An Empirical Evaluation of the Long-Run Risks Model for Asset Prices

Ravi Bansal[1], Dana Kiku[2] and Amir Yaron[3]

[1]*Fuqua School of Business, Duke University, and NBER; ravi.bansal@duke.edu.*
[2]*The Wharton School, University of Pennsylvania; kiku@wharton.upenn.edu.*
[3]*The Wharton School, University of Pennsylvania, and NBER;*
yaron@wharton.upenn.edu.

ABSTRACT

We provide an empirical evaluation of the Long-Run Risks (LRR) model, and highlight important differences in the asset pricing implications of the LRR model relative to the habit model. We feature three key results: (i) consistent with the LRR model there is considerable evidence in the data for time-varying expected consumption growth and consumption volatility, (ii) the LRR model matches the key asset markets data features, (iii) in the data and in the LRR model accordingly, lagged consumption growth does not predict the future price-dividend ratio, while in the habit-model it counterfactually predicts the future price-dividend with an R^2 of over 40%. Overall, we find considerable empirical support for the LRR model.

1   Introduction

The Long-Run Risks (LRR) model of Bansal and Yaron (2004) highlights two long risk channels to quantitatively explain a wide-range of asset pricing phenomena: (i) long-run fluctuations in expected growth and (ii) long-run fluctuations in consumption volatility. The model features an Epstein and Zin (1989) utility function with an investor preference for early resolution of uncertainty. Bansal and Yaron (2004) and Bansal *et al.* (2007a) calibrate the LRR model to annual data from 1930–2008 and document that the model can match the risk-free rate, equity premium, predictability and other important asset market data features. Beeler and Campbell (2011) carry out an empirical evaluation of the LRR model and argue that the model falls short of the data on some dimensions. In contrast, in this paper, we present evidence which shows the LRR model implications find significant support in the data. As the LRR model is a structural model, it is insightful to compare it to alternative models such as the Campbell and Cochrane (1999) habit model; we document that, along key dimensions, the LRR model fits the data very well, while the Campbell and Cochrane (1999) habit model yields sharply counterfactual implications.

The key data features we focus on in this article include: (i) consumption and dividend dynamics, (ii) mean and volatility of the market return, risk-free rate, and price-dividend ratio, (iii) consumption and return predictability, (iv) relation between consumption volatility and asset prices, (v) predictability of return volatility, and (vi) price-dividend ratio predictability by consumption growth. We discuss the empirical evidence on the magnitude of preference parameters, in particularly, the elasticity of intertemporal substitution, and model implications for the yield curve and their fit to the observed data.

In evaluating the LRR model we focus on the Bansal *et al.* (2007a) calibration since it matches a broader set of empirical facts and utilizes the improved approximate analytical solution method described in Bansal *et al.* (2007b). In their evaluation, Beeler and Campbell (2011) report numbers on two calibrations and two data sets; they consider the Bansal and Yaron (2004) and Bansal *et al.* (2007a) calibration, and the long sample of annual data from 1930–2008, as well as the post-war shorter sub-sample of quarterly data. Our empirical evaluation relies on the annual data from 1930 to 2008, as both Bansal and Yaron (2004) and Bansal *et al.* (2007a) calibrate the model to the post-1930 long sample; and therefore it seems to be the

only appropriate sample to use to critically evaluate their analysis.[1] Using the long sample of annual data is consistent with Shiller and Perron (1985), who show that longer span of data is more important for measuring low-frequency movements than a more frequently sampled shorter span of data (such as the post-war quarterly data).

We document that consumption growth is highly predictable at both short- and long-horizons in the data. A vector autoregression (VAR) based on consumption growth, price-dividend ratio, and the real risk-free rate implies consumption predictability of more than 15% at the one- and five-year horizons, which is statistically different from zero. Using a VAR framework, Hansen *et al.* (2008) also find strong evidence of predictable variations in consumption growth. The VAR-based predictability of consumption growth in the LRR model is of a similar magnitude to the data. We also document that even if one only relies on the price-dividend ratio to forecast future consumption, the regression statistics implied by the LRR model are well within the two standard-error (2-SE) from the data. This evidence shows that consumption dynamics and their predictability properties in the LRR model are consistent with the data.

As in the literature, we find that future equity returns are predictable by the current price-dividend ratio. However, it is also well recognized that the evidence for return predictability is very fragile — confidence bands for predictive R^2s include zero, suggesting lack of predictability. We show that after accounting for standard errors, the LRR model is consistent with the observed predictability of returns. Furthermore, there is a concern that the long-horizon return predictability by the price-dividend ratio is spurious, since the predictive regressor is very persistent. To account for this, we consider a modified predicting variable that is less persistent, the dividend-price ratio less the real risk-free rate. We find that in the data, return predictability based on the adjusted dividend yield declines from 31% to only about 14% at the five-year horizon.[2] In typical asset pricing models this modified predictive variable does not alter the predicability implications, suggesting

---

[1] If one were to apply the model to any sub-sample (as Beeler and Campbell (2011) do), a recalibration (or re-estimation) of the model should be employed to match the different features of the sub-sample. Bansal *et al.* (2007b) re-estimate the model on the shorter 1947–2008 quarterly sample, and show that the model implications are comparable to those of the longer 1930–2008 sample.

[2] The difference in the magnitude of R^2s from the dividend yield-based regression and the predictive regression based on the adjusted dividend-price ratio suggests that the difference is likely due to the very high persistence of the dividend yield in the data, which biases the predicability evidence upwards (see also Hodrick, 1992; Stambaugh, 1999).

that the predictability evidence of high R^2s based on the unadjusted dividend yield is suspect.

Bansal and Yaron (2004) show that, in the LRR model, consumption volatility is a source of systematic risk, as shocks to volatility carry a separate risk premium. They characterize the equilibrium stochastic discount factor and the market price of short-run, long-run and volatility risks. They also show that in the data, a rise in current consumption volatility lowers price-dividend ratios, and that future consumption volatility can be forecasted by current price-dividend ratios. Bansal *et al.* (2005) document the robustness of the negative relation between consumption volatility and asset prices, and further confirm that movements in consumption volatility are indeed an important risk channel. As highlighted in Bansal and Yaron (2004), this evidence suggests that the elasticity of intertemporal substitution (IES) is larger than one. We show that quantitatively, the LRR model matches the sign and the magnitude of the inverse relation between prices and consumption uncertainty and accounts for the observed predictable variation of the integrated volatility of asset returns.

Beeler and Campbell (2011) argue that the IES used in the LRR model is large. The literature reports a wide range of IES magnitudes. For example, Campbell (1999) argues that it is less than one; however, a very large literature (cited in Section 4.6) estimates the IES to be larger than one. Given this, an IES larger than one is well within the range of estimated magnitudes for the IES. The arguments presented in Campbell (1999) — that the IES is less than one — are based on regressing consumption growth on the risk-free rate. However, this regression can yield significantly downward biased estimates of the IES when variables exhibit stochastic volatility, as in the LRR model. Moreover, there are other (more) informative moments that can be used to measure the IES. For example, the data feature that volatility and valuation ratios are inversely related implies that the IES should be larger than one. Hansen *et al.* (2007) use alternative moments and estimate the IES using model implications which include information regarding the level of the real rate, and find that the IES is larger than one. In sum, an IES larger than one, as used in the LRR literature, is supported by the data.

The LRR model implies a downward-sloping real yield curve. Given the evidence on real rates, we view this as a strength of the model. In the UK, which has the longest sample for real yields, Evans (1998) finds a negatively sloped real term structure. We extend his sample and confirm his evidence for a downward sloping real-yield curve.

The LRR model provides an interesting contrast to the habit model in the context of predictability of price-dividend ratios. In the data, forecasting future price-dividend ratios with lagged consumption yields an R^2 close to zero. Consistent with the data, in the LRR model, lagged consumption growth rates do not predict future prices. In fact, the R^2 in the LRR model and in the data are almost identical. Asset prices in the Bansal and Yaron (2004) model are forward looking — they are determined by expectations of future growth and volatility, and therefore changes in these expectations drive movements in current price-dividend ratios. In contrast, in the habit model of Campbell and Cochrane (1999), asset prices are backward looking, as lagged consumption growth, counterfactually, forecasts future price-dividend ratios with an R^2 of more than 40%. In terms of economic differences in the two models, the LRR model would attribute a sharp decline in equity prices to a decline in future expected growth and/or a rise in volatility of future growth. The habit model, on the other hand, would attribute a decline in equity prices to past and current reductions in consumption growth and a resulting rise in risk-aversion. In all, the absence of predictability of future price-dividend ratios by lagged consumption in the data raises considerable doubts regarding the key channel featured in the Campbell and Cochrane (1999) habit model.

Overall, our results (i) support the view that there is a small long-run predictable component in consumption growth, and that consumption volatility is time-varying, (ii) confirm that the forward-looking LRR model can account for the key dynamic properties of asset market data, and (iii) suggest that there is little empirical support for the key mechanism of the backward-looking habit model — that lagged consumption growth forecasts asset prices.

Finally, it should be noted that it is relatively easy to generalize the LRR framework to address additional data features not directly discussed in this article (e.g., options). For example, Bansal *et al.* (2010) consider an augmented LRR model that includes data-consistent mean-reverting transitory jumps (disasters). Bansal and Shaliastovich (2009) and Drechsler and Yaron (2011) entertain a LRR model with two volatility processes, one volatility captures long-run movements and is very persistent, and the second volatility process captures a rapidly mean-reverting shorter-run component of consumption volatility. These generalizations of the LRR model improve the model implications for the predictability of consumption growth and excess returns by the price-dividend ratio.

The paper continues as follows. Section 2 outlines the LRR model and highlights its key features. Section 3 describes the data used in our analysis. Section 4 discusses the results of our empirical analysis. Section 5 provides concluding comments.

2 Long-Run Risks Model

In this section, we specify a model based on Bansal and Yaron (2004). The underlying environment is one with complete markets and a representative agent who has Epstein and Zin (1989) type recursive preferences and maximizes her lifetime utility,

$$V_t = \left[(1 - \delta) C_t^{\frac{1-\gamma}{\theta}} + \delta \big(E_t [V_{t+1}^{1-\gamma}] \big)^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \tag{1}$$

where C_t is consumption at time t, $0 < \delta < 1$ reflects the agent's time preference, γ is the coefficient of risk aversion, $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$, and ψ is the elasticity of intertemporal substitution (IES). Utility maximization is subject to the budget constraint,

$$W_{t+1} = (W_t - C_t) R_{c,t+1}, \tag{2}$$

where W_t is the wealth of the agent, and $R_{c,t}$ is the return on all invested wealth.

Consumption and dividends have the following joint dynamics:

$$\begin{aligned}
\Delta c_{t+1} &= \mu_c + x_t + \sigma_t \eta_{t+1} \\
x_{t+1} &= \rho x_t + \varphi_e \sigma_t e_{t+1} \\
\sigma_{t+1}^2 &= \bar{\sigma}^2 + \nu(\sigma_t^2 - \bar{\sigma}^2) + \sigma_w w_{t+1} \\
\Delta d_{t+1} &= \mu_d + \phi x_t + \pi \sigma_t \eta_{t+1} + \varphi \sigma_t u_{d,t+1},
\end{aligned} \tag{3}$$

where Δc_{t+1} and Δd_{t+1} are the growth rate of consumption and dividends, respectively. In addition, we assume that all shocks are *i.i.d* normal and are orthogonal to each other. As in the long run risks model of Bansal and Yaron (2004), $\mu_c + x_t$ is the conditional expectation of consumption growth, and x_t is a small but persistent component that captures long run risks in consumption growth. For parsimony, as in Bansal and Yaron (2004), volatility of consumption and dividends is driven by a common time-varying component. As shown in their paper, predictable variations in the conditional

second moment of growth rates lead to time-varying risk premia. Dividends have a levered exposure to the persistent component in consumption, x_t, which is captured by the parameter ϕ. In addition, we allow the *i.i.d* consumption shock η_{t+1} to influence the dividend process, and thus serve as an additional source of risk premia. The magnitude of this influence is governed by the parameter π. Save for this addition, the dynamics are similar to those in Bansal and Yaron (2004).

As in Epstein and Zin (1989), for any asset j, the first-order condition yields the following asset pricing Euler condition,

$$E_t[\exp(m_{t+1} + r_{j,t+1})] = 1, \tag{4}$$

where m_{t+1} is the log of the intertemporal marginal rate of substitution (IMRS), and $r_{j,t+1}$ is the log of the gross return on asset j. The log of the IMRS, m_{t+1}, is given by

$$m_{t+1} = \theta \log \delta - \frac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1)r_{c,t+1}, \tag{5}$$

where $r_{c,t+1}$ is the continuous return on the consumption asset. To solve for the return on wealth (the return on the consumption asset), we use the log-linear approximation for the continuous return on the wealth portfolio, namely,

$$r_{c,t+1} = \kappa_0 + \kappa_1 z_{t+1} + \Delta c_{t+1} - z_t, \tag{6}$$

where $z_t = \log(P_t/C_t)$ is the log of the price-consumption ratio (i.e., the valuation ratio corresponding to a claim that pays consumption) and κ's are log linearization constants, which are discussed in more detail below.

To derive the dynamics of asset prices we rely on approximate analytical solutions (instead of the polynomial-based numerical approximation in the original paper of Bansal and Yaron (2004)), which we find provide a more accurate solution to the model. This easy-to-implement solution technique allows us to better address certain predictability dimensions. Specifically, we conjecture that the price-consumption ratio follows,

$$z_t = A_0 + A_1 x_t + A_2 \sigma_t^2 \tag{7}$$

and solve for As using the Euler equation (4), the return equation (6) and the conjectured dynamics (7). In solving for the price-consumption ratio, we impose model consistency between its mean, $\bar{z}$, and approximation κ's, which themselves depend on the average price-consumption ratio.

This allows us to make sure that any change in the model parameters that alters $\bar{z}$ is also incorporated in the approximation constants. The model-based endogenous solution for $\bar{z}$ is thus obtained by solving the equation,

$$\bar{z} = A_0(\bar{z}) + A_2(\bar{z})\bar{\sigma}^2, \tag{8}$$

and recognizing that approximation constants that enter A's are defined by $\kappa_0 = \log(1 + \exp(\bar{z})) - \kappa_1\bar{z}$ and $\kappa_1 = \frac{\exp(\bar{z})}{1+\exp(\bar{z})}$.

The solutions for A's that describe the dynamics of the price-consumption ratio are determined by the preference and technology parameters as:

$$A_0 = \frac{1}{1 - \kappa_1}\left[\log\delta + \kappa_0 + \left(1 - \frac{1}{\psi}\right)\mu_c + \kappa_1 A_2(1 - \nu)\bar{\sigma}^2 + \frac{\theta}{2}(\kappa_1 A_2 \sigma_w)^2\right]$$

$$A_1 = \frac{1 - \frac{1}{\psi}}{1 - \kappa_1\rho} \tag{9}$$

$$A_2 = -\frac{(\gamma - 1)(1 - \frac{1}{\psi})}{2(1 - \kappa_1\nu)}\left[1 + \left(\frac{\kappa_1\varphi_e}{1 - \kappa_1\rho}\right)^2\right].$$

Bansal and Yaron (2004) show that solution (9) captures the intuition that, as long as IES is larger than one, the substitution effect dominates the wealth effect. Consequently, high expected growth raises asset valuations, while high consumption volatility lowers the price-consumption (and price-dividend) ratio. This is an important implication of the model as it may help identify the magnitude of IES in the data.

Given the solution for z_t, the innovation to the return to wealth can be derived, which in turn allows us to specify the innovations to the IMRS and facilitates the computation of the risk premia of various assets. In particular, it follows that the risk premium on the stock market portfolio is derived from three sources of risks. Specifically,

$$E_t[r_{m,t+1} - r_{f,t} + 0.5\sigma_{t,r_m}^2] = \beta_{\eta,m}\lambda_\eta\sigma_t^2 + \beta_{e,m}\lambda_e\sigma_t^2 + \beta_{w,m}\lambda_w\sigma_w^2, \tag{10}$$

where $\beta_{j,m}$, $j = \{\eta, e, w\}$ are the betas of the market return with respect to the "short-run" risk (η_t), the long-run growth risk (e_t), and the volatility risk (w_t), respectively. The market return betas are determined by the underlying preferences and cash-flow dynamics and are presented in the Appendix. λ's represent the corresponding market prices of risks that, as

shown in Bansal and Yaron (2004), are given by:

$$\lambda_\eta = \gamma$$

$$\lambda_e = (1 - \theta)\kappa_1 A_1 \varphi_e = \left(\gamma - \frac{1}{\psi}\right)\frac{\kappa_1 \varphi_e}{1 - \kappa_1 \rho} \tag{11}$$

$$\lambda_w = (1 - \theta)\kappa_1 A_2 = -(\gamma - 1)\left(\gamma - \frac{1}{\psi}\right)\frac{\kappa_1}{2\,(1 - \kappa_1 \nu)}\left[1 + \left(\frac{\kappa_1 \varphi_e}{1 - \kappa_1 \rho}\right)^2\right].$$

Note that, due to the separation between risk aversion and IES, each risk carries a separate premium. In the power utility framework, where IES equals the reciprocal of risk aversion, only short-run risks receive compensation, while long-run and volatility risks carry no separate risk premia. The market prices of risks in equation (11) show that preference for early resolution of uncertainty (i.e., γ larger than the reciprocal of IES) is required for long-run risks to earn a positive risk premium.

As discussed in the introduction, the above LRR model can be easily modified to include additional features such as jumps (e.g., disasters). Such an augmented LRR framework is presented in Bansal *et al.* (2010), who show that this additional feature, which can capture the sharp declines in consumption seen in the Great Depression, does not contribute to the risk-premium in any measurable manner. This implies that it is fairly easy to modify the LRR model to account for sharp declines in consumption without materially altering the asset pricing implications of the baseline LRR model presented above.

3 Data

Consistent with Bansal *et al.* (2007b), we use annual data on consumption and asset prices for the time period from 1930 to 2008. Consumption data are based on seasonally adjusted per-capita series on real consumption from the National Income and Product Accounts (NIPA) tables available on the Bureau of Economic Analysis, website. Aggregate consumption is defined as consumer expenditures on non-durables and services. Growth rates are constructed by taking the first difference of the corresponding log series. Our asset menu comprises the aggregate stock market portfolio on the value weighted return of the NYSE/AMEX/NASDAQ from the Center for Research

in Security Prices (CRSP) and a proxy of a risk-less asset. To construct the real risk-free rate, we regress the ex-post real three-month Treasury Bill yield on the nominal rate and past annual inflation. The fitted value from this regression is the proxy for the ex-ante real interest rate. Use of other estimates of expected inflation to construct the real rate does not lead to any significant changes in our results.

In terms of the data, Beeler and Campbell (2011) also use a sub-sample of post-war quarterly data from 1947 to 2008. However, in their analysis they continue to use the Bansal *et al.* (2007a) calibration that is based on the longer 1930–2008 sample. It is not obvious that the post-war data provide a representative sample of macroeconomic dynamics in terms of their volatility, autocorrelations, and other moments of interest. In our view, the longer sample better captures all the relevant macroeconomic outcomes. An appropriate approach to using sub-samples is to re-estimate (or re-calibrate) the model using a different sample. This approach is pursued in Bansal *et al.* (2007b) who find the LRR model estimated on the quarterly post-war data performs well along the dimensions discussed in this paper for the 1930–2008 sample.

4 Empirical Findings

4.1 Calibration and Long-Run Dynamics

In calibration and simulations, following the standard in the literature, we assume that the decision interval of the agent is monthly. To make the model-implied data comparable to the observed annual data, we appropriately aggregate the simulated monthly observations and construct annual growth rates and annual asset returns. The price-dividend ratio, as in the data, is constructed by dividing the end-of-year price by the trailing sum of 12-month dividends.

For statistical inference, as in Drechsler and Yaron (2011) and Beeler and Campbell (2011), we sample from the calibrated model and construct the finite-sample empirical distribution for various statistics of interest. Reported statistics are based on 10,000 simulated samples with 79×12 monthly observations that match the length of the actual data. We report the median and tail percentiles of the Monte-Carlo distributions. In addition, we present population values that correspond to the statistics constructed from a long-sample of 10,000 annualized observations.[3]

---

[3] Our model inferences are robust to using alternative standard methods to construct standard errors that are reported in Bansal *et al.* (2007a).

Preferences	δ	γ	ψ			
	0.9989	10	1.5			
Consumption	μ_c	ρ	ϕ_e	$\bar{\sigma}$	ν	σ_w
	0.0015	0.975	0.038	0.0072	0.999	0.0000028
Dividends	μ_d	ϕ	π	φ		
	0.0015	2.5	2.6	5.96		

Table 1. Configuration of model parameters.

Table 1 reports the configuration of investors' preferences and the time-series parameters that describe the dynamics of consumption and dividend growth rates. The model is calibrated on a monthly decision interval.

Table 1 provides the Bansal *et al.* (2007a) parameter configuration used to calibrate the model. This BKY configuration is chosen to match several key statistics of the 1930–2008 annual consumption and dividend data, and it refines the Bansal and Yaron (2004) configuration in two directions. First, the persistence of volatility shocks is assumed to be higher; second, dividend shocks are assumed to be correlated with short-run shocks in consumption growth, while in Bansal and Yaron (2004) the correlation between the two is set at zero. These changes enhance the role of the volatility channel relative to Bansal and Yaron (2004); however, low-frequency movements in expected growth are critical to magnify the role of time-varying volatility for asset prices. As we illustrate below, in the absence of the expected growth channel, the time-varying volatility channel by itself cannot account for asset prices.

Table 2 displays the model implications for the unconditional moments of consumption and dividend growth rates. The calibrated model matches closely the mean, volatility, the first to third and the fifth autocorrelations of consumption growth, though the fourth autocorrelation is slightly outside the confidence band. The model also matches quite well the volatility of dividend growth and its correlation with consumption growth.[4] It is worth noting that the first-order autocorrelation of consumption growth in the data is 0.45, which is much higher than the one implied by monthly *i.i.d* growth rates even after accounting for time-aggregation. According to

---

[4] The first five autocorrelation of dividend growth are also well within their corresponding model-implied confidence band.

Moment	Data	Model					
	Estimate	Median	2.5%	5%	95%	97.5%	Pop
$E[\Delta c]$	1.93	1.80	0.72	0.92	2.73	2.93	1.79
$\sigma(\Delta c)$	2.16	2.47	1.52	1.64	3.60	3.84	2.83
$AC1(\Delta c)$	0.45	0.39	0.15	0.19	0.57	0.60	0.45
$AC2(\Delta c)$	0.16	0.15	−0.13	−0.10	0.37	0.41	0.19
$AC3(\Delta c)$	−0.10	0.09	−0.17	−0.13	0.31	0.35	0.15
$AC4(\Delta c)$	−0.24	0.05	−0.21	−0.17	0.28	0.32	0.10
$AC5(\Delta c)$	−0.02	0.03	−0.23	−0.19	0.24	0.28	0.08
$E[\Delta d]$	1.15	1.84	−2.79	−1.96	5.64	6.52	1.45
$\sigma(\Delta d)$	11.05	14.11	8.53	9.19	20.02	21.15	15.83
$AC1(\Delta d)$	0.21	0.27	0.03	0.07	0.45	0.48	0.27
$\text{Corr}(\Delta d, \Delta c)$	0.55	0.46	0.22	0.26	0.62	0.65	0.46
$E[R]$	7.66	8.12	3.59	4.38	13.62	14.96	8.75
$\sigma(R)$	20.28	20.44	12.45	13.62	30.15	31.90	23.37
$E[p-d]$	3.36	3.14	2.79	2.85	3.29	3.32	3.07
$\sigma(p-d)$	0.45	0.18	0.11	0.11	0.28	0.30	0.26
$AC1(p-d)$	0.87	0.62	0.33	0.38	0.79	0.82	0.79
$E[R^f]$	0.57	1.24	0.10	0.31	1.78	1.86	1.05
$\sigma(R^f)$	2.86	0.94	0.54	0.59	1.46	1.59	1.22

Table 2. Dynamics of growth rates and prices.

Table 2 presents descriptive statistics for aggregate consumption growth, dividends, prices and returns of the aggregate stock market, and the risk-free rate. Data statistics along with standard deviations of bootstrap distributions (in parentheses) are reported in the "Data" panel. The data are real, sampled on an annual frequency and cover the period from 1930 to 2008. The "Model" panel presents the corresponding moments implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data. Means and volatilities of returns and growth rates are expressed in percentage terms.

the results of Working (1960), the annual autocorrelation with *i.i.d* growth rates would be only 0.25.[5]

---

[5] It would be even lower under plausible scenarios of measurement errors in monthly consumption data.

The half-life of expected growth shocks in our calibration is about 2.25 years. Beeler and Campbell (2011) report that the monthly magnitude of the change in the expected long-run level of consumption due to long-run shocks is 1.3%. Their calculations are based on the monthly model (not annual) and therefore are not particularly useful, as they cannot be directly compared to the observed annual consumption data. To understand if the model implications for the long-run consumption dynamics are plausible, we compare the impulse response functions of the *annual* consumption growth in the data and in the model. In particular, similar to Stock and Watson (1988), we estimate various univariate ARMA models of annual consumption growth in the data, and in the model using our simulated annual consumption growth. Using the fitted models, we measure the accumulated impulse response to a one-standard deviation shock in consumption growth. Table 3 reports the long-run consumption response in the data and model for the AR(1) specification. As can be seen, the model and data match very well. The estimated response in the data is 3.34%, and is in the center of the distribution of the model-based response (in the model, the median response is 3.76%). For an AR(2) specification, the consumption response is 3.02% in the data, while the median response in the model is 3.64%, and the data is well within the 2-SE bounds. We also estimate an ARMA(1,1), MA(1) and MA(2) models, and in each the long-run accumulated consumption response is also well within the model's 2-SE bounds. This evidence shows that the BKY calibrated model accounts well for the observed dynamics of consumption growth in the long-run.

Data	Model					
Estimate	Median	2.5%	5%	95%	97.5%	Pop
3.34%	3.76%	1.93%	2.14%	6.38%	7.04%	4.46%

Table 3. AR(1) based long run consumption response.

Table 3 presents the long-run response of consumption to a one-standard deviation innovation shock. The consumption response is based on fitting an AR(1) to annual data. The "Model" panel presents evidence implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data.

In all, the model and the data are a close match in terms of short, and long-run consumption dynamics. Recent work by Kaltenbrunner and Lochstoer (2010) and Croce (2005) shows that consumption and savings decisions of agents in a production economy lead to low-frequency movements in consumption growth, similar to those in the LRR model.

The LRR model calibration assumes a persistent consumption-volatility process. Earlier work (e.g., Kim and Nelson, 1999; McConnell and Perez-Quiros, 2000; Stock and Watson, 2002) documents a very persistent and low-frequency decline in the volatility of consumption and other macro series from 1930 to more recent times. Figure 1 shows the volatility of consumption by decade, and one can easily see the slow and persistent decline in volatility. The LRR model calibration of the volatility process is designed to capture these low-frequency movements in consumption volatility.[6]
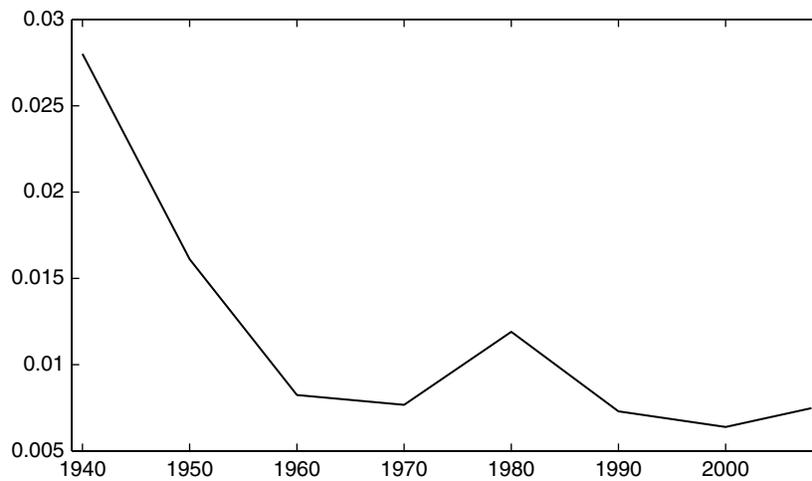
Figure 1. Volatility by decades.

Figure 1 depicts consumption volatility by decades by computing the average of the absolute value of AR(1)-filtered consumption growth within each decade.

---

[6]   Bansal and Shaliastovich (2009) and Drechsler and Yaron (2011) entertain a generalized LRR model with two volatility processes, one volatility captures long-run movements and is very persistent, and the second volatility process captures a rapidly mean-reverting shorter-run component of consumption volatility movement. The Bansal *et al.* (2007a) calibration focuses solely on low-frequency movements in volatility and highlights their importance for asset pricing.

The LRR model specification as stated in Equation (3), for analytical tractability and ease of solution, assumes that volatility shocks are normally distributed. In simulations, we replace negative realizations of σ^2 with a small positive number. We have also evaluated the approach of replacing negative volatility realizations by redrawing volatility news and found the results to be virtually identical in the two cases. Note that the standard deviation of volatility shocks (σ_w) is quite small relative to its mean. The fraction of negative realizations, therefore, is also small, averaging about 0.6% of the draws at our calibrated values; that is, the probability of a negative volatility realizations in any model draw is very small. A conceptually cleaner approach is pursued in Bansal and Shaliastovich (2009), who follow Barndorff-Nielsen and Shephard (2001) and assume that volatility shocks have a gamma distribution, which ensures positivity of the volatility process. Bansal and Shaliastovich (2009) show that the model implications in the gamma distribution case are similar to the Gaussian case presented here, and hence we continue to use the Gaussian case, particularly as the probability of a negative realization is very small.

4.2 Equity Premium and Risk-free Rate Puzzles

Table 2 also displays the model implications for the unconditional moments of the equity return, price-dividend ratio, and the risk-free rate. Overall, the model matches well the key asset price moments. Specifically, the model matches quite well the level and volatility of the equity returns and the risk-free rate. The average excess return in our data set is around 7%. For comparison, the model-implied risk premium of the stock market portfolio averages 6.9%. In the model, as in the data, the volatility of equity returns is about 20%, which is much higher than the volatility of the underlying cash-flow growth rates. Consistent with the data, the model-implied mean of the real risk-free rate is around 1% per annum.

In our calibration, the contribution to total risk premium from short-run risks is 25%, long-run growth risks is 32%, and long-run volatility risks is 43%. The two persistent sources combined account for 75% of the equity premium.

It is important to note that the long-run expected growth risk is critical for explaining the equity risk premium, as it not only accounts for a significant portion of the premium itself but also magnifies the contribution of the volatility risk. In the absence of the long-run growth risk (i.e., if the variance

of x_t is zero), the annualized equity premium is only 0.92%. In this case, the population value of the volatility of the price-dividend ratio is about 0.19. If, on the other hand, the long-run growth risk is present but the volatility channel is shut down, the annualized equity premium is 3.95% but the volatility of the price-dividend ratio drops to 0.09. Thus, the long-run growth risk is important for the level of the equity risk-premium, while the volatility channel is important for the variability of asset prices.

4.3   Consumption, Dividends and Return Predictability

In this section, we report strong evidence for consumption predictability in a multivariate VAR framework. Table 4 provides evidence on consumption predictability using a VAR with consumption growth, the real risk-free rate, and the log price-dividend ratio. The R^2 for consumption predictability starts at 23% at the one-year horizon and drops only to 15% at the five-year horizon.[7] Thus, in the data, consumption growth is strongly predictable

	Data	Model					
	Estimate	Median	2.5%	5%	95%	97.5%	Pop
1yr	0.23	0.32	0.10	0.13	0.50	0.54	0.27
3yr	0.17	0.34	0.10	0.13	0.57	0.61	0.30
5yr	0.15	0.31	0.08	0.12	0.55	0.60	0.31
10yr	0.13	0.21	0.05	0.07	0.45	0.50	0.27
15yr	0.11	0.15	0.03	0.05	0.35	0.41	0.22
20yr	0.09	0.11	0.02	0.04	0.28	0.33	0.17

Table 4. VAR-implied predictability of consumption growth.

Table 4 presents predictive R^2s for consumption growth implied by a first-order VAR model for consumption growth, price-dividend ratio of the aggregate stock market portfolio and risk-free rate. Data statistics are reported in "Data" panel. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "Model" panel presents predictability evidence implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data.

---

[7]   As in Hodrick (1992), R^2 are constructed by exploiting the dynamics of the first-order VAR specification.

at both short and long horizons, which is consistent with consumption predictability evidence reported in Hansen *et al.* (2008) and Bansal *et al.* (2007b). Table 4 further shows that the LRR model matches well the documented pattern of consumption predictability. Note that a monthly *i.i.d* consumption growth process, time-aggregated to the annual frequency, would imply an R^2 of only 6% for the first year and close to zero for the second and subsequent years. Our empirical evidence, therefore, casts doubt on the view that consumption growth is *i.i.d*, as often assumed in the literature (e.g., Campbell and Cochrane, 1999).

Panel A of Table 5 provides the evidence of consumption growth predictability using the log of the price-dividend ratio as the only regressor. Estimates of slope coefficients ($\hat{\beta}$) in these regressions for various horizons are presented in Panel B of the table. In the data, the R^2s in these regressions are 6% at the one-year horizon and close to zero at the five-year horizon. The model-implied evidence reveals a similar modest forecasting

	Data	Model					
	Estimate	Median	2.5%	5%	95%	97.5%	Pop
Panel A: Predictive R^2s							
1yr	0.06	0.14	0.00	0.01	0.36	0.41	0.07
3yr	0.01	0.10	0.00	0.00	0.37	0.42	0.05
5yr	0.00	0.08	0.00	0.00	0.37	0.43	0.04
Panel B: Predictive Slopes							
1yr	0.01	0.05	0.01	0.01	0.09	0.10	0.03
3yr	0.01	0.10	−0.03	−0.01	0.20	0.22	0.05
5yr	0.00	0.12	−0.09	−0.06	0.28	0.32	0.06

Table 5. Predictability of consumption growth by PD-ratio.

Table 5 presents R^2s and slope coefficients from projecting one-, three- and five-year consumption growth onto the lagged price-dividend ratio of the aggregate stock market portfolio. Data statistics are reported in the "Data" panel. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "Model" panel presents predictability evidence implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data.

power of the price-dividend ratio. In particular, the population R^2 in these predictive regressions is only 7% and 4% at the one- and five-year horizons, respectively. Likewise, the model-implied regression slopes, on average and in population, are close to the corresponding point estimates. Formally, both the slope and R^2s of these predictive regressions in the data are within the model-implied 2-SE confidence bands. Note that in the LRR model, variation in price-dividend ratios is driven by two state variables: the conditional mean and volatility of consumption growth. This suggests that the price-dividend ratio by itself may not forecast future growth rates in any significant manner. Consequently, univariate regressions of future consumption growth on current price-dividend ratios, emphasized in Beeler and Campbell (2011), may fail to capture all the predictable variation in consumption growth. As shown above, in the data, consumption growth is highly predictable when one relies on a multivariate regression setting and a richer information set to learn about predictable variation in expected growth rates. Therefore, the view that consumption growth is in general unpredictable is misguided, as there is strong evidence for consumption predictability.

Table 6 provides evidence on dividend predictability using a VAR with dividend growth, real risk-free rate, and the log price-dividend ratio. In the data, the VAR R^2 of predicting dividend growth starts at 16% at the one-year horizon, rises to 27% at the five-year horizon, and then gradually tapers off. The model implications for dividend growth predictability line up with the data.[8] Table 7 documents evidence on short- and long-horizon dividend predictability using only the price-dividend ratio as a regressor. The data feature modest predictability, with an R^2 in the range of 4%–9%, and the slope coefficients varying from 0.07 at the one-year horizon to 0.09 at the five-year horizon. After accounting for sampling uncertainty, the LRR model matches well both the R^2s and the estimated slopes.

Our evidence of growth rate predictability is robust to alternative measures of asset cash flows. In particular, a VAR for earnings growth, price-earnings ratio and risk-free rate yields a predictive R^2 for the earnings

---

[8] This evidence is consistent with dividend predictability documented in Bansal *et al.* (2009). They find that cash-flow growth rates of the aggregate stock market, as well as book-to-market and size sorted portfolios are strongly predictable at both short and long horizons, and highlight the importance of long-run predictable variations in asset cash flows for understanding the term structure of the risk-return trade-off.

	Data	Model					
	Estimate	Median	2.5%	5%	95%	97.5%	Pop
1yr	0.16	0.26	0.09	0.12	0.44	0.47	0.14
3yr	0.24	0.20	0.06	0.08	0.39	0.43	0.18
5yr	0.27	0.16	0.04	0.05	0.36	0.39	0.20
10yr	0.27	0.10	0.02	0.03	0.29	0.32	0.19
15yr	0.23	0.07	0.01	0.02	0.23	0.28	0.15
20yr	0.19	0.05	0.01	0.01	0.18	0.23	0.13

Table 6. VAR-implied predictability of dividend growth.

Table 6 presents predictive R^2s for dividend growth implied by a first-order VAR model for cash-flow growth, price-dividend ratio of the aggregate stock market portfolio and risk-free rate. Data statistics are reported in the "Data" panel. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "Model" panel presents predictability evidence implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data.

growth of 25% at the one-year horizon and about 44% at the five-year horizon. Bansal *et al.* (2005) and Ang and Bekaert (2007) also examine predictability of dividend and earnings growth rates in univariate and multivariate regression settings and find similar strong evidence of predictable variation in asset cash flows.

Table 8 provides evidence on predictability of multi-period excess returns by the log of the price-dividend ratio. Consistent with evidence in earlier papers, the R^2 rises with maturity, from 4% at the one-year horizon to about 31% at the five-year horizon. The model-implied predictability of equity returns is somewhat lower, but the data R^2s are well inside the 2-SE confidence bands. Return predictability is known to be highly uncertain. Not surprisingly, the model-based confidence bands for the R^2s are wide and include both zero (indicating lack of predictability) and the sample estimate. As shown in Panel B, the slope coefficients in the multi-horizon return projections implied by the model are of the right sign and magnitude compared to those in the data. Recall that variation in the risk premia in the Bansal and Yaron (2004) model is entirely due to variation in consumption

	Data	Model					
	Estimate	Median	2.5%	5%	95%	97.5%	Pop

Panel A: Predictive R^2s

	Estimate	Median	2.5%	5%	95%	97.5%	Pop
1yr	0.09	0.20	0.04	0.05	0.40	0.42	0.11
3yr	0.06	0.09	0.00	0.00	0.31	0.36	0.05
5yr	0.04	0.06	0.00	0.00	0.32	0.38	0.04

Panel B: Predictive Slopes

	Estimate	Median	2.5%	5%	95%	97.5%	Pop
1yr	0.07	0.35	0.12	0.16	0.56	0.60	0.19
3yr	0.11	0.45	−0.06	0.02	0.94	1.04	0.26
5yr	0.09	0.49	−0.33	−0.20	1.26	1.43	0.30

Table 7. Predictability of dividend growth by PD-ratio.

Table 7 presents R^2s and slope coefficients from projecting one-, three- and five-year dividends growth of the aggregate stock market portfolio onto lagged price-dividend ratio. Data statistics are reported in the "Data" panel. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "Model" panel presents predictability evidence implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data.

volatility. Shutting down the volatility channel by assuming homoscedastic growth rates will make the conditional risk premia constant.

It is well known that the return predictability evidence is quite fragile. To highlight this, in Table 9 we run the same multi-horizon return regressions as above but alter the regressor. Instead of the traditional price-dividend ratio, we use the log dividend yield minus the real risk-free rate. Econometrically, subtracting the risk-free rate from the dividend yield leads to a regressor that is not extremely persistent and is therefore not subject to potential spurious regression problems. Conceptually, subtracting the real risk-free rate from the dividend-price ratio should make virtually no difference to its predictive ability, as only short-horizon risks embodied in the risk-free rate are subtracted from the dividend yield. In the LRR model or the habit-based model of Campbell and Cochrane (1999), the implications for return predictability with the dividend-price ratio or the real-rate adjusted

	Data	Model					
	Estimate	Median	2.5%	5%	95%	97.5%	Pop
Panel A: Predictive R^2s							
1yr	0.04	0.01	0.00	0.00	0.07	0.10	0.01
3yr	0.19	0.03	0.00	0.00	0.18	0.22	0.03
5yr	0.31	0.05	0.00	0.00	0.27	0.32	0.04
Panel B: Predictive Slopes							
1yr	-0.09	−0.09	−0.38	−0.33	0.11	0.15	−0.08
3yr	-0.27	−0.26	−0.96	−0.82	0.25	0.34	−0.24
5yr	-0.43	−0.39	−1.39	−1.21	0.38	0.57	−0.38

Table 8. Predictability of excess return by PD-ratio.

Table 8 presents R^2s and slope coefficients from projecting one-, three- and five-year excess return of the aggregate stock market portfolio onto lagged price-dividend ratio. Data statistics are reported in the the "Data" panel. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "Model" panel presents predictability evidence implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data.

dividend yield are the same. In the data, however, return predictability with the adjusted dividend yield is much weaker than the one implied by the price-dividend ratio. As shown in Tables 8 and 9, once the dividend-price ratio is replaced with the adjusted dividend yield, the level of the three-year horizon R^2 drops from 19% to 7%, while the five-year horizon R^2 drops from 31% to only 14%. This evidence raises serious concerns about the magnitude of return predictability in the data. The difference in predictability evidence reported in Table 8 and Table 9 suggests that much of the ability of the dividend yield to predict future returns might be spurious and due to the very high persistence of the observed price-dividend ratio (e.g., Stambaugh, 1999). Adjusting the dividend-price ratio for the risk-free rate lowers the persistence in the predictive variable and ensures that the regressor is well behaved. This alleviates the possibility of a spurious regression and provides more reliable estimates. Therefore, the magnitude of predictability

	Data	Model					
	Estimate	Median	2.5%	5%	95%	97.5%	Pop
Panel A: Predictive R^2s							
1yr	0.03	0.01	0.00	0.00	0.07	0.09	0.01
3yr	0.07	0.03	0.00	0.00	0.18	0.21	0.03
5yr	0.14	0.04	0.00	0.00	0.26	0.31	0.04
Panel B: Predictive Slopes							
1yr	0.97	0.93	−1.64	−1.33	3.60	4.44	0.74
3yr	2.14	2.63	−4.93	−3.71	9.89	11.64	2.18
5yr	3.65	4.21	−7.83	−5.37	13.96	16.70	3.72

Table 9. Predictability of excess return by adjusted dividend yield.

Table 9 presents R^2s and slope coefficients from projecting one-, three- and five-year excess return of the aggregate stock market portfolio onto lagged dividend yield adjusted by the risk-free rate. Data statistics are reported in the "Data" panel. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "Model" panel presents predictability evidence implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data.

with the adjusted dividend yield of about 10% at long horizons, in our view, is more plausible and close to what should be considered realistic. As shown in Table 9, the LRR model matches the level of predictability and slope coefficients from the regressions based on the adjusted dividend price ratio quite well.

4.4   Forward- and Backward-Looking Models

Alternative asset pricing models generally are able to account for the equity and risk-free rate puzzles, and therefore may be hard to distinguish if focusing only on these dimensions. However, it may be possible to learn about the plausibility of different models by evaluating the link between price-dividend ratios and consumption growth. In the LRR model, current price-dividend ratios are determined by time-varying expected growth and consumption volatility. Hence, current prices anticipate the future state of

the economy: in the model, a drop in current price-dividend ratios reflects either a decline in future expected growth and/or a rise in future volatility. In this sense, the LRR model is forward-looking. In contrast, in the habit model, the shock of habit is driven by lagged consumption growth, and a reduction in growth rates raises risk aversion, the equity premium, and the discount rate, leading to a fall in the current price-dividend ratio. That is, backward consumption plays an important role in determining current prices. This important distinction between the two models provides an avenue to evaluate their plausibility in the data. To accomplish this we solve and simulate the habit model. In particular, we simulate cash-flow and asset price data from the habit model using the same calibration as in Campbell and Cochrane (1999) and relying on their numerical solution methods. As the standard sets of model implications for asset returns are already reported in their paper, for brevity, we do not repeat them here.

To highlight the key distinction between the two models, we run the following regression:

$$p_{t+1} - d_{t+1} = \alpha_0 + \sum_{j=1}^{L} \alpha_j \Delta c_{t+1-j} + u_{t+1}.$$

In the actual data and in the simulated data, we regress the log of price-dividend ratio on L lags of consumption growth.

Figure 2 reports the evidence in the data and the two models for various lag-length L. To have a uniform metric for drawing inferences and model comparison, in Figure 2, we rely on the data-based standard errors constructed using a block-bootstrap. The shaded area in the figure corresponds to the 95% confidence band around the data estimates. In the data, at all lag-lengths, predictability of the price dividend ratio by lagged consumption growth is close to zero. For example, for the five-year lag-length, lagged consumption forecasts the future price dividend ratio with an R^2 of only 3%. In the LRR model, as in the data, future price dividend ratio is predicted by lagged consumption with an R^2 that is close to zero. However, in the habit model, price dividend ratio predictability by lagged consumption is quite large — at the five-year lag-length, lagged consumption predicts future prices with an R^2 of 41%. At the ten-year horizon, the predictability, in the population, is 50%.[9] This is not surprising as prices in the Campbell

---

[9] The data R^2s are well below the 2.5-percentile of the finite-sample distributions of the habit model, for all lag lengths.
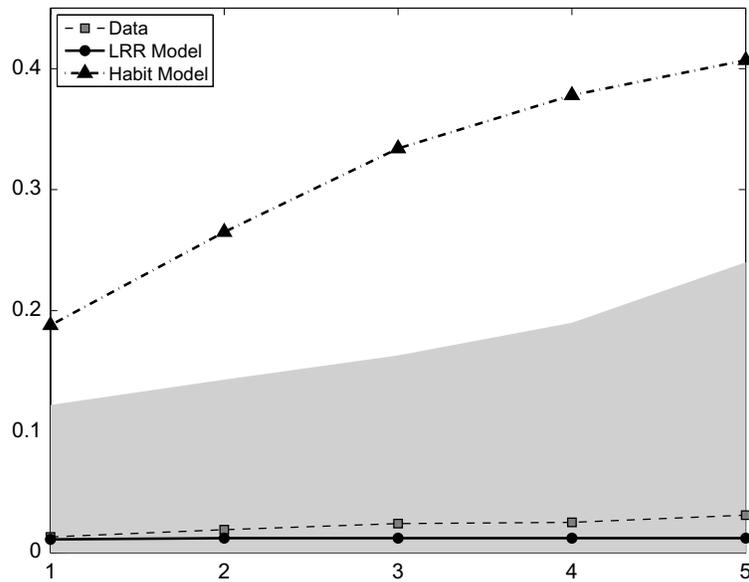
Figure 2. Price-dividend ratio and backward consumption growth.

Figure 2 plots the R^2 for regressing future log price-dividend ratio onto distributed lags of consumption growth:

$$p_{t+1} - d_{t+1} = \alpha_0 + \sum_{j=1}^{L} \alpha_j \Delta c_{t+1-j} + u_{t+1}$$

where L, the number of lags, is depicted on the x-axis. The shaded area in the figure corresponds to the 95% confidence band in which data-based standard errors are constructed using a block-bootstrap. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "LRR Model" panel presents predictability evidence implied by the Long-Run Risks model. The "Habit Model" panel shows the corresponding statistics in the habit model of Campbell and Cochrane (1999).

and Cochrane (1999) model are driven primarily by the habit stock and, hence, by movements in the lagged consumption. The lack of predictability of price-dividend ratios by past consumption growth in the data presents an important challenge for habit models, which emphasize the backward-looking consumption predictability channel for asset price determination. Related evidence regarding the predictability of price-dividend ratios in the LRR and habit model, in a set-up where dividends and consumption are cointegrated, is also presented in Bansal *et al.* (2007). Yu (2007) explores

the distinction between the forward-looking LRR model and the backward-looking habit model by looking at long-horizon correlations of returns with consumption growth and finds that the LRR model matches the data much better. More recently, Lustig *et al.* (2009) provide data-driven estimates of the wealth-consumption ratio and the risk premium on aggregate wealth and compare the LRR and habit-models; they document that their estimates and findings are quite close to the LRR model.

Beeler and Campbell (2011) regress leads and lags of consumption growth (over one, three and five years) onto the current price-dividend ratio. They write that "the price-dividend ratio in the long run risks model is just as correlated with past consumption growth as it is correlated with forward looking consumption growth" — this feature of the LRR model is indeed consistent with the data and provides additional evidence in favor of the LRR model. This should not be surprising since we have already reported that the predictability of future consumption growth by the price-dividend ratio in the LRR model is relatively modest (see Table 5), while the model's forward-looking feature produces relatively low correlations of price-dividend ratio with past consumption growth (see Figure 2). For completeness, in Figure 3 we report the R^2s from regressing leads and lags of consumption growth (for one, three, and five years) onto the price-dividend ratio for both the LRR and the Campbell and Cochrane (1999) model. Again, to facilitate a uniform inference and model comparison, the figure provides data-based standard errors constructed using a block-bootstrap. The LRR model fairs well. As in the data, the R^2s from regressing lagged (negative *j*s) accumulated consumption growth onto the current price-dividend ratio are small. Similarly, the R^2s from regressing future accumulated consumption growth onto the current price-dividend ratio are small in the data and the model. It is important to note that in all the cases the LRR model estimates are well within the data confidence band. In sharp contrast, the R^2s from the habit model are too large relative to the data and are significantly outside the data's confidence band for lagged horizons -4 to -1, for *all* horizons of consumption growth (one, three, and five year accumulated consumption growth). In sum, the habit model is significantly at odds with this important data feature, while the LRR model matches the data well.

In terms of the underlying economics of the two models and to highlight the distinction between them, consider the sharp decline in asset prices over the 2007–2008 period. According to the LRR model, the decline would be
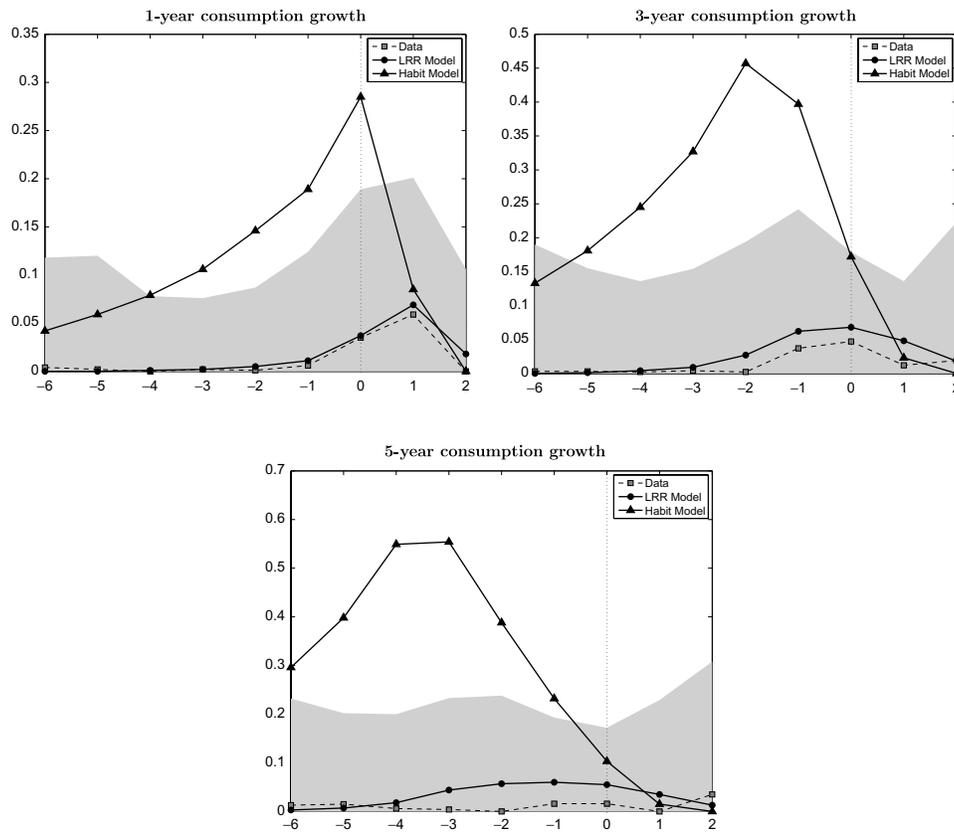
Figure 3. Leads and lags of consumption growth and price-dividend ratio.

Figure 3 plots the R^2s from regressing leads and lags of consumption growth onto the price-dividend ratio–leads (lags) are negative (positive) values on the *x*-axis. In the top, middle, and bottom figures, consumption growth is one, three, and five years, respectively. The shaded area in the figure corresponds to the 95% confidence band in which data-based standard errors are constructed using a block-bootstrap. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "LRR Model" panel presents predictability evidence implied by the Long-Run Risks model. The "Habit Model" panel shows the corresponding statistics in the habit model of Campbell and Cochrane (1999).

attributed to a decline in expected growth and/or a rise in future consumption volatility. To explain the same decline, the Campbell and Cochrane habit model would argue that a string of past and current negative consumption shocks raises risk aversion and the discount rate, leading to a

decline in asset prices. As shown in Figure 2, there is not much evidence for this channel, as lagged consumption does not forecast movements in future prices.

4.5 Volatility

As discussed above, Bansal and Yaron (2004) introduce the volatility channel and show that volatility risks are priced and contribute significantly to the equity risk premia. Fluctuations in volatility are the source of time-varying risk premia in the model (that is, risk premia varies as aggregate risk varies). An important implication of the volatility channel in the LRR set-up, with a preference for early resolution of uncertainty, is that higher volatility lowers the price-dividend ratio. Table 10 reports the evidence on the relation between asset prices and consumption volatility. The annual realized volatility of consumption is measured by fitting an

	Data	Model					
	Estimate	Median	2.5%	5%	95%	97.5%	Pop
Panel A: Predictive R^2s							
1yr	0.06	0.03	0.00	0.00	0.15	0.17	0.08
3yr	0.14	0.10	0.00	0.00	0.35	0.40	0.26
5yr	0.20	0.13	0.00	0.00	0.43	0.48	0.33
Panel B: Predictive Slopes							
1yr	−0.66	−0.98	−3.49	−3.03	0.44	0.68	−1.22
3yr	−0.57	−0.92	−2.85	−2.56	0.10	0.27	−1.18
5yr	−0.55	−0.83	−2.59	−2.27	0.09	0.24	−1.13

Table 10. Predictability of volatility of consumption growth by PD-ratio.

Table 10 presents R^2s and slope coefficients from projecting one-, three- and five-year volatility of consumption growth onto lagged price-dividend ratio of the aggregate stock market portfolio. Data statistics are reported in the "Data" panel. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "Model" panel presents predictability evidence implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data.

AR(1) process to consumption growth and taking the absolute value of the residuals. At date t, the K-horizon future realized volatility is measured by $\log \sum_{j=1}^{K} |u_{t+j}|$, where u_t is the date-t consumption residual. We regress this measure of volatility on the current price-dividend ratio to see how well current asset prices predict future consumption volatility. In the data, the predictive R^2 rises from 6% to 20%, indicating that consumption volatility is indeed predictable and time-varying. The model matches this data dimension very well — the model confidence bands include the data R^2s and, similar to the data, the magnitude of the model-implied R^2 rises with horizon. Panel B reports the slope coefficients from these regressions. In the data, the current price-dividend ratio and volatility at all horizons are negatively related. The size of the slope coefficients is quite large, and the model captures their magnitude quite well. In the data there is also pronounced negative relationship between the price-earnings ratio and future consumption volatility; the correlation between the price-earnings ratio and the five-year measure of future of consumption volatility is -0.38. Bansal *et al.* (2005) show that the negative relationship between valuation ratios and future uncertainty is robust to alternative measures of volatility, cash-flow data, and is present in other countries in addition to the US.

The LRR model also implies a negative relationship between the price-dividend ratio and future return volatility. In Table 11 we ask how much predictability does the model imply for an integrated return volatility measure? The integrated volatility is constructed by summing up the demeaned monthly squared returns. The table shows that in the data return volatility is predictable; the R^2s are 11% and 6% for the one- and five-year horizons respectively. While the median R^2s are somewhat larger than their data counterparts, the confidence bands for these R^2s contain the data magnitudes. The model matches well the negative slope coefficients. This underscores the economics in the LRR model, that when IES is larger than one, higher consumption volatility (and return volatility) are negatively related to the price-dividend ratio.[10]

Recently, Bansal *et al.* (2007b) estimate the LRR model and show that both components of the long-run risk model, long-run expected growth and

---

[10] Generalized LRR models that incorporate transitory jumps and/or two volatility components (e.g., Bansal and Shaliastovich (2009), Drechsler and Yaron (2011), Bansal *et al.* (2010) discussed earlier) further improve the model's predictability implications by reducing the predictability of future consumption growth and integrated volatility by the price-dividend ratio.

	Data	Model					
	Estimate	Median	2.5%	5%	95%	97.5%	Pop
Panel A: Predictive R^2s							
1yr	0.11	0.22	0.01	0.02	0.48	0.53	0.44
3yr	0.08	0.30	0.01	0.03	0.61	0.66	0.58
5yr	0.06	0.29	0.01	0.02	0.62	0.67	0.61
Panel B: Predictive Slopes							
1yr	−0.04	−0.05	−0.08	−0.08	−0.02	−0.01	−0.08
3yr	−0.02	−0.05	−0.08	−0.08	−0.01	−0.01	−0.08
5yr	−0.01	−0.04	−0.08	−0.08	−0.01	−0.01	−0.08

Table 11. Predictability of volatility of excess return by PD-ratio.

Table 11 presents R^2s and slope coefficients from projecting one-, three- and five-year variance of excess return of the aggregate stock market portfolio onto lagged price-dividend ratio. Data statistics are reported in the "Data" panel. The data employed in the estimation are real, compounded continuously, sampled on an annual frequency and cover the period from 1930 to 2008. The "Model" panel presents predictability evidence implied by the Long-Run Risks model. The first five columns in the right panel represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data.

long-run volatility fluctuations, contribute to the cross-sectional dispersion of expected returns. Related evidence showing that the volatility channel is important for the cross-section is also provided by Tédongap (2006).

Empirical evidence presented in Sections 4.1–4.4 is robust to alternative methods of computing standard errors. We have also evaluated the model fit using the data-based (bootstrap) confidence regions for all statistics of interest. We construct empirical distributions by re-sampling the observed data 10,000 times in blocks of eight years with replacement, and find that the inference based on the bootstrap standard errors is virtually unchanged from the one reported above.

4.6 Parameter Magnitudes

In terms of the preference parameters, the magnitude of risk aversion typically used in the LRR literature is 10 or below, which is consistent with the

magnitudes argued for in Mehra and Prescott (1985). The Campbell and Cochrane habit model, in contrast, relies on extreme risk aversion that can be as high as 250 in some states.

With respect to the IES we make the following observations: (i) in the literature, there is a wide-range of estimates of the IES and a significant number of these estimates are larger than one, (ii) the stochastic volatility in consumption and hence in returns leads to significant downward biases in the estimated IES when regressing consumption growth onto the risk free rate, (iii) when the risk-free rate is measured with error, regressions of either consumption growth onto the risk-free rate or vice-versa lead to very dispersed estimates of the IES, as found in the literature, (iv) additional informative moments useful for identifying the IES, such as moments relating volatility and valuation ratios, suggest the IES should be greater than one.

The literature examining the IES magnitude in the data leads to estimates that are both well above and below one. A large number of papers (Hansen and Singleton, 1982; Attanasio and Weber, 1989; Beaudry and van Wincoop, 1996; Vissing-Jorgensen, 2002; Attanasio and Vissing-Jorgensen, 2003; Mulligan, 2004; Gruber, 2006; Guvenen, 2006; Hansen *et al.*, 2007; Engelhardt and Kumar, 2008; Barro, 2009), show that the IES is large and indeed greater than one. Hall (1988) and Campbell (1999), however, argue that IES is small and close to zero. There seems to be little agreement on the magnitude of IES in the data, and both high and low magnitudes seem possible.

The Hall (1988) and Campbell (1999) argument for low IES is based on estimating the slope coefficient from regressing consumption growth on the real rate. In the data, this slope coefficient is small. As discussed above, it is not clear that this regression is the best way to estimate the IES. Moreover, Bansal and Yaron (2004) show that, if consumption volatility is time-varying, the slope coefficient from regressing consumption growth onto the real-risk free rate is downward biased and, hence, cannot be a guide for the true value of the IES. Beeler and Campbell (2011) question the magnitude of this bias. They report that when the population magnitude of the IES is 1.5, the finite sample estimate, using Hall's approach of regressing consumption growth on the real rate is centered around 0.93; this is a big bias in statistical and, more importantly, in economic terms, as the estimated value, in contrast to the population value, is below one. Moreover,

the 2-SE confidence bands for this regression encompass values as low as 0.28.

Beeler and Campbell also suggest using an instrumental variable approach to circumvent the bias highlighted in Bansal and Yaron (2004). While the bias in the model is smaller in this case, the IES estimate in the data that they report is negative, suggesting severe measurement error and/or lack of power. More generally, this regression approach is highly sensitive to samples, assets used, and instruments. For example, in the data and using the instrumental variable approach, Beeler and Campbell (2009) also report an IES that is larger than one. Further, when we use the risk-free rate as the asset, and an expanded instrument set that includes lagged (three years) consumption growth, price-dividend ratio, and risk-free rate the IES estimate in the data is 1.115 (1.656).

In Table 12 we demonstrate the sensitivity of the above regressions to the presence of measurement errors. Based on the model simulations, the top row of Table 12 provides the finite sample values of the IES from regressing consumption growth onto the risk-free rate, while the bottom row provides analogous output from regressing the risk-free rate onto consumption growth. For both entries, the table reports evidence for the case

Specification	Model				
	Median	2.5%	5%	95%	97.5%
Δc onto r_f	1.39	-12.26	-4.57	7.06	12.11
r_f onto Δc	0.84	-1.27	-0.82	2.33	2.70

Table 12. IES estimates when the risk-free rate is measured with error.

Table 12 presents the IES estimates from projecting the risk free rate onto consumption growth (first row) and the IES estimates from regressing consumption growth onto the risk-free rate (second row). The entries present evidence implied by the Long-Run Risks model when the risk-free rate is augmented with plausible measurement error. The magnitude of the measurement error is calibrated so the model-implied risk-free rate volatility is increased to match that of the data. The true value of the IES in the simulated data is 1.5. The first five columns represent percentiles of finite sample Monte-Carlo distributions. Population values (Pop) are computed from a very long sample of simulated data. The table shows how a small degree of measurement error can result in a very dispersed, low and large estimates of IES, as observed in the literature.

in which the risk free rate is measured with error. It is difficult to measure the real risk-free rate and the presence of measurement error is very likely. To calibrate the magnitude of the added measurement error to the risk free rate, we require the volatility of the model-implied risk-free rate to be equal to its volatility in the data, i.e., we add an *i.i.d* noise to the model-implied risk-free rate to increase its volatility to the level in data. Table 12 shows that a small amount of measurement error induces a finite sample distribution of IES estimates that, consistent with the literature, encompasses negative, small, as well as large estimates.

The central issue about the IES is its true value in the data and if it is larger than one. As mentioned, the above approach of regressing consumption growth on the risk-free rate or the market return is very sensitive to samples and instruments. A far better approach to measure the IES is to use a larger set of model-based moment restrictions, for example, some that exploit the level of the real rate, the consumption volatility effects on price-dividend ratios, and incorporate the conditional version of the Euler equation associated with the real bond. Bansal *et al.* (2007b) and Hansen *et al.* (2007) pursue the approach of using a larger set of moments and find an IES estimate larger than one. Overall, asset market data suggests that the IES is larger than one.

4.7   Yield Curve

Evaluating the model implications for the yield curve, Bansal and Yaron (2004), Piazzesi and Schneider (2007), Bansal and Shaliastovich (2009) show that the real curve is downward sloping. That is, real bonds provide insurance in the model. This implication of the model is consistent with the real yield curve data from the UK (comparable data sample for the US is not available). In particular Evans (1998) shows that the real yield curve in the UK is downward sloping; the ten minus one-year yield spread is −0.88% for the 1984.01 to 1995.08 sample. We find that for the more recent sample of 1996.07–2008.12, the ten-year minus one-year term spread is −1.92%. Therefore, a downward-sloping real yield curve is the appropriate target for models.

Piazzesi and Schneider (2007) and Bansal and Shaliastovich (2009) show that the nominal yield curve in the LRR model is upward sloping since the inflation risk-premia in the model increases with maturity. Bansal and Shaliastovich (2009) further show that the model can account for the

predictability evidence on bond returns and the violations of the expectations hypothesis documented for nominal bonds.[11]

Beeler and Campbell (2011) argue that, for some calibrations, the price of the real console (a real bond that pays one unit of consumption in every period) is infinity in the LRR model. This is hardly surprising, as even in the standard CRRA model with *i.i.d* consumption growth, the real yield curve is flat and the yield can be negative when risk aversion is sufficiently high or when consumption volatility is high — in this case, the price of a pure discount bond at infinity is infinity. The price of the console will also be infinity. This model implication for the CRRA model and for the LRR model should not be a concern, as a proper equilibrium exists and the price of the consumption claim (i.e., aggregate wealth) in the economy is finite. Also, there are no data counterparts to a real console. Nominal consoles do exist (that deliver one dollar each period) and the price of the nominal console is finite in the LRR model, as the nominal yields are positive and the nominal yield curve is upward sloping.

4.8   Additional Considerations

In addition to data features such as the equity-premium and the risk-free rate puzzles, about which we learn mostly from the time series, there are additional puzzles, which primarily focus on the cross-sectional differences in expected returns, such as the differences in returns of size sorted, book-to-market sorted, and momentum sorted portfolios (see Fama and French, 1992). These data features also help us learn about the model dynamics and the economic sources of risks. The cross-sectional differences in expected returns on these assets must reflect differences in systematic risks. To evaluate the LRR model, Bansal *et al.* (2005) measure the exposure of cash-flows to long-run consumption growth risks for 30 portfolios, sorted by size, book-to-market, and momentum. They show that these long-run cash-flow betas can explain more than 60% of the cross-sectional differences in expected returns of these 30 portfolios. At the same time, exposure to short-term consumption shocks or markets betas have almost no explanatory power in accounting for the cross-sectional differences in expected returns. The evidence in the cross-section is robust to alternative ways of measuring

---

[11] In addition, they show that the model can account for the violations of the expectations hypothesis in currency markets, that is, differences in expected returns between foreign and domestic bonds.

the exposure of cash-flows to long-run consumption shocks. Bansal *et al.* (2009) measure long-run consumption betas of the cross-section of assets by exploiting the cointegrating relation between aggregate consumption and dividends and show that this long-run cointegration-based dividend beta is critical for explaining both the cross-sectional differences in short-horizon expected returns and the long-horizon differences in risk premia. These papers underscore the importance of long-run consumption-based cash-flow risks in explaining differences in expected returns across assets. Malloy *et al.* (2009) focus on consumption of stock-holders and show that long-run risks in their consumption also accounts for the cross-section of assets returns.

Using simulations from a calibrated model, Kiku (2006) and Bansal *et al.* (2007b) show that the LRR model can simultaneously account for the differences in value and growth returns and the empirical failure of the standard CAPM betas. Santos and Veronesi (2010) evaluate the ability of the habit-based model to explain the cross-section of book-to-market returns, and show that the benchmark model of Campbell and Cochrane (1999) implies a "growth" rather than a value premium. They argue that since growth firms are characterized by a relatively long duration of their cash flows, they are more sensitive to discount rate risks than value firms and, consequently, have to carry a high risk premium inside the habit model. As a result the habit-based model cannot account for the cross-sectional differences in expected returns.

Most recent work on LRR models incorporates jumps in the expected growth and/or volatility dynamics. Eraker and Shaliastovich (2008) provide a framework for analyzing jumps in the growth rates and consumption volatility. They show that this can help account for some of the puzzling options markets features. More extensively, Drechsler and Yaron (2011) incorporate jumps in the expected growth and volatility dynamics and show that this augmented LRR framework can explain the variance premium in options markets. Drechsler (2008) highlights the effect of model uncertainty in the LRR framework with jumps and options market data. Bansal and Shaliastovich (2010) incorporate a confidence risk channel in the LRR framework that includes jumps. Shaliastovich (2008) shows that this broader LRR set-up can empirically account for several option market puzzles. These extensions open up a channel for jumps in expected returns and yield significantly higher asset excess returns and lower consumption

growth predictability by the price-dividend ratio, relative to the LRR benchmark model without jumps.

5 Conclusions

In this article we provide an empirical evaluation of the Bansal and Yaron (2004) LRR model and compare some key features of the model to the Campbell and Cochrane (1999) habit model. We show that the LRR model matches the key asset market facts, and almost all the model implications are well within the usual two-standard error confidence range from the data. We provide statistical evidence, which shows that consumption growth and consumption volatility are predictable both in the short and in the long run, and that the LRR model replicates these data features. Bansal and Yaron (2004) develop and underscore the importance of the volatility channel in their LRR model; this channel leads to volatility shocks receiving separate risk compensation in asset markets. The volatility channel of Bansal and Yaron (2004) also helps identify the intertemporal elasticity of substitution, as a large IES is needed to capture the robust evidence that higher consumption volatility lowers asset prices. We show that the data and the LRR model are quite consistent in their quantitative implications for the volatility and valuation link.

We provide an important distinction between the Campbell and Cochrane (1999) habit model and the Bansal and Yaron (2004) LRR model. Price-dividend ratios in the habit model are driven by long lags of consumption and, consequently, past consumption growth forecasts future price-dividend ratios with an R^2 of up to 40%. In the data, this predictability is close to zero. In the LRR model, the price-dividend ratio is forward-looking, as it is driven by anticipations of future growth and risk (consumption volatility). Therefore, consistent with the data, lagged consumption growth in the LRR model does not predict future price-dividend ratios. The large predictability of future price-dividend ratios with lagged consumption in the habit model raises considerable questions about its plausibility.

Our evidence calls for estimation procedures, which can incorporate a wide range of data features to evaluate the LRR, habit-based, or other models. Early approaches that use EMM and GMM to empirically test these models include Bansal *et al.* (2007) and Bansal *et al.* (2007b), respectively.

The moments used in these approaches would have to be extended to exploit all the data features discussed in this paper.

Appendix

The price-dividend ratio for the market claim to dividends, $z_{m,t} = A_{0,m} + A_{1,m}x_t + A_{2,m}\sigma_t^2$, where

$$
\begin{aligned}
A_{0,m} &= \frac{1}{1-\kappa_{1,m}}\left[\Gamma_0 + \kappa_{0,m} + \mu_d + \kappa_{1,m}A_{2,m}(1-v)\bar{\sigma}^2 \right. \\
&\quad \left. + \frac{1}{2}\left(\kappa_{1,m}A_{2,m} - \lambda_w\right)^2\sigma_w^2\right] \\
A_{1,m} &= \frac{\phi_m - \frac{1}{\psi}}{1-\kappa_{1,m}\rho} \qquad\qquad\qquad\qquad\qquad\qquad (12)\\
A_{2,m} &= \frac{1}{1-\kappa_{1,m}v}\left[\Gamma_2 + \frac{1}{2}\left((\pi - \lambda_\eta)^2 + (\kappa_{1,m}A_{1,m}\varphi_e - \lambda_\epsilon)^2\right)\right],
\end{aligned}
$$

where $\Gamma_0 = \log\delta - \frac{1}{\psi}\mu_c - (\theta-1)\left[A_2(1-v)\bar{\sigma}^2 + \frac{\theta}{2}\left(\kappa_1 A_2\sigma_w\right)^2\right]$ and $\Gamma_2 = (\theta-1)(\kappa_1 v - 1)A_2$.

The risk premium is determined by the covariation of the return innovation with the pricing kernel innovation. Thus, the risk premium for $r_{m,t+1}$ is equal to the asset's exposures to systematic risks multiplied by the corresponding risk prices,

$$
\begin{aligned}
E_t(r_{m,t+1} - r_{f,t}) + 0.5\sigma_{t,r_m}^2 &= -Cov_t(m_{t+1} - E_t(m_{t+1}), r_{m,t+1} - E_t(r_{m,t+1})) \\
&= \lambda_\eta\sigma_t^2\beta_{\eta,m} + \lambda_e\sigma_t^2\beta_{e,m} + \lambda_w\sigma_w^2\beta_{w,m},
\end{aligned}
$$

where the asset's βs are defined as,

$$
\begin{aligned}
\beta_{\eta,m} &= \pi \\
\beta_{e,m} &= \kappa_{1,m}A_{1,m}\varphi_e \\
\beta_{w,m} &= \kappa_{1,m}A_{2,m}.
\end{aligned}
$$

Acknowledgments

References

Ang, A. and G. Bekaert. 2007. "Stock Return Predictability: Is it There?" *Review of Financial Studies* 20: 651–707.

Attanasio, O. P. and A. Vissing-Jorgensen. 2003. "Stock-market Participation, Intertemporal Substitution, and Risk-aversion." *American Economic Review* 93: 383–391.

Attanasio, O. P. and G. Weber. 1989. "Intertemporal Substitution, Risk Aversion and the Euler Equation for Consumption." *Economic Journal* 99: 59–73.

Bansal, R., R. F. Dittmar, and D. Kiku. 2009. "Cointegration and Consumption Risks in Asset Returns." *Review of Financial Studies* 22: 1343–1375.

Bansal, R., R. F. Dittmar, and C. Lundblad. 2005. "Consumption, Dividends, and The Cross-section of Equity Returns." *Journal of Finance* 60: 1639–1672.

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

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

Bansal, R., D. Kiku, and A. Yaron. 2007a. "A Note on the Economics and Statistics of Predictability: A Long Run Risks Perspective." University of Pennsylvania, Working Paper.

Bansal, R., D. Kiku, and A. Yaron. 2007b. "Risks for the Long Run: Estimation and Inference." University of Pennsylvania, Working Paper.

Bansal, R., D. Kiku, and A. Yaron. 2010. "Long Run Risks, the Macroeconomy, and Asset Prices." *American Economic Review* 100(2): 542–546.

Bansal, R. and I. Shaliastovich. 2009. "A Long-run Risks Explanation of Predictability Puzzles in Bond and Currency Markets." Duke University, Working Paper.

Bansal, R. and I. Shaliastovich. 2010. "Confidence Risk and Asset Prices." *American Economic Review* 100(2): 537–541.

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

Barndorff-Nielsen, O. E. and N. Shephard. 2001. "Non-gaussian Ornstein-uhlenbeck-based Models and Some of Their Uses in Financial Economics." *Journal of Royal Statistical Society* 63: 167–241.

Barro, R. 2009. "Rare Disasters, Asset Prices, and Welfare Costs." *American Economic Review* 99(1): 243–264.

Beaudry, P. and E. van Wincoop. 1996. "The Intertemporal Elasticity Of Substitution: An Exploration Using a Panel of State Data." *Economica* 63: 495–512.

Beeler, J. and J. Y. Campbell. 2009. "The Long-run Risks Model and Aggregate Asset Prices: An Empirical Assessment." NBER Working Paper #14788.

Beeler, J. and J. Y. Campbell. 2011. "The Long-run Risks Model and Aggregate Asset Prices: An Empirical Assessment." Harvard Working Paper.

Campbell, J. Y. 1999. "Asset Prices, Consumption and the Business Cycle." In: *Handbook of Macroeconomics, Volume 1*, J. B. Taylor and M. Woodford, eds., Elsevier, The Netherlands.

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: 205–255.

Croce, M. M. 2005. "Long-run Productivity Risk: A New Hope for Production-based Asset Pricing." University of North Carolina, Working Paper.

Drechsler, I. 2008. Uncertainty, time-varying fear, and asset prices. New York University, Working Paper.

Drechsler, I. and A. Yaron. 2011. "What's vol got to do with it." *Review of Financial Studies* 24(1): 1–45.

Engelhardt, G. V. and A. Kumar. 2008. "The Elasticity of Intertemporal Substitution: New Evidence from 401(k) Participation." Federal Reserve Bank of Dallas, Working Paper #0812.

Epstein, L. G. and S. E. Zin. 1989. "Substitution, Risk Aversion, and the Intertemporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57: 937–969.

Eraker, B. and I. Shaliastovich. 2008. "An Equilibrium Guide to Designing Affine Pricing Models." *Mathematical Finance* 18: 519–543.

Evans, M. 1998. "Real Rates, Expected Inflation and Inflation Risk Premia." *Journal of Finance* 53: 187–218.

Fama, E. and K. French. 1992. "The Cross-section of Expected Stock Returns." *Journal of Finance* 59: 427–465.

Gruber, J. 2006. "A Tax-based Estimate of the Elasticity of Intertemporal Substitution." NBER Working Paper #11945.

Guvenen, F. 2006. "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective." *Journal of Monetary Economics* 53: 1451–1472.

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

Hansen, L., J. Heaton, and N. Li. 2008. "Consumption Strikes Back? Measuring Long-run Risk." *Journal of Political Economy* 116(2): 260–302.

Hansen, L. P., J. Heaton, J. Lee, and N. Roussanov. 2007. "Intertemporal Substitution and Risk Aversion." In: *Handbook of Econometrics*, J. J. Heckman and E. E. Leamer, eds., Vol. 6, Elsevier, The Netherlands.

Hansen, L. P. and K. Singleton. 1982. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica* 63: 767–804.

Hodrick, R. 1992. "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement." *Review of Financial Studies* 5–3: 357–386.

Kaltenbrunner, G. and L. Lochstoer. 2010. "Long-run Risk Through Consumption Smoothing." *Review of Financial Studies* 23: 3141–3189.

Kiku, D. 2006. "Long Run Risks and the Value Premium Puzzle." Technical Report, Duke University, Working Paper.

Kim, C.-J. and C. R. Nelson. 1999. "Has the Economy Become More Stable? A Bayesian Approach Based on a Markov-switching Model of the Business Cycle." *The Review of Economics and Statistics* 81: 608–616.

Lustig, H., S. Nieuwerburgh, and A. Verdelhan. 2009. "The Wealth-Consumption Ratio." New York Univeristy, Working Paper.

Malloy, C., T. Moskowitz, and A. Vissing-Jorgensen. 2009. "Long-run Stockholder Consumption Risk and Asset Returns." *Journal of Finance* LXIV(6): 2427–2479.

McConnell, M. and G. Perez-Quiros. 2000. "Output Fluctuations in the United States: What has Changed since the Early 1980s?" *American Economic Review* 90(5): 1464–1476.

Mehra, R. and E. C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15: 145–162.

Mulligan, C. B. 2004. "What Do Aggregate Consumption Euler Equations Say About the Capital Income Tax Burden?" NBER Working Paper #10262.

Piazzesi, M. and M. Schneider. 2007. "Equilibrium Yield Curves." In *NBER Macroeconomics Annual 2006*, D. Acemoglu, K. Rogoff, and M. Woodford, eds., The MIT Press, Cambridge, MA.

Santos, T. and P. Veronesi. 2010. "Habit Formation, the Cross Section of Stock Returns, and the Cash-Flow Risk Puzzle." *Journal of Financial Economics* 98(2): 385–413.

Shaliastovich, I. 2008. "Learning, Confidence and Option Prices." Duke University, Working Paper.

Shiller, R. and P. Perron. 1985. "Testing the Random Walk Hypothesis: Power versus Frequency of Observation." *Economic Letters* 18: 381–386.

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

Stock, J. H. and M. Watson. 1988. "Variable Trends in Economic Time Series." *Journal of Economic Perspectives* 2(3): 147–174.

Stock, J. H. and M. Watson. 2002. "Has the Business Cycle Change and Why?" In: *NBER Macroeconomics Annual 2002*, M. Gertler and K. Rogoff, eds., Cambridge, MA: The MIT Press.

Tédongap, R. 2006. "Consumption Volatility and the Cross-section of Stock Returns." Stockholm School of Economics, Working Paper.

Vissing-Jorgensen, A. 2002. " Limited Asset Market Participation and the Elasticity of Intertemporal Substitution." *Journal of Political Economy* 110: 825–853.

Working, H. 1960. "Note on the Correlation of First Differences of Averages in a Random Chain." *Econometrica* 28: 916–918.

Yu, J. 2007. "The Long and the Short of Asset Prices: Using Long Run Consumption-Return Correlations to Test Asset Pricing Models." University of Minnesota, Working Paper.