Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach

Frank Schorfheide* | Dongho Song | Amir Yaron
University of Pennsylvania | Boston College | University of Pennsylvania
CEPR and NBER | | NBER

This Version: December 1, 2015

Abstract

We develop a nonlinear state-space model that captures the joint dynamics of consumption, dividend growth, and asset returns. Our model consists of an economy containing a common predictable component for consumption and dividend growth and multiple stochastic volatility processes. The estimation is based on annual consumption data from 1929 to 1959, monthly consumption data after 1959, and monthly asset return data throughout. We maximize the span of the sample to recover the predictable component and use high-frequency data, whenever available, to efficiently identify the volatility processes. Our Bayesian estimation provides strong evidence for a small predictable component in consumption growth (even if asset return data are omitted from the estimation). Three independent volatility processes capture different frequency dynamics; our measurement error specification implies that consumption is measured much more precisely at an annual than monthly frequency; and the estimated model is able to capture key asset-pricing facts of the data.

*Correspondence: Department of Economics, 3718 Locust Walk, University of Pennsylvania, Philadelphia, PA 19104-6297. Email: schorf@ssc.upenn.edu (Frank Schorfheide). Department of Economics, Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467. Email: dongho.song@bc.edu (Dongho Song). The Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6367. Email: yaron@wharton.upenn.edu (Amir Yaron).
We thank Arthur Lewbel, Bent J. Christensen, Frank Diebold, Emily Fox, Ian Dew-Becker, Ivan Shaliastovich, Neil Shephard, Minchul Shin, and seminar participants at the 2013 SED Meetings, the 2013 SBIES Meetings, the 2014 AEA Meetings, the 2014 Aarhus Macro-Finance Symposium, the 2015 UBC Summer Finance Conference, the Board of Governors, Boston College, Columbia University, Cornell University, the European Central Bank, Universite de Toulouse, and the University of Pennsylvania for helpful comments and discussions. Schorfheide gratefully acknowledges financial support from the National Science Foundation under Grant SES 1061725. Yaron thanks the Rodney White Center for financial support.

1 Introduction

Financial economists seek to understand the sources underlying risk and return in the economy. In equilibrium models this endeavor hinges on the preference specification and the joint dynamics of cash flows, which in an endowment economy correspond to consumption and dividends. There are many equilibrium models that appeal to low-frequency components in these cash flows as well as important time variation in the fundamentals (e.g., models of long-run risks (LRR) as in Bansal and Yaron (2004), and models of rare disasters as in Barro (2009)). Identifying both of these components is challenging. To measure the small persistent component in, say, consumption and dividend growth one would want the longest span of data. On the other hand, to estimate the time variation in second moments of cash flows one would ideally like to use high-frequency data. The empirical analysis is constrained by the availability of consumption data. For the U.S., the longest span of available data for consumption growth is at the annual frequency starting in 1929. The highest-frequency consumption data is available at the monthly frequency from 1959. To exploit all the available information in mixed-frequency data, this paper develops a Bayesian state-space model that prominently features stochastic volatility and time-aggregates consumption whenever it is observed only at a low frequency.

Our state-space model is designed to capture the joint dynamics of consumption, dividend growth, and asset returns. Building on the work of Bansal and Yaron (2004), the core of our model consists of an endowment economy that is, in part, driven by a common predictable component for consumption and dividend growth. Our model distinguishes itself from the existing LRR literature in several important dimensions. First, our state-space representation contains measurement equations that time-aggregate consumption to the observed frequency, yet allow us to maintain the likelihood representation (see Bansal, Kiku, and Yaron (2012b) for a generalized methods-of-moments (GMM) approach using time aggregation). Our measurement-error specification accounts for different types of measurement errors at monthly and annual frequencies while respecting the constraint that monthly growth rates have to be consistent with annual growth rates.

Second, we generalize the volatility dynamics of Bansal and Yaron (2004)'s model specification by allowing for three separate volatility processes — one capturing long-run consumption innovations, one capturing short-run consumption innovations, and a separate process for dividend dynamics. We do so since our estimation procedure, which focuses on the joint distribution of consumption, dividends, and asset prices, requires separate stochastic volatility processes to fit the data. Third, we specify an additional process for variation in the time rate of preference as in Albuquerque, Eichenbaum, and Rebelo (2012), 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 state-space model generates several important empirical findings. First, we find strong evidence for a small predictable component in consumption growth. This evidence consists of two parts. We begin by estimating the state-space model on cash flow growth data only. Our carefully specified measurement-error model for cash flow data allows us to measure this component which otherwise is difficult to detect.[1] We then proceed by adding asset return data to the estimation and, in line with the existing LRR literature, find even stronger evidence for this predictable component. The Bayesian approach allows us to characterize the uncertainty about the persistence of the conditional mean growth process. We find that in spite of using a prior with a mean of 0.9 and a standard deviation of 0.5 our estimation yields a posterior distribution that is tightly centered around 0.99. Second, our estimated measurement errors for consumption growth are consistent with the common view (see Wilcox (1992)) that consumption growth is measured more precisely at an annual rather than monthly frequency.

Third, all three stochastic volatility processes display significant time variation yet behave distinctly over time. The volatility processes partly capture heteroskedasticity 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. As emphasized by the LRR literature, the volatility processes have to be very persistent in order to have significant quantitative effects on asset prices. 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 volatility process to account for the weak correlation between consumption growth and the risk-free rate. In the generalized specification of the model in which there are independent time rate of preference shocks, this correlation is further muted and the model fit for the dynamics of the risk-free rate is improved.

Fourth, it is worth noting that the median posterior estimate for risk aversion is 10-11 while it is around 1.5 for the intertemporal elasticity of substitution (IES). These estimates are broadly consistent with the parameter values highlighted in the LRR literature (see Bansal and Yaron (2004), Bansal, Kiku, and Yaron (2012a), and Bansal, Kiku, and Yaron (2012b)). Fifth, 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

[1] In section 4.2 we show that modeling measurement errors is important for sharp identification of our parameter estimates and that when measurement error are ignored the MLE estimates are bimodal with evidence for low or high persistent consumption growth component.

the equity premium is 6%, 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-squares from predicting returns and consumption growth by the price-dividend ratio.

Our paper also contains a number of technical innovations. First, in the specification of our state-space model we follow the stochastic volatility literature and assume that volatilities evolve according to exponential Gaussian processes that guarantee nonnegativity. 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. However, the approximation of the exponential volatility process is used only to derive the coefficients in the law of motion of the asset prices.

Second, we use a Markov chain Monte Carlo (MCMC) algorithm to generate parameter draws from the posterior distribution. This algorithm requires us to evaluate the likelihood function of our state-space model with a nonlinear filter. Due to the high-dimensional state space that arises from 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.

Our paper is related to several strands of the literature. In terms of the LRR literature, our paper is closely related to that of Bansal, Kiku, and Yaron (2012b) who utilize time aggregation and GMM to estimate the LRR model (see also Bansal, Gallant, and Tauchen (2007) for an approach using the efficient method of moments (EMM)). As noted above, our likelihood-based approach provides evidence which is broadly consistent with the results highlighted in that paper and other calibrated LRR models (see Bansal, Kiku, and Yaron (2012a)). 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 model to track cash flow and asset return time series.

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). Since 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 auxiliary particle filter of Pitt and Shephard (1999) and Chen and Liu (2000) and is recently discussed in Shephard (2013). Particle filter methods are also utilized in Johannes, Lochstoer, and Mou (2013), who estimate 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 (2013) examine the role of parameter uncertainty for asset prices, which is ignored in our analysis, they use a more restrictive version of the cash flow process and do not utilize mixed-frequency observations.

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 (2012)). 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 (2012a), and Bansal, Kiku, Shaliastovich, and Yaron (2013)). 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 model environment and describes the model solution. Section 3 presents the empirical state-space model and describes the estimation procedure. Section 4 discusses the empirical findings and Section 5 provides concluding remarks.

2 The Long-Run Risks (LRR) Model

Our baseline LRR model is described in Section 2.1. The solution of the model is outlined in Section 2.2. Section 2.3 presents a generalized version of the LRR with an exogenous shock to the time rate of preference.

2.1 Model Statement

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

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

subject to budget constraint

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

, 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.

Following Bansal and Yaron (2004),we decompose consumption growth, $g_{c,t+1}$, into a persistent component, x_t, and a transitory component, $\sigma_{c,t}\eta_{c,t+1}$. The dynamics for the persistent conditional mean follow an AR(1) with its own stochastic volatility process. Dividend streams have levered exposures to both the persistent and transitory component in consumption 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 dynamics for the cash flows are

$$\begin{aligned} g_{c,t+1} &= \mu_c + x_t + \sigma_{c,t}\eta_{c,t+1} \\ x_{t+1} &= \rho x_t + \sigma_{x,t}\eta_{x,t+1} \\ g_{d,t+1} &= \mu_d + \phi x_t + \pi\sigma_{c,t}\eta_{c,t+1} + \sigma_{d,t}\eta_{d,t+1}, \end{aligned} \tag{1}$$

where the conditional volatilities evolve according to[2]

$$\sigma_{i,t} = \varphi_i\bar{\sigma}\exp(h_{i,t}), \quad h_{i,t+1} = \rho_{h_i}h_{i,t} + \sigma_{h_i}\sqrt{1-\rho_{h_i}^2}w_{i,t+1}, \quad i = \{c,x,d\} \tag{2}$$

and the shocks are assumed to be

$$\eta_{i,t+1}, w_{i,t+1} \sim N(0,1), \quad i = \{c,x,d\}.$$

Relative to Bansal and Yaron (2004), the volatility dynamics contain three separate volatility processes. More importantly, the logarithm of the volatility process is assumed to be normal, which ensures that the standard deviation of the shocks remains positive at every point in time.

2.2 Solution

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

$$E_t\left[\exp\left(m_{t+1} + r_{i,t+1}\right)\right] = 1, \tag{3}$$

where $m_{t+1} = \theta\log\delta - \frac{\theta}{\psi}g_{c,t+1} + (\theta-1)r_{c,t+1}$ 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 dividend cash flows. Given the cash flow dynamics in (1) and

[2]Strictly speaking, to guarantee the existence of conditional moments involved in key equilibrium conditions, the exponential function needs to be spliced together with a non-exponential function, e.g., a square-root function, for volatilities exceeding some large threshold $\bar{h}_i$. See Chernov, Gallant, Ghysels, and Tauchen (2003) and Andreasen (2010).

the Euler equation (3), we derive asset prices using the approximate analytical solution described in Bansal, Kiku, and Yaron (2012a) which utilizes the Campbell and Shiller (1988a) log-linear approximation for returns.

However, since the volatility processes in (2) do not follow normal distributions, an analytical expression to (3) is infeasible. To accommodate an analytical solution, we utilize a linear approximation to (2) and express volatility in (4) as a process that follows Gaussian dynamics:

$$\begin{aligned}\sigma_{i,t}^2 - (\varphi_i\bar{\sigma})^2 &= 2(\varphi_i\bar{\sigma})^2 h_{i,t} + O(|h_{i,t}^2|), \quad h_{i,t+1} = \rho_{h_i} h_{i,t} + \sigma_{h_i}\sqrt{1-\rho_{h_i}^2}\, w_{i,t+1} \\ \sigma_{i,t+1}^2 &\approx (\varphi_i\bar{\sigma})^2(1-\rho_{h_i}) + \rho_{h_i}\sigma_{i,t}^2 + (2(\varphi_i\bar{\sigma})^2\sigma_{h_i}\sqrt{1-\rho_{h_i}^2})w_{i,t+1} \\ &= (\varphi_i\bar{\sigma})^2(1-\nu_i) + \nu_i\sigma_{i,t}^2 + \sigma_{w_i} w_{i,t+1}, \quad i = \{c,x,d\}. \end{aligned} \tag{4}$$

The analytical solution afforded via this pseudo-volatility process is important since it facilitates estimation (see details below).

The solution to the log price-consumption ratio follows, $pc_t = A_0 + A_1 x_t + A_{2,c}\sigma_{c,t}^2 + A_{2,x}\sigma_{x,t}^2$. 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. 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 agents have a preference for early resolution of uncertainty.

State prices in the economy are reflected in the innovations to the stochastic discount factor (SDF),

$$m_{t+1} - E_t\left[m_{t+1}\right] = \underbrace{\lambda_c\sigma_{c,t}\eta_{c,t+1}}_{\text{short-run consumption risk}} + \underbrace{\lambda_x\sigma_{x,t}\eta_{x,t+1}}_{\text{long-run growth risk}} + \underbrace{\lambda_{w_x}\sigma_{w_x}w_{x,t+1} + \lambda_{w_c}\sigma_{w_c}w_{c,t+1}}_{\text{volatility risks}},$$

where the derivation and λs are given in Appendix A. It is instructive to note that $\lambda_c = -\gamma$, $\lambda_x = \frac{-(\gamma-\frac{1}{\psi})\kappa_1}{1-\kappa_1\rho}$ (and λ_{w_c} and λ_{w_x}) is negative (positive) whenever preferences exhibit early resolution of uncertainty $\gamma > 1/\psi$. Furthermore the λs (except λ_c) will be zero when preferences are time separable, namely, when $\theta = 1$.

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

$$\begin{aligned} E_t(r_{m,t+1} - r_{f,t}) + \frac{1}{2}var_t(r_{m,t+1}) &= -cov_t(m_{t+1}, r_{m,t+1}) \\ &= \underbrace{\beta_{m,c}\lambda_c\sigma_{c,t}^2}_{\text{short-run risk}} + \underbrace{\beta_{m,x}\lambda_x\sigma_{x,t}^2}_{\text{long-run growth risk}} + \underbrace{\beta_{m,w_x}\lambda_{w_x}\sigma_{w_x}^2 + \beta_{m,w_c}\lambda_{w_c}\lambda_c\sigma_{w_c}^2}_{\text{volatility risks}}, \end{aligned} \tag{5}$$

where the βs are given in Appendix A and reflect the exposures of the market return to the underlying consumption risks. Equation (5) highlights that the conditional equity premium can be

attributed to (i) short-run consumption growth, (ii) long-run growth, (iii) 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 (1), the risk-free rate is affine in the state variables and follows

$$r_{f,t} = B_0 + B_1 x_t + B_{2,c}\sigma_{c,t}^2 + B_{2,x}\sigma_{x,t}^2,$$

where the Bs are derived in Appendix A. It is worth noting that $B_1 = \frac{1}{\psi} > 0$ and the risk-free rate rises with good economic prospects, while under $\psi > 1$, $\gamma > 1$ and whenever preferences exhibit early resolution of uncertainty, $B_{2,c} = -\frac{1}{2}(\frac{\gamma-1}{\psi} + \gamma)$ and $B_{2,x} = -\frac{(1-\frac{1}{\psi})(\gamma-\frac{1}{\psi})\kappa_1^2}{2(1-\kappa_1\rho)^2}$ are negative so the risk-free rate declines with a rise in economic uncertainty.

2.3 Generalized Model

In this section we augment the baseline model, as highlighted in Albuquerque, Eichenbaum, and Rebelo (2012), by allowing for a preference shock to the time rate of preference. Specifically, now the utility function contains a time rate of preference shock, λ_t, so the lifetime utility is

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

The resulting SDF equals the SDF described above plus the term $\theta x_{\lambda,t}$, where $x_{\lambda,t} = \lambda_{t+1}/\lambda_t$ is the growth rate of the preference shock. x_λ is assumed to follow an AR(1) process with persistence parameter ρ_λ and innovation shocks that are independent of other model shocks (see Appendix A.6 for derivation of this augmented SDF). Since $x_{\lambda,t}$ is known at time t, the risk-free rate will incorporate its values and consequently allow this generalized model to fit the risk-free rate dynamics better than the benchmark model.

3 State-Space Representation of the LRR Model

In order to conduct our empirical analysis, we cast the LRR model of Section 2 into state-space form. The state-space representation consists of a measurement equation that expresses the observables as a function of underlying state variables and a transition equation that describes the law of motion of the state variables. The measurement equation takes the form

$$y_{t+1} = A_{t+1}\big(D + Zs_{t+1} + Z^v s_{t+1}^v + \Sigma^u u_{t+1}\big), \quad u_{t+1} \sim iidN(0, I). \tag{6}$$

In our application, y_{t+1} consists of consumption growth, dividend growth, market returns, and the risk-free rate. The vector s_{t+1} stacks state variables that characterize the level of cash flows. The vector s^v_{t+1} is a function of the log volatilities of cash flows, h_t and h_{t+1}, in (2). Finally, 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. The solution of the LRR model sketched in Section 2.2 provides the link between the state variables and the observables y_{t+1}.

The state variables themselves follow vector autoregressive processes of the form

$$s_{t+1} = \Phi s_t + v_{t+1}(h_t), \quad h_{t+1} = \Psi h_t + \Sigma_h w_{t+1}, \quad w_{t+1} \sim iidN(0, I), \tag{7}$$

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. Roughly speaking, the vector s_{t+1} consists of the persistent cash flow component x_t (see (1)) as well as $x_{\lambda,t}$ in the generalized model of Section 2.3. However, in order to express the observables y_{t+1} as a linear function of s_{t+1} and to account 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. Since the details are cumbersome and at this stage non essential, a precise definition of s_{t+1} is relegated to the appendix.

A novel feature of our empirical analysis is the mixed-frequency approach. While dividend growth, equity return, and risk-free rate data are available at a monthly frequency from 1929 onwards, consumption data prior to 1959 are not available at a monthly frequency. Moreover, post-1959 monthly consumption growth data are subject to sizeable measurement errors, which is why many authors prefer to estimate consumption-based asset pricing models based on time-aggregated data. Our state-space approach avoids the loss of information due to time aggregation, yet we can allow for imprecisely measured consumption data at a monthly frequency. We discuss the measurement equations for consumption in Section 3.1 and the other observables in Section 3.2. Section 3.3 describes the implementation of Bayesian inference.

3.1 A Measurement Equation for Consumption

In our empirical analysis we use annual consumption growth rates prior to 1959 and monthly consumption growth rates subsequently.[3] The measurement equation for consumption in our state-space representation has to be general enough to capture two features: (i) the switch from annual to monthly observations in 1959, and (ii) measurement errors that are potentially larger at a monthly frequency than an annual frequency. To describe the measurement equation for consumption growth

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

data, we introduce some additional notation. We use the superscript o to distinguish observed consumption and consumption growth, C_t^o and $g_{c,t}^o$, from model-implied consumption and consumption growth, C_t and $g_{c,t}$. 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.

We define annual consumption as the sum of monthly consumption over the span of one year, i.e.:

$$C_{(j)}^a = \sum_{m=1}^{12} C_{12(j-1)+m}.$$

Log-linearizing this relationship around a monthly value C_* and defining lowercase c as percentage deviations from the log-linearization point, i.e., $c = \log C/C_*$, we obtain

$$c_{(j)}^a = \frac{1}{12}\sum_{m=1}^{12} c_{12(j-1)+m}.$$

Thus, monthly consumption growth rates can be defined as

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

and annual growth rates are given by

$$g_{c,(j)}^a = c_{(j)}^a - c_{(j-1)}^a = \sum_{\tau=1}^{23}\left(\frac{12-|\tau-12|}{12}\right) g_{c,12j-\tau+1}. \tag{8}$$

We assume a multiplicative *iid* measurement-error model for the level of annual consumption, which implies that, after taking log differences,

$$g_{c,(j)}^{a,o} = g_{c,(j)}^a + \sigma_\epsilon^a\big(\epsilon_{(j)}^a - \epsilon_{(j-1)}^a\big). \tag{9}$$

Moreover, consistent with the practice of the Bureau of Economic Analysis, we assume that the levels of monthly consumption are constructed by distributing annual consumption over the 12 months of a year. This distribution is based on an observed monthly proxy series z_t that is assumed to provide a noisy measure of monthly consumption. 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:

$$\begin{aligned} g_{c,12(j-1)+1}^o &= g_{c,12(j-1)+1} + \sigma_\epsilon\big(\epsilon_{12(j-1)+1} - \epsilon_{12(j-2)+12}\big) \\ &\quad -\frac{1}{12}\sum_{m=1}^{12}\sigma_\epsilon\big(\epsilon_{12(j-1)+m} - \epsilon_{12(j-2)+m}\big) + \sigma_\epsilon^a\big(\epsilon_{(j)}^a - \epsilon_{(j-1)}^a\big) \\ g_{c,12(j-1)+m}^o &= g_{c,12(j-1)+m} + \sigma_\epsilon\big(\epsilon_{12(j-1)+m} - \epsilon_{12(j-1)+m-1}\big), \quad m = 2, \ldots, 12 \end{aligned} \tag{10}$$

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 (10) 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 (8) averages out the measurement errors and yields (9).

We operate under the assumption that the agents in the model observe consumption growth, dividend growth, and asset returns in every period. As econometricians who are estimating the model, we have to rely on the statistical agency to release the consumption growth data. 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, for months $m = 1, \ldots, 11$ the vector $y_{12(j-1)+m}$ in (6) does not contain any observations on consumption growth. At the end of each year, in month $m = 12$, the vector $y_{12(j-1)+m}$ contains the 12 monthly growth rates of year j and (10) provides the portion of the measurement equation for the consumption data. The vector s_t has to contain sufficiently many lags of the model states as well as some lagged measurement errors such that it is possible to write (10) as a linear function of s_t. For the earlier part of the sample in which monthly consumption growth observations are not available, (10) is replaced by (8) and (9). The matrix A_t in (6) adapts the system to the availability of consumption data and the changing dimension of the vector y_t. Further details are provided in the appendix.

3.2 Measurement Equations for Dividend Growth and Asset Returns

It is reasonable to believe that consumption measurement errors are large, but those for financial variables (e.g., dividend streams, market returns, risk-free rates) are negligible. However, to be chary, we introduce idiosyncratic components for dividend growth, market returns, and the risk-free rate as well:

$$\begin{aligned} g^o_{d,t+1} &= g_{d,t+1} + \sigma^d_\epsilon \epsilon_{d,t+1} \\ r^o_{m,t+1} &= r_{m,t+1} + \sigma^{r_m}_\epsilon \epsilon_{r_m,t+1} \\ r^o_{f,t+1} &= r_{f,t+1} + \sigma^{r_f}_\epsilon \epsilon_{r_f,t+1}. \end{aligned} \tag{11}$$

In the subsequent empirical analysis we consider a version of the model in which only the risk-free rate is measured with error, i.e., $\sigma^d_\epsilon = 0$, $\sigma^{r_m}_\epsilon = 0$, $\sigma^{r_f}_\epsilon > 0$. We believe that aggregate dividend growth and stock market data are measured relatively cleanly and do not want to deviate too far from the existing asset pricing literature.

3.3 Bayesian Inference

Equations (6) and (7) define a nonlinear state-space system in which the size of the vector of observables y_t changes in a deterministic manner. The system matrices are functions of the parameter vector

$$\Theta = \left(\delta, \psi, \gamma, \rho, \phi, \varphi_x, \varphi_d, \bar{\sigma}, \mu, \mu_d, \pi, \sigma_\epsilon, \sigma_\epsilon^a, \rho_\lambda, \sigma_\lambda, \sigma_\epsilon^{r_f}, \rho_{h_c}, \sigma_{h_c}, \rho_{h_x}, \sigma_{h_x}, \rho_{h_d}, \sigma_{h_d}\right). \tag{12}$$

We will use a Bayesian approach to make inference about Θ and to study the implications of our model. Bayesian inference requires the specification of a prior distribution $p(\Theta)$ and the evaluation of the likelihood function $p(Y|\Theta)$. The posterior can be expressed as

$$p(\Theta|Y) = \frac{p(Y|\Theta)p(\Theta)}{p(Y)}. \tag{13}$$

We will use MCMC methods to generate a sequence of draws $\{\Theta^{(s)}\}_{s=1}^{n_{sim}}$ from the posterior distribution.

To generate the draws from the posterior distribution, we need to be able to numerically evaluate the prior density and the likelihood function $p(Y|\Theta)$. Since our state-space system is nonlinear, it is not possible to evaluate the likelihood function using the Kalman filter. Instead, we use a sequential Monte Carlo procedure also known as particle filter. The particle filter creates an approximation $\hat{p}(Y|\Theta)$ of the likelihood function $p(Y|\Theta)$. It has been shown in Andrieu, Doucet, and Holenstein (2010) 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.

Capturing the annual release schedule for the monthly consumption data described in Section 3.1 requires a high-dimensional state vector s_t. This creates a computational challenge for the evaluation of the likelihood function because the accuracy of particle filter approximations tends to decrease as the dimension of the latent state vector increases. In order to obtain a computationally efficient filter, we exploit the fact that our state-space model is linear and Gaussian conditional on the volatility states (h_t, h_{t-1}). We use a swarm of particles to represent the distribution of $(h_t, h_{t-1})|Y_{1:t}$ and employ the Kalman filter to characterize the conditional distribution of $s_t|(h_t, h_{t-1}, Y_{1:t})$. This idea has been used by Chen and Liu (2000) and more recently by Shephard (2013). A full description of the particle filter is provided in the appendix. We embed the likelihood approximation in a fairly standard random-walk Metropolis algorithm that is widely used in the DSGE model literature; see for instance Del Negro and Schorfheide (2010).

4 Empirical Results

The data set used in the empirical analysis is described in Section 4.1. The subsequent analysis is divided into two parts. In Section 4.3 we use consumption and dividend growth data to estimate the persistent components in conditional mean and volatility dynamics of cash flows. In Section 4.4 we include the market return and risk-free rate data in the estimation and analyze the asset pricing implications of our model.

4.1 Data

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 2011, quarterly from 1947:Q1 to 2011:Q4, and monthly from 1959:M1 to 2011:M12. We also use monthly observations of returns, dividends, and prices of the CRSP value-weighted portfolio of all stocks traded on the NYSE, AMEX, and NASDAQ. Price and dividend series are 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. A more detailed explanation of the data sources is provided in Appendix B. Growth rates of consumption and dividends are constructed by taking the first difference of the corresponding log series. The time-series span of the stock market data and the risk-free rate is from 1929:M1 to 2011:M12.

Table 1 presents descriptive statistics for aggregate consumption growth, dividend growth, aggregate stock market returns, the risk-free rate, and the log price-dividend ratio. The statistics are computed for a sample of annual observations from 1930 to 2011, a sample of monthly observations from 1929:M1 to 2011:M12, and a sample of monthly observations from 1959:M2 to 2011:M12. Consumption data is only available for the shorter of the two monthly samples. For our subsequent analysis, a few features of the data turn out to be important. First, the sample first autocorrelation function of monthly and annual consumption have different signs. Second, consumption and dividend growth are highly correlated at the low (annual) frequency but not at the high (monthly) frequency. Third, the sample standard deviations for the long monthly sample starting in 1929:M1 are larger than the sample standard deviations for the post-1958 sample.

Table 1: Descriptive Statistics - Data Moments

Annual Frequency: 1930 to 2011					
	Δc	Δd	r_m	r_f	pd
Mean	1.83	0.98	5.43	0.46	3.36
StdDev	2.19	11.24	19.98	2.78	0.43
AC1	0.48	0.21	0.01	0.72	0.90
AC2	0.18	-0.21	-0.16	0.40	0.81
AC3	-0.07	-0.15	0.01	0.31	0.75
Corr	1.00	0.56	0.12	-0.26	0.07
Monthly Frequency: 1929:M1 to 2011:M12					
	Δc	Δd	r_m	r_f	pd
Mean	-	0.09	0.45	0.04	3.36
StdDev	-	1.68	5.50	0.24	0.44
AC1	-	0.20	0.11	0.98	0.99
Monthly Frequency: 1959:M2 to 2011:M12					
	Δc	Δd	r_m	r_f	pd
Mean	0.16	0.11	0.43	0.10	3.57
StdDev	0.34	1.26	4.55	0.14	0.39
AC1	-0.16	-0.01	0.10	0.96	0.99
Corr	1.00	0.04	0.16	0.13	0.00

Notes: We report descriptive statistics for aggregate consumption growth (Δc), dividend growth (Δd), log returns of the aggregate stock market (r_m), the log risk-free rate (r_f), and log price-dividend ratio (pd). It shows mean, standard deviation, sample autocorrelations up to order three, and correlation with aggregate consumption growth. Means and standard deviations are expressed in percentage terms.

4.2 Estimation with Consumption Data Only

In this subsection we show the importance of accounting for measurement errors and that consumption by its own displays strong evidence for a persistent component. We focus on measurement error modeling in monthly consumption series. For simplicity, we only work with post-1959 monthly data and assume homoskedastic errors.

$$\begin{aligned} g_{c,t+1} &= \mu_c + x_t + \sigma\eta_{c,t+1} + \underbrace{\sigma_\epsilon\epsilon_{t+1} - \sigma_\epsilon\epsilon_t}_{\text{measurement error}} \\ x_{t+1} &= \rho x_t + \sqrt{1-\rho^2}(\varphi_x\sigma)\eta_{x,t+1} \end{aligned}$$

We set μ_c to the sample mean and fix σ and σ_ϵ to their maximum likelihood estimates, respectively. We vary the two parameters, ρ and φ_x, that govern the dynamics of x_t and evaluate the likelihood function with and without allowing for measurement errors.

Figure 1: Log-Likelihood Contour

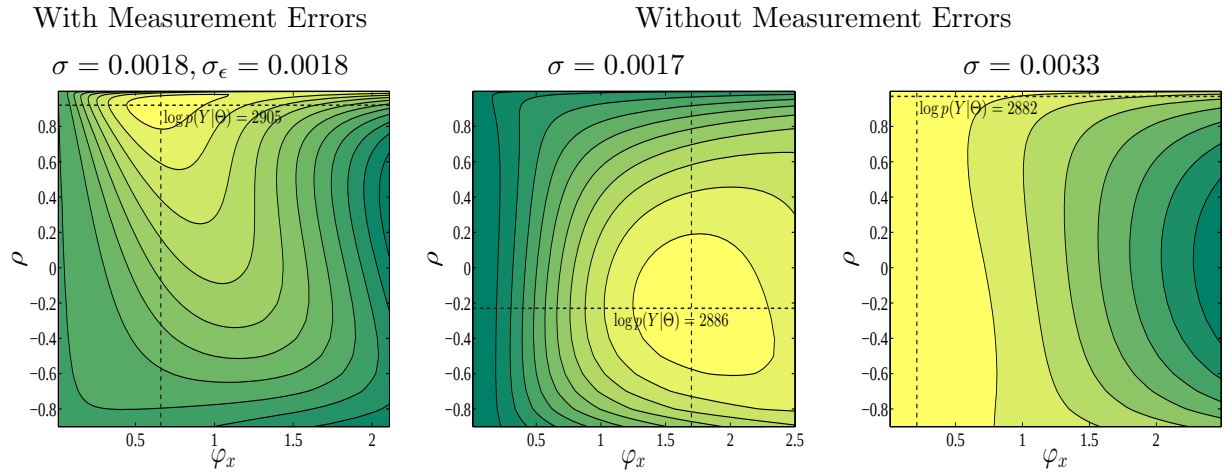

Notes: Consumption dynamics are

$$\begin{aligned} \Delta c^o_{t+1} &= \Delta c_{t+1} + \text{measurement errors} \\ \Delta c_{t+1} &= \mu_c + x_t + \sigma \eta_{c,t+1} \\ x_{t+1} &= \rho x_t + \sqrt{1-\rho^2}(\varphi_x \sigma)\eta_{x,t+1} \end{aligned}$$

where $\eta_{c,t+1}, \eta_{x,t+1} \sim N(0,1)$.

Figure 2: Sample Autocorrelation

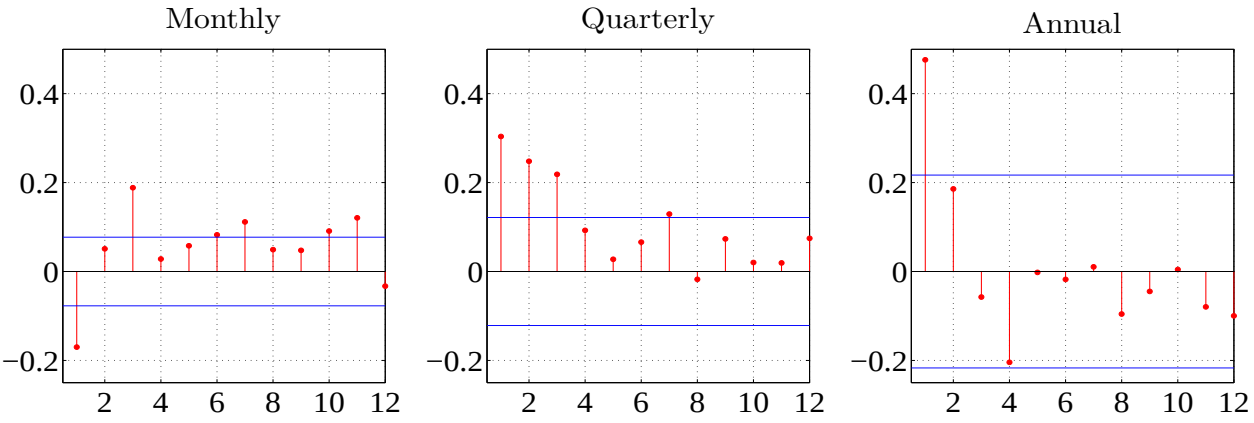

Notes: Monthly data span from 1959 to 2014, quarter from 1947 to 2014, annual from 1930 to 2014.

As we can see from Figure 1, the log-likelihood function is bimodal if measurement errors are ignored. The location of the first mode is $\rho = -0.23$ which matches the negative sample autocorrelation (see Figure 2). The location of the second mode is $\rho = 0.96$, but the log-likelihood function is flat across all values of $\rho \in (-1, 1)$. On the other-hand, when we allow for measurement errors, the log likelihood function has a very sharp peak displaying a very persistent expected consumption growth process with $\rho = 0.92$.

Next, we examine two things: (1) the informational gain achieved by using monthly as opposed to aggregated quarterly data in the estimation; and (2) the role of stochastic volatility in identifying the predictable component. Table 2, shows that sharper identification of ρ is achieved as we add in more frequent data and allow for stochastic volatility.

Table 2: Informational Gain: Posterior Distribution of ρ

	Data Frequency	Model Frequency	Measurement Errors	Stochastic Volatility	5%	50%	95%
(1)	Quarterly	Quarterly	No	No	0.421	0.649	0.817
(2)	Quarterly	Quarterly	Yes	No	0.543	0.676	0.812
(3)	Quarterly	Monthly	Yes	No	0.783	0.891	0.958
(4)	Monthly	Monthly	Yes	No	0.847	0.917	0.963
(5)	Monthly	Monthly	Yes	Yes	0.904	0.951	0.980

4.3 Estimation with Cash Flow Data Only

We begin by estimating the state-space model described in Section 3 based only on consumption and dividend growth data, dropping market returns and the risk-free rate from the measurement equation. We employ the mixed-frequency approach by utilizing annual consumption growth data from 1929 to 1959 and monthly data from 1960:M1 to 2011:M12.

Prior Distribution. We begin with a brief discussion of the prior distribution for the parameters of the cash flow process specified in (1) and (2). In general, our prior attempts 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. Percentiles of marginal prior distributions are reported in Table 3.

The prior 90% credible intervals for average annualized consumption and dividend growth range from approximately ± 7%. In view of the sample statistics reported in Table 1, this range is fairly wide and agnostic. The prior distribution for the persistence of the predictable cash flow growth component x_t is a normal distribution centered at 0.9 with a standard deviation of 0.5, truncated to the interval $(-1, 1)$. The corresponding 90% credible interval ranges from -0.1 to 0.97, encompassing values that imply *iid* cash flow growth dynamics as well as very persistent local levels. The priors for ϕ and π, parameters that determine the comovement of cash flows, are centered at zero and have large variances. $\bar{\sigma}$ is, roughly speaking, the average standard deviation of the *iid* component of consumption growth. At an annualized rate, our 90% credible interval ranges from 1.2% to 7.2%. For comparison, the sample standard deviation of annual consumption growth and annualized monthly consumption growth are approximately 2% and 4%, respectively (see Table 1).

The parameters φ_d and φ_x capture the magnitude of innovations to dividend growth and the persistent cash flow component relative to the magnitude of consumption growth innovations. The prior for the former covers the interval 0.2 to 12, whereas the prior for the latter captures the interval 0 to 0.11. Thus, *a priori* we expect dividends to be more volatile than consumption and the persistent component of cash flow growth to be much smoother than the *iid* component. Our prior interval for the persistence of the volatility processes ranges from -0.1 to 0.97 and the prior for the standard deviation of the volatility process implies that the volatility may fluctuate either relatively little, over the range of 0.67 to 1.5 times the average volatility, or substantially, over the range of 0.1 to 7 times the average volatility.

Posterior Distribution. Percentiles of the posterior distribution are also reported in Table 3. The most important result for the subsequent analysis of the asset pricing implications of the LRR model is the large estimate of ρ, the autocorrelation coefficient of the persistent cash flow component x_t. The posterior median of ρ is 0.97. Thus, according to our estimate, cash flow growth dynamics are very different from *iid* dynamics; the half-life of the persistent component is about three years; and the magnitude of the parameter estimate is quite close to the values used in the LRR literature (see Bansal, Kiku, and Yaron (2012a)).

At first glance, the large estimate of ρ may appear inconsistent with the negative sample autocorrelation of consumption growth and the near-zero autocorrelation of dividend growth at the monthly frequency reported in the third panel of Table 1. However, these sample moments confound the persistence of the "true" cash flow processes and the dynamics of the measurement errors. Our state-space framework is able to disentangle the various components of observed cash flow growth, thereby detecting a highly persistent predictable component x_t that is hidden under a layer of measurement errors. Based on our measurement-error model, we can compute the fraction

Table 3: Posterior Estimates: Cashflows Only

	Prior			Posterior				Prior			Posterior		
	Distr.	5%	95%	5%	50%	95%		Distr.	5%	95%	5%	50%	95%
	Consumption Process							Dividend Process					
μ_c	N	[-.006	.006]	.0014	.0016	.0019	μ_d	N	[-.007	.006]	.0000	.0006	.0013
ρ	N^T	[-0.08	0.97]	0.95	0.97	0.98	ϕ	N	[-13.1	13.4]	2.04	2.11	2.20
φ_x	G	[0.00	0.11]	0.17	0.20	0.22	π	N	[-1.68	1.63]	- 0.18	0.03	0.14
$\bar{\sigma}$	IG	[.001	.006]	.0021	.0024	.0026	φ_d	G	[0.22	11.9]	4.92	5.30	5.78
ρ_{h_c}	N^T	[-0.08	0.97]	.993	.995	.997	ρ_{h_d}	N^T	[-0.08	0.97]	0.83	0.89	0.94
σ_{h_c}	IG	[0.22	1.03]	0.31	0.39	0.49	σ_{h_d}	IG	[0.22	1.03]	0.47	0.53	0.61
ρ_{h_x}	N^T	[-0.08	0.97]	.979	.992	.998							
σ_{h_x}	IG	[0.22	1.03]	0.23	0.43	1.07							

Notes: We utilize the mixed-frequency approach in the estimation: For consumption we use annual data from 1929 to 1959 and monthly data from 1960:M1 to 2011:M12; we use monthly dividend growth data from 1929:M1 to 2011:M12. For consumption we adopt the measurement error model of Section 3.1. We fix φ_c in (2) at $\varphi_c = 1$. N, N^T, G, and IG denote normal, truncated (outside of the interval $(-1, 1)$) normal, gamma, and inverse gamma distributions, respectively.

of the variance of observed consumption growth that is due to measurement errors. In a constant-volatility version of our state-space model, 46% of the observed consumption growth variation at the monthly frequency is due to measurement errors. For annualized consumption growth data, this fraction drops below 1%.

The estimation results also provide strong evidence for stochastic volatility. According to the posteriors reported in Table 3, both $\sigma_{c,t}$ and $\sigma_{d,t}$ exhibit significant time variation. The posterior medians of ρ_{h_c} and ρ_{h_d} are .995 and 0.89, respectively, and the unconditional volatility standard deviations σ_{h_c} and σ_{h_d} are 0.39 and 0.53, respectively. Also, the volatility of the growth prospect component, $\sigma_{x,t}$, shows clear evidence for time variation: the posterior medians of ρ_{h_x} and σ_{h_x} are 0.992 and 0.43, respectively. It is evident that the estimation supports three independent volatility processes for consumption growth and dividend growth.

Robustness. The evidence for a persistent component in consumption and dividend growth is robust to the choice of estimation sample. We shift the beginning of our estimation sample from 1929:M1 to 1959:M1 and use only monthly data. Given that this shorter sample is dominated by the Great Moderation and does not contain the fluctuations associated with the Great Depression, this sample should be conservative in terms of providing evidence for predictable component and aggregate stochastic volatility. Interestingly, we obtain similar estimates of ρ and find that changes

in the estimates of the other parameters are generally small.[4] In all, this sample also provides strong evidence for a predictable component as well as stochastic volatility in consumption and dividends.

4.4 Estimation with Cash Flow and Asset Return Data

We now include data on market returns and the risk-free rate in the estimation of our state-space model. Recall from Section 2 that we distinguish between a benchmark model and a generalized model that allows for a shock to the time rate of preference. We will report estimates for both specifications and discuss the role played by the preference shock in fitting our observations.

Prior Distribution. The prior distribution for the parameters associated with the exogenous cash flow process are the same as the ones used in Section 4.3. Thus, we focus 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 4. 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 as reasonable in the asset pricing literature. We also use the same prior for the persistence and the innovation standard deviation of the preference shock as we did for the cash flow parameters ρ and $\bar{\sigma}$. Finally, we assume that consumption growth is measured without error at the annual frequency. We estimate measurement errors only for monthly consumption growth rates and the risk-free rates, using the same prior distributions as for $\bar{\sigma}$.

Posterior Distribution. The remaining columns of Table 4 summarize the percentiles of the posterior distribution for the parameters of the benchmark model and the generalized model. While the estimated cash flow parameters are, by and large, similar to those reported in Table 3 when asset prices are not utilized, a few noteworthy differences emerge. The estimate of ρ, the persistence of the predictable cash flow component, increases from 0.97 to 0.99 to capture part of the equity premium. The time variation in the volatility of the long-run risk innovation, σ_{h_x}, also increases, reflecting the information in asset prices about growth uncertainty. The estimate of φ_x drops from 0.20 to 0.03, which reduces the model-implied predictability of asset returns and consumption growth and brings it more in line with the data. Finally, the estimate of $\bar{\sigma}$ increases by a factor of 2 to explain the highly volatile asset prices 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

[4] For brevity these results are not displayed, but they are available upon request.

Table 4: Posterior Estimates

				Benchmark Model			Generalized Model		
		Prior		Posterior			Posterior		
	Distr.	5%	95%	5%	50%	95%	5%	50%	95%
				Preferences					
δ	B	[.9951	.9999]	.9992	.9996	.9998	.9990	.9992	.9996
ψ	G	[0.31	3.45]	1.62	1.70	1.75	1.33	1.36	1.44
γ	G	[2.74	15.45]	10.14	10.84	11.37	9.88	9.97	10.32
ρ_λ	N^T	[-0.08	0.97]	-	-	-	.935	.936	.938
σ_λ	IG	[.001	.006]	-	-	-	.0003	.0004	.0005
				Consumption					
μ_c	N	[-.006	.006]	-	.0016	-	-	.0016	-
ρ	N^T	[-0.08	0.97]	.989	.993	.994	.990	.992	.994
φ_x	G	[0.00	0.11]	0.03	0.04	0.04	0.02	0.02	0.03
$\bar{\sigma}$	IG	[.001	.006]	.004	.005	.006	.003	.004	.005
ρ_{h_c}	N^T	[-0.08	0.97]	.944	.956	.967	0.943	.946	.951
σ_{h_c}	IG	[0.22	1.03]	0.55	0.60	0.67	0.83	0.84	0.84
ρ_{h_x}	N^T	[-0.08	0.97]	.981	.990	.993	.990	.992	.994
σ_{h_x}	IG	[0.22	1.03]	0.50	0.53	0.54	0.56	0.57	0.57
				Dividend					
μ_d	N	[-.007	.006]	-	.0010	-	-	.0010	-
ϕ	N	[-13.07	13.40]	3.01	3.20	3.45	3.09	3.11	3.13
π	N	[-1.68	1.63]	1.08	1.17	1.25	1.13	1.19	1.31
φ_d	G	[0.22	11.90]	5.39	5.46	5.68	6.27	6.30	6.48
ρ_{h_d}	N^T	[-0.08	0.97]	.936	.940	.947	.939	.949	.952
σ_{h_d}	IG	[0.22	1.03]	0.44	0.45	0.46	0.56	0.57	0.57

Notes: The estimation results are based on annual consumption growth data from 1930 to 1960 and monthly consumption growth data from 1960:M1 to 2011:M12. For the other three series we use monthly data from 1929:M1 to 2011:M12. We fix $\mu_c = 0.0016$, $\mu_d = 0.0010$, and $\varphi_c = 1$ in the estimation. B, N, N^T, G, and IG are beta, normal, truncated (outside of the interval $(-1, 1)$) normal, gamma, and inverse gamma distributions, respectively.

conditional mean of cash flow growth rates as well as time variation in the volatility components. Table 4 also provides the estimated preference parameters. The IES is estimated above 1 with a relatively tight credible band. Risk aversion is estimated at 11 for the benchmark model and 10 for the generalized model.

Matching First and Second Moments. Much of the asset pricing literature, e.g., Bansal, Gallant, and Tauchen (2007), Bansal, Kiku, and Yaron (2012a), and Beeler and Campbell (2012), uses unconditional moments to calibrate or estimate model parameters and judge model fit. While these 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 conduct a posterior predictive check (see, for instance, Geweke (2005) for a textbook treatment). We use previously generated draws $\Theta^{(s)}$, $s = 1, \ldots, n_{sim}$, from the posterior distribution of the model parameters $p(\Theta|Y)$ and simulate for each $\Theta^{(s)}$ the benchmark and the generalized LLR models for 996 periods, which corresponds to the number of monthly observations in our estimation sample.[5] This leads to n_{sim} simulated trajectories, which we denote by $Y^{(s)}$. For each of these trajectories, we compute various sample moments, such as means, standard deviations, cross correlations, and autocorrelations. Suppose we denote such statistics generically by $\mathcal{S}(Y^{(s)})$. The simulations provide a characterization of the posterior predictive distribution $p(\mathcal{S}(Y^{(s)})|Y)$. Percentiles of this distribution for various sample moments are reported in Table 5. The table also lists the same moments computed from U.S. data. "Actual" sample moments that fall far into the tails of the posterior predictive distribution provide evidence for model deficiencies. The moments reported in Table 5 are computed for year-on-year cash flow growth rates. Market returns, the risk-free rate, and the price-dividend ratio are 12-month averages.

We first focus on the results from the benchmark model. Except for the first-order autocorrelation (AC1) of the risk-free rate r_f, all of the "actual" sample moments are within the 5th and the 95th percentile of the corresponding posterior predictive distribution. The variance of the posterior predictive distribution reflects the uncertainty about model parameters as well as the variability of the sample moments. The 90% credible intervals for the consumption growth and risk-free rate moments are much smaller than the intervals for the dividend growth and market-return moments, indicating that much of the uncertainty in the posterior predictive moments is due to the variability of the sample moments themselves. The high volatility of dividend growth and market returns translates into a large variability of their sample moments.

More specifically, the benchmark model replicates well the first two moments of consumption and dividend growth and their correlation. The benchmark model also generates a sizable equity risk premium with a median value of 6%. The model's return variability is about 20% with the market return being not highly autocorrelated. As in the data, the model generates both a highly variable and persistent price-dividend ratio. It is particularly noteworthy that the median and 95th

[5] To generate the simulated data, we also draw measurement errors.

Table 5: Moments of Cash Flow Growth and Asset Prices

	Data	Benchmark Model			Generalized Model		
		5%	50%	95%	5%	50%	95%
Mean (Δc)	1.83	0.91	1.89	2.79	0.88	1.93	2.86
StdDev (Δc)	2.19	1.65	2.19	2.99	1.52	2.22	3.49
AC1 (Δc)	0.48	0.09	0.32	0.56	0.08	0.33	0.57
Mean (Δd)	1.00	-2.55	1.02	4.61	-2.27	1.30	4.68
StdDev (Δd)	11.15	11.01	13.29	16.60	10.35	12.97	16.99
AC1 (Δd)	0.20	-0.19	0.03	0.23	-0.20	0.03	0.27
Corr ($\Delta c, \Delta d$)	0.55	0.12	0.32	0.51	0.13	0.34	0.56
Mean (r_m)	5.71	1.88	5.10	8.46	2.40	5.61	9.64
StdDev (r_m)	19.95	14.70	20.30	38.04	13.38	19.99	46.21
AC1 (r_m)	-0.01	-0.28	-0.06	0.17	-0.28	-0.05	0.17
Corr ($\Delta c, r_m$)	0.12	-0.03	0.18	0.39	-0.06	0.17	0.40
Mean (r_f)	0.44	-0.44	0.46	1.21	-0.39	0.67	1.49
StdDev (r_f)	2.88	2.47	2.87	3.45	1.26	1.96	4.29
AC1 (r_f)	0.64	-0.13	0.07	0.30	0.13	0.43	0.66
Mean (pd)	3.36	2.90	3.24	3.41	2.72	3.15	3.36
StdDev (pd)	0.45	0.15	0.27	0.64	0.13	0.27	0.86
AC1 (pd)	0.86	0.50	0.74	0.87	0.47	0.74	0.89

Notes: We present descriptive statistics for aggregate consumption growth (Δc), dividends growth (Δd), log returns of the aggregate stock market (r_m), 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). Cash flow growth rates are year-on-year (in percent); market returns, the risk-free rate, and the price-dividend ratio refer to 12-month averages (in percent).

percentile of the price-dividend volatility distribution are significantly larger than in other LRR calibrated models with Gaussian shocks. This feature owes in part to the fact that the models contain three volatility components with underlying log-volatility dynamics, thus accommodating some non-Gaussian features.

The sample moments implied by the generalized model are very similar to those of the benchmark model, except for the moments associated with the risk-free rate. Most notably, the benchmark model generates a slightly negative autocorrelation of the risk-free rate, whereas the generalized

Figure 3: Model-Implied Risk-Free Rate

(a) Benchmark Model

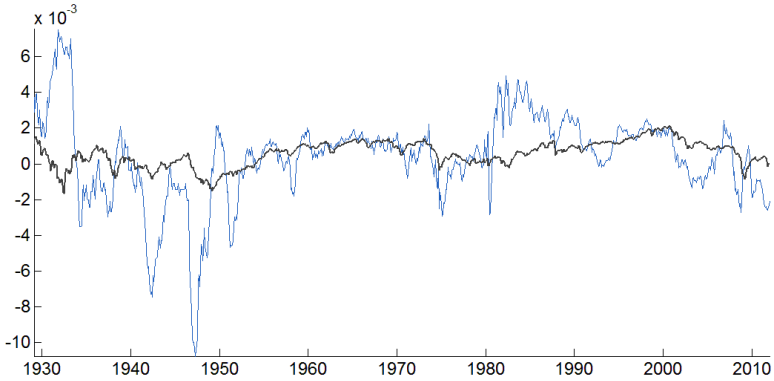

(b) Generalized Model

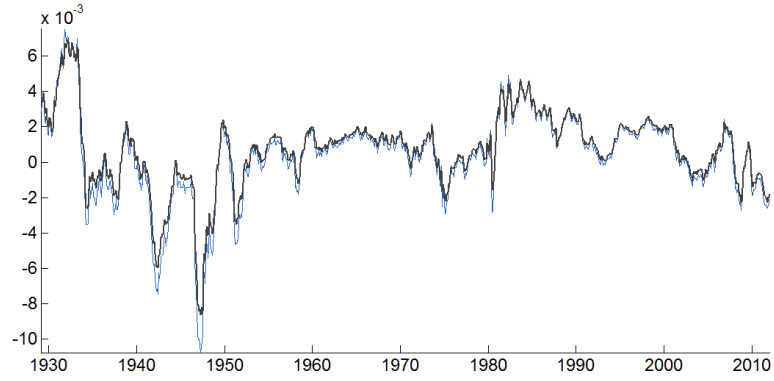

Notes: Blue lines depict the actual risk-free rate, and black lines depict the smoothed, model-implied risk-free rate without measurement errors.

model with the preference shock is able to reproduce the strongly positive serial correlation in the data.

Risk-Free Rate Dynamics. Our estimated state-space model can be used to decompose the observed risk-free rate into the "true" risk-free rate and a component that is due to measurement errors. Figure 3 overlays the actual risk rate and the smoothed "true" or model-implied risk-free rate. It is clear from the top panel of the figure that the model has difficulties generating the high volatility of the risk-free rate in the pre-1960 sample, the 1980s, and the period since 2002. The benchmark model attributes these fluctuations to measurement errors. Recall that the risk-free rate series is constructed by subtracting random-walk inflation forecasts from a nominal interest rate series, which makes the presence of measurement errors plausible. In particular, our nominal interest rate series includes several periods with negative nominal yields in the period from 1938

to 1941. The pre-1960 sample also contains periods with artificially large inflation rates, which are partly due to price adjustments following price controls after World War II. Overall, the estimated benchmark model implies that about 70-80% of the fluctuations in the risk-free rate are due to measurement errors.

The generalized model with the preference shock λ_t is able to track the risk-free rate much better than the benchmark model. By construction, λ_t generates additional fluctuations in the model-implied expected stochastic discount factor and hence the model-implied risk-free rate. The likelihood-based estimation procedure reverses this logic. Persistent movements in the observed risk-free rate suggest that λ_t fluctuated substantially between 1929 and 2011. The fraction of the fluctuations in the risk-free rate attributed to measurement errors is now much smaller. In fact, the bottom panel of Figure 3 illustrates that the difference between the observed series and the smoothed, model-implied series is now very small. This is consistent with the predictive checks reported in Table 5. Since the generalized model is more successful at tracking the observed risk-free rate, we focus on the model specification with preference shock in the remainder of this section unless otherwise noted.[6]

Smoothed Mean and Volatility States. Figure 4 depicts smoothed estimates of the predictable component x_t and the preference shock process $x_{\lambda,t}$. Since the estimate of x_t is, to a large extent, determined by the time path of consumption, the 90% credible bands 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 (see Section 4.3). 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. A visual comparison of the smoothed $x_{\lambda,t}$ process with the standardized risk-free rate in the bottom panel of Figure 4 confirms that the preference shock in the generalized model mainly helps track the observed risk-free rate.

The smoothed volatility processes are plotted in Figure 5. 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 1950 and 1965. In magnitude, this drop in volatility is much larger than a subsequent decrease around 1984, the

[6] An alternative way to interpret the preference shocks is that the model requires correlated measurement errors to capture the time series dynamics of the real risk-free rate.

Figure 4: Smoothed Mean States

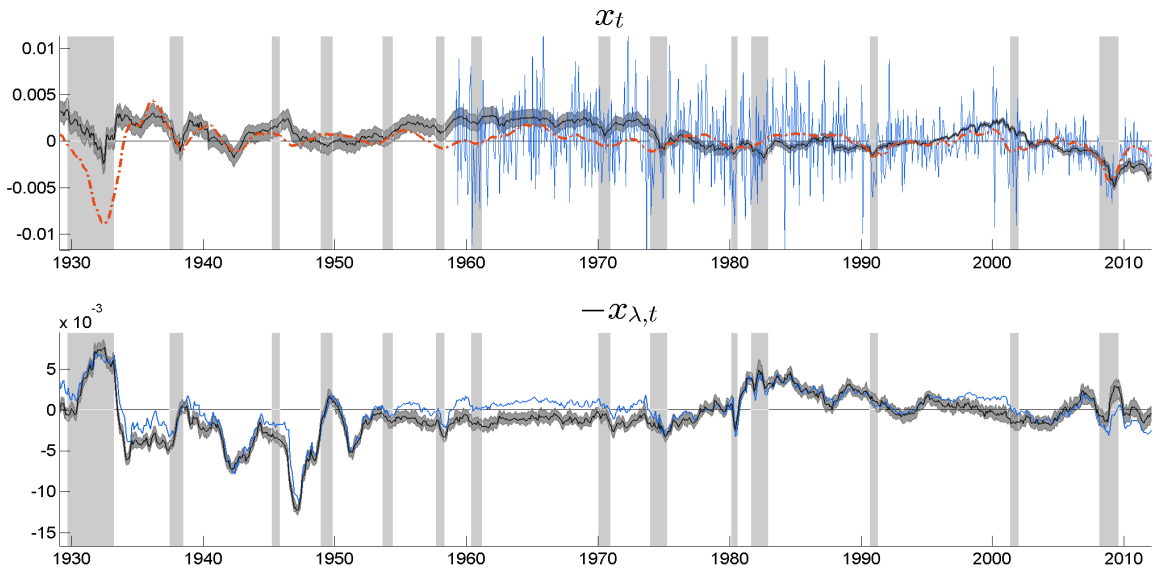

Notes: Black lines represent posterior medians of smoothed states and gray-shaded areas correspond to 90% credible intervals. In the top panel we overlay the smoothed state x_t obtained from the estimation without asset prices (red dashed line) and monthly consumption growth data (blue solid line). In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

year typically associated with the Great Moderation. The stochastic volatility process for dividend growth shows a drop around 1955, but it also features an increase in volatility starting in 2000, which is not apparent in $h_{c,t}$. Overall, the smoothed $h_{d,t}$ 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.

Determinants of Asset Price Fluctuations. After a visual inspection of the latent mean and volatility processes in Figures 4 and 5, we now examine their implications for asset prices. In equilibrium, the market returns, the risk-free rate, and the price-dividend ratios are functions of the mean and volatility states. Figure 6 depicts the contribution of various risk factors: namely, the variation in growth prospects, x_t, the preference shock, $x_{\lambda,t}$, and the conditional variability of growth prospects, $\sigma_{x,t}$, to asset price volatility. Given the posterior estimates of our state-space model, we can compute smoothed estimates of the latent asset price volatilities at every point in time. Moreover, we can also generate counterfactual volatilities by shutting down x_t, $x_{\lambda,t}$, or $\sigma_{x,t}$. The ratio of the counterfactual and the actual volatilities measures the contribution of the non-omitted risk factors. If we subtract this ratio from 1, we obtain the relative contribution of the omitted risk factor, which is shown in Figure 6.

Figure 5: Smoothed Volatility States

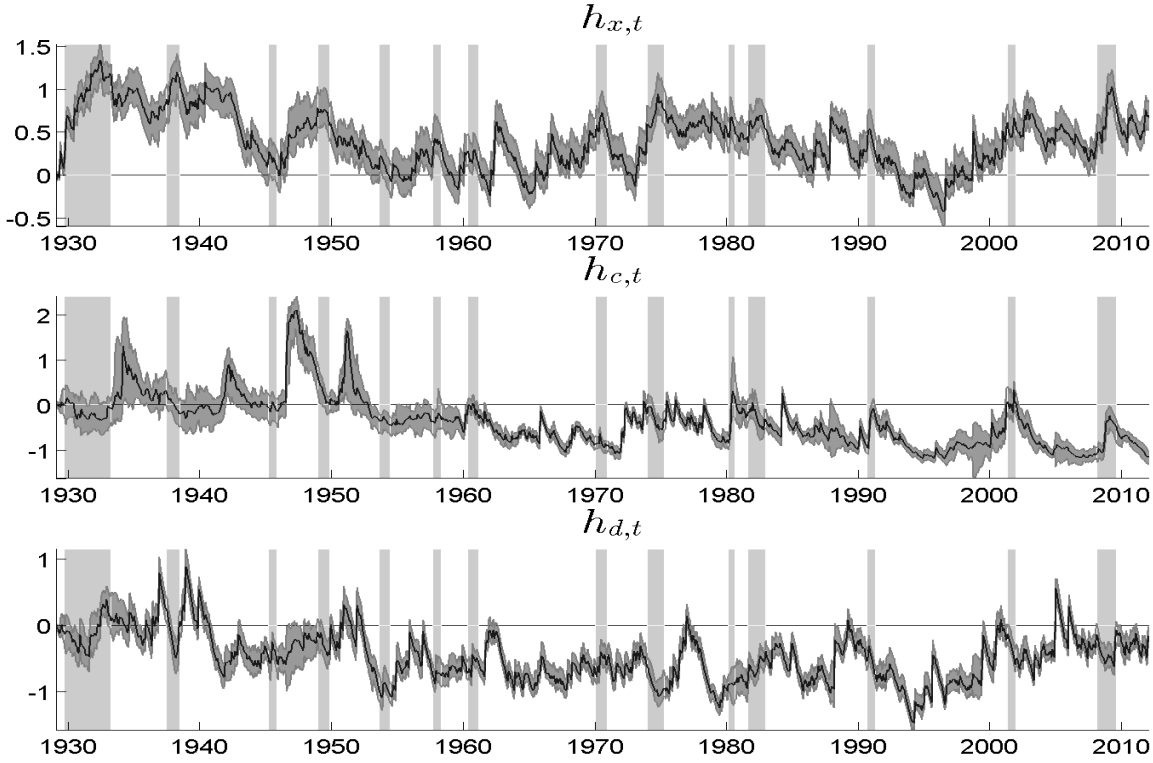

Notes: Black lines represent posterior medians of smoothed states and gray-shaded areas correspond to 90% credible intervals. Shaded bars indicate NBER recession dates.

While the preference shocks are important for the risk-free rate, they contribute very little to the variance of the price-dividend ratio and the market return. The figure shows that most of the variability of the price-dividend ratio is, in equal parts, due to the variation in x_t and $\sigma_{x,t}$. As Appendix A shows, the risk premium on the market return is barely affected by the preference shocks and consequently its variation is almost entirely attributable to the time variation in the stochastic volatility $\sigma^2_{x,t}$ and the growth prospect x_t. The remaining risk factors $\sigma^2_{c,t}$ and $\sigma^2_{d,t}$ have negligible effects (less than 1% on average) on asset price volatilities.

We assumed that in our endowment economy the preference shock is uncorrelated with 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. To assess whether our assumption of uncorrelated shocks is contradicted by the data, we computed the 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 $\Theta^{(s)}$ from the posterior distribution. The 90% posterior predictive intervals range from -0.09 to 0.03 for the correlation

Figure 6: Variance Decomposition for Market Returns and Risk-Free Rate

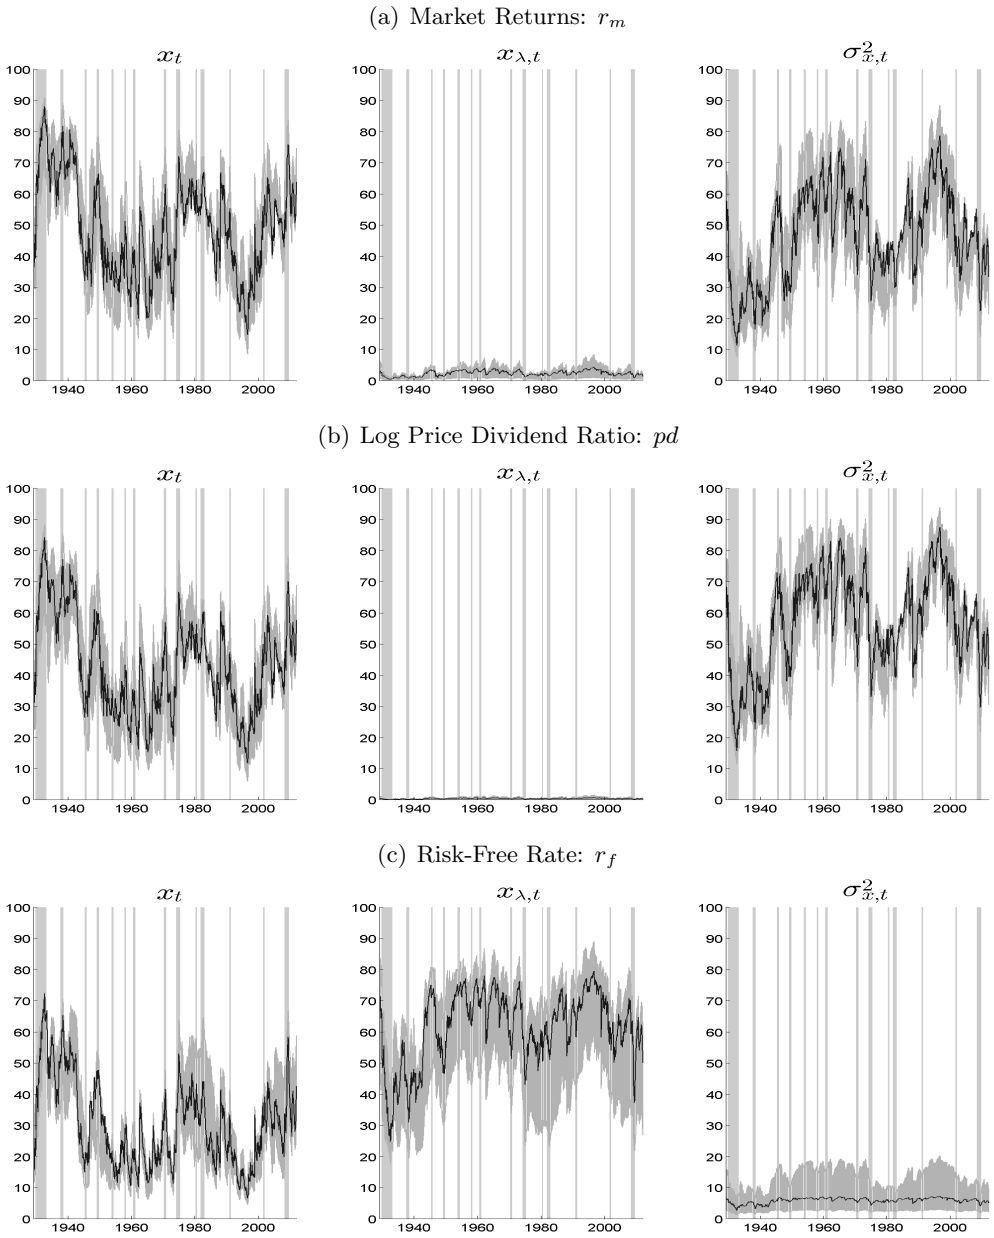

Notes: Fraction of volatility fluctuations (in percent) in the market returns, the price-dividend ratio, and the risk-free rate that is due to x_t, $x_{\lambda,t}$, and $\sigma^2_{x,t}$, respectively. We do not present the graphs for $\sigma^2_{c,t}, \sigma^2_{d,t}$ since their time-varying shares are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

between $\eta_{\lambda,t}$ and $\eta_{c,t}$ and from 0 to 0.2 for the correlation between $\eta_{\lambda,t}$ and $\eta_{x,t}$. Based on these results we conclude that there is no strong evidence that contradicts the assumption of uncorrelated preference shocks.

Figure 7: Univariate Predictability Checks

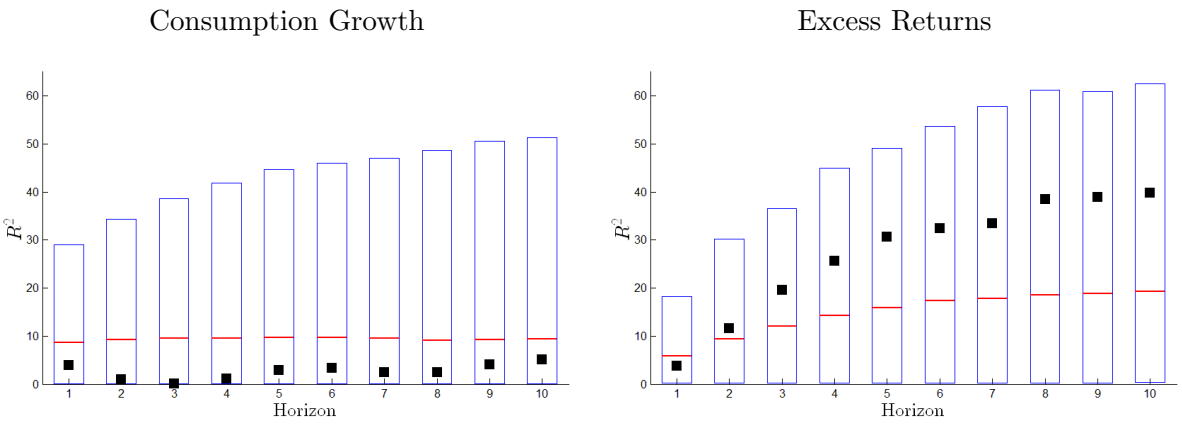

Notes: Black boxes indicate regression R^2 values from actual data. Figure also depicts medians (red lines) and 90% credible intervals (top and bottom lines of boxes) of distribution of R^2 values obtained with model-generated data.

Predictability. One aspect of the data that is often discussed in the context of asset pricing models — and in particular, in the context of models featuring long-run risks — is the low predictability of future consumption growth by the current price-dividend ratio. Another key issue in the asset pricing literature is return predictability by the price-dividend ratio (e.g., Hodrick (1992)). We address these two issues in Figure 7 where we regress cumulative consumption growth and multi-period excess returns on the current price-dividend ratio using OLS:

$$\begin{aligned} \sum_{h=1}^{H} \Delta c_{t+h} &= \alpha + \beta pd_t + \text{resid}_{t+H} \\ \sum_{h=1}^{H} (r_{m,t+h} - r_{f,t+h-1}) &= \alpha + \beta pd_t + \text{resid}_{t+H}. \end{aligned}$$

The results are presented as posterior predictive checks, similar to those in Table 5, but now depicted graphically. The statistics $\mathcal{S}(Y)$ considered are the R^2 values obtained from the two 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 contribution of parameter uncertainty to the posterior predictive distribution is negligible. The predictive intervals reflect the fact that we are repeatedly generating data from the model and computing the sample statistics $\mathcal{S}(Y)$ for each of these simulated trajectories. Finally, the small squares correspond to statistics computed from "actual" U.S. data.

The left panel of Figure 7 documents the predictability of consumption growth. While the model's median R^2 value is somewhat larger (red lines) than the corresponding data estimate, the model's

finite sample R^2 distribution for consumption growth encompasses the low data estimate. In terms of return predictability, depicted in the right panel of Figure 7, the model's median R^2 for the five-year horizon R^2 is large at 15%, with a 95 percentile value of 47% that clearly contains the data estimate. These model-implied R^2s are larger than what is typically found in models with long-run risks (e.g., Bansal, Kiku, and Yaron (2012a)) — a feature attributable to the presence of the three exponential volatility processes that allow this model specification for an improved fit.

One may wonder whether the large predictive intervals depicted in Figure 7 are an artifact of the LRR model used in our empirical analysis. To investigate this issue, we replaced the model-implied structural relationship between the price-dividend ratio and our state variables by a reduced-form relationship that was obtained by regressing the price-dividend ratio on the smoothed states using an unrestricted OLS regression. We used the Campbell-Shiller approximation to map the predicted price-dividend ratio into market returns. A predictive check based on the reduced-form relationship yields credible intervals that were of similar size as the ones depicted in Figure 7. We concluded from this exercise that it is not the LRR mechanism that generates the sampling uncertainty of the R^2 but the fact that the R^2 values are obtained from regressions in which the persistent component of the dependent variable (consumption growth or excess returns) is dominated by *iid* shocks and the right-hand-side regressor is highly persistent. In fact, Valkanov (2003) derived an asymptotic distribution of the R^2 under the assumption that the regressor follows a local-to-unity process. He shows that the goodness-of-fit measure converges to a random limit as the sample size increases. The predictive intervals in Figure 7 essentially capture the quantiles of this limit distribution because the effect of posterior parameter uncertainty on the predictive distribution is very small and our sample size is large.

It is well known that, in the data, the price-dividend ratio is very persistent, a feature that can render the aforementioned regressions spurious (see Hodrick (1992) and Stambaugh (1999)). In the model, and possibly in the data, the price-dividend ratio reflects multiple state variables. Consequently, a VAR-based predictive regression may offer a more robust characterization. As in Bansal, Kiku, and Yaron (2012a), Figure 8 displays the predictability of consumption growth and the market excess returns based on a first-order VAR that includes consumption growth, the price-dividend ratio, the real risk-free rate, and the market excess return. The first thing to note is that, with multiple predictive variables, consumption growth is *highly* predictable. The VAR provides quite a different view on consumption predictability relative to the case of using the price-dividend ratio as a univariate regressor. In particular, now consumption growth predictability at the one-year horizon is very large with an R^2 of about 55% (see also Bansal, Kiku, Shaliastovich, and Yaron (2013)). While the predictability diminishes over time, it is still nontrivial with an R^2 of 14% at the 10-year horizon. It is important to note that the model-based VAR yields very comparable

Figure 8: VAR Predictability Checks

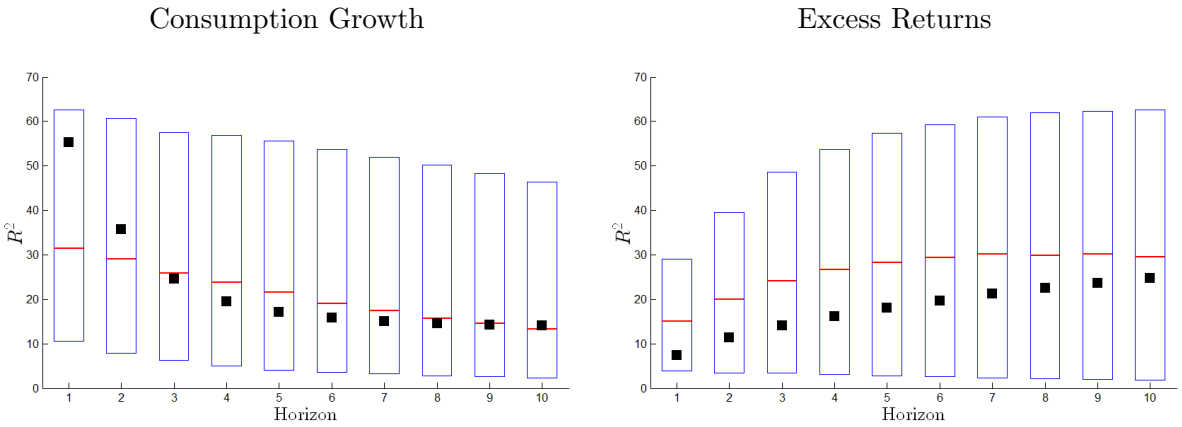

Notes: Black boxes indicate VAR R^2 values from actual data. The figure also depicts medians (red lines) and 90% credible intervals (top and bottom lines of boxes) of distribution of R^2 values obtained with model-generated data.

Figure 9: Correlation between Market Return and Growth Rates of Fundamentals

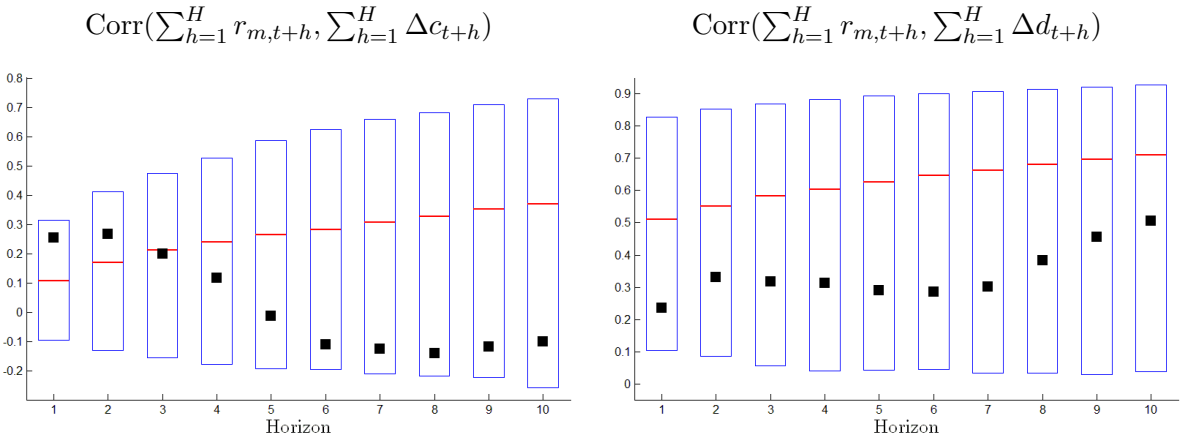

Notes: Black boxes indicate sample correlations of actual data. The figure also depicts medians (red lines) and 90% credible intervals (top and bottom lines of boxes) of distribution of sample correlations obtained with model-generated data.

results (and in fact yields a median R^2 for the one-year horizon that is somewhat *lower* than its data estimate). On the other hand, since long-horizon return predictability is highly influenced by the price-dividend ratio, the VAR-based implications for excess return predictability do not change much relative to the univariate estimates. Nonetheless, the model performs well along this dimension and its generated VAR-based R^2s are closer to their VAR data estimates, relative to the R^2s based on univariate price-dividend regressor.

Table 6: ψ (IES) from Instrumental Variables Estimation

Specification	Data	Generalized Model		
		5%	50%	95%
Δc onto r_f	-0.30	-0.45	0.56	1.50
r_f onto Δc	-0.90	-5.53	1.26	6.80

Notes: The first row provides finite sample values of the ψ from the regression $\Delta c_{t+1} = \text{constant} + \psi r_{f,t+1} + \text{resid}_{t+1}$, while the second row provides the ψ values from the regression $r_{f,t+1} = \text{constant} + \frac{1}{\psi}\Delta c_{t+1} + \text{resid}_{t+1}$. The instruments are lagged (two years) consumption growth, log price-dividend ratio, market return, and risk-free rate. The "true" ψ value in the model is 1.36 from Table 4. Regressions are implemented at an annual frequency.

One additional feature in which the generalized model performance is improved relative to the benchmark model is the correlation between long-horizon return and long-horizon dividend and consumption growth. Figure 9 presents these correlations in the data and the model. In the model, the 10-year consumption growth and 10-year return have a correlation of 0.3, but with enormous standard deviations that encompass -0.3 to 0.7, which contain the data estimate. The analogous correlations for dividend growth are 0 to 0.9 with the data at 0.5 close to the model median estimate. In the benchmark model, without the x_λ process, this last correlation will be quite a bit larger for all percentiles and would be a challenging dimension for the model.

IES Parameter. It is well understood that whether the IES parameter is above or below 1 plays a significant role in the asset pricing implication of the model. One common approach for estimating the IES has been to regress the growth rate of consumption on the risk-free rate (e.g., Hall (1988)). Bansal and Yaron (2004) and Bansal, Kiku, and Yaron (2012a) show that this regression is misspecified in the presence of stochastic volatility and leads to downward-biased estimates of the IES. Given that our estimation formally ascribes measurement errors to both consumption and the risk-free rate, we revisit the implication of this regression for inference on the IES. For completeness, we also run the reverse equation of regressing the risk-free rate on consumption growth. We use the two-year lagged consumption growth, log price-dividend ratio, market return, and risk-free rate as instrumental variables. As shown in Table 6, in both regression approaches the data based estimates are in fact negative, but lie well within the very wide 90% model-based credible band, even though in all model simulations the IES was set to its median estimate of 1.36. In totality, this evidence shows that, with the estimated levels of measurement errors, it is very difficult to precisely estimate the IES via this regression approach.

Term Structure of Real Yields. Next, we briefly discuss the model implications for the term structure of real yields. The evidence on the slope of the term structure of real bonds is mixed.

The data for the U.K., the longest sample available for real bond data, suggests the real yield curve is downward sloping (see Evans (2003)). The evidence for the U.S. suggests a slight upward sloping curve (see Gurkaynak, Sack, and Wright (2010)). However, it is well recognized that the U.S. sample is very short and that the Treasury Inflation-Protected Securities (TIPS) market is relatively illiquid. Posterior predictive distributions for average real yields are depicted in Figure 10. The right panel of the figure is based on the generalized model with preference shock, whereas we set $x_{\lambda,t} = 0$ to generate the results presented in the left panel. We show predictive distributions for the level of the 1-year yield and consider yield spreads of 5-year, 10-year, and 30-year securities relative to the 1-year bond. The benchmark model, which abstracts from preference shocks, leads to a very mild downward sloping real yield curve (model's median estimates across simulations display a decline of 50 basis points (bps) from the 1 year to the 30 year yield respectively). This feature of the long-run risk model is well recognized (see Bansal, Kiku, and Yaron (2012a)) as well as the fact that an upward nominal term structure will result if consumption growth and inflation are negatively correlated as is the case in the data (see Piazzesi and Schneider (2006) and Bansal and Shaliastovich (2013)). Our generalized model leads to a slight upward sloping real term structure (an increase of about 100 bps when comparing the 30-year to the 1-year yield).

We can use our state-space model to compute a time-path for the bond yields conditional on the smoothed states. The results for 5, 10, and 20-year yields are depicted in Figure 11. We also plot actual yields based on TIPS data obtained from the Board of Governors. Before and after the Great Recession, the model tracks the TIPS data very well, despite the fact that they have not been used to estimate the latent states. During the Great Recession, model predictions and actual data move in the opposite direction. At that time the TIPS market was very illiquid and perhaps mispriced (see, e.g. Lustig, Longstaff, and Fleckenstein (2014)), which can account for the artificial low prices and high yields.

5 Conclusion

We developed a nonlinear state-space model to capture the joint dynamics of consumption, dividend growth, and asset returns. Building on Bansal and Yaron (2004), our model consists of an economy containing a common predictable component for consumption and dividend growth and multiple stochastic volatility processes. To maximize the economic span of the data for recovering the predictable components and maximizing the frequency of data for efficiently identifying the volatility processes, we use mixed-frequency data. Our econometric framework is general enough to encompass other asset pricing models that can be written as state-space models that are linear

Figure 10: Term Structure of Real Yields

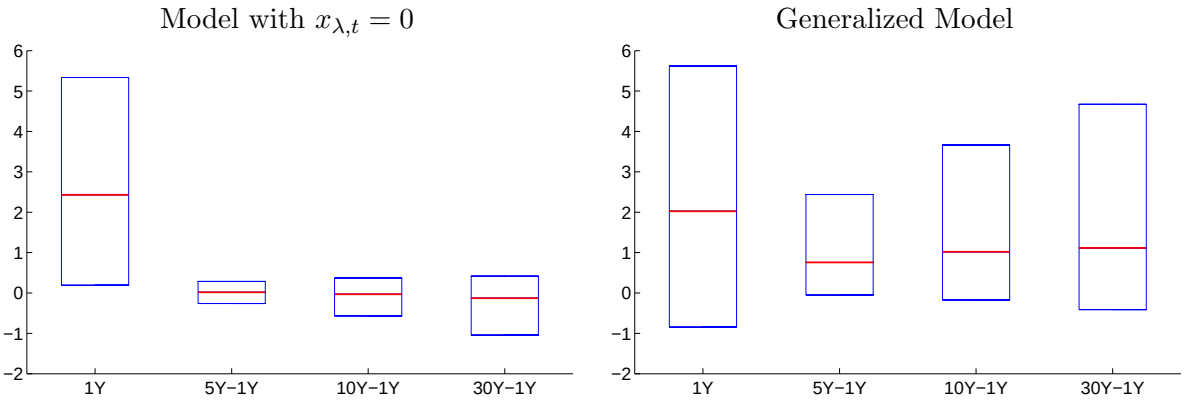

Notes: Boxes indicate 90% credible intervals for the posterior predictive distribution of the average 1-year yield and three term spreads. Horizontal red lines indicate actuals computed from TIPS data of Gurkaynak, Sack, and Wright (2010) updated by the Board of Governors.

Figure 11: Model-Implied vs. Actual Real Bond Yields

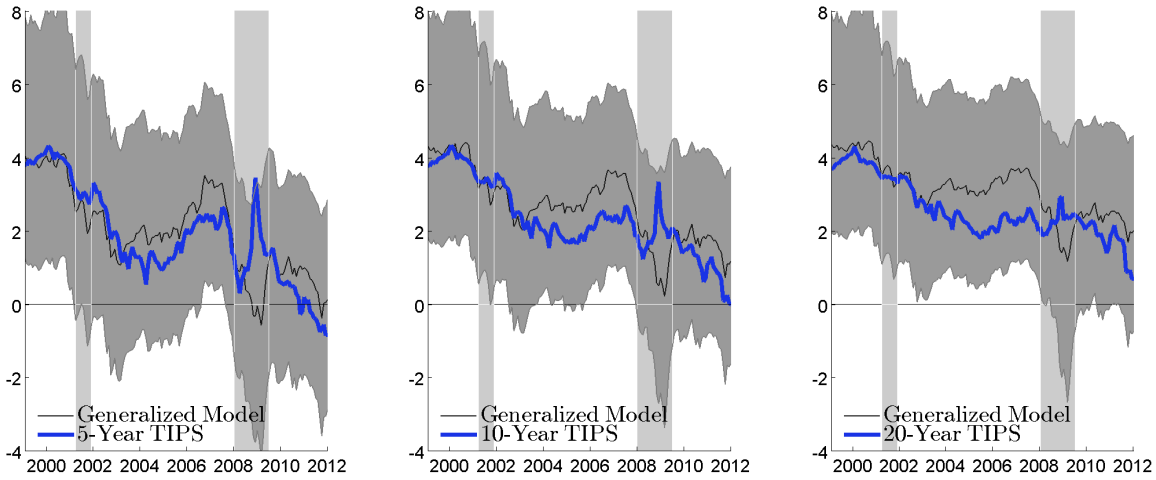

Notes: Black lines represent posterior medians of smoothed states and gray-shaded areas correspond to 90% credible intervals. Blue lines indicate TIPS data from Gurkaynak, Sack, and Wright (2010). Shaded bars indicate NBER recession dates.

conditional on the volatility states. A careful modeling of measurement errors in consumption growth reveals that the predictable cash flow component can be identified from consumption and dividend growth data only. The additional inclusion of asset prices sharpens the inference. The inclusion of two additional volatility processes improves the model fit considerably. The preference shock included in the generalized version of our model mostly captures the dynamics of the risk-free

rate, but has little effect on market returns and price-dividend ratios. Overall, the estimated model is able to capture key asset-pricing facts of the data.

References

Albuquerque, R., M. Eichenbaum, and S. Rebelo (2012): "Valuation Risk and Asset Pricing," NBER Working Paper No. 18617.

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

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, 1–33.

Aruoba, S., F. Diebold, and C. Scotti (2009): "Real-Time Measurement of Business Conditions," *Journal of Business and Economic Statistics*, 27(4), 417–427.

Bansal, R., A. Gallant, and G. Tauchen (2007): "Rational Pessimism, Rational Exuberance, and Asset Pricing Models," *Review of Economic Studies*, 74, 1005–1033.

Bansal, R., V. Khatacharian, and A. Yaron (2005): "Interpretable Asset Markets?," *European Economic Review*, 49, 531–560.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2013): "Volatility, the Macroeconomy and Asset Prices," *Journal of Finance*.

Bansal, R., D. Kiku, and A. Yaron (2012a): "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 1, 183–221.

——— (2012b): "Risks for the Long Run: Estimation with Time Aggregation," Manuscript, University of Pennsylvania and Duke University.

Bansal, R., and I. Shaliastovich (2013): "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *Review of Financial Studies*, 26, 1–33.

Bansal, R., and A. Yaron (2004): "Risks For the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.

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

Beeler, J., and J. Campbell (2012): "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 1, 141–182.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

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.

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

Chen, R., and J. Liu (2000): "Mixture Kalman Filters," *Journal of the Royal Statistical Society Series B*, 62, 493–508.

Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen (2003): "Alternative Models for Stock Price Dynamics," *Journal of Econometrics*, 116, 225–257.

Del Negro, M., and F. Schorfheide (2010): "Bayesian Macroeconometrics," in *Handbook of Bayesian Econometrics*, ed. by H. K. van Dijk, G. Koop, and J. Geweke. Oxford University Press.

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.

Evans, M. (2003): "Real Risk, Inflation Risk, and the Term Structure," *The Economic Journal*, 113, 345–389.

Fernández-Villaverde, J., and J. Rubio-Ramírez (2011): "Macroeconomics and Volatility: Data, Models, and Estimation," in D. Acemoglu, M. Arellano and E. Deckel (eds.), *Advances in Economics and Econometrics: Theory and Applications*, Tenth World Congress of the Econometric Society, Cambridge University Press.

Fernández-Villaverde, J., and J. F. Rubio-Ramírez (2007): "Estimating Macroeconomic Models: A Likelihood Approach," *Review of Economic Studies*, 74(4), 1059–1087.

Geweke, J. (2005): *Contemporary Bayesian Econometrics and Statistics*. New Jersey: John Wiley & Sons.

Gurkaynak, R., B. Sack, and J. Wright (2010): "The TIPS Yield Curve and Inflation Compensation," *American Economic Journal: Macroeconomics*, 2(1), 70–92.

Hall, R. (1988): "Intertemporal Substitution in Consumption," *Journal of Political Economy*, 96, 339–357.

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

Hodrick, R. (1992): "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies*, 5, 357–386.

Johannes, M., L. Lochstoer, and Y. Mou (2013): "Learning about Consumption Dynamics," Manuscript, Columbia Business School.

Lustig, H., F. Longstaff, and M. Fleckenstein (2014): "The TIPS-Treasury Bond Puzzle," *Journal of Finance*, forthcoming.

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.

Piazzesi, M., and M. Schneider (2006): "Equilibrium Yield Curves," *NBER Macro Annual*, pp. 389–442.

Pitt, M., and N. Shephard (1999): "Filtering via Simulation: Auxiliary Particle Filters," *Journal of the American Statistical Association*, 94, 590–599.

Schorfheide, F., and D. Song (2012): "Real-Time Forecasting with a Mixed-Frequency VAR," Federal Reserve Bank of Minneapolis Working Paper 701.

Shephard, N. (2013): "Martingale Unobserved Component Models," Manuscript, University of Oxford.

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

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

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.

Appendix

A Solving the Long-Run Risks Model

This section provides solutions for the consumption and dividend claims for the endowment process:

$$\begin{aligned}
g_{c,t+1} &= \mu_c + x_t + \sigma_{c,t}\eta_{c,t+1} \\
g_{d,t+1} &= \mu_d + \phi x_t + \pi\sigma_{c,t}\eta_{c,t+1} + \sigma_{d,t}\eta_{d,t+1} \\
x_{t+1} &= \rho x_t + \sigma_{x,t}\eta_{x,t+1} \\
x_{\lambda,t+1} &= \rho_\lambda x_{\lambda,t} + \sigma_\lambda \eta_{\lambda,t+1} \\
\sigma^2_{c,t+1} &= (1-\nu_c)(\varphi_c\bar{\sigma})^2 + \nu_c\sigma^2_{c,t} + \sigma_{w_c} w_{c,t+1} \\
\sigma^2_{x,t+1} &= (1-\nu_x)(\varphi_x\bar{\sigma})^2 + \nu_x\sigma^2_{x,t} + \sigma_{w_x} w_{x,t+1} \\
\sigma^2_{d,t+1} &= (1-\nu_d)(\varphi_d\bar{\sigma})^2 + \nu_d\sigma^2_{d,t} + \sigma_{w_d} w_{d,t+1} \\
&\quad \eta_{i,t+1}, \eta_{\lambda,t+1}, w_{i,t+1} \sim N(0,1), \quad i \in \{c,x,d\}.
\end{aligned} \tag{A.1}$$

The Euler equation for the economy is

$$E_t\left[\exp\left(m_{t+1} + r_{i,t+1}\right)\right] = 1, \quad i \in \{c, m\}, \tag{A.2}$$

where

$$m_{t+1} = \theta \log\delta + \theta x_{\lambda,t+1} - \frac{\theta}{\psi} g_{c,t+1} + (\theta - 1) r_{c,t+1} \tag{A.3}$$

is the log of the real stochastic discount factor (SDF), $r_{c,t+1}$ is the log return on the consumption claim, and $r_{m,t+1}$ is the log market return. (A.3) is derived in Section A.6 below. Returns are given by the approximation of Campbell and Shiller (1988a):

$$\begin{aligned}
r_{c,t+1} &= \kappa_0 + \kappa_1 pc_{t+1} - pc_t + g_{c,t+1} \\
r_{m,t+1} &= \kappa_{0,m} + \kappa_{1,m} pd_{t+1} - pd_t + g_{d,t+1}.
\end{aligned} \tag{A.4}$$

The risk premium on any asset is

$$E_t(r_{i,t+1} - r_{f,t}) + \frac{1}{2} Var_t(r_{i,t+1}) = -Cov_t(m_{t+1}, r_{i,t+1}). \tag{A.5}$$

In Section A.1 we solve for the law of motion for the return on the consumption claim, $r_{c,t+1}$. In Section A.2 we solve for the law of motion for the market return, $r_{m,t+1}$. The risk-free rate is derived in Section A.3. All three solutions depend on linearization parameters that are derived in Section A.4. Finally, as mentioned above, the SDF is derived in Section A.6.

A.1 Consumption Claim

In order to derive the dynamics of asset prices, we rely on approximate analytical solutions. Specifically, we conjecture that the price-consumption ratio follows

$$pc_t = A_0 + A_1 x_t + A_{1,\lambda} x_{\lambda,t} + A_{2,c}\sigma_{c,t}^2 + A_{2,x}\sigma_{x,t}^2 \tag{A.6}$$

and solve for A's using (A.1), (A.2), (A.4), and (A.6).

From (A.1), (A.4), and (A.6)

$$\begin{aligned} r_{c,t+1} &= \left\{\kappa_0 + A_0(\kappa_1 - 1) + \mu_c + \kappa_1 A_{2,x}(1-\nu_x)(\varphi_x\bar{\sigma})^2 + \kappa_1 A_{2,c}(1-\nu_c)(\varphi_c\bar{\sigma})^2\right\} \\ &+ \frac{1}{\psi}x_t + A_{1,\lambda}(\kappa_1\rho_\lambda - 1)x_{\lambda,t} + A_{2,x}(\kappa_1\nu_x - 1)\sigma_{x,t}^2 + A_{2,c}(\kappa_1\nu_c - 1)\sigma_{c,t}^2 \\ &+ \sigma_{c,t}\eta_{c,t+1} + \kappa_1 A_1\sigma_{x,t}\eta_{x,t+1} + \kappa_1 A_{1,\lambda}\sigma_\lambda\eta_{\lambda,t+1} + \kappa_1 A_{2,x}\sigma_{w_x}w_{x,t+1} + \kappa_1 A_{2,c}\sigma_{w_c}w_{c,t+1} \end{aligned} \tag{A.7}$$

and from (A.1), (A.2), (A.4), and (A.6)

$$\begin{aligned} m_{t+1} &= (\theta-1)\left\{\kappa_0 + A_0(\kappa_1 - 1) + \kappa_1 A_{2,x}(1-\nu_x)(\varphi_x\bar{\sigma})^2 + \kappa_1 A_{2,c}(1-\nu_c)(\varphi_c\bar{\sigma})^2\right\} \\ &- \gamma\mu + \theta\log\delta - \frac{1}{\psi}x_t + \rho_\lambda x_{\lambda,t} + (\theta-1)A_{2,x}(\kappa_1\nu_x - 1)\sigma_{x,t}^2 + (\theta-1)A_{2,c}(\kappa_1\nu_c - 1)\sigma_{c,t}^2 \\ &- \gamma\sigma_{c,t}\eta_{c,t+1} + (\theta-1)\kappa_1 A_1\sigma_{x,t}\eta_{x,t+1} + \left\{(\theta-1)\kappa_1 A_{1,\lambda} + \theta\right\}\sigma_\lambda\eta_{\lambda,t+1} \\ &+ (\theta-1)\kappa_1 A_{2,x}\sigma_{w_x}w_{x,t+1} + (\theta-1)\kappa_1 A_{2,c}\sigma_{w_c}w_{c,t+1}. \end{aligned} \tag{A.8}$$

The solutions for A's that describe the dynamics of the price-consumption ratio are determined from

$$E_t\left[m_{t+1} + r_{c,t+1}\right] + \frac{1}{2}Var_t\left[m_{t+1} + r_{c,t+1}\right] = 0$$

and they are

$$A_1 = \frac{1-\frac{1}{\psi}}{1-\kappa_1\rho}, \quad A_{1,\lambda} = \frac{\rho_\lambda}{1-\kappa_1\rho_\lambda}, \quad A_{2,x} = \frac{\frac{\theta}{2}(\kappa_1 A_1)^2}{1-\kappa_1\nu_x}, \quad A_{2,c} = \frac{\frac{\theta}{2}(1-\frac{1}{\psi})^2}{1-\kappa_1\nu_c} \tag{A.9}$$

and $A_0 = \frac{A_0^1 + A_0^2}{1-\kappa_1}$, where

$$\begin{aligned} A_0^1 &= \log\delta + \kappa_0 + \mu(1-\frac{1}{\psi}) + \kappa_1 A_{2,x}(1-\nu_x)(\varphi_x\bar{\sigma})^2 + \kappa_1 A_{2,c}(1-\nu_c)(\varphi_c\bar{\sigma})^2 \\ A_0^2 &= \frac{\theta}{2}\left\{(\kappa_1 A_{1,\lambda} + 1)^2\sigma_\lambda^2 + (\kappa_1 A_{2,x}\sigma_{w_x})^2 + (\kappa_1 A_{2,c}\sigma_{w_c})^2\right\}. \end{aligned}$$

For convenience, (A.8) can be rewritten as

$$m_{t+1} - E_t[m_{t+1}] = \lambda_c\sigma_{c,t}\eta_{c,t+1} + \lambda_x\sigma_{x,t}\eta_{x,t+1} + \lambda_\lambda\sigma_\lambda\eta_{\lambda,t+1} + \lambda_{w_x}\sigma_{w_x}w_{x,t+1} + \lambda_{w_c}\sigma_{w_c}w_{c,t+1}.$$

Note that λs represent the market price of risk for each source of risk. To be specific,

$$\begin{aligned}\lambda_c &= -\gamma, \quad \lambda_x = -(\gamma - \frac{1}{\psi})\frac{\kappa_1}{1-\kappa_1\rho}, \quad \lambda_\lambda = \frac{\theta - \kappa_1\rho_\lambda}{1-\kappa_1\rho_\lambda}, \\ \lambda_{w_x} &= -\frac{\theta(\gamma-\frac{1}{\psi})(1-\frac{1}{\psi})\kappa_1}{2(1-\kappa_1\nu_x)}(\frac{\kappa_1}{1-\kappa_1\rho})^2, \quad \lambda_{w_c} = -\frac{\theta(\gamma-\frac{1}{\psi})(1-\frac{1}{\psi})\kappa_1}{2(1-\kappa_1\nu_c)}.\end{aligned} \tag{A.10}$$

Similarly, rewrite (A.7) as

$$r_{c,t+1} - E_t[r_{c,t+1}] = -\beta_{c,c}\sigma_{c,t}\eta_{c,t+1} - \beta_{c,x}\sigma_{x,t}\eta_{x,t+1} - \beta_{c,\lambda}\sigma_\lambda\eta_{\lambda,t+1} - \beta_{c,w_x}\sigma_{w_x}w_{x,t+1} - \beta_{c,w_c}\sigma_{w_c}w_{c,t+1}$$

where

$$\beta_{c,c} = -1, \quad \beta_{c,x} = -\kappa_1 A_1, \quad \beta_{c,\lambda} = -\kappa_1 A_{1,\lambda}, \quad \beta_{c,w_x} = -\kappa_1 A_{2,x}, \quad \beta_{c,w_c} = -\kappa_1 A_{2,c}. \tag{A.11}$$

The risk premium for the consumption claim is

$$\begin{aligned}E_t(r_{c,t+1} - r_{f,t}) + \frac{1}{2}Var_t(r_{c,t+1}) &= -Cov_t(m_{t+1}, r_{c,t+1}) \\ &= \beta_{c,x}\lambda_x\sigma^2_{x,t} + \beta_{c,c}\lambda_c\sigma^2_{c,t} + \beta_{c,\lambda}\lambda_\lambda\sigma^2_\lambda + \beta_{c,w_x}\lambda_{w_x}\sigma^2_{w_x} + \beta_{c,w_c}\lambda_{w_c}\sigma^2_{w_c}.\end{aligned} \tag{A.12}$$

A.2 Market Return

Similarly, using the conjectured solution to the price-dividend ratio

$$pd_t = A_{0,m} + A_{1,m}x_t + A_{1,\lambda,m}x_{\lambda,t} + A_{2,x,m}\sigma^2_{x,t} + A_{2,c,m}\sigma^2_{c,t} + A_{2,d,m}\sigma^2_{d,t} \tag{A.13}$$

the market return can be expressed as

$$\begin{aligned}r_{m,t+1} &= \kappa_{0,m} + A_{0,m}(\kappa_{1,m}-1) + \mu_d + \kappa_{1,m}A_{2,x,m}(1-\nu_x)(\varphi_x\bar{\sigma})^2 \\ &+ \kappa_{1,m}A_{2,c,m}(1-\nu_c)(\varphi_c\bar{\sigma})^2 + \kappa_{1,m}A_{2,d,m}(1-\nu_d)(\varphi_d\bar{\sigma})^2 + \{\phi + A_{1,m}(\kappa_{1,m}\rho - 1)\}x_t \\ &+ (\kappa_{1,m}\rho_\lambda - 1)A_{1,\lambda,m}x_{\lambda,t} + A_{2,x,m}(\kappa_{1,m}\nu_x - 1)\sigma^2_{x,t} + A_{2,c,m}(\kappa_{1,m}\nu_c - 1)\sigma^2_{c,t} \\ &+ A_{2,d,m}(\kappa_{1,m}\nu_d - 1)\sigma^2_{d,t} + \pi\sigma_{c,t}\eta_{c,t+1} + \sigma_{d,t}\eta_{d,t+1} + \kappa_{1,m}A_{1,m}\sigma_{x,t}\eta_{x,t+1} + \kappa_{1,m}A_{1,\lambda,m}\sigma_\lambda\eta_{\lambda,t+1} \\ &+ \kappa_{1,m}A_{2,x,m}\sigma_{w_x}w_{x,t+1} + \kappa_{1,m}A_{2,c,m}\sigma_{w_c}w_{c,t+1} + \kappa_{1,m}A_{2,d,m}\sigma_{w_d}w_{d,t+1}.\end{aligned} \tag{A.14}$$

Given the solution for A's, A_m's can be derived as follows:

$$\begin{aligned}
A_{0,m} &= \frac{A_{0,m}^{1st} + A_{0,m}^{2nd}}{1-\kappa_{1,m}} \\
A_{1,m} &= \frac{\phi - \frac{1}{\psi}}{1-\kappa_{1,m}\rho} \\
A_{1,\lambda,m} &= \frac{\rho_\lambda}{1-\kappa_{1,m}\rho_\lambda} \\
A_{2,x,m} &= \frac{\frac{1}{2}\left\{(\theta-1)\kappa_1 A_1 + \kappa_{1,m}A_{1,m}\right\}^2 + (\theta-1)(\kappa_1\nu_x - 1)A_{2,x}}{1-\kappa_{1,m}\nu_x} \\
A_{2,c,m} &= \frac{\frac{1}{2}(\pi-\gamma)^2 + (\theta-1)(\kappa_1\nu_c - 1)A_{2,c}}{1-\kappa_{1,m}\nu_c} \\
A_{2,d,m} &= \frac{\frac{1}{2}}{1-\kappa_{1,m}\nu_d},
\end{aligned} \tag{A.15}$$

where

$$\begin{aligned}
A_{0,m}^{1st} &= \theta\log\delta + (\theta-1)\left\{\kappa_0 + A_0(\kappa_1-1) + \kappa_1 A_{2,x}(1-\nu_x)(\varphi_x\bar{\sigma})^2 + \kappa_1 A_{2,c}(1-\nu_c)(\varphi_c\bar{\sigma})^2\right\} \\
&- \gamma\mu + \kappa_{0,m} + \mu_d + \kappa_{1,m}A_{2,x,m}(1-\nu_x)(\varphi_x\bar{\sigma})^2 + \kappa_{1,m}A_{2,c,m}(1-\nu_c)(\varphi_c\bar{\sigma})^2 \\
&+ \kappa_{1,m}A_{2,d,m}(1-\nu_d)(\varphi_d\bar{\sigma})^2 \\
A_{0,m}^{2nd} &= \frac{1}{2}\left(\kappa_{1,m}A_{2,x,m}\sigma_{w_x} + (\theta-1)\kappa_1 A_{2,x}\sigma_{w_x}\right)^2 + \frac{1}{2}\left(\kappa_{1,m}A_{2,c,m}\sigma_{w_c} + (\theta-1)\kappa_1 A_{2,c}\sigma_{w_c}\right)^2 \\
&+ \frac{1}{2}\left(\kappa_{1,m}A_{2,d,m}\sigma_{w_d}\right)^2 + \frac{1}{2}\left(\kappa_{1,m}A_{1,\lambda,m}\sigma_\lambda + (\theta-1)\kappa_1 A_{1,\lambda}\sigma_\lambda + \theta\sigma_\lambda\right)^2.
\end{aligned}$$

Rewrite market-return equation (A.14) as

$$r_{m,t+1} - E_t[r_{m,t+1}] = -\beta_{m,c}\sigma_{c,t}\eta_{c,t+1} - \beta_{m,x}\sigma_{x,t}\eta_{x,t+1} - \beta_{m,\lambda}\sigma_\lambda\eta_{\lambda,t+1} - \beta_{m,w_x}\sigma_{w_x}w_{x,t+1} - \beta_{m,w_c}\sigma_{w_c}w_{c,t+1},$$

where

$$\begin{aligned}
\beta_{m,c} &= -\pi, \quad \beta_{m,x} = -\kappa_{1,m}A_{1,m}, \quad \beta_{m,\lambda} = -\kappa_{1,m}A_{1,\lambda,m}, \\
\beta_{m,w_x} &= -\kappa_{1,m}A_{2,x,m}, \quad \beta_{m,w_c} = -\kappa_{1,m}A_{2,c,m}.
\end{aligned} \tag{A.16}$$

The risk premium for the dividend claim is

$$\begin{aligned}
E_t(r_{m,t+1} - r_{f,t}) + \frac{1}{2}Var_t(r_{m,t+1}) &= -Cov_t(m_{t+1}, r_{m,t+1}) \\
&= \beta_{m,x}\lambda_x\sigma_{x,t}^2 + \beta_{m,c}\lambda_c\sigma_{c,t}^2 + \beta_{m,\lambda}\lambda_\lambda\sigma_\lambda^2 + \beta_{m,w_x}\lambda_{w_x}\sigma_{w_x}^2 + \beta_{m,w_c}\lambda_{w_c}\sigma_{w_c}^2.
\end{aligned} \tag{A.17}$$

A.3 Risk-Free Rate

The model-driven equation for the risk-free rate is

$$\begin{aligned} r_{f,t} &= -E_t\left[m_{t+1}\right] - \frac{1}{2}var_t\left[m_{t+1}\right] \\ &= -\theta\log\delta - E_t\left[x_{\lambda,t+1}\right] + \frac{\theta}{\psi}E_t\left[g_{c,t+1}\right] + (1-\theta)E_t\left[r_{c,t+1}\right] - \frac{1}{2}var_t\left[m_{t+1}\right]. \end{aligned} \tag{A.18}$$

Subtract $(1-\theta)r_{f,t}$ from both sides and divide by θ,

$$r_{f,t} = -\log\delta - \frac{1}{\theta}E_t\left[x_{\lambda,t+1}\right] + \frac{1}{\psi}E_t\left[g_{c,t+1}\right] + \frac{(1-\theta)}{\theta}E_t\left[r_{c,t+1} - r_{f,t}\right] - \frac{1}{2\theta}var_t\left[m_{t+1}\right] \tag{A.19}$$

From (A.1) and (A.8)

$$r_{f,t} = B_0 + B_1 x_t + B_{1,\lambda}x_{\lambda,t} + B_{2,x}\sigma^2_{x,t} + B_{2,c}\sigma^2_{c,t},$$

where

$$B_1 = \frac{1}{\psi}, \quad B_{1,\lambda} = -\rho_\lambda, \quad B_{2,x} = -\frac{(1-\frac{1}{\psi})(\gamma-\frac{1}{\psi})\kappa_1^2}{2(1-\kappa_1\rho)^2}, \quad B_{2,c} = -\frac{1}{2}(\frac{\gamma-1}{\psi}+\gamma) \tag{A.20}$$

and

$$\begin{aligned} B_0 &= -\theta\log\delta - (\theta-1)\left\{\kappa_0 + (\kappa_1-1)A_0 + \kappa_1 A_{2,x}(1-\nu_x)(\varphi_x\bar{\sigma})^2 + \kappa_1 A_{2,c}(1-\nu_c)(\varphi_c\bar{\sigma})^2\right\} \\ &+ \gamma\mu - \frac{1}{2}\left\{(\theta-1)\kappa_1 A_{2,x}\sigma_{w_x}\right\}^2 - \frac{1}{2}\left\{(\theta-1)\kappa_1 A_{2,c}\sigma_{w_c}\right\}^2 - \frac{1}{2}\left\{((\theta-1)\kappa_1 A_{1,\lambda}+\theta)^2\sigma^2_\lambda\right\}. \end{aligned}$$

A.4 Linearization Parameters

For any asset, the linearization parameters are determined endogenously by the following system of equations:

$$\begin{aligned} \bar{pd}_i &= A_{0,i}(\bar{pd}_i) + \sum_{j\in\{c,x,d\}} A_{2,i,j}(\bar{pd}_i)\times(\varphi_j\bar{\sigma})^2 \\ \kappa_{1,i} &= \frac{\exp(\bar{pd}_i)}{1+\exp(\bar{pd}_i)} \\ \kappa_{0,i} &= \log(1+\exp(\bar{pd}_i)) - \kappa_{1,i}\bar{pd}_i. \end{aligned}$$

The solution is determined numerically by iteration until reaching a fixed point of $\bar{pd}_i$.

A.5 Zero-Coupon Real Bonds

Let $p_{n,t}$ be the log t-price of an n-period zero-coupon real bond. Conjecture that $p_{n,t}$ is a linear function of state variables

$$p_{n,t} = C_{n,0} + C_{n,1}x_t + C_{n,1\lambda}x_{\lambda,t} + C_{n,2x}\sigma^2_{x,t} + C_{n,2c}\sigma^2_{c,t}. \tag{A.21}$$

The pricing equation implies

$$p_{n,t} = E_t\left[p_{n-1,t+1} + m_{t+1}\right] + \frac{1}{2}Var_t\left[p_{n-1,t+1} + m_{t+1}\right]. \tag{A.22}$$

The coefficients of the pricing equation are expressed recursively as

$$\begin{aligned}
C_{n,1} &= C_{n-1,1}\rho - \frac{1}{\psi} \qquad (A.23)\\
C_{n,1\lambda} &= C_{n-1,1\lambda}\rho_\lambda + \rho_\lambda\\
C_{n,2x} &= C_{n-1,2x}\nu_x + (\theta-1)A_{2,x}(\kappa_1\nu_x - 1) + \frac{1}{2}\left\{C_{n-1,1} + (\theta-1)\kappa_1 A_1\right\}^2\\
C_{n,2c} &= C_{n-1,2c}\nu_c + (\theta-1)A_{2,c}(\kappa_1\nu_c - 1) + \frac{1}{2}\gamma^2\\
C_{n,0} &= \theta\log\delta + (\theta-1)\left\{\kappa_0 + (\kappa_1-1)A_0 + \kappa_1 A_{2,x}(1-\nu_x)(\varphi_x\bar{\sigma})^2 + \kappa_1 A_{2,c}(1-\nu_c)(\varphi_c\bar{\sigma})^2\right\}\\
&\quad -\gamma\mu + C_{n-1,0} + C_{n-1,2x}(1-\nu_x)(\varphi_x\bar{\sigma})^2 + C_{n-1,2c}(1-\nu_c)(\varphi_c\bar{\sigma})^2\\
&\quad +\frac{1}{2}\left(C_{n-1,1\lambda} + \left\{(\theta-1)\kappa_1 A_{1,\lambda} + \theta\right\}\right)^2\sigma^2_\lambda + \frac{1}{2}\left((\theta-1)\kappa_1 A_{2,x} + C_{n-1,2x}\right)^2\sigma^2_{w_x}\\
&\quad +\frac{1}{2}\left((\theta-1)\kappa_1 A_{2,c} + C_{n-1,2c}\right)^2\sigma^2_{w_c}
\end{aligned}$$

with initial conditions that $C_{0,1} = C_{0,1\lambda} = C_{0,2x} = C_{0,2c} = C_{0,0} = 0$. However, in order to develop economic intuition, it is useful to express them in a non-recursive fashion:

$$\begin{aligned}
C_{n,1} &= -\frac{1}{\psi}\frac{(1-\rho^n)}{(1-\rho)}, \quad n \geq 1\\
C_{n,1\lambda} &= \rho_\lambda\frac{(1-\rho_\lambda^n)}{(1-\rho_\lambda)}, \quad n \geq 1\\
C_{n,2x} &= \left((\theta-1)A_{2,x}(\kappa_1\nu_x - 1) + \frac{1}{2}\left\{-\frac{1}{\psi}\frac{(1-\rho^{n-1})}{(1-\rho)} + (\theta-1)\kappa_1 A_1\right\}^2\right)\frac{(1-\nu_x^n)}{(1-\nu_x)}, \quad n \geq 2\\
C_{n,2c} &= \left((\theta-1)A_{2,c}(\kappa_1\nu_c - 1) + \frac{1}{2}\gamma^2\right)\frac{(1-\nu_c^n)}{(1-\nu_c)}, \quad n \geq 1.
\end{aligned}$$

Define return on an n-period zero-coupon bond as

$$r_{n,t+1} = p_{n-1,t+1} - p_{n,t}.$$

The risk premium for the bond return is

$$\begin{aligned} & E_t(r_{n,t+1} - r_{f,t}) + \frac{1}{2} Var_t(r_{n,t+1}) \qquad (A.24) \\ = \; & -cov_t(m_{t+1}, r_{n,t+1}) \\ = \; & -(\theta-1)\kappa_1 A_{2,c} C_{n-1,2c} \sigma_{w_c}^2 - (\theta-1)\kappa_1 A_{2,x} C_{n-1,2x} \sigma_{w_x}^2 - \left\{(\theta-1)\kappa_1 A_{1,\lambda} + \theta\right\} C_{n-1,1\lambda} \sigma_\lambda^2 \\ & -(\theta-1)\kappa_1 A_1 C_{n-1,1} \sigma_{x,t}^2. \end{aligned}$$

A.6 Deriving the Intertemporal Marginal Rate of Substitution (MRS)

We consider a representative-agent endowment economy modified to allow for time-preference shocks. The representative agent has Epstein and Zin (1989) recursive preferences and maximizes her lifetime utility

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

subject to budget constraint

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

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. The ratio $\frac{\lambda_{t+1}}{\lambda_t}$ determines how agents trade off current versus future utility and is referred to as the time-preference shock (see Albuquerque, Eichenbaum, and Rebelo (2012)).

First conjecture a solution for $V_t = \phi_t W_t$. The value function is homogenous of degree 1 in wealth; it can now be written as

$$\phi_t W_t = \max_{C_t} \left[(1-\delta)\lambda_t C_t^{\frac{1-\gamma}{\theta}} + \delta \big(\mathbb{E}_t [(\phi_{t+1} W_{t+1})^{1-\gamma}] \big)^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}} \qquad (A.25)$$

subject to

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

Epstein and Zin (1989) show that the above dynamic program has a maximum.

Using the dynamics of the wealth equation, we substitute W_{t+1} into (A.25) to derive

$$\phi_t W_t = \max_{C_t} \left[(1-\delta)\lambda_t C_t^{\frac{1-\gamma}{\theta}} + \delta (W_t - C_t)^{\frac{1-\gamma}{\theta}} \big(\mathbb{E}_t [(\phi_{t+1} R_{c,t+1})^{1-\gamma}] \big)^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}. \qquad (A.26)$$

At the optimum, $C_t = b_t W_t$, where b_t is the consumption-wealth ratio. Using (A.26) and shifting the exponent on the braces to the left-hand side, and dividing by W_t, yields

$$\phi_t^{\frac{1-\gamma}{\theta}} = (1-\delta)\lambda_t \left(\frac{C_t}{W_t} \right)^{\frac{1-\gamma}{\theta}} + \delta \left(1 - \frac{C_t}{W_t} \right)^{\frac{1-\gamma}{\theta}} \big(\mathbb{E}_t [(\phi_{t+1} R_{c,t+1})^{1-\gamma}] \big)^{\frac{1}{\theta}} \qquad (A.27)$$

or simply

$$\phi_t^{\frac{1-\gamma}{\theta}} = (1-\delta)\lambda_t b_t^{\frac{1-\gamma}{\theta}} + \delta(1-b_t)^{\frac{1-\gamma}{\theta}}\left(\mathbb{E}_t[(\phi_{t+1}R_{c,t+1})^{1-\gamma}]\right)^{\frac{1}{\theta}}. \tag{A.28}$$

The first-order condition with respect to the consumption choice yields

$$(1-\delta)\lambda_t b_t^{\frac{1-\gamma}{\theta}-1} = \delta(1-b_t)^{\frac{1-\gamma}{\theta}-1}\left(\mathbb{E}_t[(\phi_{t+1}R_{c,t+1})^{1-\gamma}]\right)^{\frac{1}{\theta}}. \tag{A.29}$$

Plugging (A.29) into (A.28) yields

$$\phi_t = (1-\delta)^{\frac{\theta}{1-\gamma}}\lambda_t^{\frac{\theta}{1-\gamma}}\left(\frac{C_t}{W_t}\right)^{\frac{1-\gamma-\theta}{1-\gamma}} = (1-\delta)^{\frac{\psi}{\psi-1}}\lambda_t^{\frac{\psi}{\psi-1}}\left(\frac{C_t}{W_t}\right)^{\frac{1}{1-\psi}}. \tag{A.30}$$

The lifetime value function is $\phi_t W_t$, with the solution to ϕ_t stated above. This expression for ϕ_t is important: It states that the maximized lifetime utility is determined by the consumption-wealth ratio.

(A.29) can be rewritten as

$$(1-\delta)^{\theta}\lambda_t^{\theta}\left(\frac{b_t}{1-b_t}\right)^{-\frac{\theta}{\psi}} = \delta^{\theta}\mathbb{E}_t[(\phi_{t+1}R_{c,t+1})^{1-\gamma}]. \tag{A.31}$$

Consider the term $\phi_{t+1}R_{c,t+1}$:

$$\phi_{t+1}R_{c,t+1} = (1-\delta)^{\frac{\psi}{\psi-1}}\lambda_{t+1}^{\frac{\psi}{\psi-1}}\left(\frac{C_{t+1}}{W_{t+1}}\right)^{\frac{1}{1-\psi}}R_{c,t+1}. \tag{A.32}$$

After substituting the wealth constraint, $\frac{C_{t+1}}{W_{t+1}} = \frac{C_{t+1}/C_t}{W_t/C_t - 1}\cdot\frac{1}{R_{c,t+1}} = \frac{G_{t+1}}{R_{c,t+1}}\cdot\frac{b_t}{1-b_t}$, into the above expression, it follows that

$$\phi_{t+1}R_{c,t+1} = (1-\delta)^{\frac{\psi}{\psi-1}}\lambda_{t+1}^{\frac{\psi}{\psi-1}}\left(\frac{b_t}{1-b_t}\right)^{\frac{1}{1-\psi}}\left(\frac{G_{t+1}}{R_{c,t+1}}\right)^{\frac{1}{1-\psi}}R_{c,t+1}. \tag{A.33}$$

After some intermediate tedious manipulations,

$$\delta^{\theta}(\phi_{t+1}R_{c,t+1})^{1-\gamma} = \delta^{\theta}(1-\delta)^{\theta}\lambda_{t+1}^{\theta}\left(\frac{b_t}{1-b_t}\right)^{-\frac{\theta}{\psi}}G_{t+1}^{-\frac{\theta}{\psi}}R_{c,t+1}^{\theta}. \tag{A.34}$$

Taking expectations and substituting the last expression into (A.31) yields

$$\delta^{\theta}\mathbb{E}_t[\left(\frac{\lambda_{t+1}}{\lambda_t}\right)^{\theta}G_{t+1}^{-\frac{\theta}{\psi}}R_{c,t+1}^{\theta-1}R_{c,t+1}] = 1. \tag{A.35}$$

From here we see that the MRS in terms of observables is

$$M_{t+1} = \delta^{\theta}\left(\frac{\lambda_{t+1}}{\lambda_t}\right)^{\theta}G_{t+1}^{-\frac{\theta}{\psi}}R_{c,t+1}^{\theta-1}. \tag{A.36}$$

The log of MRS is

$$m_{t+1} = \theta\log\delta + \theta x_{\lambda,t+1} - \frac{\theta}{\psi}g_{t+1} + (\theta-1)r_{c,t+1}, \tag{A.37}$$

where $x_{\lambda,t+1} = \log(\frac{\lambda_{t+1}}{\lambda_t})$.

B Data Source

B.1 Nominal PCE

We download seasonally adjusted data for nominal PCE from NIPA Tables 2.3.5 and 2.8.5. We then compute within-quarter averages of monthly observations and within-year averages of quarterly observations.

B.2 Real PCE

We use Table 2.3.3., Real Personal Consumption Expenditures by Major Type of Product, Quantity Indexes (A:1929-2011)(Q:1947:Q1-2011:Q4) to extend Table 2.3.6., Real Personal Consumption Expenditures by Major Type of Product, Chained Dollars (A:1995-2011) (Q:1995:Q1-2011:Q4). Monthly data are constructed analogously using Table 2.8.3. and Table 2.8.6.

B.3 Real Per Capita PCE: ND+S

The LRR model defines consumption as per capita consumer expenditures on nondurables and services. We download mid-month population data from NIPA Table 7.1.(A:1929-2011)(Q:1947:Q1-2011:Q4) and from Federal Reserve Bank of St. Louis' FRED database (M:1959:M1-2011:M12). We convert consumption to per capita terms.

B.4 Dividend and Market Returns Data

Data are from the Center for Research in Security Prices (CRSP). The three monthly series from CRSP are the value-weighted with-, RN_t, and without-dividend nominal returns, RX_t, of CRSP stock market indexes (NYSE/AMEX/NASDAQ/ARCA), and the CPI inflation rates, π_t. The sample period is from 1928:M1 to 2011:M12. The monthly real dividend series are constructed as in Hodrick (1992):

1. A normalized nominal value-weighted price series is produced by initializing $P_0 = 1$ and recursively setting $P_t = (1 + RX_t)P_{t-1}$.
2. A normalized nominal divided series, d_t, is obtained by recognizing that $d_t = (RN_t - RX_t)P_{t-1}$.
3. The annualized dividend is $D_t = \sum_{j=0}^{11} d_{t-j}$, which sums the previous 11 months of dividends with the current dividend. The first observation is 1928:M12.

Both dividend growth, $\log(\frac{D_{t+1}}{D_t})$, and market returns, RN_{t+1}, are converted from nominal to real terms using the CPI inflation rates, which are denoted by $g_{d,t+1}$ and $r_{m,t+1}$ respectively. They are available from 1929:M1 to 2011:M12.

B.5 Ex Ante Risk-Free Rate

The ex ante risk-free rate is constructed as in the online appendix of Beeler and Campbell (2012). Nominal yields to calculate risk-free rates are the CRSP Fama Risk Free Rates. Even though our model runs in monthly frequencies, we use the three-month yield because of the larger volume and higher reliability. We subtract annualized three-month inflation, $\pi_{t,t+3}$, from the nominal yield, $i_{f,t}$, to form a measure of the ex post (annualized) real three-month interest rate. The ex ante real risk-free rate, $r_{f,t}$, is constructed as a fitted value from a projection of the ex post real rate on the current nominal yield, $i_{f,t}$, and inflation over the previous year, $\pi_{t-12,t}$:

$$\begin{aligned} i_{f,t} - \pi_{t,t+3} &= \beta_0 + \beta_1 i_{f,t} + \beta_2 \pi_{t-12,t} + \varepsilon_{t+3} \\ r_{f,t} &= \hat{\beta}_0 + \hat{\beta}_1 i_{f,t} + \hat{\beta}_2 \pi_{t-12,t}. \end{aligned}$$

The ex ante real risk-free rates are available from 1929:M1 to 2011:M12.

C The State-Space Representation of the LRR Model

C.1 Measurement Equations

In order to capture the correlation structure between the measurement errors at monthly frequency, we assumed in the main text that 12 months of consumption growth data are released at the end of each year. We will now present the resulting measurement equation. To simplify the exposition, we assume that the monthly consumption data are released at the end of the quarter (rather than at the end of the year). In the main text, the measurement equation is written as

$$y_{t+1} = A_{t+1}\Big(D + Zs_{t+1} + Z^v s^v_{t+1} + \Sigma^u u_{t+1}\Big), \quad u_{t+1} \sim N(0, I). \tag{A.38}$$

The selection matrix A_{t+1} accounts for the deterministic changes in the vector of observables, y_{t+1}. Recall that monthly observations are available only starting in 1959:M1. For the sake of exposition, suppose prior to 1959:M1 consumption growth was available at a quarterly frequency. Then:

1. Prior to 1959:M1:

   (a) If $t+1$ is the last month of the quarter:

$$y_{t+1} = \begin{bmatrix} g^q_{c,t+1} \\ g_{d,t+1} \\ r_{m,t+1} \\ r_{f,t} \end{bmatrix}, \quad A_{t+1} = \begin{bmatrix} \frac{1}{3} & \frac{2}{3} & 1 & \frac{2}{3} & \frac{1}{3} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

   (b) If $t+1$ is not the last month of the quarter:

$$y_{t+1} = \begin{bmatrix} g_{d,t+1} \\ r_{m,t+1} \\ r_{f,t} \end{bmatrix}, \quad A_{t+1} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

2. From 1959:M1 to present:

   (a) If $t+1$ is the last month of the quarter:

$$y_{t+1} = \begin{bmatrix} g_{c,t+1} \\ g_{c,t} \\ g_{c,t-1} \\ g_{d,t+1} \\ r_{m,t+1} \\ r_{f,t} \end{bmatrix}, \quad A_{t+1} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

(b) If $t+1$ is not the last month of the quarter:

$$y_{t+1} = \begin{bmatrix} g_{d,t+1} \\ r_{m,t+1} \\ r_{f,t} \end{bmatrix}, \quad A_{t+1} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

The relationship between observations and states (ignoring the measurement errors) is given by the approximate analytical solution of the LRR model described in Section A:

$$\begin{aligned} g_{c,t+1} &= \mu_c + x_t + \sigma_{c,t}\eta_{c,t+1} \\ g_{d,t+1} &= \mu_d + \phi x_t + \pi\sigma_{c,t}\eta_{c,t+1} + \sigma_{d,t}\eta_{d,t+1} \\ r_{m,t+1} &= \{\kappa_{0,m} + (\kappa_{1,m}-1)A_{0,m} + \mu_d\} \\ &+ (\kappa_{1,m}A_{1,m})x_{t+1} + (\phi - A_{1,m})x_t + (\kappa_{1,m}A_{1,\lambda,m})x_{\lambda,t+1} - A_{1,\lambda,m}x_{\lambda,t} + \pi\sigma_{c,t}\eta_{c,t+1} + \sigma_{d,t}\eta_{d,t+1} \\ &+ (\kappa_{1,m}A_{2,x,m})\sigma^2_{x,t+1} - A_{2,x,m}\sigma^2_{x,t} + (\kappa_{1,m}A_{2,c,m})\sigma^2_{c,t+1} - A_{2,c,m}\sigma^2_{c,t} + (\kappa_{1,m}A_{2,d,m})\sigma^2_{d,t+1} - A_{2,d,m}\sigma^2_{d,t} \\ r_{f,t} &= B_0 + B_1 x_t + B_{1,\lambda}x_{\lambda,t} + B_{2,x}\sigma^2_{x,t} + B_{2,c}\sigma^2_{c,t} \\ &\eta_{i,t+1}, \eta_{\lambda,t+1}, w_{i,t+1} \sim N(0,1), \quad i \in \{c,x,d\}. \end{aligned} \tag{A.39}$$

In order to reproduce (A.39) and the measurement-error structure described in Sections 3.1 and 3.2, we define the vectors of states s_{t+1} and s^v_{t+1} as

$$s_{t+1} = \begin{bmatrix} x_{t+1} \\ x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ x_{t-4} \\ \sigma_{c,t}\eta_{c,t+1} \\ \sigma_{c,t-1}\eta_{c,t} \\ \sigma_{c,t-2}\eta_{c,t-1} \\ \sigma_{c,t-3}\eta_{c,t-2} \\ \sigma_{c,t-4}\eta_{c,t-3} \\ \sigma_\epsilon\epsilon_{t+1} \\ \sigma_\epsilon\epsilon_t \\ \sigma_\epsilon\epsilon_{t-1} \\ \sigma_\epsilon\epsilon_{t-2} \\ \sigma_\epsilon\epsilon_{t-3} \\ \sigma_\epsilon\epsilon_{t-4} \\ \sigma_\epsilon^q\epsilon_{t+1}^q \\ \sigma_\epsilon^q\epsilon_t^q \\ \sigma_\epsilon^q\epsilon_{t-1}^q \\ \sigma_\epsilon^q\epsilon_{t-2}^q \\ \sigma_{d,t}\eta_{d,t+1} \\ x_{\lambda,t+1} \\ x_{\lambda,t} \end{bmatrix}, \quad s_{t+1}^v = \begin{bmatrix} \sigma_{x,t+1}^2 \\ \sigma_{x,t}^2 \\ \sigma_{c,t+1}^2 \\ \sigma_{c,t}^2 \\ \sigma_{d,t+1}^2 \\ \sigma_{d,t}^2 \end{bmatrix}. \tag{A.40}$$

It can be verified that the coefficient matrices D, Z, Z^v, and Σ^e are given by

$$Z = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & -\frac{1}{3} & -\frac{1}{3} & \frac{2}{3} & -\frac{2}{3} & \frac{1}{3} & \frac{1}{3} & 1 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \phi & 0 & 0 & 0 & 0 & \pi & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ \mu_{r,1} & \mu_{r,2} & 0 & 0 & 0 & 0 & \mu_{r,3} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{r,4} & \mu_{r,5} & \mu_{r,6} \\ 0 & B_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & B_{1,\lambda} \end{bmatrix}$$

$$Z^v = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \mu_{r,7} & \mu_{r,8} & \mu_{r,9} & \mu_{r,10} & \mu_{r,11} & \mu_{r,12} \\ 0 & B_{2,x} & 0 & B_{2,c} & 0 & 0 \end{bmatrix}, \quad D = \begin{bmatrix} \mu \\ \mu \\ \mu \\ \mu \\ \mu \\ \mu_d \\ \mu_{r,0} \\ B_0 \end{bmatrix}, \quad \Sigma^u = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \sigma_\epsilon^{rf} \end{bmatrix}.$$

The coefficients $\mu_{r,0}$ to $\mu_{r,12}$ are obtained from the solution of the LRR model:

$$\begin{bmatrix} \mu_{r,0} \\ \mu_{r,1} \\ \mu_{r,2} \\ \mu_{r,3} \\ \mu_{r,4} \\ \mu_{r,5} \\ \mu_{r,6} \end{bmatrix} = \begin{bmatrix} \kappa_{0,m} + A_{0,m}(\kappa_{1,m} - 1) + \mu_d \\ \kappa_{1,m} A_{1,m} \\ \phi - A_{1,m} \\ \pi \\ 1 \\ \kappa_{1,m} A_{1,\lambda,m} \\ -A_{1,\lambda,m} \end{bmatrix}, \quad \begin{bmatrix} \mu_{r,7} \\ \mu_{r,8} \\ \mu_{r,9} \\ \mu_{r,10} \\ \mu_{r,11} \\ \mu_{r,12} \end{bmatrix} = \begin{bmatrix} \kappa_{1,m} A_{2,x,m} \\ -A_{2,x,m} \\ \kappa_{1,m} A_{2,c,m} \\ -A_{2,c,m} \\ \kappa_{1,m} A_{2,d,m} \\ -A_{2,d,m} \end{bmatrix}.$$

C.2 State Transition Equations

Using the definition of s_{t+1} in (A.40), we write the state-transition equation as

$$s_{t+1} = \Phi s_t + v_{t+1}(h_t). \tag{A.41}$$

Conditional on the volatilities h_t, this equation reproduces the law of motion of the two persistent conditional mean processes

$$\begin{aligned} x_{t+1} &= \rho x_t + \sigma_{x,t} \eta_{x,t+1} \\ x_{\lambda,t+1} &= \rho_\lambda x_{\lambda,t} + \sigma_\lambda \eta_{\lambda,t+1} \end{aligned} \tag{A.42}$$

and it contains some trivial relationships among the measurement-error states. The matrices Φ and $v_{t+1}(h_t)$ are defined as

$$\Phi = \begin{bmatrix}
\rho & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \rho_\lambda & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0
\end{bmatrix}$$

and

$$v_{t+1}(h_t) = \begin{bmatrix} \sigma_{x,t}\eta_{x,t+1} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \sigma_{c,t}\eta_{c,t+1} \\ 0 \\ 0 \\ 0 \\ 0 \\ \sigma_\epsilon \epsilon_{t+1} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \sigma_\epsilon^q \epsilon_{t+1}^q \\ 0 \\ 0 \\ 0 \\ \sigma_{d,t}\eta_{d,t+1} \\ \sigma_\lambda \eta_{\lambda,t+1} \\ 0 \end{bmatrix}.$$

The law of motion of the three persistent conditional log volatility processes is given by

$$h_{t+1} = \Psi h_t + \Sigma_h w_{t+1}, \tag{A.43}$$

where

$$h_{t+1} = \begin{bmatrix} h_{x,t+1} \\ h_{c,t+1} \\ h_{d,t+1} \end{bmatrix}, \quad \Psi = \begin{bmatrix} \rho_{h_x} & 0 & 0 \\ 0 & \rho_{h_c} & 0 \\ 0 & 0 & \rho_{h_d} \end{bmatrix}$$

$$\Sigma_h = \begin{bmatrix} \sigma_{h_x}\sqrt{1-\rho_{h_x}^2} & 0 & 0 \\ 0 & \sigma_{h_c}\sqrt{1-\rho_{h_c}^2} & 0 \\ 0 & 0 & \sigma_{h_d}\sqrt{1-\rho_{h_d}^2} \end{bmatrix}, \quad w_{t+1} = \begin{bmatrix} w_{x,t+1} \\ w_{c,t+1} \\ w_{d,t+1} \end{bmatrix}.$$

We express

$$\sigma_{x,t} = \varphi_x \bar{\sigma} \exp(h_{x,t}), \quad \sigma_{c,t} = \varphi_c \bar{\sigma} \exp(h_{c,t}), \quad \sigma_{d,t} = \varphi_d \bar{\sigma} \exp(h_{d,t}),$$

which delivers the dependence on h_t in the above definition of $v_{t+1}(\cdot)$.

D Posterior Inference

As discussed in the main text, we use a particle-filter approximation of the likelihood function and embed this approximation into a fairly standard random walk Metropolis algorithm.

D.1 Particle Filter

Our state-space representation, given by equations (A.38), (A.41), and (A.43), is linear conditional on the volatility states h_t. Thus, following Chen and Liu (2000), we update s_{t+1} conditional on h_t using Kalman filter iterations, which improves the efficiency of the filter substantially. In the subsequent exposition we omit the dependence of all densities on the parameter vector Θ. The particle filter approximates the sequence of distributions $\{p(z_t|Y_{1:t})\}_{t=1}^T$ by a set of pairs $\left\{z_t^{(i)}, \pi_t^{(i)}\right\}_{i=1}^N$, where $z_t^{(i)}$ is the i'th particle vector, $\pi_t^{(i)}$ is its weight, and N is the number of particles. As a by-product, the filter produces a sequence of likelihood approximations $\hat{p}(y_t|Y_{1:t-1})$, $t = 1, \ldots, T$.

- **Initialization:** We generate the particle values $z_0^{(i)}$ by drawing the volatilities (h_0, h_{-1}) from the unconditional distribution associated with (A.43). Conditional on the volatility state $(h_0^{(i)}, h_{-1}^{(i)})$, $s_0^{(i)}$ is generated from the unconditional distribution associated with (A.41). We set $\pi_0^{(i)} = 1/N$ for each i.

- **Propagation of particles:** We simulate (A.43) forward to generate $(h_t^{(i)}, h_{t-1}^{(i)})$ conditional on $(h_{t-1}^{(i)}, h_{t-2}^{(i)})$. Taking $s_{t-1}^{(i)}$ and $(h_t^{(i)}, h_{t-1}^{(i)})$ as given, for each particle we run one iteration of the Kalman filter based on the linear state-space system comprised of (A.38)and (A.41) to determine $p(s_t|y_t, s_{t-1}^{(i)}, h_t^{(i)}, h_{t-1}^{(i)})$. This distribution is normal with mean $s_{t|t}^{(i)}$ and $P_{t|t}^{(i)}$. We sample $s_t^{(i)}$ from $N\big(s_{t|t}^{(i)}, P_{t|t}^{(i)}\big)$. We use $q(z_t|z_{t-1}^{(i)}, y_t)$ to represent the distribution from which we draw $z_t^{(i)}$.

- **Correction of particle weights:** Define the unnormalized particle weights for period t as

$$\begin{aligned}
\tilde{\pi}_t^{(i)} &= \pi_{t-1}^{(i)} \times \frac{p(y_t|z_t^{(i)})p(z_t^{(i)}|z_{t-1}^{(i)})}{q(z_t^{(i)}|z_{t-1}^{(i)}, y_t)} \\
&= \pi_{t-1}^{(i)} \times \frac{p(y_t|z_t^{(i)})p(z_t^{(i)}|z_{t-1}^{(i)})}{p(s_t^{(i)}|h_t^{(i)}, h_{t-1}^{(i)}, s_{t-1}^{(i)}, y_t)q(h_t^{(i)}|h_{t-1}^{(i)})} \\
&= \pi_{t-1}^{(i)} \times \frac{p(y_t|z_t^{(i)})p(s_t^{(i)}|h_t^{(i)}, h_{t-1}^{(i)}, s_{t-1}^{(i)})p(h_t^{(i)}|h_{t-1}^{(i)})}{p(s_t^{(i)}|h_t^{(i)}, h_{t-1}^{(i)}, s_{t-1}^{(i)}, y_t)q(h_t^{(i)}|h_{t-1}^{(i)})} \\
&= \pi_{t-1}^{(i)} \times \frac{p(y_t|z_t^{(i)})p(s_t^{(i)}|h_t^{(i)}, h_{t-1}^{(i)}, s_{t-1}^{(i)})}{p(s_t^{(i)}|h_t^{(i)}, h_{t-1}^{(i)}, s_{t-1}^{(i)}, y_t)}.
\end{aligned} \tag{A.44}$$

The term $\pi_{t-1}^{(i)}$ is the initial particle weight and the ratio $p(y_t|z_t^{(i)})p(z_t^{(i)}|z_{t-1}^{(i)})/q(z_t^{(i)}|z_{t-1}^{(i)},y_t)$ is the importance weight of the particle. The second equality is obtained by factorizing $q(z_t^{(i)}|z_{t-1}^{(i)},y_t)$ into the density of $h_t^{(i)}$ associated with the forward simulation of the volatility states, and the conditional density of $s_t|(y_t,s_{t-1}^{(i)},h_t^{(i)},h_{t-1}^{(i)})$ is obtained from the Kalman filter updating step. The third equality is obtained by factorizing the joint density of $(s_t^{(i)},h_t^{(i)})$, $p(z_t^{(i)}|z_{t-1}^{(i)})$, into a marginal density for $h_t^{(i)}$ and a conditional density for $s_t^{(i)}|h_t^{(i)}$. The last equality follows from the fact that we chose $q(h_t^{(i)}|h_{t-1}^{(i)}) = p(h_t^{(i)}|h_{t-1}^{(i)})$. We further simplify the expression in the last line of (A.44) in the next subsection.

The log likelihood function approximation is given by

$$\log \hat{p}(y_t|Y_{1:t-1}) = \log \hat{p}(y_{t-1}|Y_{1:t-2}) + \log\left(\sum_{i=1}^{N} \tilde{\pi}_t^{(i)}\right).$$

- **Resampling:** Define the normalized weights

$$\pi_t^{(i)} = \frac{\tilde{\pi}_t^{(i)}}{\sum_{j=1}^{N} \tilde{\pi}_t^{(j)}}$$

and generate N draws from the distribution $\{s_t^{(i)},\pi_t^{(i)}\}_{i=1}^N$ using multinomial resampling. In slight abuse of notation, we denote the resampled particles and their weights also by $s_t^{(i)}$ and $\pi_t^{(i)}$, where $\pi_t^{(i)} = 1/N$.

D.2 Further Details on the Correction and Updating Step

We now derive the density $p(s_t|y_t,s_{t-1}^{(i)},h_t^{(i)},h_{t-1}^{(i)})$ as well as a simplified expression for the density ratio in the last line of (A.44). Recall that, conditional on the volatilities (h_t,h_{t-1}), the state-space representation of our model takes the form

$$y_t = A_t\big(D + Zs_t + Z^v s_t^v + \Sigma^u u_t\big), \quad u_t \sim N(0,I) \tag{A.45}$$

$$s_t = \Phi s_{t-1} + v_t(h_{t-1}). \tag{A.46}$$

We now proceed with a Kalman filter forecasting and updating step. Conditional on $(s_{t-1}^{(i)},h_t^{(i)},h_{t-1}^{(i)})$, the state-transition equation can be used to forecast s_t:

$$s_t\big|(s_{t-1}^{(i)},h_t^{(i)},h_{t-1}^{(i)}) \sim N\big(s_{t|t-1}^{(i)},P_{t|t-1}^{(i)}\big),$$

where

$$s_{t|t-1}^{(i)} = \Phi s_{t-1}^{(i)}, \quad P_{t|t-1}^{(i)} = \mathbb{E}[v_t(h_{t-1}^{(i)})v_t'(h_{t-1}^{(i)})].$$

Using the measurement equation we can forecast y_t, conditional on $(s_{t-1}^{(i)}, h_t^{(i)}, h_{t-1}^{(i)})$, as follows:

$$y_t\big|(s_{t-1}^{(i)}, h_t^{(i)}, h_{t-1}^{(i)}) \sim N\big(\hat{y}_{t|t-1}^{(i)}, F_{t|t-1}^{(i)}\big), \tag{A.47}$$

where

$$\hat{y}_{t|t-1}^{(i)} = A_t\big(D + Zs_{t|t-1}^{(i)} + Z^v s_t^v(h_t^{(i)}, h_{t-1}^{(i)})\big), \quad F_{t|t-1}^{(i)} = (A_tZ)P_{t|t-1}^{(i)}(A_tZ)' + (A_t\Sigma^u)(A_t\Sigma^u)'.$$

Finally, we can apply the Kalman filter updating step to obtain

$$s_t\big|(y_t, s_{t-1}^{(i)}, h_t^{(i)}, h_{t-1}^{(i)}) \sim N\big(s_{t|t}^{(i)}, P_{t|t}^{(i)}\big), \tag{A.48}$$

where

$$\begin{aligned} s_{t|t}^{(i)} &= s_{t|t-1}^{(i)} + (A_tZP_{t|t-1}^{(i)})'(F_{t|t-1}^{(i)})^{-1}(y_t - \hat{y}_{t|t-1}^{(i)}) \\ P_{t|t}^{(i)} &= P_{t|t-1}^{(i)} - (A_tZP_{t|t-1}^{(i)})'(F_{t|t-1}^{(i)})^{-1}(A_tZP_{t|t-1}^{(i)}). \end{aligned}$$

Define $\mathcal{F}^{(i)} = \{h_t^{(i)}, h_{t-1}^{(i)}, s_{t-1}^{(i)}\}$ and consider the density ratio used to update the particle weights:

$$\begin{aligned} \frac{p(y_t|z_t^{(i)})p(s_t^{(i)}|h_t^{(i)}, h_{t-1}^{(i)}, s_{t-1}^{(i)})}{p(s_t^{(i)}|h_t^{(i)}, h_{t-1}^{(i)}, s_{t-1}^{(i)}, y_t)} &= \frac{p(y_t|s_t^{(i)}, \mathcal{F}^{(i)})p(s_t^{(i)}|\mathcal{F}^{(i)})}{p(s_t^{(i)}|y_t, \mathcal{F}^{(i)})} \\ &= \frac{p(s_t^{(i)}|y_t, \mathcal{F}^{(i)})p(y_t|\mathcal{F}^{(i)})}{p(s_t^{(i)}|y_t, \mathcal{F}^{(i)})} \\ &= p(y_t|\mathcal{F}^{(i)}). \end{aligned} \tag{A.49}$$

The first equality in (A.49) follows from

$$p(y_t|z_t^{(i)}) = p(y_t|s_t^{(i)}, h_t^{(i)}, h_{t-1}^{(i)}) = p(y_t|s_t^{(i)}, h_t^{(i)}, h_{t-1}^{(i)}, s_{t-1}^{(i)})$$

and the second equality in (A.49) is an application of Bayes' Theorem. The expression for $p(y_t|\mathcal{F}^{(i)})$ was previously derived in (A.47).

E The Measurement-Error Model for Consumption

E.1 Monthly Interpolation and Adjustment of Consumption

For expositional purposes, we assume that the accurately measured low-frequency observations are available at quarterly frequency (instead of annual frequency as in the main text). Correspondingly, we define the time subscript $t = 3(j-1)+m$, where month $m = 1, 2, 3$ and quarter $j = 1, \ldots$. We use uppercase C to denote the level of consumption and lowercase c to denote percentage deviations from some log-linearization point. Growth rates are approximated as log differences and we use a superscript o to distinguish observed from "true" values.

The measurement-error model presented in the main text can be justified by assuming that the statistical agency uses a high-frequency proxy series to determine monthly consumption growth rates. We use $Z_{3(j-1)+m}$ to denote the monthly value of the proxy series and $Z^q_{(j)}$ the quarterly aggregate. Suppose the proxy variable provides a noisy measure of monthly consumption. More specifically, we consider a multiplicative error model of the form

$$Z_{3(j-1)+m} = C_{3(j-1)+m} \exp(\epsilon_{3(j-1)+m}). \tag{A.50}$$

The interpolation is executed in two steps. In the first step we construct a series $\tilde{C}^o_{3(j-1)+m}$, and in the second step we rescale the series to ensure that the reported monthly consumption data add up to the reported quarterly consumption data within the period. In Step 1, we start from the level of consumption in quarter $j-1$, $C^q_{(j-1)}$, and define

$$\begin{aligned}
\tilde{C}^o_{3(j-1)+1} &= C^{q,o}_{(j-1)} \left(\frac{Z_{3(j-1)+1}}{Z^q_{(j-1)}} \right) \\
\tilde{C}^o_{3(j-1)+2} &= C^{q,o}_{(j-1)} \left(\frac{Z_{3(j-1)+1}}{Z^q_{(j-1)}} \right) \left(\frac{Z_{3(j-1)+2}}{Z_{3(j-1)+1}} \right) = C^{q,o}_{(j-1)} \left(\frac{Z_{3(j-1)+2}}{Z^q_{(j-1)}} \right) \\
\tilde{C}^o_{3(j-1)+3} &= C^{q,o}_{(j-1)} \left(\frac{Z_{3(j-1)+1}}{Z^q_{(j-1)}} \right) \left(\frac{Z_{3(j-1)+2}}{Z_{3(j-1)+1}} \right) \left(\frac{Z_{3(j-1)+3}}{Z_{3(j-1)+2}} \right) = C^{q,o}_{(j-1)} \left(\frac{Z_{3(j-1)+3}}{Z^q_{(j-1)}} \right).
\end{aligned} \tag{A.51}$$

Thus, the growth rates of the proxy series are used to generate monthly consumption data for quarter q. Summing over the quarter yields

$$\tilde{C}^{q,o}_{(j)} = \sum_{m=1}^{3} \tilde{C}^o_{3(j-1)+m} = C^{q,o}_{(j-1)} \left[\frac{Z_{3(j-1)+1}}{Z^q_{(j-1)}} + \frac{Z_{3(j-1)+2}}{Z^q_{(j-1)}} + \frac{Z_{3(j-1)+3}}{Z^q_{(j-1)}} \right] = C^{q,o}_{(j-1)} \frac{Z^q_{(j)}}{Z^q_{(j-1)}}. \tag{A.52}$$

In Step 2, we adjust the monthly estimates $\tilde{C}^o_{3(j-1)+m}$ by the factor $C^{q,o}_{(j)}/\tilde{C}^{q,o}_{(j)}$, which leads to

$$\begin{aligned} C^o_{3(j-1)+1} &= \tilde{C}^o_{3(j-1)+1}\left(\frac{C^{q,o}_{(j)}}{\tilde{C}^{q,o}_{(j)}}\right) = C^{q,o}_{(j)}\frac{Z_{3(j-1)+1}}{Z^q_{(j)}} \\ C^o_{3(j-1)+2} &= \tilde{C}^o_{3(j-1)+2}\left(\frac{C^{q,o}_{(j)}}{\tilde{C}^{q,o}_{(j)}}\right) = C^{q,o}_{(j)}\frac{Z_{3(j-1)+2}}{Z^q_{(j)}} \\ C^o_{3(j-1)+3} &= \tilde{C}^o_{3(j-1)+3}\left(\frac{C^{q,o}_{(j)}}{\tilde{C}^{q,o}_{(j)}}\right) = C^{q,o}_{(j)}\frac{Z_{3(j-1)+3}}{Z^q_{(j)}} \end{aligned} \tag{A.53}$$

and guarantees that

$$C^{q,o}_{(j)} = \sum_{m=1}^{3} C^o_{3(j-1)+m}.$$

We now define the growth rates $g^o_{c,t} = \log C^o_t - \log C^o_{t-1}$ and $g_{c,t} = \log C_t - \log C_{t-1}$. By taking logarithmic transformations of (A.50) and (A.53) and combining the resulting equations, we can deduce that the growth rates for the second and third month of quarter q are given by

$$\begin{aligned} g^o_{c,3(j-1)+2} &= g_{c,3(j-1)+2} + \epsilon_{3(j-1)+2} - \epsilon_{3(j-1)+1} \\ g^o_{c,3(j-1)+3} &= g_{c,3(j-1)+3} + \epsilon_{3(j-1)+3} - \epsilon_{3(j-1)+2}. \end{aligned} \tag{A.54}$$

The derivation of the growth rate between the third month of quarter $j-1$ and the first month of quarter j is a bit more cumbersome. Using (A.53), we can write the growth rate as

$$\begin{aligned} g^o_{c,3(j-1)+1} &= \log C^{q,o}_{(j)} + \log Z_{3(j-1)+1} - \log Z^q_{(j)} \\ &\quad - \log C^{q,o}_{(j-1)} - \log Z_{3(j-2)+3} + \log Z^q_{(j-1)}. \end{aligned} \tag{A.55}$$

To simplify (A.55) further, we are using a log-linear approximation. Suppose we log-linearize an equation of the form

$$X^q_{(j)} = X_{3(j-1)+1} + X_{3(j-1)+2} + X_{3(j-1)+3}$$

around X^q_* and $X_* = X^q_*/3$, using lowercase variables to denote percentage deviations from the log-linearization point. Then,

$$x^q_{(j)} \approx \frac{1}{3}(x_{3(j-1)+1} + x_{3(j-1)+2} + x_{3(j-1)+3}).$$

Using (A.50) and the definition of quarterly variables as sums of monthly variables, we can apply the log-linearization as follows:

$$\log C^{q,o}_{(j)} - \log Z^q_{(j)} = \log(C^q_*/Z^q_*) + \epsilon^q_{(j)} - \frac{1}{3}\big(\epsilon_{3(j-1)+1} + \epsilon_{3(j-1)+2} + \epsilon_{3(j-1)+3}\big). \tag{A.56}$$

Substituting (A.56) into (A.55) yields

$$\begin{aligned} g^o_{c,3(j-1)+1} &= g_{c,3(j-1)+1} + \epsilon_{3(j-1)+1} - \epsilon_{3(j-2)+3} + \epsilon^q_{(j)} - \epsilon^q_{(j-1)} \\ &\quad -\frac{1}{3}\big(\epsilon_{3(j-1)+1} + \epsilon_{3(j-1)+2} + \epsilon_{3(j-1)+3}\big) + \frac{1}{3}\big(\epsilon_{3(j-2)+1} + \epsilon_{3(j-2)+2} + \epsilon_{3(j-2)+3}\big). \end{aligned} \tag{A.57}$$

An "annual" version of this equation appears in the main text.