016): "Risks for the Long Run: Estimation With Time Aggregation," *Journal of Monetary Economics*, 82, 52–69. [620]

BARRO, R. (2009): "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99, 243–264. [618]

BAUER, M., AND J. HAMILTON (2017): "Robust Bond Risk Premia," *Review of Financial Studies,* 31 (2), 399–448. [650]

BEELER, J., AND J. CAMPBELL (2012): "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 1, 141–182. [648]

BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685. [618,621]

CAMPBELL, J., AND J. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251. [617]

CAMPBELL, J., AND R. SHILLER (1988a): "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1, 195–227. [637,650]

——— (1988b): "Stock Prices, Earnings, and Expected Dividends," *Journal of Finance*, 43, 661–676. [632,640]

CARTER, C. K., AND R. KOHN (1994): "On Gibbs Sampling for State Space Models," *Biometrika*, 81 (3), 541–553. [624]

CHEN, H., W. D. DOU, AND L. KOGAN (2015): "Measuring the "Dark Matter" in Asset Pricing Models," Manuscript, MIT. [647]

CHEN, R., AND J. LIU (2000): "Mixture Kalman Filters," *Journal of the Royal Statistical Society Series B*, 62, 493–508. [621,640]

CHERNOV, M., R. GALLANT, E. GHYSELS, AND G. TAUCHEN (2003): "Alternative Models for Stock Price Dynamics," *Journal of Econometrics*, 116, 225–257. [636]

COCHRANE, J. (2011): "Presidential Address: Discount Rates," *Journal of Finance*, 66 (4), 1047–1108. [650]

CREAL, D. D., AND J. C. WU (2015): "Bond Risk Premia in Consumption-Based Models," Manuscript, Chicago Booth. [621]

CROCE, M. (2014): "Long-run Productivity Risk: A New Hope for Production-Based Asset Pricing?" *Journal of Monetary Economics*, 66, 13–31. [632]

DOH, T., AND S. WU (2015): "Cash Flow and Risk Premium Dynamics in an Equilibrium Asset Pricing Model With Recursive Preferences," FRB Kansas City Research Working Paper, 15-12. [621]

DROST, F., AND T. NIJMAN (1993): "Temporal Aggregation of Garch Processes," *Econometrica*, 61, 909–927. [618]

EPSTEIN, L., AND S. ZIN (1989): "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969. [635]

FERNÁNDEZ-VILLAVERDE, J., AND J. F. RUBIO-RAMÍREZ (2007): "Estimating Macroeconomic Models: A Likelihood Approach," *Review of Economic Studies*, 74 (4), 1059–1087. [621]

——— (2011): "Macroeconomics and Volatility: Data, Models, and Estimation," in *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society*, ed. by D. Acemoglu, M. Arellano and E. Deckel. Cambridge University Press. [618,621]

GEWEKE, J. (2005): *Contemporary Bayesian Econometrics and Statistics*. New Jersey: John Wiley & Sons. [646]

GOURIO, F. (2012): "Disaster Risk and Business Cycles," *American Economic Review*, 102 (6), 2734–2766. [618]

HALL, R. (1978): "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86 (6), 971–986. [617]

HANSEN, L. (2007): "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk," *American Economic Review*, 97 (2), 1–30. [628]

HANSEN, L., AND T. SARGENT (2007): *Robustness*. Princeton University Press. [618]

HANSEN, L. P., J. C. HEATON, AND N. LI (2008): "Consumption Strikes Back? Measuring Long-Run Risk," *Journal of Political Economy*, 116 (2), 260–302. [617,627]

HARVEY, A. (1989): *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press. [621]

HODRICK, R. (1992): "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies*, 5, 357–386. [632,640,649-651]

JAGANNATHAN, R., AND B. LIU (2016): "Dividend Dynamics, Learning, and Expected Stock Index Returns," Manuscript, Northwestern University. [650]

JOHANNES, M., L. LOCHSTOER, AND Y. MOU (2016): "Learning About Consumption Dynamics," *Journal of Finance*, 71 (2), 551–600. [621]

KIM, S., N. SHEPHARD, AND S. CHIB (1998): "Stochastic Volatility: Likelihood Inference and Comparison With ARCH Models," *Review of Economic Studies*, 65, 361–393. [624]

KOIJEN, R., AND J. VAN BINSBERGEN (2010): "Predictive Regressions: A Present-Value Approach," *Journal of Finance*, 65 (4), 1439–1471. [650]

LETTAU, M., AND S. LUDVIGSON (2005): "Expected Returns and Expected Dividend Growth," *Journal of Financial Economics*, 76 (3), 583–626. [650]

LUCAS, R. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46 (6), 1429–1445. [636]

MARIANO, R., AND Y. MURASAWA (2003): "A New Coincident Index of Business Cycles Based on Monthly and Quarterly Series," *Journal of Applied Econometrics*, 18, 427–443. [621]

NAKAMURA, E., D. SERGEYEV, AND J. STEINSSON (2017): "Growth-Rate and Uncertainty Shocks in Consumption: Cross-Country Evidence," *American Economic Journal: Macroeconomics*, 9 (1), 1–39. [621]

POHL, W., K. SCHMEDDERS, AND O. WILMS (2017): "Higher-Order Effects in Asset-Pricing Models With Long-Run Risks," *Journal of Finance* (forthcoming). [638]

ROMER, C. (1986): "New Estimates of Prewar Gross National Product and Unemployment," *The Journal of Economic History*, 46 (2), 341–352. [629]

——— (1989): "The Prewar Business Cycle Reconsidered: New Estimates of Gross National Product, 1869–1908," *Journal of Political Economy*, 97 (1), 1–37. [629]

SCHORFHEIDE, F., AND D. SONG (2015): "Real-Time Forecasting With a Mixed-Frequency VAR," *Journal of Business and Economic Statistics*, 33 (3), 366–380. [621]

SCHORFHEIDE, F., D. SONG, AND A. YARON (2018): "Supplement to 'Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach'," *Econometrica Supplemental Material*, 86, http://dx.doi.org/10.3982/ECTA14308. [621,638]

SLESNICK, D. (1998): "Are our Data Relevant to the Theory? The Case of Aggregate Consumption," *Journal of Business and Economic Statistics*, 16 (1), 52–61. [618]

STAMBAUGH, R. F. (1999): "Predictive Regressions," *Journal of Financial Economics*, 54, 375–421. [650]

VALKANOV, R. (2003): "Long-Horizon Regressions: Theoretical Results and Applications," *Journal of Financial Economics*, 68, 201–232. [650]

WILCOX, D. (1992): "The Construction of the U.S. Consumption Data: Some Facts and Their Implications for Empirical Work," *American Economic Review*, 82, 922–941. [618,627,635]

Co-editor Giovanni L. Violante handled this manuscript.

Manuscript received 11 April, 2016; final version accepted 24 September, 2017; available online 26 September, 2017.