

Growth-Rate and Uncertainty Shocks in Consumption: Cross-Country Evidence

Emi Nakamura Dmitriy Sergeyev Jón Steinsson*

Columbia University

August 4, 2012

Abstract

We quantify the importance of long-run risks—persistent shocks to growth rates and uncertainty—in a panel of long-term aggregate consumption data for developed countries. We identify sizable and highly persistent world growth-rate shocks as well as less persistent country-specific growth rate shocks. The world growth-rate shocks capture the productivity speed-up and slow-down many countries experienced in the second half of the 20th century. We also identify large and persistent world shocks to uncertainty. Our world uncertainty process captures the large but uneven rise and fall of volatility that occurred over the course of the 20th century. We find that negative shocks to growth rates are correlated with shocks that increase uncertainty. Our estimates based on macroeconomic data alone line up well with earlier calibrations of the long-run risks model designed to match asset pricing data. We document how these dynamics, combined with Epstein-Zin-Weil preferences, help explain a number of asset pricing puzzles.

Keywords: Long-run risks, Uncertainty shocks, Equity premium puzzle.

JEL Classification: E21, G12

*We would like to thank Mariana Garcia and Channing Verbeck for excellent research assistance. We would like to thank Andrew Ang, Ravi Bansal, Geert Bekaert, Jaroslav Borovicka, John Campbell, John Cochrane, Lars Hansen, Ralph Koijen, Lars Lochstoer, Stavros Panageas, Bernard Salanie and seminar participants at various institutions for valuable comments and discussions. We thank the Columbia University Center for International Business Education and Research for financial support.

1 Introduction

A large recent literature has emphasized the importance of long-run risks—persistent shocks to growth rates and uncertainty—for explaining a variety of asset market phenomena. Bansal and Yaron (2004) demonstrated the importance of these features for explaining the high equity premium, high volatility of stock returns, low and stable risk-free rate and predictability of stock returns. Subsequent work has used these shocks to explain failures of the expectations hypothesis of the term structure and uncovered interest rate parity, the return premium on value stocks and small stocks, the term structure of equity returns, and the volatility of the real exchange rate.¹ A comparably large recent literature has focused on the macroeconomic consequences of these same types of shocks—news shocks about future growth rates and uncertainty shocks. Beaudry and Portier (2006) argue that news shocks about future growth rates are an important driver of business cycles.² Bloom (2009) highlights the role of uncertainty shocks in generating recessions.³

Bansal and Yaron (2004) propose the following time-series model of consumption growth:

$$\begin{aligned}\Delta c_{t+1} &= \mu + x_t + \chi\sigma_t\eta_{t+1}, \\ x_{t+1} &= \rho x_t + \sigma_t\epsilon_{t+1}, \\ \sigma_{t+1}^2 &= \sigma^2 + \gamma(\sigma_t^2 - \sigma^2) + \sigma_\omega\omega_{t+1}.\end{aligned}\tag{1}$$

Relative to a simple, random-walk model for consumption, this model adds two novel features: 1) consumption growth is affected by a persistent process x_t , 2) the uncertainty about consumption growth varies over time in a persistent manner. A difficulty with empirically evaluating this model is that certain key parameters—e.g., the persistence of x_t and σ_t^2 and the volatility of the innovations to σ_t^2 —are difficult to estimate with 80 years of consumption data from a single country. This has led authors in the asset pricing literature to focus on calibrations of the long-run risks model designed to match asset pricing data (Bansal and Yaron, 2004; Bansal et al., 2012).⁴ A concern with this approach is that the asset pricing data may be driven by other factors such as habits, rare disasters and heterogeneous agents.⁵ More direct evidence for the mechanisms that the long-run

¹Important papers include Bansal and Shaliastovich (2010), Bansal, Dittmar, and Lundblad (2005), Hansen, Heaton, and Li (2008), Bonomo et al. (2011), Malloy, Moskowitz, and Vissing-Jorgensen (2009), Croce, Lettau, and Ludvigson (2010), and Colacito and Croce (2011). See Bansal, Kiku, and Yaron (2012) for a more comprehensive review of this literature.

²See also Aguiar and Gopinath (2007), Jaimovich and Rebelo (2009), Barsky and Sims (2010), Schmitt-Grohe and Uribe (2010), and Blanchard, L’Huillier, and Lorenzoni (2011).

³See also Bloom et al., 2011, Fernandez-Villaverde et al., 2011, and Basu and Bundick (2011).

⁴Several papers have also used a combination of macroeconomic and asset pricing data to estimate the parameters of the long-run risks model (e.g., Bansal, Kiku, and Yaron, 2007; Constantinides and Ghosh, 2009).

⁵See Campbell and Cochrane (1999), Barro (2006) and Constantinides and Duffie (1996) for influential asset pricing models based on these features.

risks model is based on would, therefore, strengthen the case for this model.

We quantify the importance of growth-rate and uncertainty shocks using recently assembled data on aggregate consumption for a panel of 16 developed countries. We assume that certain features of consumption dynamics are common across countries. This allows us to estimate key parameters more accurately. An advantage of our approach is that our estimates are based purely on macroeconomic data. We therefore avoid the concern that our estimates of growth-rate and uncertainty shocks are engineered to fit the asset pricing data, as opposed to being a fundamental feature of the aggregate consumption data.

Our empirical model augments Bansal and Yaron’s model to allow for common variation in growth-rates and uncertainty across countries as well as country-specific shocks to growth rates and uncertainty. We identify a substantial common component to expected growth rates in our panel of developed countries. This common variation in growth rates is highly persistent. It captures the productivity speed-up and slow-down in the second half of the 20th century as well as several world recessions, such as those of 1979-82, 1990 and 2008. The country-specific growth-rate processes we identify are less persistent, but nevertheless yields movements in consumption that differ substantially from a random walk.⁶

We also identify large and highly persistent common shocks to macroeconomic uncertainty. Our world uncertainty process captures the large but uneven rise and fall of volatility that occurred over the course of the 20th century. The “Great Moderation” identified by McConnell and Perez-Quiros (2000) is evident in our estimates. But we uncover several additional sharp swings in volatility, most recently a large increase associated with the “Great Recession.” We estimate substantial variation across countries in the timing and direction of uncertainty shocks. For example, uncertainty rose for several decades after World War II (WWII) in the U.K., while it fell in most countries over this period.

Another novel feature of our empirical model relative to earlier work is that we allow the growth-rate and uncertainty shocks to be correlated. We find that they are in fact substantially negatively correlated. In other words, negative shocks to growth rates tend to be associated with shocks that increase uncertainty. The 1960’s were both a period of high growth and low volatility, while in the 1970’s growth fell and uncertainty rose. More recently, during the recessions of 1990 and particularly 2008 growth fell and our estimates of uncertainty shot up.

⁶These findings line up well with those of Cogley (1990), who finds that long-run growth rates of output are more highly correlated across countries than one-year growth rates for nine of the countries we study.

Overall our empirical results based on macroeconomic data alone yield parameter values that are quite consistent with calibrations of the long-run risks model designed to match key asset pricing moments (Bansal and Yaron, 2004; Bansal, Kiku, and Yaron, 2012). We analyze the asset pricing implications of our estimated model of consumption dynamics within the context of a representative agent endowment economy—following Lucas (1978) and Mehra and Prescott (1985)—and assume that agents have Epstein-Zin-Weil preferences (Epstein and Zin, 1989; Weil, 1990). Our model can match the observed equity premium and risk-free rate if agents have a coefficient of relative risk aversion (CRRA) of roughly 6.5 and an intertemporal elasticity of substitution (IES) of 1.5. For the same utility function parameters, the model without growth-rate and uncertainty shocks generates an equity premium that is more than an order of magnitude smaller. Bansal and Yaron (2004) match the equity premium with a CRRA of 10. On this metric, our estimates, thus, yield more long-run risk than their original calibration. Our model also does well when it comes to other key asset pricing moments such as the volatility of stock returns, the volatility of the risk-free rate, the Sharpe ratio for equity, the volatility and persistence of the price-dividend ratio on stocks and predictability of stock returns based on the price-dividend ratio on stocks.

Uncertainty shocks play an important role in generating movements in asset prices in our model. Shocks that raise expected future uncertainty lead stock prices to fall. And expected returns are predictably high following stock market declines provoked by such uncertainty shocks. Through this mechanism, our model is able to help explain the long-term predictability of stock returns (Campbell and Shiller, 1988; Fama and French, 1988; Hodrick, 1992; Cochrane, 2008; Binsbergen and Koijen, 2010). Our model also implies that price-dividend ratios should forecast volatility and consumption growth. We show that price-dividend ratios on stocks have substantial predictive power for future realized volatility of consumption growth in our sample of countries—extending earlier evidence by Bansal et al. (2005). We also extend related work by Lettau et al. (2004, 2008) that suggests that changes in macroeconomic volatility can explain a substantial fraction of low-frequency movements in price-dividend ratios on stocks. In the data, consumption growth is not forecastable by the price-dividend ratio (Beeler and Campbell, 2012).⁷ (Croce, Lettau, and Ludvigson, 2010) show that consumption is less forecastable in a long-run risks model in which investors don't have had full knowledge of the variation in growth prospects in real time .

We analyze the quantitative implications of growth-rate and uncertainty shocks under the as-

⁷However, Bansal, Kiku, and Yaron (2012) show that consumption growth is substantially forecastable in a VAR with the price-dividend ratio, the risk-free rate and consumption growth.

sumption that the CRRA is 6.5. This value is substantially lower than the standard parameterization in the long-run risks literature of $\text{CRRA} = 10$. However, this degree of risk aversion is high relative to the values typically estimated in the microeconomics literature (Barsky et al., 1997; Chetty, 2006; Paravisini et al. 2010).⁸ Our findings, thus, leave ample “room” for additional factors, such as habit, heterogeneous agents, and rare disasters to play an important role in explaining the level and volatility of asset returns.

In addition to the work discussed above, our paper is related to several strands of work in macroeconomics and finance. A large body of work in macroeconomics has studied the long-run properties of output (Nelson and Plosser, 1982; Campbell and Mankiw, 1989; Cochrane, 1988; Cogley, 1990) and variation in the volatility of output growth (McConnell and Perez-Quiros, 2000; Blanchard and Simon, 2001; Stock and Watson, 2002; Ursua, 2010). Our paper builds heavily on the large and growing literature on long-run risks as a framework for asset pricing pioneered by Kandel and Stambaugh (1990). We consider a simple representative agent asset pricing framework with known parameter values, taking the consumption process as given. Several theoretical papers extend on this framework, studying the production-based microfoundations for long run risks (e.g., Kaltenbrunner and Lochstoer, 2010; Kung and Schmid, 2011), the asset pricing implications of parameter learning (e.g., Collin-Dufresne, Johannes, and Lochstoer, 2012), deviations from the representative agent framework (e.g., Garleanu and Panageas, 2010), and frameworks where utility depends on more than just consumption (e.g., Uhlig, 2007).

The paper proceeds as follows. Section 2 discusses the data we use. Section 3 presents the empirical model. Section 4 discusses our estimation strategy. Section 5 presents our empirical estimates. Section 6 studies the asset-pricing implications of our model. Section 7 concludes.

2 Data

We estimate our model using a dataset on long-term consumer expenditures recently constructed by Robert Barro and Jose Ursua, and described in detail in Barro and Ursua (2008).⁹ Our sample

⁸In a static context, an agent with a CRRA of 6.5 would turn down a 50-50 gamble that either raised consumption by a factor of 1 million or lowered it by 12%. An agent with a CRRA of 10 would turn down a 50-50 gamble that either raised consumption by a factor of 1 million or lowered it by 8%.

⁹One limitation of the Barro-Ursua data set is that it does not allow us to distinguish between expenditures on non-durables and services versus durables. Unfortunately, separate data on durable and non-durable consumption are not available for most of the countries and time periods we study. For the U.S., non-durables and services are about 70% as volatile as total consumer expenditures over the time period when both series are available. One way of adjusting our results would therefore be to scale down the volatility of the shocks we estimate by 0.7. Whether this adjustment is appropriate depends on the extent to which non-durables and services are less volatile at the longer

includes 16 countries: Australia, Belgium, Canada, Denmark, Finland, France, Germany, Italy, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom, United States.¹⁰ Our consumption data is an unbalanced panel with data for each country starting between 1890 and 1914. Our sample period ends in 2009. Figure 1 plots our data series for France. We have drawn a trend line through the pre-WWII period and extended this line to the present. The figure strongly suggests that France has experienced very persistent swings in growth over the last 120 year. In analyzing the asset pricing implications of our model, we also make use of total returns data on stocks and bills as well as dividend yields on stocks from Global Financial Data (GFD) and data on inflation from Barro and Ursua (2008).

3 An Empirical Model of Growth-Rate and Uncertainty Shocks

Building on the work of Bansal and Yaron (2004), we model the logarithm of the permanent component of per capita consumption in country i at time $t + 1$ —denoted $\tilde{c}_{i,t+1}$ —as evolving in the following way:

$$\begin{aligned}\Delta\tilde{c}_{i,t+1} &= \mu_i + x_{i,t} + \xi_i x_{W,t} + \eta_{i,t+1}, \\ x_{i,t+1} &= \rho x_{i,t} + \epsilon_{i,t+1}, \\ x_{W,t+1} &= \rho_W x_{W,t} + \epsilon_{W,t+1}.\end{aligned}\tag{2}$$

Permanent consumption growth is governed by three shocks: a random-walk shock ($\eta_{i,t+1}$), and two shocks that have persistent effects on the growth rate of consumption—one of which is country specific ($\epsilon_{i,t+1}$) and one of which is common across all countries ($\epsilon_{W,t+1}$). The persistence of the effects of the last two of these shocks to consumption growth is governed by AR(1) processes $x_{i,t+1}$ and $x_{W,t+1}$, respectively. We allow the different countries in our sample to differ in their sensitivity to the world growth rate process. This differing sensitivity is governed by the parameter ξ_i .

The volatility of the three shocks affecting permanent consumption growth is time varying and governed by two AR(1) stochastic volatility processes:

$$\sigma_{i,t+1}^2 = \sigma_i^2 + \gamma(\sigma_{i,t}^2 - \sigma_i^2) + \omega_{i,t+1},\tag{3}$$

$$\sigma_{W,t+1}^2 = \sigma_W^2 + \gamma(\sigma_{W,t}^2 - \sigma_W^2) + \omega_{W,t+1},\tag{4}$$

horizons over which our long-run risks shocks are most important. For example, if durables and non-durables are cointegrated, the adjustment is likely to be smaller. The adjustment is also likely to be smaller for earlier points in our sample, when the role of durables in total consumer expenditures was much smaller.

¹⁰We exclude countries in Southeast Asia and Latin America from our sample. Including these countries raises our estimates of the importance of long-run risks.

where $\sigma_{i,t+1}^2$ is a country-specific component of stochastic volatility and $\sigma_{W,t+1}^2$ is a common component of stochastic volatility. We refer to the innovations to these stochastic volatility processes— $\omega_{i,t+1}$ and $\omega_{W,t+1}$ —as uncertainty shocks.¹¹

The common component of stochastic volatility $\sigma_{W,t+1}^2$ affects the volatility of all three of the shocks to permanent consumption. The idea here is that when world uncertainty rises this affects the volatility of all shocks to permanent consumption. The country specific component of stochastic volatility $\sigma_{i,t+1}^2$, however, only affects the country specific shocks. Variation in this component, therefore, represents deviations in the uncertainty faced by a particular country from that faced by countries on average. More specifically, we assume that $\text{var}_t(\epsilon_{W,t+1}) = \sigma_{W,t}^2$, $\text{var}_t(\epsilon_{i,t+1}) = \sigma_{i,t}^2 + \sigma_{W,t}^2$, and $\text{var}_t(\eta_{i,t+1}) = \chi_i^2(\sigma_{i,t}^2 + \sigma_{W,t}^2)$, where χ_i governs the relative volatility of the two country specific shocks, $\epsilon_{i,t+1}$ and $\eta_{i,t+1}$.

We allow for correlation between the growth-rate shocks and the uncertainty shocks. This is meant to capture the possibility that times of high uncertainty may also tend to be times of low growth. Specifically, we allow the country-specific growth-rate shock $\epsilon_{i,t+1}$ and the country-specific uncertainty shock $\omega_{i,t+1}$ to be correlated with a correlation coefficient of λ . We also allow the world growth-rate shock $\epsilon_{W,t+1}$ and the world uncertainty shocks $\omega_{W,t+1}$ to be correlated with a correlation coefficient of λ_W .

To summarize, we assume the following distributions for the random-walk, growth-rate and uncertainty shocks:

$$\eta_{i,t+1} \sim \text{N}(0, \chi_i^2(\sigma_{i,t}^2 + \sigma_{W,t}^2)), \quad (5)$$

$$\begin{bmatrix} \epsilon_{i,t+1} \\ \omega_{i,t+1} \end{bmatrix} \sim \text{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{i,t}^2 + \sigma_{W,t}^2 & \lambda\sigma_\omega\sqrt{\sigma_{i,t}^2 + \sigma_{W,t}^2} \\ \lambda\sigma_\omega\sqrt{\sigma_{i,t}^2 + \sigma_{W,t}^2} & \sigma_\omega^2 \end{bmatrix} \right), \quad (6)$$

$$\begin{bmatrix} \epsilon_{W,t+1} \\ \omega_{W,t+1} \end{bmatrix} \sim \text{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{W,t}^2 & \lambda_W\sigma_{W,t}\sigma_{\omega,W} \\ \lambda_W\sigma_{W,t}\sigma_{\omega,W} & \sigma_{\omega,W}^2 \end{bmatrix} \right). \quad (7)$$

To avoid negative variances, we truncate the process for $\sigma_{W,t+1}^2$ at a small positive value ζ and we truncate the process for $\sigma_{i,t+1}^2$ such that $\sigma_{i,t+1}^2 > \zeta - \sigma_{W,t}^2$.¹²

¹¹We could alternatively model $\log \sigma_{i,t+1}^2$ and $\log \sigma_{W,t+1}^2$ as following AR(1) processes. This would allow us to avoid truncating the uncertainty shocks (see below). We have experimented with this specification. However, with this specification, the volatility of σ^2 drops to very low levels when σ^2 is small implying that σ^2 can “get stuck” close to zero for a very long time. It is not clear to us that the data support this feature. Also, our MCMC estimation algorithm runs into trouble in this case since the likelihood function is very flat when $\log \sigma^2$ becomes sufficiently negative (σ^2 sufficiently small). In this region very large movements in $\log \sigma^2$ correspond to tiny movements in σ^2 . This leads the MCMC algorithm to get stuck.

¹²For world stochastic volatility, this means that when an $\omega_{W,t+1}$ is drawn that would yield a value of $\sigma_{W,t+1}^2 < \zeta$, we set $\sigma_{W,t+1}^2 = \zeta$. This implies that the innovations to the $\sigma_{W,t+1}^2$ have a positive mean when $\sigma_{W,t+1}^2$ is close to ζ .

We allow parameters to vary across countries whenever our data contains enough information to make this feasible. For example, we allow σ_i^2 to differ across countries. This allows some countries to have permanently higher or lower volatility of macroeconomic shocks than others. However, as Bansal and Yaron (2004) emphasize, some of the key parameters of the long-run risks model are difficult to estimate precisely using data from a single country, even with over 100 years of data. For these parameters, we rely on the panel structure of the data set and assume that they are equal for all countries in our data set. The parameters we make this pooling assumption for are: the persistence of the growth-rate components ρ and ρ_W , the persistence of the stochastic volatility processes γ , the volatility of the uncertainty shocks σ_ω^2 and $\sigma_{W,\omega}^2$, the average volatility of the world stochastic volatility process σ_W^2 , and the correlations between the growth-rate and uncertainty shocks λ and λ_W .¹³

We allow measured consumption—denoted $c_{i,t}$ —to differ from permanent consumption $\tilde{c}_{i,t}$ because of two transitory shocks:

$$c_{i,t+1} = \tilde{c}_{i,t+1} + \nu_{i,t+1} + I_{i,t+1}^d \psi_{i,t+1}^d. \quad (8)$$

The first of these shocks $\nu_{i,t+1}$ is mainly meant to capture measurement error. We assume that this shock is distributed $N(0, \sigma_{i,t,\nu}^2)$, where the volatility of this shock is allowed to differ before and after 1945. By incorporating this break in the volatility of $\nu_{i,t+1}$ we can capture potential changes in national accounts measurement around this time (Romer, 1986; Balke and Gordon, 1989). This is empirically important since it avoids the possibility that our estimates of the high persistence of macroeconomic uncertainty arise spuriously from these changes in measurement procedures.

The second shock $I_{i,t+1}^d \psi_{i,t+1}^d$ captures transitory variation in consumption due to disasters.¹⁴ We do not estimate the timing of disasters in this paper. Instead, the dummy variable $I_{i,t}^d$ is set equal to one in periods identified as disaster periods by Nakamura et al. (2010) and during a two year recovery period after each such episode and zero otherwise.¹⁵ The disaster shock $\psi_{i,t}^d$ is distributed $N(\mu_d, 1)$. We fix the variance of $\psi_{i,t}^d$ at 1 (a large value), to ensure that this shock

For the estimated values of the parameters of our model (baseline estimation), $\sigma_{W,t+1}^2 = \zeta$ about 9.2% of the time. We incorporate this truncation in our asset pricing analysis in section 6.

¹³Notice also, that we assume that the same parameter (γ) governs the persistence of both the common and country-specific components of stochastic volatility. We do this because there is insufficient information in our dataset to estimate a separate parameter for the persistence of world volatility.

¹⁴The permanent effects of disasters are captured by $\eta_{i,t+1}$, $\epsilon_{i,t+1}$, and $\epsilon_{W,t+1}$.

¹⁵Nakamura et al. (2010)'s results indicate that there is unusually high growth after disasters—i.e., recoveries—but that this unusually high growth dies out rapidly—it has a half-life of 1 year. By allowing for a two year recovery period after disasters, we allow the disaster shocks in our model to capture the bulk of the unusually high growth after disasters and avoid having this growth variation inflate our estimates of long-run risks.

“soaks up” all transitory variation in consumption during the disaster periods. Were we to exclude the disaster shock, we would estimate substantially higher volatilities of the stochastic volatility processes $\sigma_{i,t+1}^2$ and $\sigma_{W,t+1}^2$.

4 Estimation

The model presented in section 3 contains a large number of unobserved state variables, since it decomposes consumption into several unobserved components. We estimate the model using Bayesian MCMC methods.¹⁶ To carry out our Bayesian estimation we need to specify a set of priors on the parameters of the model. We choose highly dispersed priors for all the main parameters of the model to minimize their effect on our inference. The full set of priors we use is:

$$\begin{aligned}
\rho &\sim \text{U}(0.005, 0.995), & \rho_W &\sim \text{U}(0.005, 0.995), \\
\gamma &\sim \text{U}(0.005, 0.98), & \sigma_W^2 &\sim \text{U}(10^{-8}, 10^{-2}), \\
\sigma_\omega^2 &\sim \text{U}(10^{-10}, 10^{-6}), & \sigma_{W,\omega}^2 &\sim \text{U}(10^{-10}, 10^{-6}), \\
\lambda &\sim \text{U}(-0.995, 0.995), & \lambda_W &\sim \text{U}(-0.995, 0.995), \\
\xi_i &\sim \text{U}(10^{-4}, 10), & \chi_i^2 &\sim \text{U}(10^{-4}, 25), \\
\sigma_i^2 &\sim \text{U}(10^{-8}, 10^{-2}), & \sigma_{\nu,i}^2 &\sim \text{U}(10^{-8}, 10^{-2}), \\
\mu_i &\sim \text{N}(0.015, 1), & \mu_d &\sim \text{N}(0, 1),
\end{aligned}$$

except that we normalize $\xi_{US} = 1$ to identify the scale of the world stochastic volatility process. We assume that the initial values of $x_{i,t}$, $x_{W,t}$, $\sigma_{i,t}$ and $\sigma_{W,t}$ are drawn from their unconditional distributions. We assume that the initial value of \tilde{c}_{it} for each country is drawn from a highly dispersed normal distribution centered on the initial observation for $c_{i,t}$. It can be shown that the model is formally identified except for a few special cases in which multiple shocks have zero variance.

5 Empirical Results

Our baseline empirical results are for the full model described in section 3 for the full sample period 1890-2009. We also report results for a shorter post-WWII sample period and for a simplified

¹⁶Our algorithm samples from the posterior distributions of the parameters and unobserved states using a Gibbs sampler augmented with Metropolis steps when needed. This algorithm is described in greater detail in appendix A. The estimates discussed in section 5 for the three versions of the model, are based on four independent Markov chains each with 5 million draws or more with the first 450,000 draws from each chain dropped as “burn-in”. To assess convergence, we employ Gelman and Rubin’s (1992) approach to monitoring convergence based on parallel chains with “over-dispersed starting points” (see also Gelman, 2004, ch. 11).

version of the model in which we shut down the world growth-rate and volatility components as well as the correlation between the country-specific growth-rate and volatility shocks. We refer to this latter model as the “simple model.” Tables 1-3 present parameter estimates for these three cases. For each parameter, we present the prior and posterior mean and standard deviation. We refer to the posterior mean of each parameter as our point estimate for that parameter.

We estimate a highly persistent world growth-rate process in our baseline model. The autoregressive coefficient for the world growth-rate component is estimated to be $\rho = 0.83$, implying a half-life of 3.8 years. The country specific growth-rate process is estimated to be less persistent. The autoregressive coefficient for the country-specific growth-rate component is estimated to be $\rho = 0.56$, implying a half-life of 1.2 years. Table 2 compares these estimates to the calibration of the persistence of the growth-rate process in Bansal and Yaron (2004) and Bansal et al. (2012). The persistence of the growth-rate process in these papers is in-between that of our world and country-specific growth-rate processes. In the simple model, the persistence of the (country-specific) growth-rate component is estimated to be $\rho = 0.68$, which implies a half-life of 1.8 years. This illustrates that allowing for a world growth-rate component is important in capturing the low-frequency variation in growth in our dataset.

Figure 2 plots the impulse response of consumption to our estimated growth-rate processes as well as to the random-walk shock. The figure shows clearly that despite the relatively modest half-lives of the growth-rate shocks, their effects on output are very different from those of the random-walk shock. After a country-specific growth-rate shock, consumption continues to grow for several periods and eventually rises by more than two times the initial size of the shock with the bulk of the growth occurring in the first 5 years. After a world growth-rate shock, continuing growth in subsequent periods leads the eventual impact of the shock on consumption to be six times its initial impact with roughly a third of that growth occurring more than 5 years after the shock.

Figure 3 presents our estimate of the world growth-rate process. The most striking feature of this process is its high values in the 1950’s, 60’s and early 70’s. This captures the persistently high growth seen in many countries in our sample in the 3rd quarter of the 20th century.¹⁷ The world growth-rate process also captures several major recessions such as the 1979-82 recession following the spike in oil prices that accompanied the Iranian Revolution as well as the tightening of U.S.

¹⁷It is intriguing that this growth spurt so closely followed World War II. It is tempting to infer that this high growth is due to post-war reconstruction. However, for most countries, the vast majority of the unusually high growth during this period occurred in years when consumption (and output) had surpassed its pre-WWII trend-adjusted level (see, e.g., Figure 1).

monetary policy, the recession of 1990 following, among other events, the Persian Gulf War, the unification of Germany, and the accompanying tightening of German monetary policy, and the 2008 recession following the sharp fall in house prices in several countries, associated collapse of major financial institutions and turmoil in financial markets.

We estimate large and very persistent shocks to economic uncertainty. Table 1 reports that our estimate of the autoregressive coefficient for the uncertainty processes in the baseline estimation is $\gamma = 0.970$. This implies that uncertainty shocks in the baseline case have a half-life of 22.8 years (Table 2). This estimate lies between the 4.4 year calibration of Bansal and Yaron (2004) and the 57.7 year calibration of Bansal et al. (2012). Uncertainty shocks are also estimated to be highly persistent in the simple model and in the post-WWII sample. For these cases, we estimate half-lives of 13.5 years and 18.2 years, respectively.

Figure 4 presents our estimates of the evolution of the world stochastic volatility process ($\sigma_{W,t}$). We estimate that world volatility was high in the early post-WWII period and has been on an uneven downward trend since then. World volatility fell a great deal in the 1960's, but was high again in the 1970's and early 1980's. It fell sharply in the mid-to-late 1980's but was relatively high in the early 1990's. From 1995 to 2007 the world experienced a long period of relative tranquility with volatility falling sharply towards the end of this period to record lows. At the end of our sample period, world volatility rose sharply once again. In studying this figure, it is important to keep in mind that our model attributes much of the volatility in the first half of our sample to our disaster and measurement error shocks.

Comparing Figures 3 and 4, it is evident that the world growth-rate process and the world stochastic volatility process are negatively correlated. Our model allows explicitly for a correlation between shocks to these processes (λ_W). Table 1 reports that our estimate of this correlation is -0.25. We also estimate a common correlation between the country-specific growth-rate and uncertainty shocks in our data and find this correlation to also be -0.40. Our estimates, thus, strongly suggest that periods of high volatility are also periods of low growth.

We estimate a substantial amount of heterogeneity in the evolution of volatility across countries. Figure 5 presents our estimates of the evolution of the volatility process for the U.S., the U.K. and Canada— $(\sigma_{i,t}^2 + \sigma_{W,t}^2)^{1/2}$ in our notation. For the United States our results reflect the “long and large” decline in macroeconomic volatility documented by Blanchard and Simon (2001) and well as the rather abrupt decline in volatility in the mid-1980's documented by McConnell and Perez-Quiros (2000) and Stock and Watson (2002). The experience of the U.K. is quite different. Volatility in the

U.K. was lower in the early part of the 20th century (excluding disasters), but then rose substantially over the first three decades after WWII. Volatility in the U.K. began falling only around the time Margaret Thatcher came to power and has remained elevated relative to volatility in the U.S. ever since 1960. In contrast, volatility in Canada fell much more abruptly in the 1950's and early 1960's than volatility in the U.S. and was substantially below U.S. volatility in the 1960's, 1970's and early 80's at which point U.S. volatility converged down to similarly low levels.

One feature of our results that differs markedly from the calibrations of the long-run risks model used in Bansal and Yaron (2004) and Bansal et al. (2012) is that the growth-rate shocks we estimate are substantially more volatile. Recall that the parameter χ_i governs the relative volatility of the random-walk shock ($\eta_{i,t}$) and the growth-rate shock ($\epsilon_{i,t}$). Estimates for this parameter as well as other country-specific parameters are reported in Table 3. For the median country, we estimate χ_i to be 0.81, while we estimate a value of 1.16 for the United States.¹⁸ Our estimates thus imply that the growth-rate shocks and the random-walk shocks are roughly equally volatile. Bansal and Yaron (2004) and Bansal et al. (2012) calibrate the growth-rate shock to be only about 5% as volatile as the random-walk shock.

We allow countries to differ in their sensitivity to the world growth-rate process. The parameter ξ_i governs this sensitivity. We fix $\xi_{US} = 1$, implying that for other countries this parameter can be interpreted as their sensitivity to world shocks relative to the sensitivity of the U.S. to these shocks. For the median country, our estimate of $\xi_i = 1.51$. In particular, many continental European countries have values of ξ_i that are substantially larger than one (see Table A.1). This heterogeneity in sensitivity to the world-growth rate shock is one source of heterogeneity in risk-premia across countries in our asset-pricing calculations in section 6. We estimate a substantial decline in the volatility of transitory shocks $\sigma_{\nu,i}$ after 1945 in most countries. This change likely reflects in part changes in national accounts measurement, as we discuss in section 3.¹⁹

One potential concern with our results is that they might be influenced by our treatment of disasters in the early part of our sample. Another potential concern is that the quality of the data for the period before World War II may be lower than for the more recent period. To address these concerns, we estimate our model on data starting in 1950. This yields results that are very similar to our baseline estimation along most dimensions. The main deviation is that in this case we estimate

¹⁸Estimates for all 16 countries for our baseline case are presented in the appendix (Table A.1).

¹⁹Ursua (2010) argues—based on methods developed by Romer (1986)—that this change also reflects changes in macroeconomic fundamentals. Since transitory shocks turn out to be relatively unimportant for asset pricing, the choice of whether to treat this change as a consequence of measurement or fundamental shocks plays a small role in our asset pricing analysis.

a smaller and less volatile world stochastic volatility process and larger values of the sensitivity to the world growth-rate shock for most countries. Also, the posterior standard deviation of several key parameters increases substantially—in particular, the standard deviation of the sensitivity to the world growth-rate—reflecting the much smaller sample. For the median country, the degree to which consumption growth is driven by the world growth-rate shock rises since the increase in the sensitivity to the world growth-rate shock is larger than the decrease in the volatility of the world growth-rate shocks.

5.1 Autocorrelations, Cross-Country Correlations and Variance Ratios

Given that we estimate large persistent components to consumption growth, one might worry that our estimated model implies too much autocorrelation of consumption growth relative to the data. Table 4 presents estimates of autocorrelations, cross-country correlations and variance ratios in the data and in the model. We report statistics for the median country in our dataset and for the United States. Both for the data and the model, we exclude transitory variation in consumption due to disasters.²⁰

Consider first the autocorrelations of consumption growth. In the data, the autocorrelations for the median country are positive but small at all horizons; for the US, they oscillate around zero. The model also generates modest autocorrelations at all horizons. The 95% probability intervals generated by the model contain the corresponding empirical statistics in almost all cases.²¹ Despite assigning an important role to long-run risks, our estimated model yields modest short-term autocorrelation in the growth rate of consumption because the model also features transitory shocks to the level of consumption, which generate an offsetting negative correlation in short-term growth rates.

The cross-country correlation of consumption growth for the median country is estimated to be substantial and to grow with the horizon of the growth rates. The median one-year cross-country correlation is 0.23, while it is 0.44 at the five year horizon and 0.56 at the ten year horizon. The model is able to capture both the magnitude and the increasing pattern of this cross-country correlation through the world growth-rate process. The correlation of the U.S. with other countries in our sample is somewhat smaller than for the median country both in the data and in the model.

²⁰For the real world data, we do this by subtracting from the raw data our estimate of the transitory disaster shock. This yields series for consumption that smoothly “interpolate” through disasters. For the simulated data from our model, we simulate the model without the transitory disaster shock.

²¹Estimated on the post-WWII sample, the autocorrelations for the U.S. oscillate less and are slightly negative at horizons longer than one year.

Table 4 also reports estimates of variance ratios for consumption growth and the volatility of consumption growth at the 15 year horizon for the median country and for the United States. Variance ratios above one indicate reduced form evidence for positive autocorrelation of consumption growth and volatility. The definition and intuition for these statistics is discussed in more detail in appendix B. In the data, the variance ratio for consumption growth for the median country is 1.69, substantially above one. The average across countries is even higher at 2.28. For the U.S. it is somewhat smaller but still above one.²² These high variance ratios provide reduced form evidence for positive autocorrelation of growth rates. Our model captures this well. For the median country, the model generates a 15-year variance ratio of 2.69. The variance ratio of realized volatility is substantially larger than one both in the median country and in the United States. Again, our model is able to capture this feature of the data well.

6 Asset Pricing

We analyze the asset pricing implications of the model of aggregate consumption described in section 3 within the context of a representative consumer endowment economy with Epstein-Zin-Weil preferences (Epstein and Zin, 1989; Weil, 1990). For this preference specification, Epstein and Zin (1989) show that the return on an arbitrary cash flow is given by the solution to the following equation:

$$E_t \left[\beta^\theta \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{(-\theta/\psi)} R_{c,t,t+1}^{-(1-\theta)} R_{i,t,t+1} \right] = 1, \quad (9)$$

where $R_{i,t,t+1}$ denotes the gross return on an arbitrary asset in country i from period t to period $t+1$, $R_{c,t,t+1}$ denotes the gross return on the agent's wealth, which in our model equals the endowment stream. The parameter β represents the subjective discount factor of the representative consumer. The parameter $\theta = \frac{1-\gamma}{1-1/\psi}$, where γ is the coefficient of relative risk aversion (CRRA) and ψ is the intertemporal elasticity of substitution (IES), which governs the agent's desire to smooth consumption over time.

We begin by calculating asset prices for two assets: a risk-free one-period bond and a risky asset we will use to represent equity. The risk-free one-period bond has a certain pay-off of one unit of consumption in the next period. We follow Bansal, Kiku, and Yaron (2012) in modeling equity as

²²We have also calculated these variance ratios including disasters and they are lower both in the data and in the model. Excluding disasters raises the variance ratio of consumption growth because disasters are typically followed by significant recoveries (Kilian and Ohanian, 2002; Nakamura et al., 2010). Ursua (2010) presents a related analysis. Rather than filtering the data the way we do, he excludes "outlier" growth observations. This simpler procedure also yields substantially larger variance ratios than raw consumption growth in his broader sample.

having a levered exposure to the stochastic component of permanent consumption. Specifically, the growth rate of dividends for our equity claim is

$$\Delta d_{t+1} = \mu + \phi(x_{i,t} + \xi_i x_{W,t} + \eta_{i,t}), \quad (10)$$

where ϕ is the leverage ratio on expected consumption growth (Abel, 1999). We base our analysis on the posterior mean estimates for the baseline case from section 5. We therefore abstract from learning, doubt and fragile beliefs (Timmermann, 1993; Pastor and Veronesi, 2009; Hansen, 2007; Hansen and Sargent, 2010). These issues are potentially important in our context, given the difficulty of estimating long-run risks, both for the econometrician, and the economic agent (see, e.g., Croce et al., 2010).

The asset-pricing implications of our model with Epstein-Zin-Weil (EZW) preferences cannot be derived analytically. We solve for asset prices in our model using standard grid-based numerical methods of the type used, e.g., by Campbell and Cochrane (1999) and Wachter (2005).²³ We choose a subjective discount factor of $\beta = 0.990$ to fit the observed average risk-free rate in our baseline specification. We choose a CRRA of $\gamma = 6.5$ to match the U.S. equity premium in our baseline specification. We follow the long-run risks literature in choosing an IES of $\psi = 1.5$ (Bansal and Yaron, 2004; Bansal et al., 2012).²⁴ We follow Bansal and Yaron (2004) in setting leverage of $\phi = 3$. We calculate asset prices for a consumption process that ignores the transitory disaster shock in our model.²⁵ We do this to focus attention on the asset-pricing implications of long-run risks. Allowing for transitory drops in consumption due to disasters would further raise the equity premium we estimate (or equivalently allow us to match the equity premium with a lower value of the CRRA) but at considerable cost in terms of computational complexity. The asset pricing implications of disaster risk have been the focus of a large recent literature (see, e.g., Barro, 2006, and Nakamura, et al., 2010). We present asset pricing results for the post-WWII estimation of our model—a sample without major disasters in our sample of countries—in an appendix (Table A.2).

²³We solve the integral in equation (9) on a grid. Specifically, we start by solving for the price-dividend ratio for a consumption claim. In this case we can rewrite equation (9) as $PDR_t^C = E_t[f(\Delta C_{t+1}, PDR_{t+1}^C)]$, where PDR_t^C denotes the price dividend ratio of the consumption claim. We specify a grid for PDR_t^C over the state space. We then solve numerically for a fixed point for PDR_t^C as a function of the state of the economy on the grid. We can then rewrite equation (9) for other assets as $PDR_t = E_t[f(\Delta C_{t+1}, \Delta D_{t+1}, PDR_{t+1}^C, PDR_{t+1})]$, where PDR_t denotes the price dividend ratio of the asset in question and ΔD_{t+1} denotes the growth rate of its dividend. Given that we have already solved for PDR_t^C , we can solve numerically for a fixed point for PDR_t for any other asset as a function of the state of the economy on the grid.

²⁴There is little agreement in the macroeconomics and finance literatures on the appropriate value for the IES. Hall (1988) and Campbell (1999) estimate the IES to be close to zero. However, Hansen and Singleton (1982), Bansal and Yaron (2004), Gruber (2006), Hansen et al. (2007) and Nakamura et al. (2010) argue for values of the IES above one.

²⁵Recall that the permanent effects of disasters on consumption are captured by the random-walk and growth rate shocks in our model.

6.1 The Effects of Long-Run Risks on Asset Prices

Figure 6 presents impulse responses for the return on equity and the risk-free rate to a world growth-rate shock. A positive world growth-rate shock yields a large positive return on equity on impact. This positive return reflects the balance of two opposing forces. On the one hand, the shock raises expected future dividends on equity, which pushes up stock prices. On the other hand, since consumption growth is expected to be high for some time, agents' desire to save falls, which pushes down all asset prices. If agents are sufficiently willing to substitute consumption over time ($IES > 1$), as we assume, the first of these effects is stronger than the second for equity and the price of equity rises on impact. In the periods after the shock, returns on equity and the risk-free rate are higher than average because of agents' reduced desire to save.

Figure 7 presents impulse responses for the return on equity and the risk-free rate to an uncertainty shock. A shock that increases economic uncertainty yields a large negative return on equity on impact. As with the growth-rate shock, there are two opposing forces that together determine the response of stock prices. The increase in economic uncertainty makes stocks riskier, which raises the equity premium. This tends to depress their value. However, the increase in uncertainty also increases the desire of agents to save. This tends to raise the price of all assets. With $CRRA > 1$ and $IES > 1$, the first force is stronger than the second and the price of stocks falls on impact (Campbell, 1993). In the periods after the shock, the equity premium remains elevated because uncertainty has risen. A one standard deviation shock to $\omega_{W,t}$ raises the equity premium by roughly 0.6% in the period after the shock.

Notice that in our model neither the growth-rate shock nor the uncertainty shock affect consumption growth on impact. For an agent with power utility, these shocks would therefore not affect marginal utility on impact. This implies that agents with power utility would not demand a risk premium on stocks as compensation for exposure to these shocks. With EZW utility, however, marginal utility depends not only on current consumption but also on news about future consumption.²⁶ In equation (9), this is captured by the presence of the return on wealth— $R_{c,t+1}$. Since negative growth-rate shocks and shocks that increase uncertainty imply negative returns on wealth on impact, they increase marginal utility. Households are, thus, willing to pay a premium for assets that provide insurance against growth-rate and uncertainty shocks. Conversely, they demand a risk premium for holding assets that expose them to these shocks.

²⁶This implication of EZW preferences is illustrated elegantly by the decomposition developed by Borovicka et al. (2011).

6.2 Risk-Premia and Return Volatility

Two key features of the asset pricing data are the equity premium and the large volatility of equity returns. Long-run risks shocks make the world a riskier place, leading both the level and the volatility of equity returns to rise relative to the risk free rate. Table 5 presents key asset pricing statistics in the data and for our baseline specification of the model. The table presents results for the U.S. and for the median country in our sample.

Our model matches the observed equity premium for the United States with a CRRA of 6.5, a value that is an order of magnitude lower than in a model without long-run risks (Mehra and Prescott, 1985; Tallarini, 2000). On this metric, we find more long-run risks than the original calibrations of the long-run risks model, which require a CRRA of 10 to fit the equity premium. Our model also generates highly volatile returns on equity. The standard deviation of equity returns for the U.S. is 18% in the model versus 17% in the data. Finally, the model generates large and persistent swings in the price-dividend ratio, similar to those observed in the data. For the U.S., the standard deviation of the price-dividend ratio in the model is 0.3 and its first-order autocorrelation is 0.85, while these statistics are 0.4 and 0.9 in the data, respectively.²⁷

One might worry that the model would generate counterfactually large variation in the risk-free rate owing to fluctuations in households' desire to save associated with the long-run risk shocks. This is, however, not the case. The standard deviation of the risk-free rate generated by our model is 1.6%. The standard deviation of ex post real returns on U.S. T-bills, our empirical measure of this statistic, is 3.3%. Since the model does not incorporate inflation risk, it is appropriate that the model yields a lower number than the data along this dimension.

Roughly half of the increase in the equity premium in our model results from the growth-rate shocks and the other half from the uncertainty shocks. This can be seen from Table 6. The table presents results on the equity premium and the risk free rate for all 16 countries in our sample. The full model generates equity premia ranging from 8-23% with an average equity premium of 13.7%. The second case preserves the baseline parameter values of the full model, but turns off the uncertainty shocks. This “constant volatility” model yields equity premia that are roughly half as large as the full model. The third case eliminates all long-run risks and re-calibrates the volatility

²⁷Table A.2 presents analogous results to Table 5 for our two alternative specifications. The simple model (without world components) generates a slightly smaller equity premium than the baseline case—roughly 4%. This specification matches the equity premium for a CRRA of 10. A higher CRRA is required because the simple model doesn't capture the persistent world component of consumption growth. The post-WWII case generates very similar results for the U.S. but a larger equity premium for the median country. This arises because the median country becomes more sensitive to the world growth-rate shock in this specification.

of the random-walk shocks to match the volatility of $\Delta\tilde{c}_{i,t}$. This case corresponds closely to the model considered by Mehra and Prescott (1985). It generates equity premia that are more than an order of magnitude smaller than the full model.

Our estimated model tends to generate a higher equity premium in countries in which the equity premium in the data has been higher. The correlation between the equity premium in our model and the equity premium in the data is 0.25. Spain is an outlier and if we exclude it this correlation rises to 0.50. In our model, cross-country variation in the equity premium is driven by variation in the extent to which growth-rates in different countries load on the world growth-rate process (ξ_{it}) and also because of variation in the average level of volatility across countries (σ_i^2).

We estimate a negative correlation between growth-rate and uncertainty shocks—i.e., negative growth-rate shocks tend to be associated with shocks that raise economic uncertainty. Since negative growth-rate shocks and shocks that increase uncertainty both raise marginal utility, being hit by both at the same time is particularly painful for the representative agent. This implies that the negative correlation between these two shocks contributes positively to the equity premium in our model. We have calculated asset prices for a case with $\lambda = \lambda_W = 0$ but keeping other parameters unchanged. This yields an equity premium that is 0.8 percentage points smaller for the U.S. than our baseline case.

Finally, we analyze the term structure implications of our model. We approximate long-term bonds by a perpetuity with coupon payments that decline over time by 10% per year. This yields a bond with a duration similar to that of 10-year coupon bonds. In our model, the term-premium for this real long-term bond is -2.0%. Piazzesi and Schneider (2006) document that the real yield curve in the United Kingdom has been downward sloping, while it has been mostly upward sloping in the United States. They caution, however, that this evidence is hard to assess because of the short sample and poor liquidity in the U.S. TIPS market.²⁸

²⁸Building on Alvarez and Jermann’s (2005) analysis of the implication of the term structure for the properties of the stochastic discount factor, Kojen et al. (2010) emphasize that the positive autocorrelation of growth rates in the long-run risk model implies that the model has a downward sloping term structure of real bond yields. Binsbergen et al. (2010a,b) show that short term dividend strips on the aggregate stock market have substantially higher expected returns than the stock market as a whole. (The price of a k-year dividend strip is the present value of the dividend paid in k years.) They point out that this fact is difficult to match using the original calibration of the long-run risk model proposed by Bansal and Yaron (2004). Croce, Lettau, and Ludvigson (2010) show that a model with long-run risk shocks that agents do not observe directly but must instead learn about over time can generate high excess returns on short-term assets relative to long-term assets.

6.3 Predictability of Returns, Consumption and Volatility

A large literature in finance has argued that a high price-dividend ratio predicts low stock returns (Campbell and Shiller, 1988; Fama and French, 1988; Hodrick, 1992; Cochrane, 2008; Binsbergen and Koijen, 2010).²⁹ Leading asset pricing models differ in their implications about return predictability. In the long-run risks model, uncertainty shocks cause variation in the price-dividend ratio on stocks that forecasts stock returns. More generally, variation in the price-dividend ratio on stocks comes from two sources in the long-run risks model: growth-rate shocks and uncertainty shocks. Consequently, the price-dividend ratio on stocks should forecast not only future returns on stocks but also future volatility and future consumption growth.

Table 7 presents results on the predictability of five-year excess returns on equity, realized volatility and consumption growth in our estimated models. We estimate equations of the following form

$$y_{i,t+5} = \alpha_i + \beta_i pd_{i,t} + \epsilon_{i,t+5}, \quad (11)$$

where $pd_{i,t}$ denotes the logarithm of the price-dividend ratio on equity and $y_{i,t+5}$ is one of three things: five-year excess returns on stocks, five-year realized volatility or five-year consumption growth.³⁰ We estimate these regressions in the data for the countries in our sample, and we run the same regressions on simulated datasets of the same length (120 years) from our baseline estimation and our simple model. We report the median from 1000 such simulations, as well as the 2.5% and 97.5% quantiles. For comparison, Table 7 also presents the degree of predictability of these variables in the models of Bansal and Yaron (2004) and Bansal et al. (2012).

The first panel of Table 7 presents results on the predictability of excess returns. Our point estimates imply a large degree of predictability of returns in the U.S. data. The regression coefficient on the price-dividend ratio is -0.41 and the R-squared of the regression is 0.24. We estimate less predictability of returns for the median country in our sample—the regression coefficient is -0.30 and the R-squared is 0.11. Our baseline case generates a median regression coefficient of -0.40 and a median R-squared of 0.10. The simple model yields similar results. Our model can thus account for a large fraction of the predictability of excess 5-year stock returns seen in the data. Our estimated model generates more predictability of excess stock returns than do the calibrations of the long-run

²⁹The statistical significance of return predictability has been hotly debated (see, e.g., Stambaugh, 1999; Ang and Bekaert, 2007). Recent work by Lewellen (2004) and Cochrane (2008) has exploited the stationarity of price-dividend ratios and the lack of predictability of dividend growth to develop more powerful tests of return predictability. These tests reject the null of no predictability of returns at the 1-2% level.

³⁰We follow Bansal et al. (2005) in using the absolute value of the residual from an AR(1) regression for consumption growth as our measure of realized volatility and summing this over five years.

risks model in Bansal and Yaron (2004) and Bansal et al. (2012).

The second panel of Table 7 presents results on the predictability of volatility. We find that the price-dividend ratio on stocks has substantial predictive power for realized volatility of consumption growth. For the U.S., the regression coefficient is -0.81 and the R-squared 0.32. For the median country in our sample, predictability of volatility is smaller, but nevertheless substantial—the regression coefficient is -0.38 and the R-squared is 0.19. These results are in line with earlier results by Bansal et al. (2005). Our model generates predictability of volatility that lines up well with the data. The regression coefficients for our baseline case is -0.37 and the R-squared is 0.07, while for the simple model we get a regression coefficient of -0.91 and an R-squared of 0.12. The values for the U.S. and for the median country are well within the 95% probability interval we construct.

Our model also implies a low frequency link between stock prices and macroeconomic uncertainty. Figure 8 plots our estimate of the evolution of economic uncertainty in the U.S. along with the dividend-price ratio on stocks. The figure illustrates the comovement between economic uncertainty and the value of the stock market emphasized by Lettau et al. (2008). Figure A.1 in the appendix presents analogous plots for all countries in our sample. This extends the results of Lettau et al. (2004) by including more countries and longer sample periods for several countries. The comovement of economic uncertainty and stock prices varies across countries and time. It is not very strong for most countries before 1970, but is stronger after this.

The third panel of Table 7 presents results on the predictability of consumption growth. The price-dividend ratio on stocks has little predictive power for consumption growth both in the U.S. or in the median country in our sample. These results extend earlier work by Beeler and Campbell (2012). Our estimated version of the long-run risks model generates somewhat more predictability of consumption growth than we see in the data. In the data, the regression coefficients for this regression is less than 0.05. In the model, the median regression coefficient across model runs is 0.19 in our baseline case. However, the empirical value lies within the 95% probability interval generated by our model. Our estimated model generates a degree of predictability of consumption growth that is intermediate between that in Bansal and Yaron (2004) and Bansal et al. (2012).

6.4 The Volatility of Real Exchange Rates

An important finding from our empirical analysis is that there is a large amount of comovement of growth-rates and uncertainty across countries. This has important implications for real exchange rates. In a world with complete markets, the log change in the real exchange rate between two

countries is

$$\Delta e_t = m_t^* - m_t, \tag{12}$$

where e_t denotes the log real exchange rate (home goods price of foreign goods), and m_t and m_t^f are the logarithm of the home and foreign stochastic discount factors, respectively. The annual standard deviation of changes in real exchange rates has been roughly 10% in the post-Bretton Woods period (see Table 8). However, Hansen and Jagannathan (1991) show that $\sigma(M_t)R_t^f \geq E(R_t^e)/\sigma(R_t^e)$, where M_t is the level of the stochastic discount factor and R^e is the excess return on the stock market. From Table 5 we can see that $R^f \simeq 1.01$, $E(R_t^e) \simeq 7\%$, and $\sigma(R_t^e) \simeq 18\%$, which implies $\sigma(M_t) \geq 40\%$. Brandt, Cochrane, and Santa-Clara (2006) point out that this logic combined with equation (12) implies that either m_t and m_t^* are highly correlated—i.e., there is a high degree of international risk sharing—or exchange rates are not as volatile as the theory predicts. In addition, the low degree of comovement of consumption growth across countries at short horizons suggests that stochastic discount factors are not highly correlated. Colacito and Croce (2011) refer to this as the international equity premium puzzle.

The common components of growth-rates and uncertainty that we estimate have the potential to resolve this puzzle. They generate comovement in the stochastic discount factors across countries that is not evident from the short-run comovement of consumption growth. Table 8 presents the standard deviation implied by our estimated model of annual changes in the bilateral real exchange rate versus the United State for each country in our sample. The table also presents a counterfactual for this statistic based on the same simulated data from our estimated model but ignoring the correlation between the stochastic discount factors of each country and the United States that is implied by our model—i.e., simply adding the variances of the two stochastic discount factors and taking a square root. We see that the presence of common long-run risk shocks in our model lowers the volatility of the real exchange rate by roughly a factor of two relative to what it would be if the stochastic discount factors were uncorrelated. Our model can therefore account for a large part of the discrepancy between the observed volatility of the real exchange rate and the volatility implied by a model in which marginal utility across countries is uncorrelated. Our results complement those of Colacito and Croce (2011), who carry out a related exercise for the exchange rate of the U.S. versus the U.K.

7 Conclusion

The long-run risks model is one of the leading frameworks of consumption-based asset pricing. It is difficult to obtain precise estimates of the key parameters of this model using even 100 years of macroeconomic data from a single country. As a consequence, previous work has used a combination of macroeconomic and asset price data to estimate the model. Our model of consumption dynamics allows for country-specific variation in the average level of volatility across countries, but pools across countries in estimating the persistence of growth-rate and uncertainty shocks as well as the volatility of shocks to uncertainty. This allows us to estimate long-run risk parameters using macroeconomic data alone. We can thereby avoid relying on a particular asset pricing model, and the concern that our estimates derive from a need to fit the asset pricing data.

Our estimates suggest that growth-rate and uncertainty shocks play an important role in asset pricing. We identify a large and persistent world growth-rate component and a less persistent country-specific growth-rate process. Shocks to uncertainty are highly persistent and yield substantial variation in uncertainty over time. With EZW preferences, current marginal utility depends not only on current consumption growth but also on news about future growth and uncertainty. With a $CRRA > 1$ and $IES > 1$, shocks that lower future expected growth or raise future economic uncertainty raise current marginal utility and cause stock prices to fall. This generates a substantial equity premium, high volatility of equity returns, and predictability of returns based on the price-dividend ratio.

A Model Estimation

We employ a Bayesian MCMC algorithm to estimate our model. More specifically, we employ a Metropolized Gibbs sampling algorithm to sample from the joint posterior distribution of the unknown parameters and variables conditional on the data. The full probability model we employ may be denoted by

$$f(Y, X, \Theta) = f(Y, X|\Theta)f(\Theta),$$

where $Y \in \{c_{i,t}, I_{i,t+1}^d\}$ is the set of observable variables for which we have data,

$$X \in \{z_{i,t}, x_{i,t}, x_{W,t}, \sigma_{i,t+1}^2, \sigma_{W,t+1}^2\}$$

is the set of unobservable variables, and

$$\Theta \in \{\rho, \rho_W, \gamma, \sigma_W^2, \sigma_\omega^2, \sigma_{W,\omega}^2, \lambda, \lambda_W, \xi_i, \chi_i, \sigma_i^2, \sigma_{\nu,i}^2, \mu_i, \mu_d, \}$$

is the set of parameters. From a Bayesian perspective, there is no real importance to the distinction between X and Θ . The only important distinction is between variables that are observed and those that are not. The function $f(Y, X|\Theta)$ is often referred to as the likelihood function of the model, while $f(\Theta)$ is often referred to as the prior distribution. Both $f(Y, X|\Theta)$ and $f(\Theta)$ are fully specified in sections 3 and 4 of the paper. The likelihood function may be constructed by combining equations (2)-(4) and (8), the distributional assumptions for the shocks in these equations detailed in section 3 and the assumptions about the distributions of $z_{i,t}$, $x_{i,t}$, $x_{W,t}$, $\sigma_{i,t}$, and $\sigma_{W,t}$ for the initial period for each country that are detailed in section 4. The prior distributions are described in detail in section 4.

The object of interest in our study is the distribution $f(X, \Theta|Y)$, i.e., the joint distribution of the unobservables conditional on the observed values of the observables. For expositional simplicity, let $\Phi = (X, \Theta)$. Using this notation, the object of interest is $f(\Phi|Y)$. The Gibbs sampler algorithm produces a sample from the joint distribution by breaking the vector of unknown variables into subsets and sampling each subvector sequentially conditional on the value of all the other unknown variables (see, e.g., Gelman et al., 2004, and Geweke, 2005). In our case we implement the Gibbs sampler as follows.

1. We derive the conditional distribution of each element of Φ conditional on all the other elements and conditional on the observables. For the i th element of Φ , we can denote this conditional distribution as $f(\Phi_i|\Phi_{-i}, Y)$, where Φ_i denotes the i th element of Φ and Φ_{-i}

denotes all but the i th element of Φ . In most cases, $f(\Phi_i|\Phi_{-i}, Y)$ are common distributions such as normal distributions or gamma distributions for which samples can be drawn in a computationally efficient manner. In cases where the Gibbs sampler cannot be applied, we use the Metropolis algorithm to sample values of $f(\Phi_i|\Phi_{-i}, Y)$.³¹

2. We propose initial values for all the unknown variables Φ . Let Φ^0 denote these initial values.
3. We cycle through Φ sampling Φ_i^t from the distribution $f(\Phi_i|\Phi_{-i}^{t-1}, Y)$ where

$$\Phi_{-i}^{t-1} = (\Phi_1^t, \dots, \Phi_{i-1}^t, \Phi_{i+1}^{t-1}, \dots, \Phi_d^{t-1})$$

and d denotes the number of elements in Φ . At the end of each cycle, we have a new draw Φ^t . We repeat this step N times to get a sample of N draws for Φ .

4. It has been shown that samples drawn in this way converge to the distribution $f(\Phi|Y)$ under very general conditions (see, e.g., Geweke, 2005). We assess convergence and throw away an appropriate burn-in sample.

In practice, we run four such “chains” starting two from one set of initial values and two from another set of initial values. We choose starting values that are far apart in the following way: For one chain, we set the initial values of $x_{i,t} = 0$ for all i and t . For the other chain, we set the initial values of $x_{i,t} = \Delta c_{i,t}$ for all i and t .

Given a sample from the joint distribution $f(\Phi|Y)$ of the unobserved variables conditional on the observed data, we can calculate any statistic of interest that involves Φ . For example, we can calculate the mean of any element of Φ by calculating the sample analogue of the integral

$$\int \Phi_i f(\Phi_i|\Phi_{-i}^{t-1}, Y) d\Phi_i.$$

³¹The Metropolis algorithm samples a proposal Φ_i^* from a proposal distribution $J_t(\Phi_i^*|\Phi_i^{t-1})$. This proposal distribution must be symmetric, i.e., $J_t(x_a|x_b) = J_t(x_b|x_a)$. The proposal is accepted with probability $\min(r, 1)$ where $r = f(\Phi_i^*|\Phi_{-i}, Y)/f(\Phi_i^{t-1}|\Phi_{-i}, Y)$. If the proposal is accepted, $\Phi_i^t = \Phi_i^*$. Otherwise $\Phi_i^t = \Phi_i^{t-1}$. Using the Metropolis algorithm to sample from $f(\Phi_i|\Phi_{-i}, Y)$ is much less efficient than the standard algorithms used to sample from known distributions such as the normal distribution in most software packages. Intuitively, this is because it is difficult to come up with an efficient proposal distribution. The proposal distribution we use is a normal distribution centered at Φ_i^{t-1} .

B Variance Ratios

Variance ratios are a simple tool to quantify the persistence of shocks to aggregate consumption (Cochrane, 1988). The k -period variance ratio for consumption growth is defined as the ratio of the variance of k -period consumption growth and 1-period consumption growth divided by k :

$$\text{VR}_{i,k} = \frac{1}{k} \frac{\text{var}\left(\sum_{j=0}^{k-1} \Delta c_{i,t-j}\right)}{\text{var}(\Delta c_{i,t})}. \quad (13)$$

The intuition for this statistic comes from the fact that for a simple random-walk process $\text{var}(c_{i,t} - c_{i,t-k})$ is equal to k times $\text{var}(c_{i,t} - c_{i,t-1})$, implying that the variance ratio for such a process is equal to one for all k . For a trend-stationary process, the variance ratio is less than one and falls toward zero as k increases. However, for a process that has persistent growth-rate shocks—i.e., positively autocorrelated growth rates—the variance ratio is larger than one.

Bansal and Yaron (2004) introduce a variance ratio statistic for assessing the persistence of shocks to volatility. They first compute the innovations to consumption growth $u_{i,t}$ as the residuals from an AR(5) regression and use the absolute value of these innovations $|u_{i,t}|$ as a measure of realized volatility of consumption growth. They then construct variance ratios for $|u_{i,t}|$,

$$\text{VR}_{i,k}^u = \frac{1}{k} \frac{\text{var}\left(\sum_{j=0}^{k-1} |u_{i,t-j}|\right)}{\text{var}(|u_{i,t}|)}. \quad (14)$$

This statistic provides a rough measure of the persistence of stochastic volatility. As with the variance ratio for consumption growth, if this variance ratio is above one, it indicates that uncertainty shocks have persistent effects on volatility—i.e., high volatility periods are “bunched together” leading to a high value of the variance in the numerator.

References

- ABEL, A. B. (1999): “Risk Premia and Term premia in General Equilibrium,” *Journal of Monetary Economics*, 43, 3–33.
- AGUIAR, M., AND G. GOPINATH (2007): “Emerging Market Business Cycles: The Cycle Is the Trend,” *Journal of Political Economy*, 115(1), 69–102.
- ALVAREZ, F., AND U. J. JERMANN (2005): “Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth,” *Econometrica*, 73(6), 1977–2016.
- ANG, A., AND G. BEKAERT (2007): “Stock Return Predictability: Is it There?,” *Review of Financial Statistics*, 20(3), 651–707.
- BALKE, N. S., AND R. J. GORDON (1989): “The Estimation of Prewar Gross National Product: Methodology and New Evidence,” *Journal of Political Economy*, 97, 38–92.
- BANSAL, R., R. F. DITTMAR, AND C. T. LUNDBLAD (2005): “Consumption, Dividends, and the Cross Section of Equity Returns,” *Journal of Finance*, 60(4), 1639–1672.
- BANSAL, R., V. KHATCHATRIAN, AND A. YARON (2005): “Interpretable Asset Markets?,” *European Economic Review*, 49, 531–560.
- BANSAL, R., D. KIKU, AND A. YARON (2007): “Risks For the Long Run: Estimation and Inference,” Working Paper, Duke University.
- (2012): “An Empirical Evaluation of the Long-Run Risks Model for Asset Prices,” *Critical Finance Review*, 1, 183–221.
- BANSAL, R., AND I. SHALIASTOVICH (2010): “A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets,” Working Paper, Duke University.
- BANSAL, R., AND A. YARON (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 59(4), 1481–1509.
- BARRO, R. (2006): “Rare Disasters and Asset Markets in the Twentieth Century,” *Quarterly Journal of Economics*, 121(3), 832–866.
- BARRO, R. J., AND J. F. URSUA (2008): “Macroeconomic Crises since 1870,” *Brookings Papers on Economic Activity*, 2008, 255–350.
- BARSKY, R. B., F. T. JUSTER, M. S. KIMBALL, AND M. D. SHAPIRO (1997): “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study,” *Quarterly Journal of Economics*, 112(2), 537–579.
- BARSKY, R. B., AND E. R. SIMS (2010): “News Shocks and Business Cycles,” Working Paper.
- BASU, S., AND B. BUNDICK (2011): “Uncertainty Shocks in a Model of Effective Demand,” Working Paper, Boston College.
- BEAUDRY, P., AND F. PORTIER (2006): “Stock Prices, News and Economic Fluctuations,” *American Economic Review*, 96(4), 1293–1307.

- BEELER, J., AND J. Y. CAMPBELL (2012): “The Long-Run Risk Model and Aggregate Asset Prices: An Empirical Assessment,” *Critical Finance Review*, 1, 141–182.
- BINSBERGEN, J. H., M. W. BRANDT, AND R. S. KOIJEN (2010a): “On the Timing and Pricing of Dividends,” Working Paper, Northwestern University.
- BINSBERGEN, J. H., W. HUESKES, R. S. KOIJEN, AND E. VRUGT (2010b): “Asset Pricing Puzzles: Measuring the Importance of Short-Term and Long-Term Risk,” Working Paper, Northwestern University.
- BINSBERGEN, J. H., AND R. S. KOIJEN (2010): “Predictive Regressions: A Present-Value Approach,” *Journal of Finance*, 65(4), 1439–1471.
- BLANCHARD, O., AND J. SIMON (2001): “The Long and Large Decline in U.S. Output Volatility,” *Brookings Papers on Economic Activity*, 2001(1), 135–164.
- BLANCHARD, O. J., J.-P. L’HUILIER, AND G. LORENZONI (2011): “New, Noise, and Fluctuations: An Empirical Exploration,” Working Paper, MIT.
- BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77(3), 623–685.
- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. TERRY (2011): “Really Uncertain Business Cycles,” Working Paper, Stanford University.
- BONOMO, M., R. GARCIA, N. MEDDAHI, AND R. TEDONGAP (2011): “Generalized Disappointment Aversion, Long-run Volatility Risk, and Asset Prices,” *Review of Financial Studies*, 24(1), 83–122.
- BOROVICKA, J., L. P. HANSEN, M. HENDRICKS, AND J. A. SCHEINKMAN (2011): “Risk-Price Dynamics,” *Journal of Financial Econometrics*, 9(1), 3–65.
- BRANDT, M. W., J. H. COCHRANE, AND P. SANTA-CLARA (2006): “International Risk Sharing Is Better than You Think, or Exchange Rates Are too Smooth,” *Journal of Monetary Economics*, 53, 671–698.
- CAMPBELL, J. Y. (1993): “Intertemporal Asset Pricing without Consumption Data,” *American Economic Review*, 83(3), 487–512.
- (1999): “Asset Prices, Consumption and the Business Cycle,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, Amsterdam, Holland. Elsevier.
- CAMPBELL, J. Y., AND J. H. COCHRANE (1999): “By force of habit: A consumption-based explanation of aggregate stock market behavior,” *Journal of Political Economy*, 107, 205–251.
- CAMPBELL, J. Y., AND N. G. MANKIW (1989): “International Evidence on the Persistence of Economic Fluctuations,” *Journal of Monetary Economics*, 23, 319–333.
- CAMPBELL, J. Y., AND R. J. SHILLER (1988): “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *The Review of Financial Studies*, 1(3), 195–228.
- CHETTY, R. (2006): “A New Method of Estimating Risk Aversion,” *American Economic Review*, 96(5), 1821–1834.

- COCHRANE, J. H. (1988): “How Big Is the Random Walk in GNP?,” *Journal of Political Economy*, 96(5), 893–920.
- (2008): “The Dog That Did Not Bark: A Defense of Return Predictability,” *Review of Financial Studies*, 21(4), 1533–1575.
- COGLEY, T. (1990): “International Evidence on the Size of the Random Walk in Output,” *Journal of Political Economy*, 98(3), 501–518.
- COLACITO, R., AND M. M. CROCE (2011): “Risks for the Long-Run and the Real Exchange Rate,” *Journal of Political Economy*, 119(1), 153–181.
- COLLIN-DUFRESNE, P., M. JOHANNES, AND L. A. LOCHSTOER (2012): “Parameter Learning in General Equilibrium: The Asset Pricing Implications,” Working Paper, Columbia University.
- CONSTANTINIDES, G., AND D. DUFFIE (1996): “Asset Pricing with heterogeneous consumers,” *Journal of Political Economy*, 104(2), 219–240.
- CONSTANTINIDES, G. M., AND A. GHOSH (2009): “Asset Pricing Tests with Long-Run Risks in Consumption Growth,” Working Paper, University of Chicago.
- CROCE, M. M., M. LETTAU, AND S. C. LUDVIGSON (2010): “Investor Information, Long-Run Risk, and the Duration of Risky Cash Flows,” Working Paper, UNC, Chapel Hill.
- EPSTEIN, L. G., AND S. E. ZIN (1989): “Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns,” *Econometrica*, 57, 937–969.
- FAMA, E. F., AND K. R. FRENCH (1988): “Dividend Yields and Expected Stock Returns,” *Journal of Financial Economics*, 22, 3–27.
- FERNANDEZ-VILLAYERDE, J., P. GUERRON-QUINTANA, J. RUBIO-RAMIREZ, AND M. URIBE (2011): “Risk Matters: The Real Effects of Volatility Shocks,” *American Economic Review*, 101, 2530–2561.
- GARLEANU, N., AND S. PANAGEAS (2010): “Young, Old, Conservative and Bold. The Implications of Finite Lives and Heterogeneity for Asset Pricing,” Working Paper, University of Chicago.
- GELMAN, A., J. B. CARLIN, H. S. STERN, AND D. B. RUBIN (2004): *Bayesian Data Analysis*. John Wiley and Sons, Hoboken, New Jersey.
- GELMAN, A., AND D. B. RUBIN (1992): “Inference from Iterative Simulation Using Multiple Sequences,” *Statistical Science*, 7, 457–511.
- GEWEKE, J. (2005): *Contemporary Bayesian Econometrics and Statistics*. Chapman & Hall/CRC, Boca Raton, Florida.
- GRUBER, J. (2006): “A Tax-Based Estimate of the Elasticity of Intertemporal Substitution,” NBER Working Paper No. 11945.
- HALL, R. E. (1988): “Intertemporal Substitution in Consumption,” *Journal of Political Economy*, 96(2), 339–357.

- HANSEN, L. P. (2007): “Beliefs, Doubts and Learning: Valuing Macroeconomic Risk,” *American Economic Review*, 97(2), 1–30.
- HANSEN, L. P., J. HEATON, J. LEE, AND N. ROUSSANOV (2007): “Intertemporal Substitution and Risk Aversion,” in *Handbook of Econometrics*, ed. by J. J. Heckman, and E. E. Leamer, pp. 3968–4056, Amsterdam, Holland. Elsevier.
- HANSEN, L. P., J. C. HEATON, AND N. LI (2008): “Consumption Strikes Back? Measuring Long-Run Risk,” *Journal of Political Economy*, 116(2), 260–302.
- HANSEN, L. P., AND R. JAGANNATHAN (1991): “Implications of Security Market Data for Models of Dynamic Economies,” *Journal of Political Economy*, 99, 225–262.
- HANSEN, L. P., AND T. J. SARGENT (2010): “Fragile Beliefs and the Price of Uncertainty,” *Quantitative Economics*, 1(1), 129–162.
- HANSEN, L. P., AND K. J. SINGLETON (1982): “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models,” *Econometrica*, 50, 1269–1288.
- HODRICK, R. J. (1992): “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement,” *Review of Financial Studies*, 5(3), 357–386.
- JAIMOVICH, N., AND S. REBELO (2009): “Can News about the Future Drive the Business Cycle?,” *American Economic Review*, 99(4), 1097–1118.
- KALTENBRUNNER, G., AND L. A. LOCHSTOER (2010): “Long-Run Risk through Consumption Smoothing,” *Review of Financial Studies*, 23, 3141–3189.
- KANDEL, S., AND R. F. STAMBAUGH (1990): “Expectations and Volatility of Consumption and Asset Returns,” *Review of Financial Studies*, 3(2), 207–232.
- KILIAN, L., AND L. E. OHANIAN (2002): “Unit Roots, Trend Breaks, and Transitory Dynamics: A Macroeconomic Perspective,” *Macroeconomic Dynamics*, 6(5), 614–632.
- KOIJEN, R. S., H. LUSTIG, S. V. NIEUWERBURGH, AND A. VERDELHAN (2010): “Long Run Risk, the Wealth-Consumption Ratio, and the Temporal Pricing of Risk,” *American Economic Review*, 100(2), 552–556.
- KUNG, H., AND L. SCHMID (2011): “Innovation, Growth, and Asset Pricing,” Working Paper, Duke University.
- LETTAU, M., S. C. LUDVIGSON, AND J. A. WACHTER (2004): “The Declining Equity Premium: What Role Does Macroeconomic Risk Play,” NBER Working Paper No. 10270.
- (2008): “The Declining Equity Premium: What Role Does Macroeconomic Risk Play,” *The Review of Financial Studies*, 21(4), 1653–1687.
- LEWELLEN, J. W. (2004): “Predicting Returns with Financial Ratios,” *Journal of Financial Economics*, 74, 209–235.
- LUCAS, R. E. (1978): “Asset Prices in an Exchange Economy,” *Econometrica*, 46(6), 1429–1445.

- MALLOY, C. J., T. J. MOSKOWITZ, AND A. VISSING-JORGENSEN (2009): “Long-Run Stockholder Consumption Risk and Asset Returns,” *Journal of Finance*, 64(6), 2427–2479.
- MCCONNELL, M. M., AND G. PEREZ-QUIROS (2000): “Output Fluctuations in the United States: What Has Changed since the Early 1980’s,” *American Economic Review*, 90(5), 1464–1476.
- MEHRA, R., AND E. C. PRESCOTT (1985): “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 15, 145–161.
- NAKAMURA, E., J. STEINSSON, R. BARRO, AND J. URSUA (2010): “Crises and Recoveries in an Empirical Model of Consumption Disasters,” Working Paper, Columbia University.
- NELSON, C. R., AND C. I. PLOSSER (1982): “Trends and Random Walk in Macroeconomic Time Series: Some Evidence and Implications,” *Journal of Monetary Economics*, 10, 139–162.
- PARAVISINI, D., V. RAPPOPORT, AND E. RAVINA (2010): “Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios,” NBER Working Paper No. 16063.
- PASTOR, L., AND P. VERONESI (2009): “Learning in Financial Markets,” *Annual Review of Financial Economics*, 1, 361–381.
- PIAZZESI, M., AND M. SCHNEIDER (2006): “Equilibrium Yield Curves,” in *NBER Macroeconomics Annual*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, pp. 389–442, Cambridge, Ma. MIT Press.
- ROMER, C. D. (1986): “Is the Stabilization of the Postwar Economy a Figment of the Data,” *American Economic Review*, 76(3), 314–334.
- SCHMITT-GROHE, S., AND M. URIBE (2010): “What’s News in Business Cycles,” Working Paper, Columbia University.
- STAMBAUGH, R. F. (1999): “Predictive Regressions,” *Journal of Financial Economics*, 54, 375–421.
- STOCK, J. H., AND M. W. WATSON (2002): “Has the Business Cycle Changed and Why?,” in *NBER Macroeconomics Annual*, ed. by M. Gertler, and K. Rogoff, pp. 159–218, Cambridge, MA. MIT Press.
- TALLARINI, T. D. (2000): “Risk-sensitive real business cycles,” *Journal of Monetary Economics*, 45(3), 507–532.
- TIMMERMANN, A. G. (1993): “How Learning in Financial Markets Generates Excess Volatility and Predictability of Stock Returns,” *Quarterly Journal of Economics*, 108, 1135–1145.
- UHLIG, H. (2007): “Leisure, Growth and Long Run Risk,” Working Paper, University of Chicago.
- URSUA, J. (2010): “Long-Run Volatility,” Working Paper, Harvard University.
- WACHTER, J. A. (2005): “A Consumption-Based Model of the Term Structure of Interest Rates,” *Journal of Financial Economics*, 79, 365–399.
- WEIL, P. (1990): “Unexpected Utility in Macroeconomics,” *Quarterly Journal of Economics*, 105(1), 29–42.

TABLE I
Estimates for Pooled Parameters

	Prior	Baseline	Simple Model	Post-WWII
<i>Persistence:</i>				
Country-Specific Growth-Rate Shocks (ρ)	0.500 (0.286)	0.565 (0.046)	0.682 (0.038)	0.622 (0.060)
World Growth-Rate Shocks (ρ_W)	0.500 (0.286)	0.832 (0.077)	--	0.832 (0.093)
Stochastic Volatility (γ)	0.493 (0.281)	0.970 (0.011)	0.950 (0.028)	0.963 (0.024)
<i>Standard Deviations:</i>				
Mean of World Stoch. Vol. Process (σ_W)	0.0667 (0.0236)	0.0053 (0.0028)	--	0.0032 (0.0025)
Country-Specific Stoch. Vol. Shock (σ_ω)	0.000667 (0.000236)	0.000028 (0.000007)	0.000054 (0.000015)	0.000037 (0.000011)
World Stoch. Vol. Shock ($\sigma_{\omega,W}$)	0.000667 (0.000236)	0.000024 (0.000009)	--	0.000007 (0.000007)
<i>Correlations:</i>				
Country-Specific (λ)	0.00 (0.57)	-0.40 (0.17)	--	-0.34 (0.20)
World (λ_W)	0.00 (0.57)	-0.25 (0.28)	--	-0.32 (0.32)

The table reports prior and posterior means of the parameters with prior and posterior standard deviations in parentheses. The "Baseline" case is for our full model estimated on data from 1890-2009. The "Simple Model" case is for our simple model estimated on data from 1890-2009. The "Post-WWII" case is for our full model estimated on data from 1950-2009.

TABLE II
Half-Life of Growth-Rate and Uncertainty Shocks

	Half-Life in Years		
	Growth-Rate Process		Uncertainty Process
	Country-Specific ($x_{i,t}$)	World ($x_{W,t}$)	($\sigma^2_{i,t}$ and $\sigma^2_{W,t}$)
Baseline	1.2	3.8	22.8
Simple Model	1.8	--	13.5
Post-WWII	1.5	3.8	18.2
Bansal and Yaron (2004)	2.7	--	4.4
Bansal, Kiku and Yaron (2012)	2.3	--	57.7

TABLE III
Estimates for Country-Specific Parameters

	Prior	Baseline		Simple Model		Post-WWII	
		Median	U.S.	Median	U.S.	Median	U.S.
Rel. St. Dev. of Random Walk Shock (χ_i)	3.38 (1.18)	0.81 (0.45)	1.16 (0.44)	0.87 (0.45)	0.99 (0.50)	1.09 (0.60)	1.09 (0.60)
Sensitivity to Common Shocks (ξ_i)	5.00 (2.89)	1.51 (0.55)	1.00 (0.00)	-- --	-- --	4.65 (1.62)	1.00 (0.00)
Average Growth (μ_i)	0.015 (1.00)	0.016 (0.008)	0.018 (0.006)	0.016 (0.004)	0.017 (0.004)	0.022 (0.015)	0.021 (0.006)
<i>Standard Deviations:</i>							
Average Stochastic Volatility (σ_i)	0.067 (0.024)	0.009 (0.004)	0.009 (0.004)	0.013 (0.005)	0.012 (0.005)	0.011 (0.004)	0.011 (0.004)
Post-1945 Transitory Shock (σ_{vi})	0.067 (0.024)	0.004 (0.002)	0.003 (0.001)	0.004 (0.002)	0.004 (0.002)	0.003 (0.002)	0.003 (0.002)
Pre-1945 Transitory Shock (σ_{vi})	0.067 (0.024)	0.024 (0.005)	0.023 (0.005)	0.022 (0.005)	0.023 (0.005)	-- --	-- --

The table reports prior and posterior means of the parameters with prior and posterior standard deviations in parentheses. The "Baseline" case is for our full model estimated on data from 1890-2009. The "Simple Model" case is for our simple model estimated on data from 1890-2009. The "Post-WWII" case is for our full model estimated on data from 1950-2009. "Median" refers to the median country. In other words, we report the value of each statistic - both means and standard deviations - for the country that has the median value of that statistic.

TABLE IV
Properties of Consumption Growth

	Median Country			United States		
	Data	Model Median	Model [2.5%, 97.5%]	Data	Model Median	Model [2.5%, 97.5%]
AC(1)	0.13	0.25	[0.01,0.51]	-0.08	0.07	[-0.21,0.35]
AC(2)	0.14	0.28	[0.12,0.50]	0.15	0.18	[-0.06,0.40]
AC(3)	0.04	0.21	[0.06,0.42]	-0.21	0.12	[-0.11,0.35]
AC(4)	0.07	0.15	[0.01,0.36]	0.28	0.09	[-0.15,0.31]
AC(5)	0.00	0.11	[-0.02,0.32]	-0.09	0.07	[-0.16,0.28]
AC(10)	0.12	0.01	[-0.13,0.17]	0.12	0.01	[-0.21,0.22]
CrossC(1)	0.23	0.34	[0.17,0.56]	0.18	0.25	[0.08,0.47]
CrossC(5)	0.44	0.65	[0.37,0.84]	0.43	0.54	[0.18,0.79]
CrossC(10)	0.56	0.73	[0.42,0.90]	0.54	0.64	[0.20,0.87]
VR(15) ΔC	1.62	2.57	[1.18, 5.17]	1.29	1.70	[0.56, 4.14]
VR(15) Vol	2.14	1.76	[1.01, 3.04]	1.80	1.95	[0.68, 4.48]

The Table reports autocorrelations, cross-country correlations and variance ratios for the real-world data and simulated data from the model (excluding disasters in both cases). The first through sixth rows present the autocorrelation of one year through five year and ten year consumption growth. The next three rows present cross-country correlations of one, five and ten year consumption growth. The last two rows present the fifteen year variance ratio of consumption growth and the realized volatility of consumption growth. For the cross-country correlations, the median country results are the median of the 120 cross-country correlations across our 16 countries. For the results based on data from the model, we simulate 1000 datasets from the model of the same size as the actual data. For each such simulation we calculate the median across countries as well as the value for the U.S. for each statistic. We then report the median along with the 2.5% and 97.5% quantiles across simulations for each of these statistic.

TABLE V
Asset Pricing Statistics

	Data		Baseline Model	
	Median	U.S.	Median	U.S.
$E(R_m - R_f)$	6.87	7.10	11.07	7.41
$\sigma(R_m - R_f)$	21.82	17.37	21.69	17.88
$E(R_m - R_f)/\sigma(R_m - R_f)$	0.32	0.41	0.51	0.41
$E(R_m)$	9.10	8.23	11.83	8.83
$\sigma(R_m)$	21.99	17.89	21.65	17.84
$E(R_f)$	1.43	1.13	0.66	1.41
$\sigma(R_f)$	4.57	3.33	2.26	1.74
$E(p-d)$	3.30	3.30	2.58	2.94
$\sigma(p-d)$	0.41	0.40	0.36	0.30
AC1(p-d)	0.85	0.90	0.85	0.85

Columns labeled as "Median" report the result for the median country for each statistic. Columns labeled as "U.S." report these statistics for the United States. For returns the statistics we report are the unconditional average of the level of the ex-post real net return in percentage points (i.e., multiplied by 100). R_m denotes the return on equity (the market), while R_f denotes the return on a short term nominal government bond (risk-free rate). The last three rows report statistics for the logarithm of the price-dividend ratio on equity. For the model, these results are for a CRRA = 6.5, IES = 1.5, and subjective discount factor of $\beta = 0.99$, and are calculated using a sample of length 1 million years.

TABLE VI
The Equity Premium and Risk-Free Rate Across Countries and Models

	Equity Premium				Risk-Free Rate	
	Data	Full Model	Constant Volatility	Mehra-Prescott	Data	Full Model
Australia	0.090	0.083	0.036	0.005	0.007	0.010
Belgium	0.068	0.137	0.064	0.006	0.012	0.000
Canada	0.061	0.098	0.045	0.008	0.013	0.009
Denmark	0.043	0.094	0.045	0.005	0.028	0.010
Finland	0.128	0.193	0.105	0.014	-0.001	-0.006
France	0.078	0.123	0.056	0.006	-0.018	0.005
Germany	0.101	0.110	0.053	0.008	-0.027	0.007
Italy	0.061	0.153	0.088	0.008	-0.008	0.000
Netherlands	0.067	0.147	0.072	0.007	0.007	0.001
Norway	0.058	0.111	0.052	0.007	0.013	0.008
Portugal	0.089	0.187	0.094	0.016	0.001	-0.005
Spain	0.051	0.221	0.116	0.011	0.006	-0.011
Sweden	0.073	0.099	0.046	0.004	0.019	0.011
Switzerland	0.056	0.084	0.037	0.002	0.009	0.009
United Kingdom	0.054	0.104	0.048	0.005	0.013	0.007
United States	0.075	0.074	0.033	0.005	0.009	0.014
Average	0.072	0.126	0.062	0.007	0.005	0.004
Median	0.067	0.111	0.053	0.006	0.008	0.007

The table presents asset pricing statistics based on simulated data from our model as well as from the historical data. The "Constant Volatility" model is a version of the full model where we "turn off" the stochastic volatility by setting the volatility of the uncertainty shocks ω and ω_w to zero but keep other parameters at their estimated values for the full model. For the "Mehra-Prescott" model we "turn off" both the stochastic volatility and the growth-rate shocks and then we recalibrate the random-walk shocks based on the volatility of permanent consumption in the full model. These results are for a CRRA = 6.5, IES = 1.5 and subjective discount factor of $\beta = 0.99$.

TABLE VII
Predictability Regressions

	Data		Baseline (U.S.)		Simple Model (U.S.)		BY	BKY
	Median	U.S.	Median	95% Prob. Int.	Median	95% Prob. Int.	Median	Median
<u>5 Year Excess Returns on Price Dividend Ratio</u>								
β	-0.30	-0.41	-0.40	[-1.01, 0.18]	-0.44	[-1.06, 0.19]	-0.23	-0.39
R^2	0.11	0.24	0.10	[0.00, 0.39]	0.09	[0.00, 0.40]	0.03	0.05
<u>5 Year Realized Volatility on Price-Dividend Ratio</u>								
β	-0.38	-0.81	-0.37	[-1.35, 0.48]	-0.91	[-2.25, 0.25]	-0.10	-0.83
R^2	0.19	0.32	0.06	[0.00, 0.35]	0.12	[0.00, 0.44]	0.02	0.13
<u>5 Year Consumption Growth on Price-Dividend Ratio</u>								
β	0.03	0.02	0.19	[0.01, 0.36]	0.18	[-0.03, 0.39]	0.35	0.12
R^2	0.04	0.02	0.26	[0.00, 0.64]	0.18	[0.00, 0.55]	0.32	0.08

The table reports results from regressions of excess returns, consumption growth and realized volatility at a 1, 3 and 5 year horizon on the price-dividend ratio. Our measure of realized volatility is the absolute value of the residual from an AR(1) model for consumption growth. The first two columns report results using data from our 16 country sample and the U.S., respectively. The first column is the median across countries of the statistic in question. The next two columns report results from our model. The last two columns report results for the models of Bansal and Yaron (2004) and Bansal, Kiku and Yaron (2012). The results for the Bansal-Yaron model are taken from Beeler and Campbell (2009). We use the end of year convention for the timing of consumption, whereby time t consumption is assumed to occur at the end of year t .

TABLE VIII
World Long-Run Risks and Real Exchange Rate Volatility

	Exchange Rate Volatility		
	Data	Baseline Estimation	Ignoring Correlation
Australia	0.09	0.37	0.79
Belgium	0.11	0.51	1.06
Canada	0.05	0.40	0.85
Denmark	0.10	0.38	0.87
Finland	0.10	0.71	1.22
France	0.10	0.45	1.00
Germany	0.10	0.42	0.94
Italy	0.10	0.58	1.13
Netherlands	0.10	0.56	1.11
Norway	0.08	0.42	0.94
Portugal	0.10	0.69	1.21
Spain	0.11	0.88	1.43
Sweden	0.11	0.38	0.89
Switzerland	0.11	0.34	0.82
United Kingdom	0.09	0.39	0.91
Average	0.10	0.50	1.01
Median	0.10	0.42	0.94

The table presents the standard deviation of the log change in the real exchange rate of each country with the United States. First, it presents results based on historical data from 1975-2009. Second, it presents results based on simulated data from our baseline estimates. The last column calculates counterfactual exchange rates based on the simulated data from our estimated model but ignoring the correlation between the stochastic discount factors of the two countries in question.

TABLE A.1
Estimates of Country-Specific Parameters

	Rel. St. Dev. Random Walk Shock (χ_i)		Sensitivity to Common Shocks (ξ_i)		Average St. Dev. Stoch. Vol. (σ_i)		St. Dev. Transitory Shock (σ_{vi})				Average Growth (μ_i)	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.	post-1945		pre-1945		Mean	St. Dev.
							Mean	St. Dev.	Mean	St. Dev.		
Australia	1.73	0.61	1.05	0.42	0.008	0.004	0.004	0.002	0.035	0.009	0.015	0.006
Belgium	1.01	0.50	1.95	0.56	0.007	0.003	0.004	0.002	0.020	0.008	0.012	0.009
Canada	1.96	0.63	1.17	0.41	0.009	0.004	0.003	0.001	0.028	0.009	0.018	0.006
Denmark	0.79	0.41	1.29	0.58	0.012	0.004	0.007	0.003	0.012	0.003	0.016	0.007
Finland	2.00	1.23	2.30	0.83	0.012	0.006	0.004	0.002	0.024	0.007	0.022	0.012
France	0.83	0.42	1.73	0.46	0.007	0.004	0.002	0.001	0.027	0.004	0.015	0.008
Germany	0.63	0.34	1.54	0.53	0.012	0.004	0.003	0.001	0.012	0.004	0.015	0.008
Italy	0.58	0.30	2.16	0.65	0.011	0.004	0.003	0.002	0.015	0.003	0.017	0.011
Netherlands	0.59	0.32	2.09	0.63	0.010	0.004	0.003	0.002	0.023	0.004	0.016	0.010
Norway	1.17	0.57	1.49	0.60	0.009	0.004	0.006	0.003	0.006	0.003	0.019	0.008
Portugal	2.59	0.80	2.27	0.68	0.008	0.004	0.005	0.003	0.030	0.009	0.021	0.011
Spain	0.76	0.45	3.24	0.86	0.010	0.004	0.002	0.001	0.048	0.008	0.019	0.016
Sweden	0.72	0.47	1.36	0.52	0.010	0.004	0.004	0.002	0.024	0.005	0.018	0.007
Switzerland	0.64	0.44	1.21	0.42	0.009	0.004	0.001	0.001	0.043	0.006	0.011	0.006
United Kingdom	0.63	0.30	1.48	0.49	0.010	0.004	0.004	0.002	0.006	0.002	0.013	0.008
United States	1.16	0.44	1.00	0.00	0.009	0.004	0.003	0.001	0.023	0.005	0.018	0.006
Average	1.11	0.51	1.71	0.54	0.010	0.004	0.004	0.002	0.024	0.006	0.017	0.009
Median	0.81	0.45	1.51	0.55	0.009	0.004	0.004	0.002	0.024	0.005	0.016	0.008

The table presents our estimates of the posterior mean and standard deviation of the country-specific parameters in our full model.

TABLE A.2
Asset Pricing Statistics

	Data		Baseline		Simple Model		Post-WWII	
	Median	U.S.	Median	U.S.	Median	U.S.	Median	U.S.
$E(R_m - R_f)$	6.87	7.10	11.07	7.41	4.13	3.98	20.74	6.39
$\sigma(R_m - R_f)$	21.82	17.37	21.69	17.88	13.76	13.53	31.38	16.99
$E(R_m - R_f) / \sigma(R_m - R_f)$	0.32	0.41	0.51	0.41	0.30	0.29	0.66	0.38
$E(R_m)$	9.10	8.23	11.83	8.83	5.68	5.68	19.94	8.13
$\sigma(R_m)$	21.99	17.89	21.65	17.84	13.81	13.55	31.40	16.86
$E(R_f)$	1.43	1.13	0.66	1.41	1.54	1.71	-0.77	1.74
$\sigma(R_f)$	4.57	3.33	2.26	1.74	1.41	1.34	3.51	1.50
$E(p-d)$	3.30	3.30	2.58	2.94	3.45	3.52	2.08	3.10
$\sigma(p-d)$	0.41	0.40	0.36	0.30	0.20	0.19	0.48	0.26
AC1(p-d)	0.85	0.90	0.85	0.85	0.79	0.79	0.84	0.81

Columns labeled as "Median" report the result for the median country for each statistic. Columns labeled as "U.S." report these statistics for the United States. For returns the statistics we report are the unconditional average of the level of the ex-post real net return in percentage points (i.e., multiplied by 100). R_m denotes the return on equity (the market), while R_f denotes the return on a short term nominal government bond (risk-free rate). The last three rows report statistics for the logarithm of the price-dividend ratio on equity. For the model, these results are for a CRRA = 6.5, IES = 1.5, and subjective discount factor of $\beta = 0.99$, and are calculated using a sample of length 1 million years.

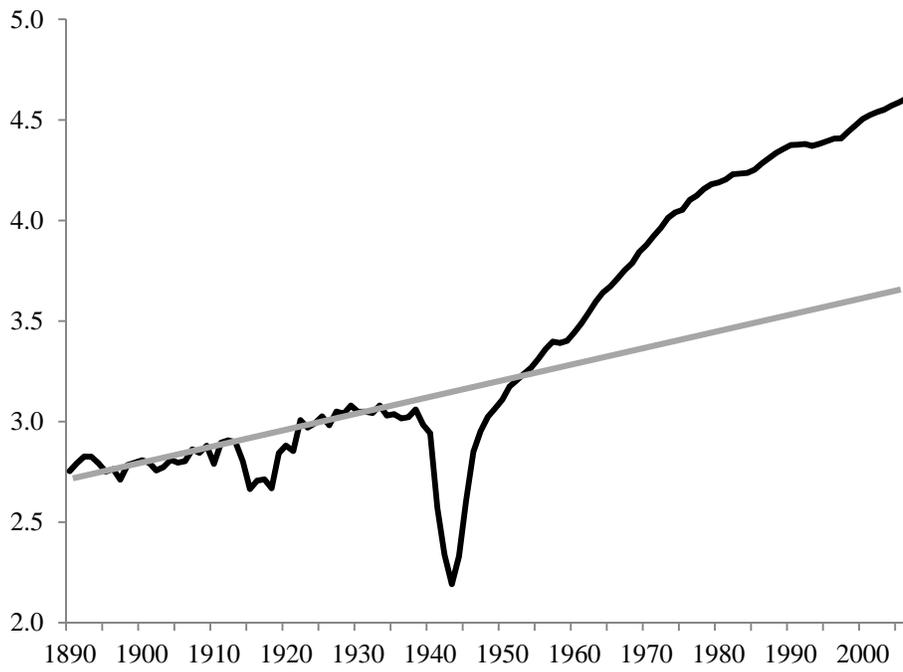


FIGURE I
Log per Capita Consumption in France

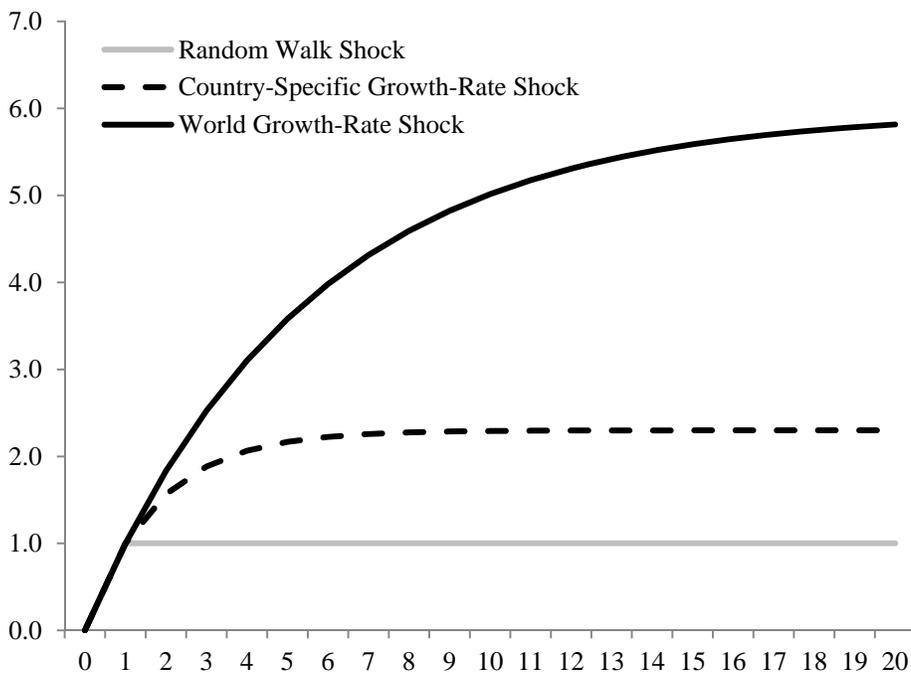


FIGURE II
Response of Consumption to Growth-Rate and Random-Walk Shocks

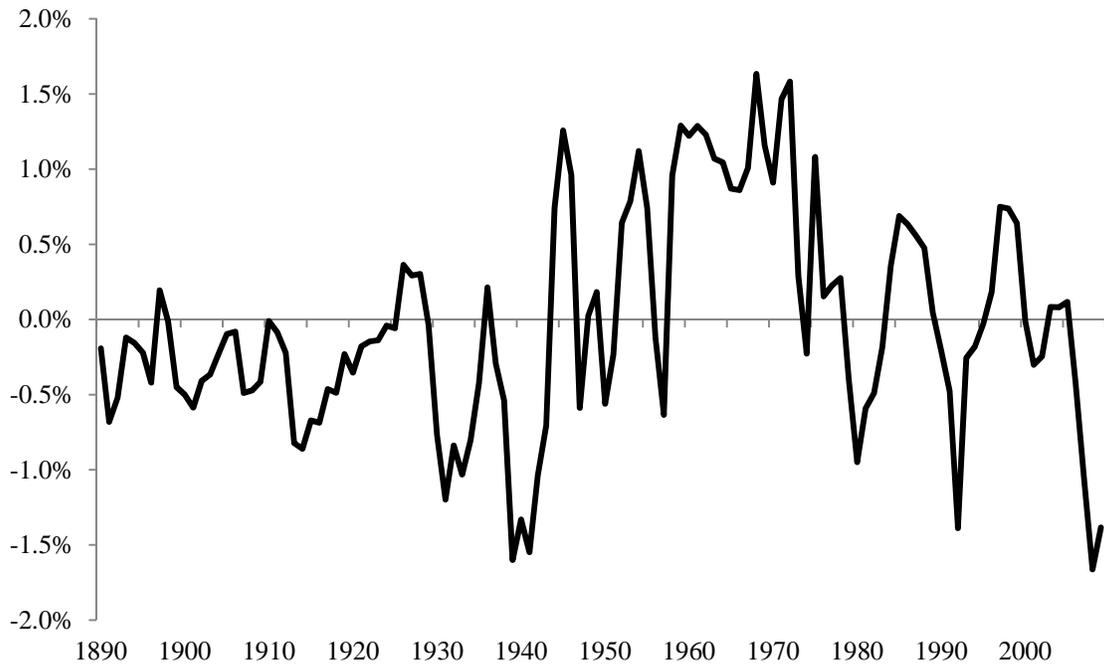


FIGURE III
 The World Growth-Rate Process
 The figure plots the posterior mean value of $x_{w,t}$ for each year in our sample.

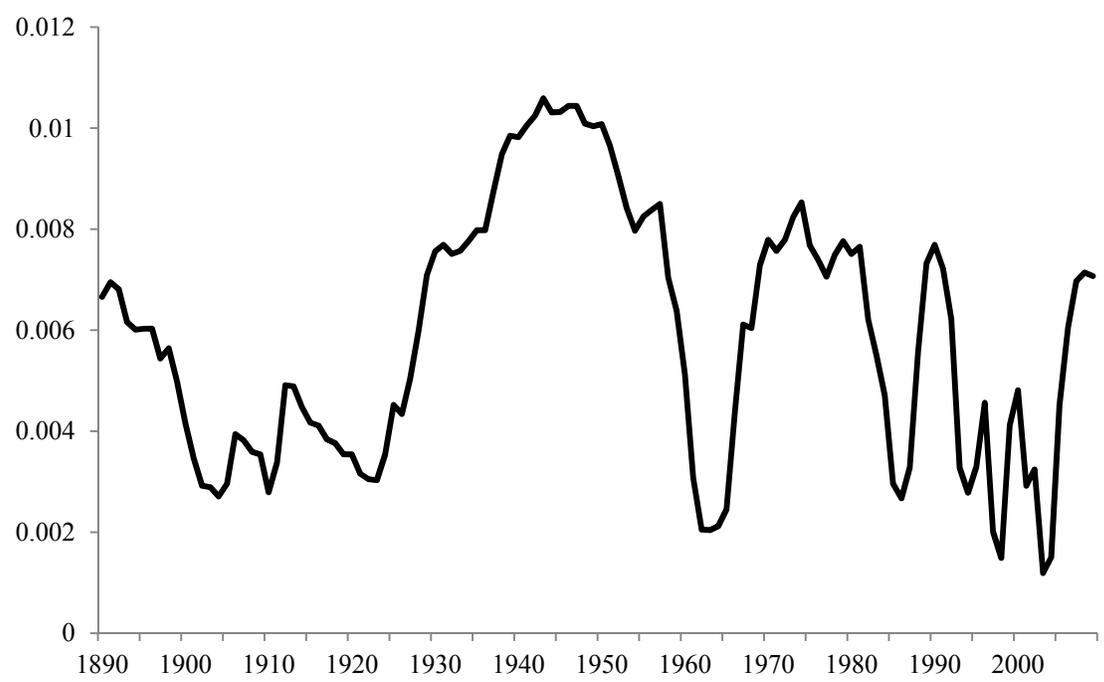


FIGURE IV
 World Stochastic Volatility
 The figure plots the posterior mean value of $\sigma_{w,t}$ for each year in our sample.

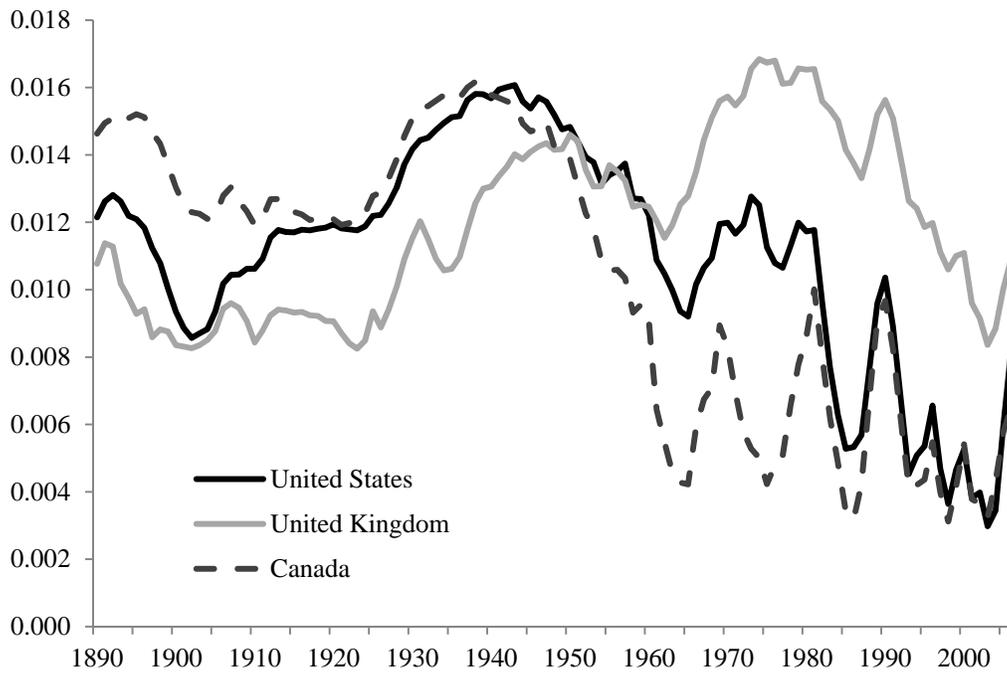


FIGURE V
Stochastic Volatility for the United States, the United Kingdom and Canada

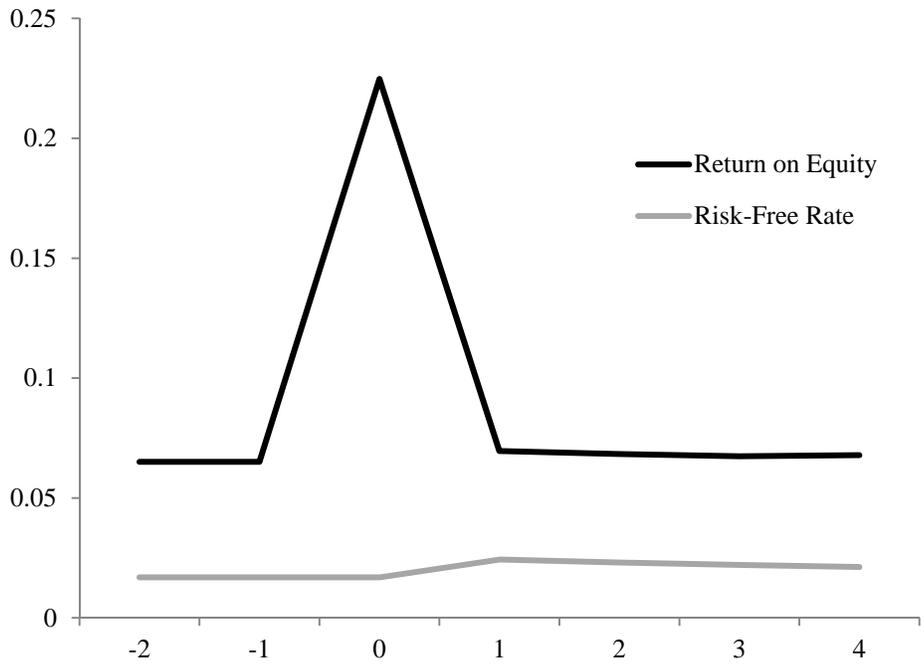


FIGURE VI

Asset Returns in Response to a World Growth-Rate Shock

Response of asset returns to a one standard deviation shock in ε_{Wt} starting from the model's steady state

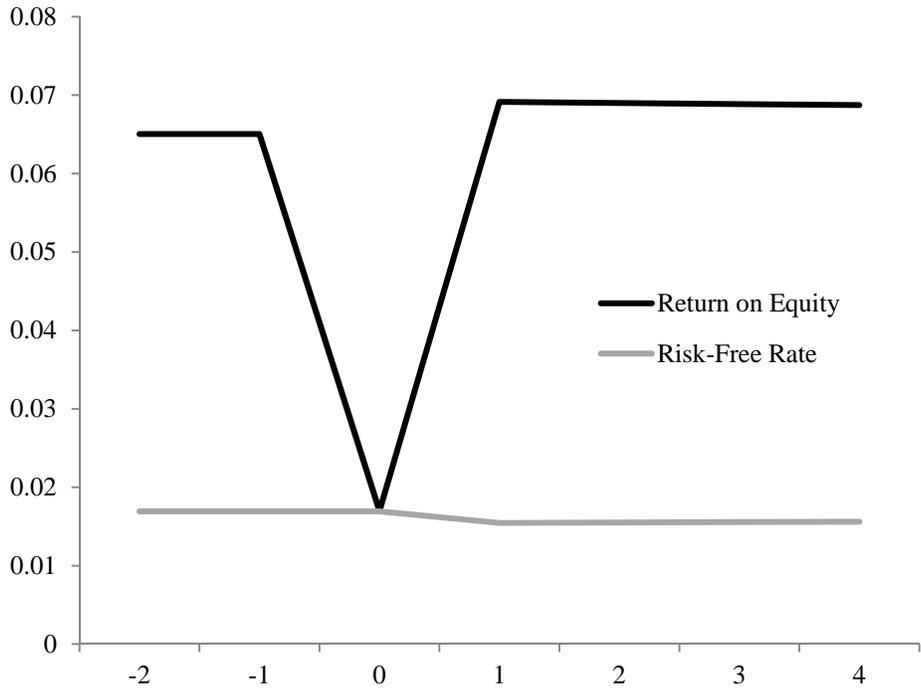


FIGURE VII

Asset Returns in Response to a World Uncertainty Shock

Response of asset returns to a one standard deviation shock in ω_{Wt} starting from the model's steady state

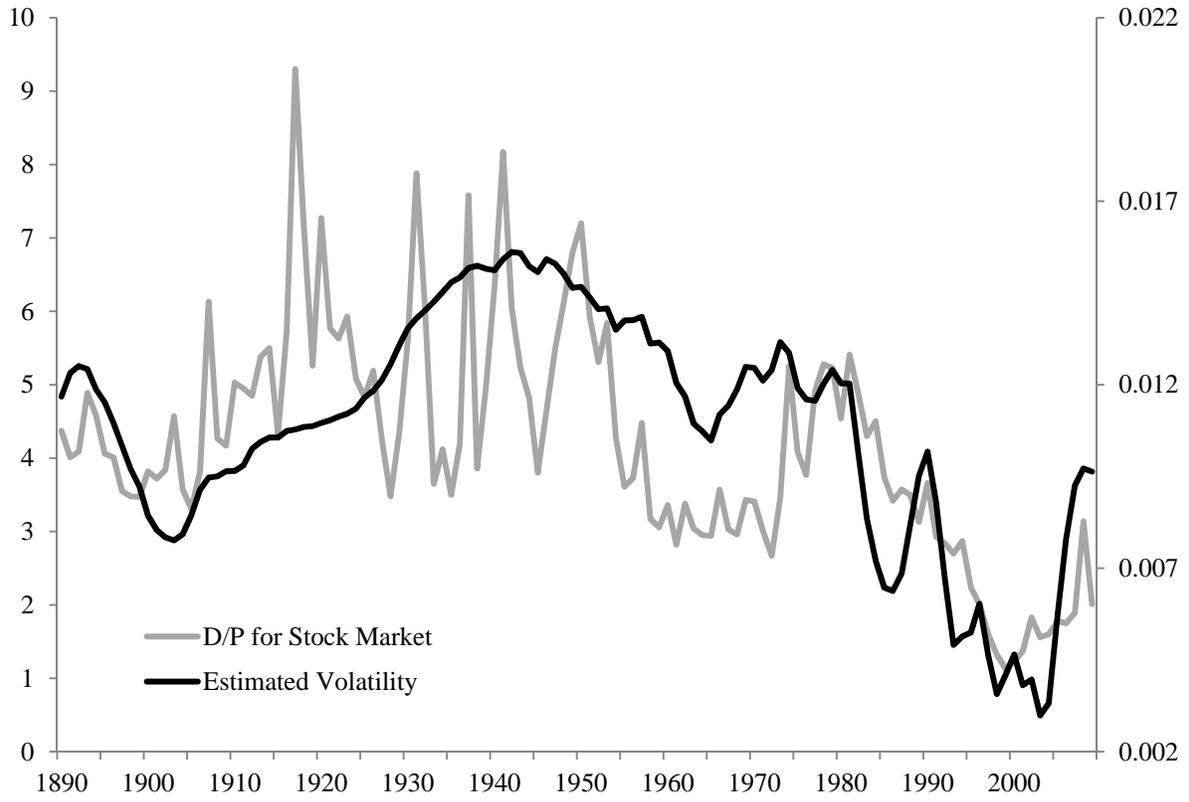


FIGURE VIII
Stock Prices and Economic Uncertainty for the United States

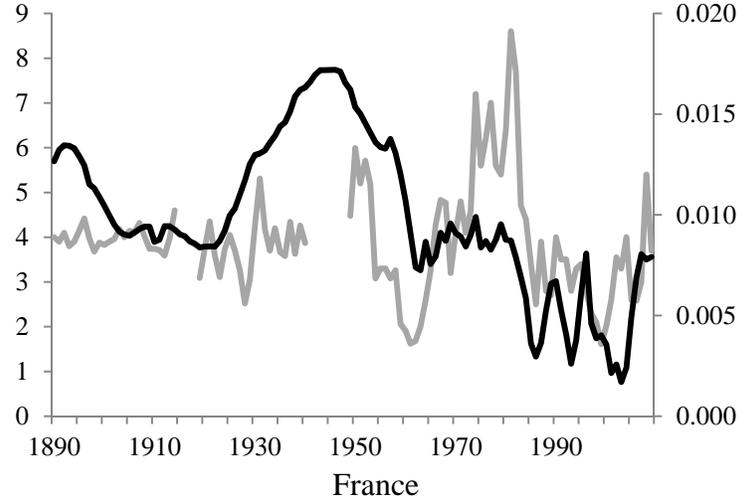
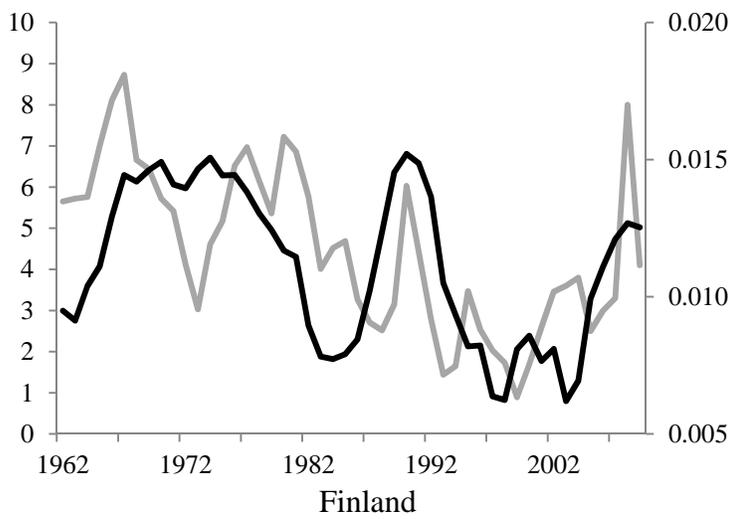
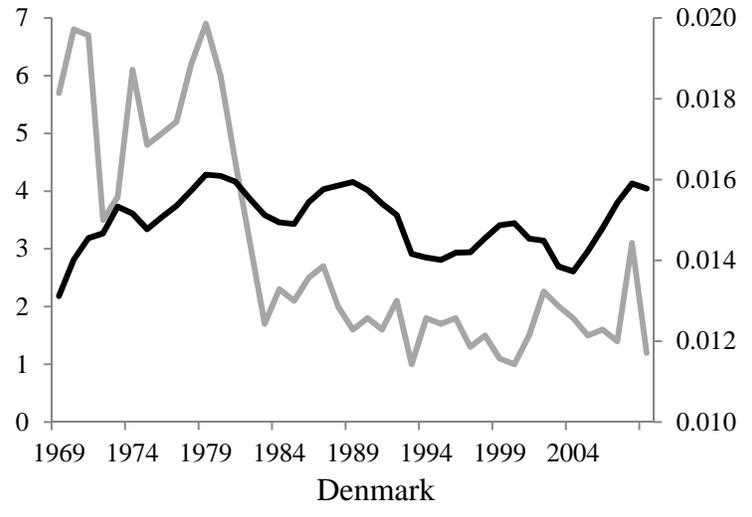
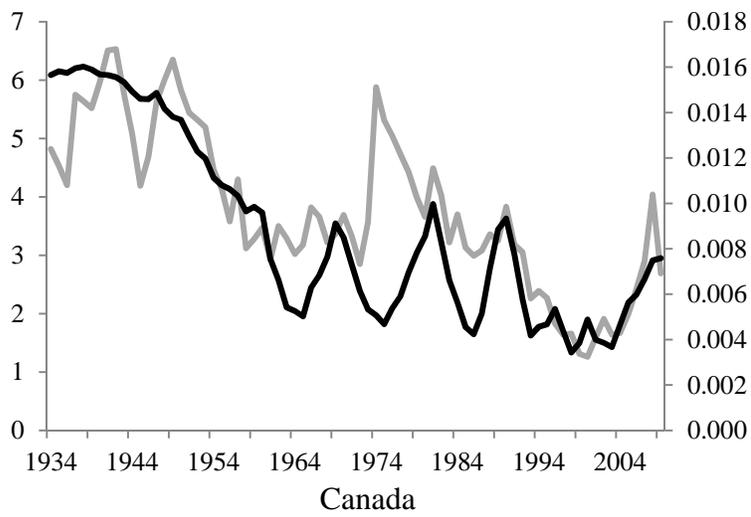
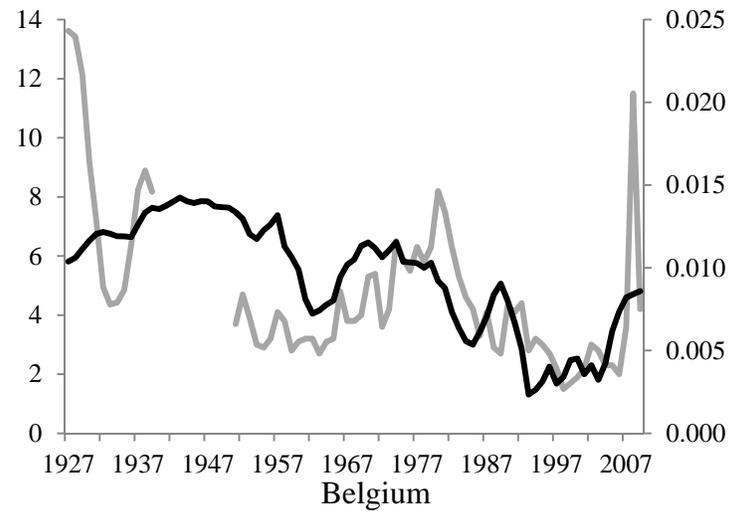
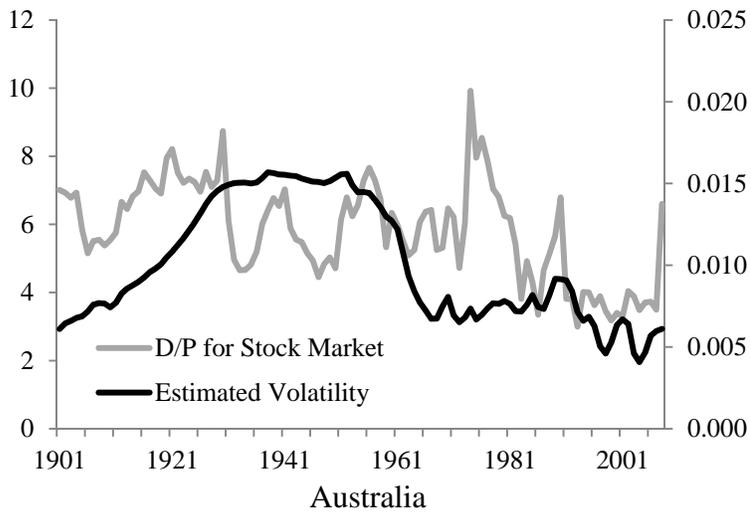


FIGURE A.1
Dividend-Price Ratio for Stocks and Economic Uncertainty

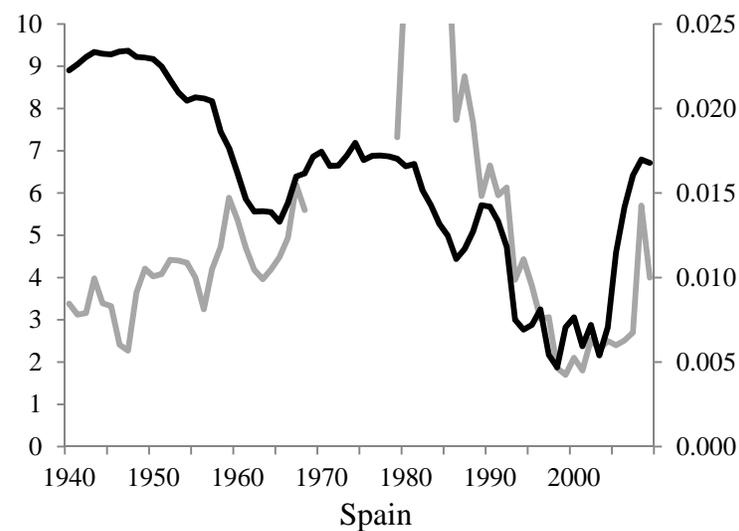
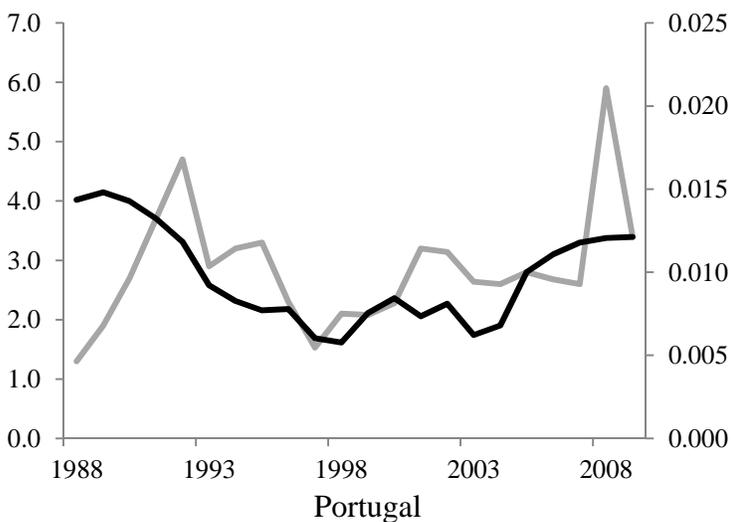
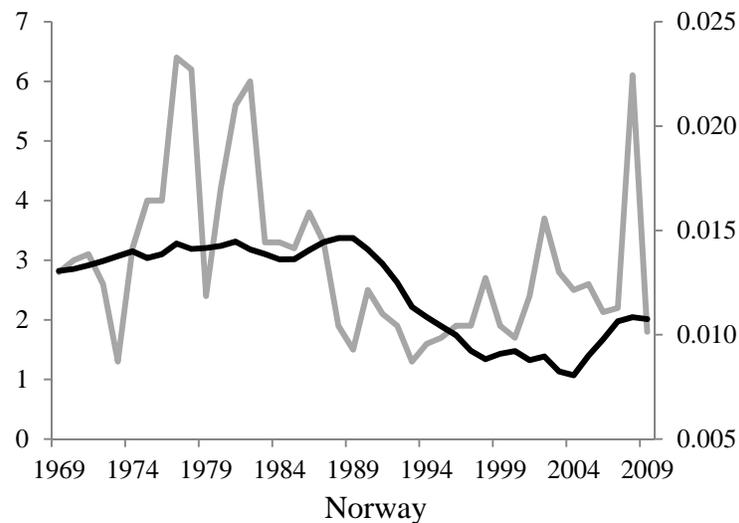
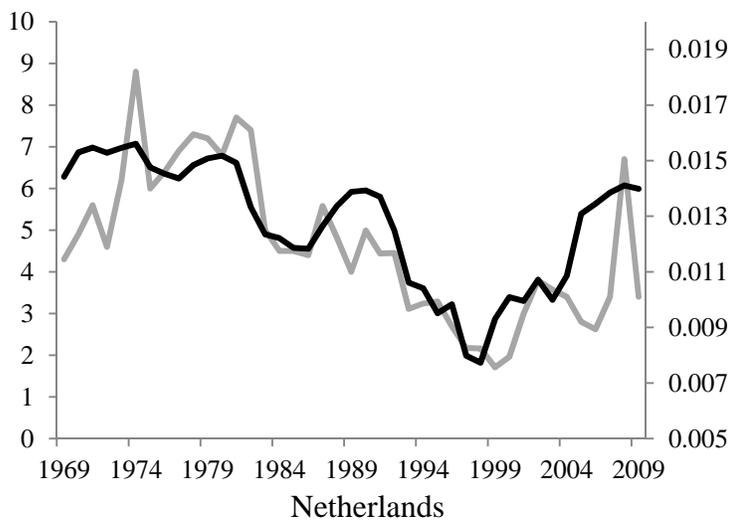
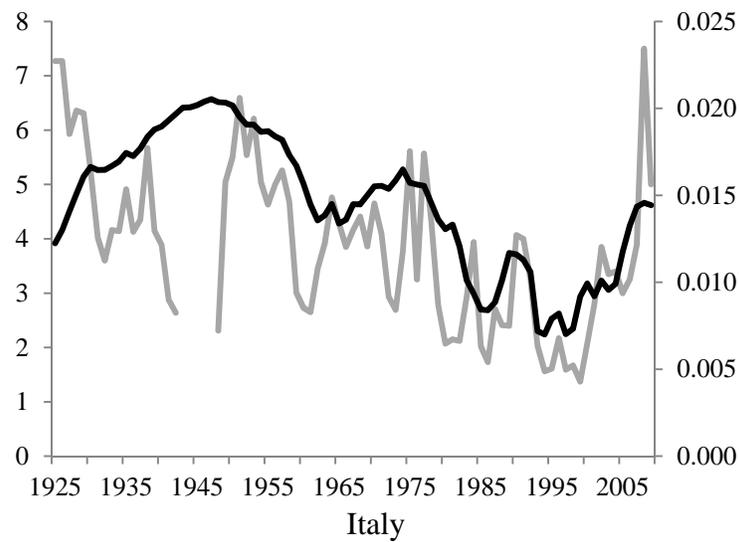
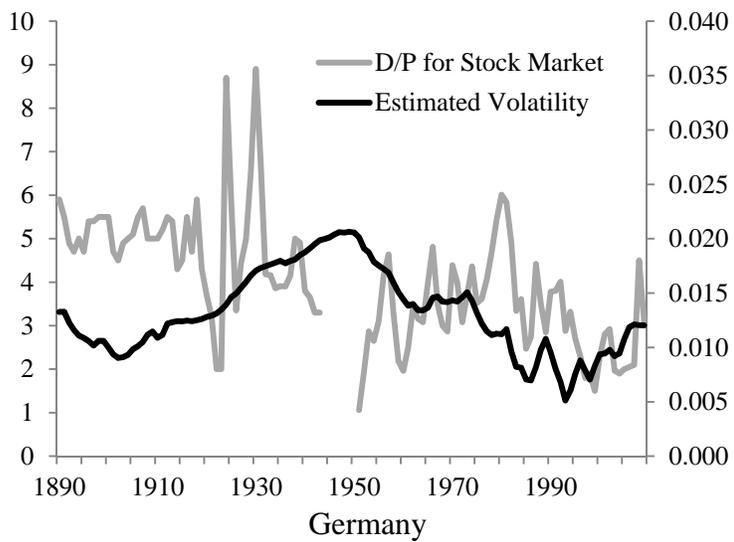


FIGURE A.1 (cont.)
Dividend-Price Ratio for Stocks and Economic Uncertainty

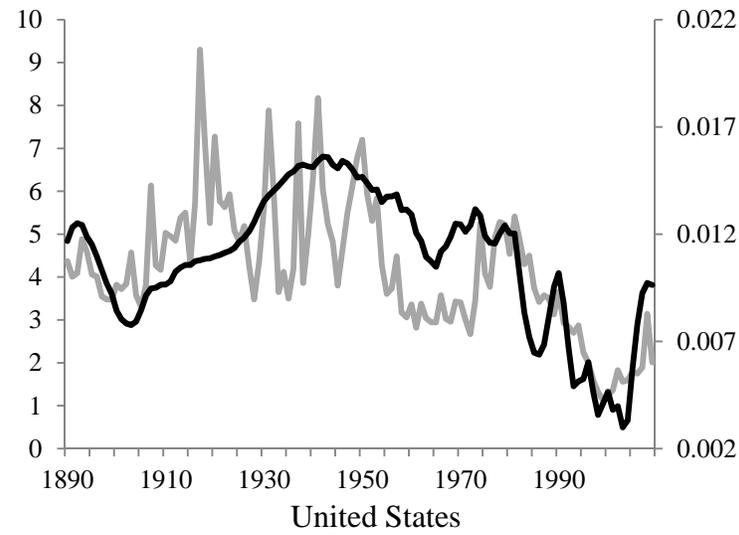
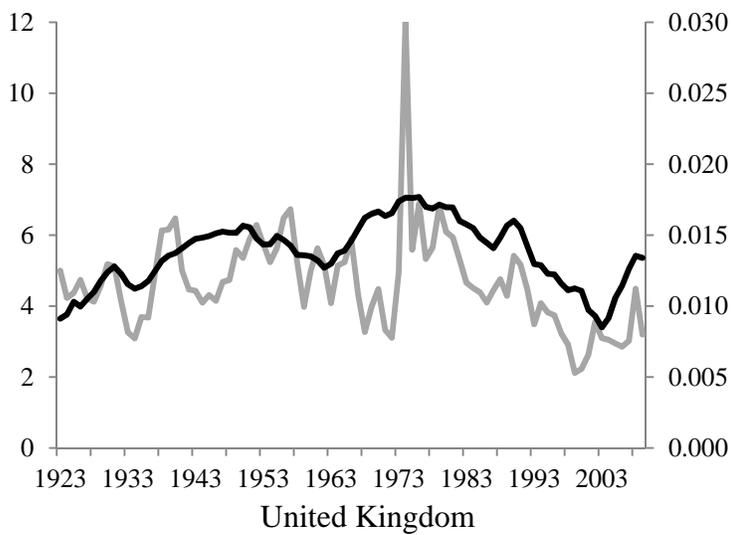
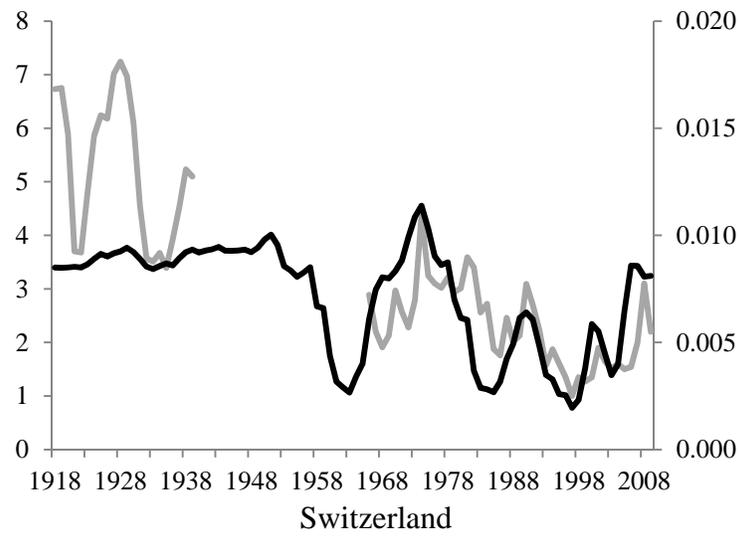
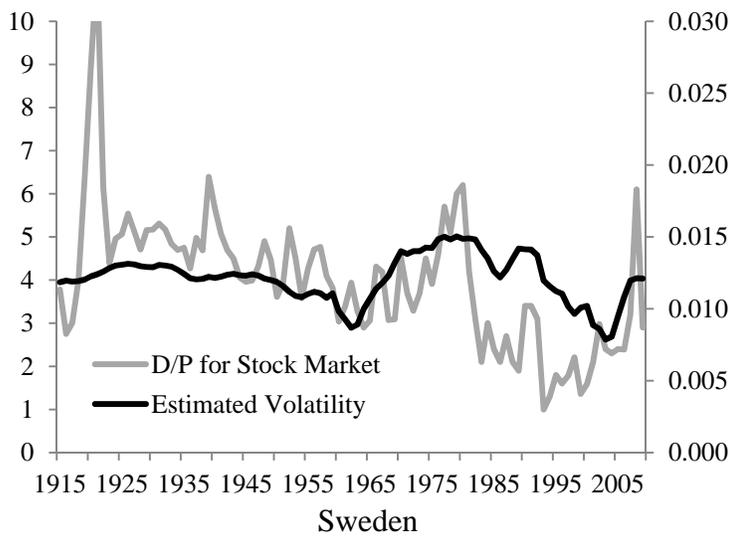


FIGURE A.1 (cont.)
Dividend-Price Ratio for Stocks and Economic Uncertainty