Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles

RAVI BANSAL and AMIR YARON*

ABSTRACT

We model consumption and dividend growth rates as containing (1) a small long-run predictable component, and (2) fluctuating economic uncertainty (consumption volatility). These dynamics, for which we provide empirical support, in conjunction with Epstein and Zin's (1989) preferences, can explain key asset markets phenomena. In our economy, financial markets dislike economic uncertainty and better long-run growth prospects raise equity prices. The model can justify the equity premium, the risk-free rate, and the volatility of the market return, risk-free rate, and the price–dividend ratio. As in the data, dividend yields predict returns and the volatility of returns is time-varying.

SEVERAL KEY ASPECTS OF ASSET MARKET DATA pose a serious challenge to economic models.[1] It is difficult to justify the 6% equity premium and the low risk-free rate (see Mehra and Prescott (1985), Weil (1989), and Hansen and Jagannathan (1991)). The literature on variance bounds highlights the difficulty in justifying the market volatility of 19% per annum (see Shiller (1981) and LeRoy and Porter (1981)). The conditional variance of the market return, as shown in Bollerslev, Engle, and Wooldridge (1988), fluctuates across time and is very persistent. Price–dividend ratios seem to predict long-horizon equity returns (see Campbell and Shiller (1988)). In addition, as documented in this paper, consumption volatility and future price–dividend ratios are significantly negatively correlated—a rise in consumption volatility lowers asset prices.

[1] Notable papers addressing asset market anomalies include Abel (1990, 1999), Bansal and Coleman (1997), Barberis, Huang, and Santos (2001), Campbell and Cochrane (1999), Cecchetti, Lam, and Mark (1990), Chapman (2002), Constantinides (1990), Constantinides and Duffie (1996), Hansen, Sargent, and Tallarini (1999), Heaton (1995), Heaton and Lucas (1996), and Kandel and Stambaugh (1991).

We present a model that helps explain the above features of asset market data. There are two main ingredients in the model. First, we rely on the standard Epstein and Zin (1989) preferences, which allow for a separation between the intertemporal elasticity of substitution (IES) and risk aversion, and consequently permit both parameters to be simultaneously larger than 1. Second, we model consumption and dividend growth rates as containing (1) a small persistent expected growth rate component, and (2) fluctuating volatility, which captures time-varying economic uncertainty. We show that this specification for consumption and dividends is consistent with observed annual consumption and dividend data. In our economy, when the IES is larger than 1, agents demand large equity risk premia because they fear that a reduction in economic growth prospects or a rise in economic uncertainty will lower asset prices. Our results show that risks related to varying growth prospects and fluctuating economic uncertainty can quantitatively justify many of the observed features of asset market data.

Why is persistence in the growth prospects important? In a partial equilibrium model, Barsky and DeLong (1993) and Bansal and Lundblad (2002) show that persistence in expected dividend growth rates is an important source of volatility in price–dividend ratios. In our equilibrium model, the degree of persistence in expected growth rate news affects the volatility of the price–dividend ratio and also determines the risk premium on the asset. News regarding future expected growth rates leads to large reactions in the price–dividend ratio and the ex post equity return; these reactions positively covary with the marginal rate of substitution of the representative agent, and hence lead to large equity risk premia. The dividend elasticity of asset prices and the risk premia on assets rise as the degree of permanence of expected dividend growth rates increases. We formalize this intuition in Section I with a simple version of the model that incorporates only fluctuations in growth prospects.

To allow for time-varying risk premia, we incorporate changes in the conditional volatility of future growth rates. Fluctuating economic uncertainty (conditional volatility of consumption) directly affects price–dividend ratios, and a rise in economic uncertainty leads to a fall in asset prices. In our model, shocks to consumption volatility carry a positive risk premium. The consumption volatility channel is important for capturing the volatility feedback effect; that is, return news and news about return volatility are negatively correlated. About half of the volatility of price–dividend ratios in the model can be attributed to variation in expected growth rates, and the remaining can be attributed to variation in economic uncertainty. This is distinct from models where growth rates are *i.i.d.*, and consequently, all the variation in price–dividend ratio is attributed to the changing cost of capital.

Our specification for growth rates emphasizes persistent movements in expected growth rates and fluctuations in economic uncertainty. For these channels to have a significant quantitative impact on the risk premium and volatility of asset prices, the persistence in expected growth rate has to be quite large,

close to 0.98.[2] A pertinent question is whether this is consistent with growth rate data, as observed autocorrelations in realized growth rates of consumption and dividends are small. Shephard and Harvey (1990) show that in finite samples, it is very difficult to distinguish between a purely *i.i.d.* process and one which incorporates a small persistent component. While it is hard to distinguish econometrically between the two alternative processes, the asset pricing implications across them are very different. We show that our specification for the consumption and dividend growth rates, which incorporates the persistent component, is consistent with the growth rate data and helps justify several puzzling aspects of asset market data.

We provide direct empirical evidence for fluctuating consumption volatility, which motivates our time-varying economic uncertainty channel. The variance ratios of realized consumption volatility increase up to 10 years. If residuals of consumption growth were *i.i.d.*, then the variance ratio of the absolute value of these residuals would be flat across different horizons. Evidence presented below and explored further in Bansal, Khatchatrian, and Yaron (2002) shows that realized consumption volatility predicts and is predicted by the price–dividend ratio. This again corroborates the view that consumption volatility is time-varying.

In terms of preferences, our main results are based on a risk aversion of 10 and an IES of 1.5. There is considerable debate about what are reasonable magnitudes for these parameters. Mehra and Prescott (1985) argue that a risk aversion of 10 and below seems reasonable. Our value for the IES is consistent with the findings of Hansen and Singleton (1982) and many other authors. Moreover, as established below, an IES greater than 1 is critical for capturing the observed negative correlation between consumption volatility and price–dividend ratios. Further, we show that the presence of fluctuating consumption volatility leads to a serious downward bias in the estimates for the IES using the regression approach pursued in Hall (1988). This bias may help interpret Hall's small estimates of the IES.

The remainder of the paper is organized as follows. In Section I we formalize this intuition and present the economics behind our model. The data and the model's quantitative implications are described in Section II. The last Section provides concluding comments.

I. An Economic Model for Asset Markets

Consider a representative agent with the Epstein and Zin (1989) and Weil (1989) recursive preferences. For these preferences, Epstein and Zin (1989) show that the asset pricing restrictions for gross return $R_{i,t+1}$ satisfy

$$E_t\left[\delta^\theta G_{t+1}^{-\frac{\theta}{\psi}} R_{a,t+1}^{-(1-\theta)} R_{i,t+1}\right] = 1, \tag{1}$$

---

[2] Barsky and DeLong (1993) choose a value of 1. Our choice ensures that the growth rate process is stationary.

where G_{t+1} is the aggregate gross growth rate of consumption and $R_{a,t+1}$ is the gross return on an asset that delivers aggregate consumption as its dividends each period. The parameter $0 < \delta < 1$ is the time discount factor. The parameter $\theta \equiv \frac{1-\gamma}{1-\frac{1}{\psi}}$, with $\gamma \geq 0$ being the risk-aversion parameter and $\psi \geq 0$ the IES parameter. The sign of θ is determined by the magnitudes of the risk aversion and the elasticity of substitution.[3]

We distinguish between the *unobservable* return on a claim to aggregate consumption, $R_{a,t+1}$, and the *observable* return on the market portfolio $R_{m,t+1}$; the latter is the return on the aggregate dividend claim. As in Campbell (1996), we model aggregate consumption and aggregate dividends as two separate processes; the agent is implicitly assumed to have access to labor income.

Although we solve our model numerically, we demonstrate the mechanisms working in our model via approximate analytical solutions. To derive these solutions for the model, we use the standard approximations utilized in Campbell and Shiller (1988),

$$r_{a,t+1} = \kappa_0 + \kappa_1 z_{t+1} - z_t + g_{t+1}, \tag{2}$$

where lowercase letters refer to logs, so that $r_{a,t+1} = \log(R_{a,t+1})$ is the continuous return, $z_t = \log(P_t/C_t)$ is the log price–consumption ratio, and κ_0 and κ_1 are approximating constants that both depend only on the average level of z.[4] Analogously, $r_{m,t+1}$ and $z_{m,t}$ correspond to the market return and its log price–dividend ratio.

The logarithm of the intertemporal marginal rate of substitution (IMRS) is

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

It follows that the innovation in m_{t+1} is driven by the innovations in g_{t+1} and $r_{a,t+1}$. Covariation with the innovation in m_{t+1} determines the risk premium for any asset. When θ equals 1, the above IMRS collapses to the usual case of power utility. To present the intuition of our model in a simple manner, we first discuss the case (Case I) in which there are fluctuations only in the expected growth rates. Subsequently, we present the complete model (Case II), which also includes fluctuating economic uncertainty.

A. Case I: Fluctuating Expected Growth Rates

We first solve for the consumption return $r_{a,t+1}$, as this determines the pricing kernel and consequently risk premia on the market portfolio, $r_{m,t+1}$, as well as all other assets. To do so we first specify the dynamics for consumption and dividend growth rates. We model consumption and dividend growth rates, g_{t+1}

---

[3] In particular, if $\psi > 1$ and $\gamma > 1$ then θ will be negative. Note that when $\theta = 1$, that is, $\gamma = (1/\psi)$, the above recursive preferences collapse to the standard case of expected utility. Further, when $\theta = 1$ and in addition $\gamma = 1$, we get the standard case of log utility.

[4] Note that $\kappa_1 = \exp(\bar{z})/(1 + \exp(\bar{z})) \cdot \kappa_1$ is approximately 0.997, which is consistent with the magnitude of $\bar{z}$ in our sample and with magnitudes used in Campbell and Shiller (1988).

and $g_{d,t+1}$, respectively, as containing a small persistent predictable component x_t, which determines the conditional expectation of consumption growth,

$$
\begin{aligned}
x_{t+1} &= \rho x_t + \varphi_e \sigma e_{t+1} \\
g_{t+1} &= \mu + x_t + \sigma \eta_{t+1} \\
g_{d,t+1} &= \mu_d + \phi x_t + \varphi_d \sigma u_{t+1} \\
e_{t+1}, &u_{t+1}, \eta_{t+1} \sim N.i.i.d.(0,1),
\end{aligned}
\tag{4}
$$

with the three shocks, e_{t+1}, u_{t+1}, and η_{t+1} being mutually independent.[5] Two additional parameters, $\phi > 1$ and $\varphi_d > 1$, allow us to calibrate the overall volatility of dividends (which in the data are significantly larger than that of consumption) and its correlation with consumption. The parameter ϕ, as in Abel (1999), can be interpreted as the leverage ratio on expected consumption growth.[6] It is straightforward to allow the three shocks to be correlated; however, to maintain parsimony in the number of parameters, we have assumed they are independent.

The parameter ρ determines the persistence of the expected growth rate process. First, note that when $\varphi_e = 0$, the processes g_t and $g_{d,t+1}$ are *i.i.d.* Second, if $e_{t+1} = \eta_{t+1}$, the process for consumption is the ARMA(1,1) used in Bansal and Yaron (2000). Additionally, if $\varphi_e = \rho$, then consumption growth is an AR(1) process, as in Mehra and Prescott (1985).

Since g and g_d are exogenous processes, a solution for the log price–consumption ratio z_t and the log price–dividend ratio $z_{m,t}$ leads to a complete characterization of the returns $r_{a,t+1}$ and $r_{m,t+1}$ (using equation (2)). The relevant state variable for deriving the solution for z_t and $z_{m,t}$ is the expected growth rate of consumption x_t. Exploiting the Euler equation (1), the solution for the log price–consumption z_t has the form $z_t = A_0 + A_1 x_t$. An analogous expression holds for the log price–dividend ratio $z_{m,t}$. Details of both derivations are provided in the Appendix.

The solution coefficients for the effect of expected growth rate x_t on the price–consumption ratio, A_1, and the price–dividend ratio, $A_{1,m}$, respectively, are

$$
A_1 = \frac{1 - \dfrac{1}{\psi}}{1 - \kappa_1 \rho}, \quad A_{1,m} = \frac{\phi - \dfrac{1}{\psi}}{1 - \kappa_{1,m} \rho}.
\tag{5}
$$

It immediately follows that A_1 is positive if the IES, ψ, is greater than 1. In this case the intertemporal substitution effect dominates the wealth effect. In

---

[5] Similar growth rate dynamics (see equation (4)) are also considered in Campbell (1999), Cecchetti, Lam, and Mark (1993), and Wachter (2002) to model the consumption growth rate.

[6] The above specification models the growth rates of consumption (nondurables plus services) and dividends. Consequently, as in many other papers (e.g., Campbell and Cochrane (1999)), consumption and dividends are not cointegrated. It is an empirical issue if these series are cointegrated or not. Additionally, these growth-rate focused models also do not consider the implications for the ratio of dividends to consumption. It is possible that confronting the model specification for consumption and dividends with these additional issues may provide further insights regarding the appropriate time-series model for them—we leave this for future research.

response to higher expected growth (higher expected rates of return), agents buy more assets, and consequently the wealth-to-consumption ratio rises. In the standard power utility model, the need to have risk aversion larger than 1 also implies that $\psi < 1$, and hence A_1 is negative. Consequently, the wealth effect dominates the substitution effect.[7] In addition, note that $A_{1,m} > A_1$ when $\phi > 1$; consequently, expected growth rate news leads to a larger reaction in the price of the dividend claim than in the price of the consumption claim.

Substituting the equilibrium return for $r_{a,t+1}$ into the IMRS, it is straightforward to show that the innovation to the pricing kernel is (see equation (A10) in the Appendix)

$$
\begin{aligned}
m_{t+1} - E_t(m_{t+1}) &= \left[-\frac{\theta}{\psi} + \theta - 1 \right] \sigma \eta_{t+1} \\
&\quad - (1 - \theta) \left[\kappa_1 \left(1 - \frac{1}{\psi} \right) \frac{\varphi_e}{1 - \kappa_1 \rho} \right] \sigma e_{t+1} \\
&= \lambda_{m,\eta} \sigma \eta_{t+1} - \lambda_{m,e} \sigma e_{t+1}.
\end{aligned}
\tag{6}
$$

The expressions $\lambda_{m,e}$ and $\lambda_{m,\eta}$ capture the pricing kernel's exposure to the expected growth rate and the independent consumption shocks, η_{t+1}. The key observation is that the exposure to expected growth rate shocks $\lambda_{m,e}$ rises as the permanence parameter ρ rises. The conditional volatility of the pricing kernel is constant, as all risk sources have constant conditional variances.

As asset returns and the pricing kernel in this model economy are conditionally log-normal, the continuous risk premium on any asset i is $E_t[r_{i,t+1} - r_{f,t}] = -\text{cov}_t(m_{t+1}, r_{i,t+1}) - 0.5\sigma_{r_i,t}^2$. Given the solutions for A_1 and $A_{1,m}$, it is straightforward to derive the equity premium on the market portfolio (see Sec. A.4 in the Appendix),

$$
E(r_{m,t+1} - r_{f,t}) = \beta_{m,e} \lambda_{m,e} \sigma^2 - 0.5 \text{var}(r_{m,t}),
\tag{7}
$$

where $\beta_{m,e} \equiv [\kappa_{1,m}(\phi - \frac{1}{\psi}) \frac{\varphi_e}{1 - \kappa_{1,m}\rho}]$ and $\text{var}_t(r_{m,t+1}) = [\beta_{m,e}^2 + \varphi_d^2]\sigma^2$. The exposure of the market return to expected growth rate news is $\beta_{m,e}$, and the price of expected growth risk is determined by $\lambda_{m,e}$. The expressions for these parameters reveal that a rise in ρ increases both $\beta_{m,e}$ and $\lambda_{m,e}$. Consequently, the risk premium on the asset also increases with ρ. Similarly, the volatility of the market return also increases with ρ (see equation (A22) in the Appendix).

Because of our assumption of a constant σ, the conditional risk premium on the market portfolio in (7) is constant, and so is its conditional volatility. Hence, the ratio of the two, namely the Sharpe ratio, is also constant. In order to address issues that pertain to time-varying risk premia and predictability of risk premia, we augment our model in the next section and introduce time-varying economic uncertainty.

---

[7] An alternative interpretation with the power utility model is that higher expected growth rates increase the risk-free rate to an extent that discounting dominates the effects of higher expected growth rates. This leads to a fall in asset prices.

B. Case II: Incorporating Fluctuating Economic Uncertainty

We model fluctuating economic uncertainty as time-varying volatility of consumption growth. The dynamics for the system (4) that incorporate stochastic volatility are

$$x_{t+1} = \rho x_t + \varphi_e \sigma_t e_{t+1}$$
$$g_{t+1} = \mu + x_t + \sigma_t \eta_{t+1}$$
$$g_{d,t+1} = \mu_d + \phi x_t + \varphi_d \sigma_t u_{t+1} \tag{8}$$
$$\sigma_{t+1}^2 = \sigma^2 + v_1(\sigma_t^2 - \sigma^2) + \sigma_w w_{t+1}$$
$$e_{t+1}, u_{t+1}, \eta_{t+1}, w_{t+1} \sim N.i.i.d.(0, 1),$$

where σ_{t+1} represents the time-varying economic uncertainty incorporated in consumption growth rate and σ^2 is its unconditional mean. To maintain parsimony, we assume that the shocks are uncorrelated, and allow for only one source of economic uncertainty to affect consumption and dividends.

The relevant state variables in solving for the equilibrium price–consumption (and price–dividend) ratio are now x_t and σ_t^2. Thus, the approximate solution for the price–consumption ratio is $z_t = A_0 + A_1 x_t + A_2 \sigma_t^2$. The solution for A_1 is unchanged (equation (5)). The solution coefficient A_2 for measuring the sensitivity of price–consumption ratios to volatility fluctuations is

$$A_2 = \frac{0.5\left[\left(\theta - \dfrac{\theta}{\psi}\right)^2 + (\theta A_1 \kappa_1 \varphi_e)^2\right]}{\theta(1 - \kappa_1 v_1)}. \tag{9}$$

An analogous coefficient for the price–dividend ratio, $A_{2,m}$, is derived in the Appendix and has a similar form. Two features of this model specification are noteworthy. First, if the IES and risk aversion are larger than 1, then θ is negative, and a rise in volatility lowers the price–consumption ratio. Similarly, an increase in economic uncertainty raises risk premia and lowers the market price–dividend ratio. This highlights that an IES larger than 1 is critical for capturing the negative correlation between price–dividend ratios and consumption volatility. Second, an increase in the permanence of volatility shocks, that is v_1, magnifies the effects of volatility shocks on valuation ratios, as changes in economic uncertainty are perceived as being long-lasting.

As the price–consumption ratio is affected by volatility shocks, so is the return $r_{a,t+1}$. Consequently, the pricing kernel (IMRS) is also affected by volatility shocks. Specifically, the innovation in the pricing kernel is now:

$$m_{t+1} - E_t(m_{t+1}) = \lambda_{m,\eta} \sigma_t \eta_{t+1} - \lambda_{m,e} \sigma_t e_{t+1} - \lambda_{m,w} \sigma_w w_{t+1}, \tag{10}$$

where $\lambda_{m,w} \equiv (1 - \theta)A_2\kappa_1$, while $\lambda_{m,\eta}$ and $\lambda_{m,e}$ are defined in equation (6). This expression is similar to the earlier model (see equation (6)) save for the inclusion of w_{t+1}: shocks to consumption volatility. In the special case of power utility,

where $\theta = 1$, these volatility innovations are not reflected in the innovation of the pricing kernel, as $\lambda_{m,w}$ equals zero.[8]

The equation for the equity premium will now have two sources of systematic risk. The first, as before, relates to fluctuations in expected consumption growth, and the second to fluctuations in consumption volatility. The equity premium in the presence of time-varying economic uncertainty is

$$E_t(r_{m,t+1} - r_{f,t}) = \beta_{m,e}\lambda_{m,e}\sigma_t^2 + \beta_{m,w}\lambda_{m,w}\sigma_w^2 - 0.5\text{var}_t(r_{m,t+1}), \qquad (11)$$

where $\beta_{m,w} \equiv \kappa_{1,m}A_{2,m}$ and $\text{var}_t(r_{m,t+1}) = \{\beta_{m,e}^2\sigma_t^2 + \varphi_d^2\sigma_t^2 + \beta_{m,w}^2\sigma_w^2\}$.

The market compensation for stochastic volatility risk in consumption is determined by $\lambda_{m,w}$. The risk premium on the market portfolio is time-varying as σ_t fluctuates. The ratio of the conditional risk premium to the conditional volatility of the market portfolio fluctuates with σ_t, and hence the Sharpe ratio is time-varying. The maximal Sharpe ratio in this model economy, which approximately equals the conditional volatility of the pricing kernel innovation (equation (10)), also varies with σ_t.[9] This means that during periods of high economic uncertainty, risk premia will rise. For further discussion on the specialization of the risk premia under expected utility see Bansal and Yaron (2000).

The first-order effects on the level of the risk-free rate (see equation (A26) in the Appendix) are the rate of time preference and the average consumption growth rate, divided by the IES. Increasing the IES keeps the level low. In addition, the variance of the risk-free rate is primarily determined by the volatility of expected consumption growth rate and the IES. Increasing the IES lowers the volatility of the risk-free rate.

II. Data and Model Implications

To derive asset market implications from the model described in (8), we calibrate the model at the monthly frequency, such that its time-aggregated annual growth rates of consumption and dividends match salient features of observed annual data, and at the same time allow the model to reproduce many observed asset pricing features. Following Campbell and Cochrane (1999), Kandel and Stambaugh (1991), and many others, we assume that the decision interval of the agent is monthly but the targeted data to match are annual.[10]

Our choices of the time series and preference parameters are designed to simultaneously match observed growth rate data and asset market data. In

---

[8] Recall that in our specification the conditional volatility and expected growth rate processes are independent. With power utility, the volatility shocks will not be reflected in the innovations of the IMRS. With the Epstein and Zin (1989) preferences, in spite of this independence, volatility shocks influence the innovations in the IMRS.

[9] As in Campbell and Cochrane (1999), given the normality of the growth rate dynamics, the maximal Sharpe ratio is simply given by the standard deviation of the log pricing kernel.

[10] The evidence regarding the model is based on numerical solutions using standard polynomial-based projection methods discussed in Judd (1998). The numerical results are quite close to those based on the approximate analytical solutions.

order to isolate the economic effects of persistent expected growth rates from those of fluctuating economic uncertainty, we report our results first for Case I, where fluctuating economic uncertainty has been shut off (σ_w is set to zero), and then consider the model specification where both channels are operational.

A. Persistent Expected Growth

In Table I we display the time-series properties of the model given in (4). The specific parameters are given below the table. In spite of a persistent growth component, the model's implied time-series properties are largely consistent with the data.

Barsky and DeLong (1993) rely on a persistence parameter ρ equal to 1. We calibrate ρ at 0.979; this ensures that expected consumption growth rates are stationary and permits the possibility of large dividend elasticity of equity prices and equity risk premia. Our choice of φ_e and σ is motivated to ensure that we match the unconditional variance and the autocorrelation function of annual consumption growth. The standard deviation of the one-step ahead innovation in consumption, that is σ, equals 0.0078. This parameter configuration implies that the predictable variation in monthly consumption growth, that is, the R^2, is

Table I
Annualized Time-Averaged Growth Rates

The model parameters are based on the process given in equation (4). The parameters are $\mu = \mu_d = 0.0015$, $\rho = 0.979$, $\sigma = 0.0078$, $\phi = 3$, $\varphi_e = 0.044$, and $\varphi_d = 4.5$. The statistics for the data are based on annual observations from 1929 to 1998. Consumption is real nondurables and services (BEA); dividends are from the CRSP value-weighted return. The expression $AC(j)$ is the jth autocorrelation, $VR(j)$ is the jth variance ratio, and *corr* denotes the correlation. Standard errors are Newey and West (1987) corrected using 10 lags. The statistics for the model are based on 1,000 simulations each with 840 monthly observations that are time-aggregated to an annual frequency. The *mean* displays the mean across the simulations. The 95% and 5% columns display the estimated percentiles of the simulated distribution. The *p*-val column denotes the number of times in the simulation the parameter of interest was larger than the corresponding estimate in the data. The *Pop* column refers to population value.

	Data		Model				
Variable	Estimate	SE	Mean	95%	5%	*p*-Val	*Pop*
$\sigma(g)$	2.93	(0.69)	2.72	3.80	2.01	0.37	2.88
$AC(1)$	0.49	(0.14)	0.48	0.65	0.21	0.53	0.53
$AC(2)$	0.15	(0.22)	0.23	0.50	-0.17	0.70	0.27
$AC(5)$	-0.08	(0.10)	0.13	0.46	-0.13	0.93	0.09
$AC(10)$	0.05	(0.09)	0.01	0.32	-0.24	0.80	0.01
$VR(2)$	1.61	(0.34)	1.47	1.69	1.22	0.17	1.53
$VR(5)$	2.01	(1.23)	2.26	3.78	0.79	0.63	2.36
$VR(10)$	1.57	(2.07)	3.00	6.51	0.76	0.77	2.96
$\sigma(g_d)$	11.49	(1.98)	10.96	15.47	7.79	0.43	11.27
$AC(1)$	0.21	(0.13)	0.33	0.57	0.09	0.53	0.39
$corr(g, g_d)$	0.55	(0.34)	0.31	0.60	-0.03	0.07	0.35

only 4.4%. Our choice of ϕ is very similar to that in Abel (1999) and captures the "levered" nature of dividends. The standard deviation of the monthly innovation in dividends, $\varphi_d\sigma$, is 0.0351. This parameter configuration allows us to match the unconditional variance of dividend growth and its annual correlation with consumption.

Since our model emphasizes the long-horizon implications of the predictable component x_t, we first demonstrate that our proposed process for consumption is consistent with annual consumption data along a variety of dimensions. We use BEA data on real per-capita annual consumption growth of nondurables and services for the period 1929–1998. This is the longest single source of consumption data. Dividends and the value-weighted market return data are taken from CRSP. All nominal quantities are deflated using the CPI. To facilitate comparisons between the model, which is calibrated to a monthly decision interval, and the annual data, we time-aggregate our monthly model and report its annual statistics. As there is considerable evidence for small sample biases in estimating autoregression coefficients and variance ratios (see Hurwicz (1950), Ansley and Newbold (1980)), we report statistics based on 1,000 Monte Carlo experiments, each with 840 monthly observations—each experiment corresponding to the 70 annual observations available in our data set. Increasing the size of the Monte Carlo makes little difference in the results.

The annualized real per-capita consumption growth mean is 1.8% and its standard deviation is about 2.9%. Note that this volatility is somewhat lower for our sample than for the period considered in Mehra and Prescott (1985), Kandel and Stambaugh (1991), and Abel (1999). Table I shows that, in the data, consumption growth has a large first-order autocorrelation coefficient and a small second-order one. The standard errors in the data for these autocorrelations are sizeable. An alternative way to view the long-horizon properties of the model is to use variance ratios that are themselves determined by the autocorrelations (see Cochrane (1988)). In the data the variance ratios first rise significantly and at about 7 years out start to decline. The standard errors on these variance ratios, not surprisingly, are quite substantial.

The mean (across simulations) of the model's implied first-order autocorrelation is similar to that in the data. The second- and tenth-order autocorrelations are within one standard error of the data. The fifth-order autocorrelation is slightly above the two standard error range of the data. The empirical distribution of these estimates across the simulations as depicted by the 5[th] and 95[th] percentiles is wide and contains the point estimates from the data. The model's variance ratios mimic the pattern in the data. The point estimates are slightly larger than the data, but they are well within one standard error of the data. The point estimates from the data are clearly contained in the 5% confidence interval based on the empirical distribution of the simulated variance ratios. The unconditional volatility of consumption and dividend growth closely matches that in the data. In addition, the correlation of dividends with consumption of about 0.3 is somewhat lower, but within one standard error of its estimate in the data. This lower correlation is a conservative estimate, and increasing it helps the model generate a higher risk premium. Overall, Table I shows that

allowing for a persistent predictable component produces consumption and dividend moments that are largely consistent with the data.

It is often argued that consumption growth is close to being *i.i.d.* As shown in Table I, the consumption dynamics, which contain a persistent but small predictable component, are also largely consistent with the data. This evidence is consistent with Shephard and Harvey (1990), Barsky and DeLong (1993), and Bansal and Lundblad (2002), who show that in finite samples, discrimination across the *i.i.d.* growth rate model and the one considered above is extremely difficult. While the financial market data are hard to interpret from the perspective of the *i.i.d.* dynamics, they are, as shown below, interpretable from the perspective of the growth rate dynamics considered above.

Before we discuss the asset pricing implications we highlight two additional issues related to the data. First, data for consumption, dividends, and asset returns pertain to the long sample from 1929. Clearly, moments of these data will differ across subsamples. Our choice of the long sample is similar to Mehra and Prescott (1985), Kandel and Stambaugh (1991), and Abel (1999) and is motivated to keep the estimation error on the moments small. The annual autocorrelations of consumption growth for our model are well within standard error bounds, even when compared to those in the post-war annual consumption data.[11] Second, our dividend model is calibrated to cash dividends; this is similar to that used by many earlier studies. While it is common to use cash dividends, this measure of dividends may mismeasure total payouts, as it ignores other forms of payments made by corporations. Given the difficulties in accurately measuring total payouts of corporations and to maintain comparability with earlier work, we have focused on cash dividends as well. Jagannathan, McGrattan, and Scherbina (2000) provide evidence pertaining to the issue of dividends, and show that alternative measures of dividends have even higher volatility.

A.1. Case I: Asset Pricing Implications

In Table II, we display the asset pricing implications of the model for a variety of risk aversion and IES configurations. In Panel A, we use the time-series parameters from Table I. In Panel B we increase ϕ, the dividend leverage parameter, to 3.5, and in Panel C we analyze the implications of an *i.i.d.* process. The table intentionally concentrates on a relatively narrow set of asset pricing moments, namely the mean risk-free rate, equity premium, the market and risk-free rate volatility, and the volatility of the log price–dividend ratio. These moments are the main focus of many asset pricing models. In Section II.C, we discuss additional model implications.

---

[11] The first-order autocorrelations for annual consumption growth in 1951–1999 and 1961–1999 are 0.38 and 0.44, respectively—hence the consumption growth autocorrelations vary with samples. Based on Table I, both estimates are well within the model-based 5% confidence interval for the first-order autocorrelation. We have focused on annual data (consumption and dividends) to avoid dealing with seasonalities and other measurement problems discussed in Wilcox (1992).

Table II
Asset Pricing Implications—Case I

This table provides information regarding the model without fluctuating economic uncertainty (i.e., Case I, where $\sigma_w = 0$). All entries are based on $\delta = 0.998$. In Panel A the parameter configuration follows that in Table I, that is, $\mu = \mu_d = 0.0015$, $\rho = 0.979$, $\sigma = 0.0078$, $\phi = 3$, $\varphi_e = 0.044$, and $\varphi_d = 4.5$. Panels B and C describe the changes in the relevant parameters. The expressions $E(R_m - R_f)$ and $E(R_f)$ are, respectively, the annualized equity premium and mean risk-free rate. The expressions $\sigma(R_m)$, $\sigma(R_f)$, and $\sigma(p - d)$ are the annualized volatilities of the market return, risk-free rate, and the log price-dividend, respectively.

γ	ψ	$E(R_m - R_f)$	$E(R_f)$	$\sigma(R_m)$	$\sigma(R_f)$	$\sigma(p - d)$
\multicolumn{7}{c}{Panel A: $\phi = 3.0, \rho = 0.979$}						
7.5	0.5	0.55	4.80	13.11	1.17	0.07
7.5	1.5	2.71	1.61	16.21	0.39	0.16
10.0	0.5	1.19	4.89	13.11	1.17	0.07
10.0	1.5	4.20	1.34	16.21	0.39	0.16
\multicolumn{7}{c}{Panel B: $\phi = 3.5, \rho = 0.979$}						
7.5	0.5	1.11	4.80	14.17	1.17	0.10
7.5	1.5	3.29	1.61	18.23	0.39	0.19
10.0	0.5	2.07	4.89	14.17	1.17	0.10
10.0	1.5	5.10	1.34	18.23	0.39	0.19
\multicolumn{7}{c}{Panel C: $\phi = 3.0, \rho = \varphi_e = 0$}						
7.5	0.5	−0.74	4.02	12.15	0.00	0.00
7.5	1.5	−0.74	1.93	12.15	0.00	0.00
10.0	0.5	−0.74	3.75	12.15	0.00	0.00
10.0	1.5	−0.74	1.78	12.15	0.00	0.00

Our choice of parameters attempts to take economic considerations into account. In particular $\delta < 1$, and the risk-aversion parameter γ is either 7.5 or 10. Mehra and Prescott (1985) argue that a reasonable upper bound for risk aversion is around 10. In this sense, our choice for risk aversion is reasonable. The magnitude for the IES that we focus on is 1.5. Hansen and Singleton (1982) and Attanasio and Weber (1989) estimate the IES to be well in excess of 1.5. More recently, Vissing-Jorgensen (2002) and Guvenen (2001) also argue that the IES is well over 1. However, Hall (1988) and Campbell (1999) estimate the IES to be well below 1. Their results are based on a model without fluctuating economic uncertainty. In Section II.C.4, we show that ignoring the effects of time-varying consumption volatility leads to a serious downward bias in the estimates of the IES. To highlight the role of the IES, we choose one value of the IES less than 1 (IES = 0.5) and another larger than 1 (IES = 1.5).

Table II shows that the model with persistent expected growth is able to generate sizeable risk premia, market volatility, and fluctuations in price–dividend ratios. Larger risk aversion clearly increases the equity premium; changing risk aversion mainly affects this dimension of the model. To qualitatively match key

features of the data, it is important for the IES to be larger than 1. Lowering the IES lowers $A_{1,m}$, the dividend elasticity of asset prices, and the risk premia on the asset. As the IES rises, the volatility of the price–dividend ratio and asset returns rise along with $A_{1,m}$. At very low values of the IES, $A_{1,m}$ can become negative, which would imply that a rise in dividends' growth rate expectations will lower asset prices (see the discussion in Sec. I). In addition, note that if the leverage parameter ϕ is increased, it increases the riskiness of dividends, and $A_{1,m}$ rises. The price–dividend ratio becomes more volatile, and the equity premium rises.

As discussed earlier we assumed that u_t, e_t, and η_t are independent. To give a sense of how the results change if we allow for correlations in the various shocks, consider the case with the IES at 1.5 and a risk aversion of 10. When we assume that the correlation between u_t and η_t is 0.25 and all other innovations are set at zero, then the equity premium rises to 5.02%. If the correlation between u_t and e_t is assumed to be 0.25, then the equity premium and the market return volatility rise to 5.21% and 17.22% respectively. There are virtually no other changes. As stated earlier, in Table II, we have made the conservative assumption of zero correlations to maintain parsimony in the parameters that we have to calibrate.

It is also interesting to consider the case where consumption and dividend growth rates are assumed to be *i.i.d.*, that is, $\varphi_e = 0$. In this case, the equity premium for the market is $E_t(r_{m,t+1} - r_{f,t}) = \gamma \operatorname{cov}(g_{t+1}, g_{d,t+1}) - 0.5 \operatorname{var}(r_{m,t+1})$. In our baseline model, dividend innovations are independent of consumption innovations; hence, with *i.i.d.* growth rates, $\operatorname{cov}(g_{t+1}, g_{d,t+1})$ equals zero, and the market equity premium is $-0.5 \operatorname{var}(r_{m,t+1})$; this explains the negative equity premium in the *i.i.d.* case reported in Panel C of Table II. If we assume that the correlation between monthly consumption and dividend growth is 0.25, then the equity premium is 0.08% per annum. This is similar to the evidence documented in Weil (1989) and Mehra and Prescott (1985). For comparable IES and risk-aversion values, shifting from the persistent growth rate process to *i.i.d.* growth rates lowers the volatility of the equity returns. In all, this evidence highlights the fact that although the time-series dynamics of the model with small persistent expected growth are difficult to distinguish from a pure *i.i.d.* model, its asset pricing implications are vastly different from those in the *i.i.d.* model. In what follows we use the parameters in Panel A, with an IES of 1.5 as our preferred configuration, and display the implications of adding fluctuating economic uncertainty.

B. Fluctuating Economic Uncertainty

Before displaying the asset pricing implications of adding fluctuating economic uncertainty, we first briefly discuss evidence for the presence of fluctuating economic uncertainty.

Panel A of Table III documents that the variance ratios of the absolute value of residuals from regressing current consumption growth on five lags increase gradually out to 10 years. This suggests slow-moving predictable variation in

Table III
Properties of Consumption Volatility

The entries in Panel A are the variance ratios ($VR(j)$) for $|\epsilon_{g^a,t}|$, which is the absolute value of the residual from the regression $g_t^a = \sum_{j=1}^{5} A_j g_{t-j}^a + \epsilon_{g^a,t}$, where g_t^a denotes annual consumption growth rate. Panel B provides regression results for $|\epsilon_{g^a,t+j}| = \alpha + B(j)(p_t - d_t) + v_{t+j}$, and j indicates the forecast horizon in years. The statistics are based on annual observations from 1929 to 1998 of real nondurables and services consumption (BEA). The price–dividend ratio is based on the CRSP value-weighted return. Standard errors are Newey and West (1987) corrected using 10 lags.

	Panel A: Variance Ratios		Panel B: Predicting $|\epsilon_{g^a,t+j}|$		
Horizon	$VR(j)$	SE	$B(j)$	SE	R^2
2	0.95	(0.38)	−0.11	(0.04)	0.06
5	1.26	(1.09)	−0.10	(0.05)	0.04
10	1.75	(2.46)	−0.08	(0.08)	0.03

this measure of realized volatility. Note that if realized volatility were *i.i.d.*, these variance ratios would be flat.[12]

In Panel B of Table III we provide evidence that future realized consumption volatility is predicted by current price–dividend ratios. The current price–dividend ratio predicts future realized volatility with negative coefficients, with robust *t*-statistics around 2 and R^2s around 5% (for horizons of up to 5 years). If consumption volatility were not time-varying, the slope coefficient on the price-dividend ratio would be zero. As suggested by our theoretical model, this evidence indicates that information regarding persistent fluctuations in economic uncertainty is contained in asset prices. Overall, the evidence in Table III lends support to the view that the conditional volatility of consumption is time-varying. Bansal, et al. (2002) extensively document the evidence in favor of time-varying consumption volatility and show that this feature holds up quite well across different samples and economies.

Given the evidence above, a large value of ν_1, the parameter governing the persistence of conditional volatility, allows the model to capture the slow-moving fluctuations in economic uncertainty. In Table IV we provide the asset pricing implications based on the system (8), when in addition to the parameters given in Table I, we activate the volatility parameters (given below the table). It is important to note that the time-series properties displayed in Table I are virtually unaltered once we introduce the fluctuations in economic uncertainty.

Table IV provides statistics for the asset market data and for the model that incorporates fluctuating economic uncertainty (i.e., Case II). Columns 2 and 3 provide the statistics and their respective standard errors for our data sample. Columns 4 and 5 provide the model's corresponding statistics for risk aversion

---

[12] Also note that it is difficult to detect high-frequency time-varying volatility (e.g., GARCH) effects once the data are time-aggregated (see Nelson (1991), Drost and Nijman (1993)).

Table IV
Asset Pricing Implications—Case II

The entries are model population values of asset prices. The model incorporates fluctuating economic uncertainty (i.e., Case II) using the process in equation (8). In addition to the parameter values given in Panel A of Table II ($\delta = 0.998$, $\mu = \mu_d = 0.0015$, $\rho = 0.979$, $\sigma = 0.0078$, $\phi = 3$, $\varphi_e = 0.044$, and $\varphi_d = 4.5$), the parameters of the stochastic volatility process are $\nu_1 = 0.987$ and $\sigma_w = 0.23 \times 10^{-5}$. The predictable variation of realized volatility is 5.5%. The expressions $E(R_m - R_f)$ and $E(R_f)$ are, respectively, the annualized equity premium and mean risk-free rate. The expressions $\sigma(R_m)$, $\sigma(R_f)$, and $\sigma(p - d)$ are the annualized volatilities of the market return, risk-free rate, and the log price-dividend, respectively. The expressions $AC1$ and $AC2$ denote, respectively, the first and second autocorrelation. Standard errors are Newey and West (1987) corrected using 10 lags.

	Data		Model	
Variable	Estimate	SE	$\gamma = 7.5$	$\gamma = 10$
Returns				
$E(r_m - r_f)$	6.33	(2.15)	4.01	6.84
$E(r_f)$	0.86	(0.42)	1.44	0.93
$\sigma(r_m)$	19.42	(3.07)	17.81	18.65
$\sigma(r_f)$	0.97	(0.28)	0.44	0.57
Price Dividend				
$E(\exp(p - d))$	26.56	(2.53)	25.02	19.98
$\sigma(p - d)$	0.29	(0.04)	0.18	0.21
$AC1(p - d)$	0.81	(0.09)	0.80	0.82
$AC2(p - d)$	0.64	(0.15)	0.65	0.67

of 7.5 and 10, respectively. In this table the IES is always set at 1.5 and ϕ is set at 3.

Column 5 of Table IV shows that with $\gamma = 10$, the model generates an equity premium that is comparable to that in the data.[13] The mean of the risk-free rate, and the volatilities of the market return and of the risk-free rate, are by and large consistent with the data. The model essentially duplicates the volatility and persistence of the observed log price–dividend ratio. Comparing columns 4 and 5 provides sensitivity of the results to the level of risk aversion. Not surprisingly, higher risk aversion increases the equity premium and aligns the model closer to the data. A comparison of Table IV with Table II shows that when risk aversion is 10, the equity risk premium is about 2.5% higher—this additional premium reflects the premium associated with fluctuating economic uncertainty as derived in equation (11). One could, as discussed earlier, modify the above model and also include correlation between the different shocks. The inclusion of these correlations as documented above typically helps to

---

[13] To derive analytical expressions we have assumed that the volatility process is conditionally normal. When we solve the model numerically we ensure that the volatility is positive by replacing negative realizations with a very small number. This happens for about 5% of the realizations; hence, the possibility that volatility in equation (8) can become negative is primarily a technical issue.

increase the equity premium. Hence, it would seem that these correlations would help the model generate the same equity premium with a lower risk-aversion parameter.

Weil (1989) and Kandel and Stambaugh (1991) also explore the implications of the Epstein and Zin (1989) preferences for asset market data. However, these papers find it difficult to quantitatively explain the aforementioned asset market features at our configuration of preference parameters. Why, then, do we succeed in capturing these asset market features with Epstein and Zin preferences? Weil uses *i.i.d.* consumption growth rates. As discussed earlier, with *i.i.d.* consumption and dividend growth rates, the risks associated with fluctuating expected growth and economic uncertainty are absent. Consequently, the model has great difficulty in explaining the asset market data.

Kandel and Stambaugh (1991) consider a model in which there is predictable variation in consumption growth rates and volatility. However, at our preference parameters, the persistence in the expected growth and conditional volatility in their specification is not large enough to permit significant response of asset prices to news regarding expected consumption growth and volatility. In addition, Kandel and Stambaugh primarily focus on the case in which the IES is close to zero. At very low values of the IES, $\lambda_{m,e}$ and $\beta_{m,e}$ are negative (see equations (6) and (7)). This may still imply a sizeable equity premium. However, a parameter configuration with an IES less than 1 and a moderate level of risk aversion (e.g., 10 or less) leads to high levels of the risk-free rate and/or its volatility. In contrast, our IES, which is greater than 1, ensures that the level and volatility of the risk-free rate are low and comparable to those in the data. Hence, with moderate levels of risk aversion, both the high persistence and an IES greater than 1 are important in order to capture key aspects of asset market data.

C. Additional Asset Pricing Implications

As noted earlier, in the model where we shut off fluctuating economic uncertainty (Case I), both risk premia and Sharpe ratios are constant—hence, this simple specification cannot address issues regarding predictability of risk premia. The model that incorporates fluctuating economic uncertainty (Case II) does permit risk premia to fluctuate. Henceforth, we focus entirely on this model specification with the parameter configuration stated in Table IV with $\gamma = 10$.

C.1. Variability of the Pricing Kernel

The maximal Sharpe ratio, as shown in Hansen and Jagannathan (1991), is determined by the conditional volatility of the pricing kernel. This maximal Sharpe ratio for our model is the volatility of the pricing kernel innovation defined in equation (10). In Table V, we quantify the contributions of different shocks to the variance of the pricing kernel innovations (see equation (10)). The maximal annualized Sharpe ratio for our model economy is 0.73, which is

Table V
Decomposing the Variance of the Pricing Kernel

Entries are the relative variance of different shocks to the variance of the pricing kernel. The entries are based on the model configuration described in Table IV with $\gamma = 10$. The volatility of the maximal Sharpe ratio is annualized in order to make it comparable to the Sharpe ratio on annualized returns.

	Relative Variance of Shocks		
Volatility of Pricing Kernel	Independent Consumption	Expected Growth Rate	Fluctuating Economic Uncertainty
0.73	14%	47%	39%

quite large. The maximal Sharpe ratio with *i.i.d.* growth rates is $\gamma\sigma$, and with our parameter configuration its annualized value equals 0.27. Consequently, the Epstein and Zin preferences and the departure from *i.i.d.* growth rates are responsible for this larger maximal Sharpe ratio. Additionally, for our model, the maximal Sharpe ratio exceeds that of the market return, which is 0.33. The sources of risk in order of importance are shocks to the expected growth rate (i.e., e_{t+1}), followed by that of fluctuating economic uncertainty (i.e., w_{t+1}). While the variance of these shocks in themselves is small, their effects on the pricing kernel get magnified because of the long-lasting nature of these shocks (see discussion in Sec. I). Finally, the variance of high-frequency consumption news, η_{t+1}, is relatively large, but this risk source contributes little to the pricing kernel variability, as this shock is not long-lasting.

C.2. Predictability of Returns, Growth Rates, and Price–Dividend Ratios

Dividend yields seem to predict multi-horizon returns. A rise in the current dividend yield predicts a rise in future expected returns. Our model performs quite well in capturing this feature of the data. However, it is important to recognize that these predictability results are quite sensitive to changing samples, estimation techniques, and data sets (see Hodrick (1992), and Goyal and Welch (1999)). Further, most dimensions of the evidence related to predictability (be it growth rates or returns) are estimated with considerable sampling error. This, in conjunction with the rather high persistence in the price–dividend ratio, suggests that considerable caution should be exercised in interpreting the evidence regarding predictability based on price–dividend ratios.

In Panel A of Table VI, we report the predictability regressions of future excess returns for horizons of 1, 3, and 5 years for our sample data. In Column 4 we report the corresponding evidence from the perspective of the model. The model captures the positive relationship between expected returns and dividend yields. The absolute value of the slope coefficients and the corresponding R^2s rise with the return horizon, as in the data. The predictive slope coefficients and the R^2s in the model are somewhat lower than those in the data; however,

Table VI
Predictability of Returns, Growth Rates, and Price–Dividend Ratios

This table provides evidence on predictability of future excess returns and growth rates by price–dividend ratios, and the predictability of price–dividend ratios by consumption volatility. The entries in Panel A correspond to regressing $r_{t+1}^e + r_{t+2}^e \cdots + \cdots r_{t+j}^e = \alpha(j) + B(j)\log(P_t/D_t) + v_{t+j}$, where r_{t+1}^e is the excess return, and j denotes the forecast horizon in years. The entries in Panel B correspond to regressing $g_{t+1}^a + g_{t+2}^a \cdots + \cdots g_{t+j}^a = \alpha(j) + B(j)\log(P_t/D_t) + v_{t+j}$, and g^a is annualized consumption growth. The entries in Panel C correspond to $\log(P_{t+j}/D_{t+j}) = \alpha(j) + B(j)|\epsilon_{g^a,t}| + v_{t+j}$, where $|\epsilon_{g^a,t}|$ is the volatility of consumption defined as the absolute value of the residual from regressing $g_t^a = \sum_{j=1}^{5} A_j g_{t-j}^a + \epsilon_{g^a,t}$. The model is based on the process in equation (8), with parameter configuration given in Table IV and $\gamma = 10$. The entries for the model are based on 1,000 simulations each with 840 monthly observations that are time-aggregated to an annual frequency. Standard errors are Newey and West (1987) corrected using 10 lags.

Variable	Panel A: Excess Returns			Panel B: Growth Rates			Panel C: Volatility		
	Data	SE	Model	Data	SE	Model	Data	SE	Model
$B(1)$	−0.08	(0.07)	−0.18	0.04	(0.03)	0.06	−8.78	(3.58)	−3.74
$B(3)$	−0.37	(0.16)	−0.47	0.03	(0.05)	0.12	−8.32	(2.81)	−2.54
$B(5)$	−0.66	(0.21)	−0.66	0.02	(0.04)	0.15	−8.65	(2.67)	−1.56
$R^2(1)$	0.02	(0.04)	0.05	0.13	(0.09)	0.10	0.12	(0.05)	0.14
$R^2(3)$	0.19	(0.13)	0.10	0.02	(0.05)	0.12	0.11	(0.04)	0.08
$R^2(5)$	0.37	(0.15)	0.16	0.01	(0.02)	0.11	0.12	(0.04)	0.05

the model's slope coefficients are within two standard errors of the estimated coefficients in the data.[14]

In Panel B of Table VI, we provide regression results where the dependent variable is the sum of annual consumption growth rates. In the data it seems that price–dividend ratios have little predictive power, particularly at longer horizons. The slope coefficients and R^2s of these regressions are quite low both in the data and the model. The R^2s are relatively small in the model for two reasons. First, price–dividend ratios are determined by expected growth rates, and the variation in expected growth rates is quite small. Recall that the monthly R^2 for consumption dynamics is less than 5%. Second, price–dividend ratios are also affected by independent movements in economic uncertainty, which lowers their ability to predict future growth rates. Overall, the model, like the data, suggests that growth rates at long horizons are not predicted by price–dividend ratios in any economically sizeable manner.[15]

---

[14] Consistent with Lettau and Ludvigson (2001), predictability coefficients and R^2s based on the wealth–consumption ratio follow the same pattern and are slightly larger than those based on price–dividend ratios.

[15] Our model can be easily modified to further lower the predictability of growth rates. Consider an augmented model (as in Cecchetti et al. (1993)) that allows for additional predictable movements in dividend growth rates that are unrelated to consumption. This will not affect the risk-free rate and the risk premia in the model, but will additionally lower the ability of price–dividend ratios to predict future consumption growth rates.

In Panel C of Table VI, we report how well current realized consumption volatility predicts future price–dividend ratios. First, note that there is strong evidence in the data for this relationship. The regression coefficients for predicting future price–dividend ratios with current volatility for 1, 3, and 5 years are all negative, have robust t-statistics that are well above 2, and have R^2s of about 10%. The model produces similar negative coefficients, albeit in absolute terms they are slightly smaller. The R^2s are within two standard errors of the data. Taken together with the results in Panel B of Table III, the evidence is consistent with the economics of the model; fluctuating economic uncertainty, captured via realized consumption volatility, predicts future price–dividend ratios and is predicted by lagged price–dividend ratios. The empirical evidence shows that asset markets dislike economic uncertainty—a feature that our model is capable of reproducing. Using alternative measures of consumption volatility, Bansal et al. (2002) show that this evidence is robust across many samples and frequencies, and is consistently found in many developed economies.

Some caution should be exercised in interpreting the links between dividend growth rates and price–dividend ratios. Evidence from other papers (see Ang and Bekaert (2001) and Bansal et al. (2002)) indicates that alternative measures of cash flows, such as earnings, are well predicted by valuation ratios. Cash dividends, as discussed earlier, may not accurately measure the total payouts to equity holders and hence may distort the link between growth rates and asset valuations. However, given the practical difficulties in measuring the appropriate payouts, and to maintain comparability with other papers in the literature, we, like others, continue to use cash dividends. With this caveat in mind, we also explore the model's implications by exploring how much of the variation in the price–dividend ratio is from growth rates and what part is due to variation in expected returns.

In the data, the majority of the variation in price–dividend ratios seems to be due to variation in expected returns. For our sample the point estimate for the percentage of the variation in price–dividend ratio due to return fluctuations is 108%, with a standard error of 42%, while dividends' growth rates account for -6%, with a standard error of 31%.[16] Our model produces population estimates that attribute about 52 percent of the variation in price–dividend ratios to returns and 54% to fluctuations in expected dividend growth. Note that the standard errors of the point estimates of this decomposition in the data are very large. To account for any finite sample biases, we also conducted a Monte Carlo exercise using simulations from our model of sample sizes comparable to our data. This Monte Carlo evidence implies that in our model, the returns account for about 70% of the variation in price–dividend ratio, thus aligning the model closer to the data. Given the large sampling variation in measuring these

---

[16] For explicit details of this decomposition, see Cochrane (1992). Specifically, these represent the percentage of $\text{var}(p-d)$ accounted for by returns and dividend growth rates: $100 * \sum_{j=1}^{15} \Omega^j \frac{\text{cov}(p_t - d_t, x_{t+j})}{\text{var}(p_t - d_t)}$, where $x = -r$ and g_d respectively, and $\Omega = 1/(1 + E(r))$.

quantities in the data using cash dividends, and the sharp differences in predictability implications using alternative cash flow measures, makes economic inference based on this decomposition quite difficult.

Two additional features of the model are worth highlighting. First, in the data the contemporaneous correlation between equity return and consumption is very small at the monthly frequency and is about 0.20 at the annual frequency. Our model produces comparable magnitudes, with correlations of 0.04 and 0.15 for the monthly and annual frequencies, respectively. Second, the term premium on nominal bonds, the average one-period excess return on an n-period discount bond, is small. This suggests that the equity premium in the data is not driven by a large term premium. The term premium (which in our model is on real bonds) is in fact small and slightly negative. Hence the large equity premium in the model is not a by-product of a large positive term premium.[17] In totality, the above evidence, in conjunction with the results pertaining to predictability, suggest that the model is capable of capturing several key aspects of asset markets data.

C.3. Conditional Volatility and the Feedback Effect

A large literature documents that market return volatility is very persistent (see, e.g., Bollerslev, Engle, and Wooldridge (1988)). This feature of the data is easily reproduced in our model. The market volatility process, as described in equation (A13) in the Appendix, is a linear affine function of the conditional variance of the consumption growth rate process σ_t. As the conditional variance of the consumption growth rate process is an AR(1) process, it follows that the market volatility inherits this property. Note that the coefficient on the conditional variance of consumption in the market volatility process is quite large. This magnifies the conditional variance of the market portfolio relative to consumption volatility. The persistence in market volatility coincides with the persistence in the consumption volatility process. In the monthly market return data, this persistence parameter is about 0.986 (see Bollerslev et al. (1998)), and in the model it equals ν_1, 0.987. As consumption volatility is high during recessions, this implies that the market volatility also rises during recessions. Also note that during periods of high consumption volatility (e.g., recessions), in the model the equity premium also rises. This implication of the model is consistent with the evidence provided in Fama and French (1989) that risk premia are countercyclical.

Campbell and Hentschel (1992), Glosten, Jagannathan, and Runkle (1993), and others document what is known as the volatility feedback effect. That is, return innovations are negatively correlated with innovations in market

---

[17] The explicit formulas for the real term structure and the term premia are presented in Bansal and Yaron (2000). The negative real term premia of our model are consistent with the evidence provided in Evans (1998), who documents that for inflation-indexed bonds in the United Kingdom (1983–1995) the term premia are significantly negative (less than -2% at the 1-year horizon), while the term premia for nominal bonds are very slightly positive.

volatility. The model is capable of reproducing this negative correlation. The feedback effect arises within the model in spite of the fact that the volatility innovations are independent of the expected consumption growth process. The key feature that allows the model to capture this dimension is the Epstein–Zin preferences in which volatility risk is priced (see the discussion in Sec. I.B). Using the analytical expressions for the innovation in the market return (see equation (A12) in the Appendix) and the expression for the innovation in the market volatility, it is straightforward to show that the conditional covariance

$$\text{cov}_t((r_{m,t+1} - E_t r_{m,t+1}), \text{var}_{t+1}(r_{m,t+2}) - E_t[\text{var}_{t+1}(r_{m,t+2})])$$
$$= \beta_{m,w}(\beta_{m,e}^2 + \varphi_d^2)\sigma_w^2, \tag{12}$$

where $\beta_{m,w} \equiv \kappa_1 A_{2,m} < 0$ as $A_{2,m}$ is negative. The correlation between market return innovations and market volatility innovations for our model is -0.32.

An additional issue pertains to the relation between the expected return on the market portfolio and the market volatility. Glosten et al. (1993) and Whitelaw (1994) document that the expected market return and the market volatility are negatively related. French, Schwert, and Stambaugh (1987) and Campbell and Hentschel (1992) argue that this relation is likely to be positive. In our model, theoretically, the relation between expected market return and market volatility is positive, and is not consistent with the negative relation between expected returns and market volatility. Whitelaw (2000) shows that a standard power utility model with regime shifts in consumption growth can accommodate the negative relation between expected returns and market volatility. The unconditional correlation in our model between ex post excess returns on the market and the ex ante market volatility is a small positive number, 0.04. The model cannot generate the negative relation between expected returns and market volatility. To do so, we conjecture, will require significant changes, perhaps along the lines pursued in Whitelaw. This departure is well outside the scope of this paper, and we leave this exploration for future work.

C.4. Bias in Estimating the Intertemporal Elasticity of Substitution

As in Hall (1988), the IES is typically measured by the slope coefficient from regressing date $t + 1$ consumption growth rate on the date t risk-free rate. This projection would indeed recover the IES, if no fluctuating uncertainty affected the risk-free rate. However, the risk-free rate in our model fluctuates as a result of both changing expected growth rate and independent fluctuations in the volatility of consumption. Thus, the above projection is misspecified and creates a downward bias. This bias is quite significant, as inside our model, where the value of the IES is set at 1.5, Hall's regression would estimate the IES parameter to be 0.62. Our model is a simple one, and there may be alternative instrumental variable approaches to undo this bias. However, we view this result of the downward bias as suggestive of the difficulties in accurately pinning down the IES. As discussed in Section II.A, several papers report an estimated IES that is well over 1. This evidence, along with the potential

downward bias in estimating the IES, makes our choice of an IES larger than 1 quite reasonable.

III. Conclusions

In this paper, we explore the idea that news about growth rates and economic uncertainty (i.e., consumption volatility) alters perceptions regarding long-term expected growth rates and economic uncertainty and that this channel is important for explaining various asset market phenomena. If indeed news about consumption has a nontrivial impact on long-term expected growth rates or economic uncertainty, then asset prices will be fairly sensitive to small growth rate and consumption volatility news. We develop a model for growth rates that captures this intuition. Anderson, Hansen, and Sargent (2002) utilize features of our growth rate dynamics to motivate economic models that incorporate robust control with respect to the small long-run components in growth rates.

We provide empirical support for aggregate consumption and dividend growth processes that contain a small persistent expected growth rate component and a conditional volatility component. These growth rate dynamics, in conjunction with the Epstein and Zin (1989) and Weil (1989) preferences, can help explain many asset market puzzles. In our model, at plausible values for the preference parameters, a reduction in economic uncertainty or better long-run growth prospects leads to a rise in the wealth–consumption and the price–dividend ratios.

The model is capable of justifying the observed magnitudes of the equity premium, the risk-free rate, and the volatility of the market return, dividend-yield, and the risk-free rate. Further, it captures the volatility feedback effect, that is, the negative correlation between return news and return volatility news. As in the data, dividend yields predict future returns and the volatility of returns is time-varying. Evidence provided in this paper and Bansal et al. (2002) shows that there is a significant negative correlation between price–dividend ratios and consumption volatility. The model captures this dimension of the data as well. A feature of the model is that about half of the variability in equity prices is due to fluctuations in expected growth rates, and the remainder is due to fluctuations in the cost of capital.

Appendix

The consumption and dividend processes given in (8) are

$$g_{t+1} = \mu + x_t + \sigma_t \eta_{t+1}$$
$$x_{t+1} = \rho x_t + \varphi_e \sigma_t e_{t+1}$$
$$\sigma_{t+1}^2 = \sigma^2 + v_1(\sigma_t^2 - \sigma^2) + \sigma_w w_{t+1} \qquad (A1)$$
$$g_{d,t+1} = \mu_d + \phi x_t + \varphi_d \sigma_t u_{t+1}$$
$$w_{t+1}, e_{t+1}, u_{t+1}, \eta_{t+1} \sim N.i.i.d.(0, 1).$$

The IMRS (Intertemporal Marginal Rate of Substitution) for this economy is given by

$$\ln M_{t+1} = \theta \ln \delta - \frac{\theta}{\psi} g_{t+1} + (\theta - 1)r_{a,t+1}. \tag{A2}$$

We derive asset prices using this IMRS and the standard asset pricing condition $E_t[M_{t+1}R_{i,t+1}] = 1$, so that

$$E_t\left[\exp\left(\theta \ln \delta - \frac{\theta}{\psi} g_{t+1} + (\theta - 1)r_{a,t+1} + r_{i,t+1} \right) \right] = 1 \tag{A3}$$

for any asset $r_{i,t+1} \equiv \log(R_{i,t+1})$. We first start by solving the special case where $r_{i,t+1}$ is $r_{a,t+1}$—the return on the aggregate consumption claim, and then solve for market return $r_{m,t+1}$, and the risk-free rate r_f.

A. The Return on the Consumption Claim Asset, $r_{a,t+1}$

We conjecture that the log price–consumption ratio follows, $z_t = A_0 + A_1 x_t + A_2\sigma_t^2$. Armed with the endogenous variable z_t, we substitute the approximation $r_{a,t+1} = \kappa_0 + \kappa_1 z_{t+1} - z_t + g_{t+1}$ into the Euler equation (A3).

Since g, x, and σ_t^2 are conditionally normal, $r_{a,t+1}$ and $\ln M_{t+1}$ are also normal. Exploiting the normality of $r_{a,t+1}$ and $\ln M_{t+1}$, we can write down the Euler equation (A3) in terms of the state variables x_t and σ_t. As the Euler condition must hold for all values of the state variables, it follows that all terms involving x_t must satisfy the following:

$$-\frac{\theta}{\psi}x_t + \theta[\kappa_1 A_1 \rho x_t - A_1 x_t + x_t] = 0. \tag{A4}$$

It immediately follows that

$$A_1 = \frac{1 - \dfrac{1}{\psi}}{1 - \kappa_1 \rho}, \tag{A5}$$

which is (5) in the main text. Similarly, collecting all the σ_t^2 terms leads to the solution for A_2,

$$\theta\left[\kappa_1 \nu_1 A_2 \sigma_t^2 - A_2 \sigma_t^2\right] + \frac{1}{2}\left[\left(\theta - \frac{\theta}{\psi}\right)^2 + (\theta A_1 \kappa_1 \varphi_e)^2\right]\sigma_t^2 = 0, \tag{A6}$$

which implies that

$$A_2 = \frac{0.5\left[\left(\theta - \dfrac{\theta}{\psi}\right)^2 + (\theta A_1 \kappa_1 \varphi_e)^2\right]}{\theta(1 - \kappa_1 \nu_1)}, \tag{A7}$$

the solution given in (9).

Given the solution above for z_t, it is possible to derive the innovation to the return r_a as a function of the evolution of the state variables and the parameters of the model:

$$r_{a,t+1} - E_t(r_{a,t+1}) = \sigma_t \eta_{t+1} + B\sigma_t e_{t+1} + A_2 \kappa_1 \sigma_w w_{t+1}, \tag{A8}$$

where $B = \kappa_1 A_1 \varphi_e = \kappa_1 \frac{\varphi_e}{1-\kappa_1\rho}(1 - \frac{1}{\psi})$. Further, it follows that the conditional variance of $r_{a,t+1}$ is

$$\text{var}_t(r_{a,t+1}) = (1 + B^2)\sigma_t^2 + (A_2 \kappa_1)^2 \sigma_w^2. \tag{A9}$$

A.1. Intertemporal Marginal Rate of Substitution

Substituting for $r_{a,t+1}$ and the dynamics of g_{t+1}, we can rewrite the IMRS in terms of the state variables—referring to this as the pricing kernel. Suppressing all the constants in the pricing kernel,

$$m_{t+1} \equiv \ln M_{t+1} = \theta \ln \delta - \frac{\theta}{\psi} g_{t+1} + (\theta - 1)r_{a,t+1}$$

$$E_t[m_{t+1}] = m_0 - \frac{x_t}{\psi} + A_2(\kappa_1 \nu_1 - 1)(\theta - 1)\sigma_t^2$$

$$m_{t+1} - E_t(m_{t+1}) = \left(-\frac{\theta}{\psi} + \theta - 1 \right)\sigma_t \eta_{t+1} + (\theta - 1)(A_1 \kappa_1 \varphi_e)\sigma_t e_{t+1}$$

$$+ (\theta - 1)A_2 \kappa_1 \sigma_w w_{t+1}$$

$$= \lambda_{m,\eta}\sigma_t \eta_{t+1} - \lambda_{m,e}\sigma_t e_{t+1} - \lambda_{m,w}\sigma_w w_{t+1}, \tag{A10}$$

where $\lambda_{m,\eta} \equiv [-\frac{\theta}{\psi} + (\theta - 1)] = -\gamma$, $\lambda_{m,e} \equiv (1 - \theta)B$, $\lambda_{m,w} \equiv (1 - \theta)A_2\kappa_1$, and B and A_2 are defined above. Note that the λ's represent the market price of risk for each source of risk, namely η_{t+1}, e_{t+1}, and w_{t+1}.

A.2. Risk Premia for $r_{a,t+1}$

The risk premium for any asset is determined by the conditional covariance between the return and m_{t+1}. Thus, the risk premium for $r_{a,t+1}$ is equal to $E_t(r_{a,t+1} - r_{f,t}) = -\text{cov}_t[m_{t+1} - E_t(m_{t+1}), r_{a,t+1} - E_t(r_{a,t+1})] - 0.5\text{var}_t(r_{a,t+1})$. Exploiting the innovations in (A8) and (A10), it follows that

$$E_t[r_{a,t+1} - r_{f,t}] = -\lambda_{m,\eta}\sigma_t^2 + \lambda_{m,e}B\sigma_t^2 + \kappa_1 A_2 \lambda_{m,w}\sigma_w^2 - 0.5\text{var}_t(r_{a,t+1}), \tag{A11}$$

where $\text{var}_t(r_{a,t+1})$ is defined in equation (A9).

A.3. Equity Premium and Market Return Volatility

The risk premium for any asset is determined by the conditional covariance between the return and m_{t+1}. Thus the risk premium for the market portfolio

$r_{m,t+1}$ is equal to $E_t(r_{m,t+1} - r_{f,t}) = -\text{cov}_t[m_{t+1} - E_t(m_{t+1}), r_{m,t+1} - E_t(r_{m,t+1})] - 0.5\text{var}_t(r_{m,t+1})$.

Equation (A10) already provides the innovation in m_{t+1}. We now proceed to derive the innovation in the market return. The price–dividend ratio for the claim on dividends is $z_{m,t} = A_{0,m} + A_{1,m}x_t + A_{2,m}\sigma_t^2$. It follows that

$$r_{m,t+1} = g_{d,t+1} + \kappa_1 A_{1,m}x_{t+1} - A_{1,m}x_t + \kappa_1 A_{2,m}\sigma_{t+1}^2 - A_{2,m}\sigma_t^2$$

$$r_{m,t+1} - E_t(r_{m,t+1}) = \varphi_d \sigma_t u_{t+1} + \kappa_1 A_{1,m}\varphi_e \sigma_t e_{t+1} + \kappa_1 A_{2,m}\sigma_w w_{t+1}$$

$$= \varphi_d \sigma_t u_{t+1} + \beta_{m,e}\sigma_t e_{t+1} + \beta_{m,w}\sigma_w w_{t+1}, \tag{A12}$$

where $\beta_{m,e} \equiv \kappa_{1,m}A_{1,m}\varphi_e$, and $\beta_{m,w} \equiv \kappa_1 A_{2,m}$. Moreover, it follows that

$$\text{var}_t(r_{m,t+1}) = \left(\beta_{m,e}^2 + \varphi_d^2\right)\sigma_t^2 + \beta_{m,w}^2\sigma_w^2. \tag{A13}$$

Using the innovations in the market return and the pricing kernel, the expression for the equity premium is

$$E_t(r_{m,t+1} - r_{f,t}) = \beta_{m,e}\lambda_{m,e}\sigma_t^2 + \beta_{m,w}\lambda_{m,w}\sigma_w^2 - 0.5\text{var}_t(r_{m,t+1}), \tag{A14}$$

where $\text{var}_t(r_{m,t+1})$ is defined in equation (A13).

To derive the expressions for $A_{1,m}$ and $A_{2,m}$, we exploit the Euler condition $E_t[\exp(m_{t+1} + r_{m,t+1})] = 1$. Collecting all the x_t terms, we find that

$$-\frac{x}{\psi} + x\kappa_{1,m}A_{1,m}\rho - A_{1,m}x + \phi x = 0, \tag{A15}$$

which implies that

$$A_{1,m} = \frac{\phi - \dfrac{1}{\psi}}{1 - \kappa_{1,m}\rho}. \tag{A16}$$

The solution for $A_{2,m}$ follows from exploiting the asset pricing condition,

$$\exp\{E_t(m_{t+1}) + E_t(r_{m,t+1}) + 0.5\text{var}_t(m_{t+1} + r_{m,t+1})\} = 1, \tag{A17}$$

and collecting all σ_t terms. Note that $\text{var}_t(m_{t+1} + r_{m,t+1})$ equals

$$\text{var}_t[\lambda_{m,\eta}\sigma_t \eta_{t+1} - \lambda_{m,w}\sigma_w w_{t+1} - \lambda_{m,e}\sigma_t e_{t+1} + \beta_{m,e}\sigma_t e_{t+1}$$

$$+ \varphi_d \sigma_t u_{t+1} + \beta_{m,w}\sigma_w w_{t+1}]$$

$$= H_m \sigma_t^2 + [-\lambda_{m,w} + \beta_{m,w}]^2 \sigma_w^2, \tag{A18}$$

where $H_m \equiv [\lambda_{m,\eta}^2 + (-\lambda_{m,e} + \beta_{m,e})^2 + \varphi_d^2]$. Now collect all the σ_t^2 terms in equation (A17), and note that σ_t appears in $E_t(r_{m,t+1})$ as well as $E_t(m_{t+1})$. This leads to the following restriction,

$$(\theta - 1)A_2(\kappa_1 \nu_1 - 1) + A_{2,m}(\kappa_{1,m}\nu_1 - 1) + \frac{H_m}{2} = 0, \tag{A19}$$

which implies that

$$A_{2,m} = \frac{(1-\theta)A_2(1-\kappa_1 v_1) + 0.5H_m}{(1-\kappa_{1,m}v_1)}. \tag{A20}$$

To derive the unconditional variance of the market return, note that

$$r_{m,t+1} - E(r_{m,t+1}) = -\frac{x_t}{\psi} + \beta_{m,e}\sigma_t e_{t+1} + \varphi_d \sigma_t u_{t+1}$$

$$+ A_{2,m}(v_1\kappa_1 - 1)[\sigma_t^2 - E(\sigma_t^2)] + \beta_{m,w}\sigma_w w_{t+1}. \tag{A21}$$

Hence, the unconditional variance is

$$\text{var}(r_m) = \frac{\sigma_x^2}{\psi^2} + [\beta_{m,e}^2 + \varphi_d^2]\sigma^2 + [A_{2,m}(v_1\kappa_1 - 1)]^2 \text{var}(\sigma_t^2) + \beta_{m,w}^2\sigma_w^2. \tag{A22}$$

The unconditional variance of $z_{m,t}$, the price–dividend ratio for the market portfolio, can be derived as follows:

$$\text{var}(z_{m,t}) = A_{1,m}^2 \text{var}(x_t) + A_{2,m}^2 \text{var}(\sigma_t^2). \tag{A23}$$

Finally, note that the innovation to the market return volatility follows from equation (A12) and is

$$\text{var}_{t+1}(r_{m,t+2}) - E_t[\text{var}_{t+1}(r_{m,t+2})] = (\beta_{m,e}^2 + \varphi_d^2)\sigma_w w_{t+1}. \tag{A24}$$

B. The Risk-Free Rate and Its Volatility

To derive the risk-free rate, start with (A3) and plug in the risk-free rate for r_i:

$$r_{f,t} = -\theta \log(\delta) + \frac{\theta}{\psi}E_t[g_{t+1}] + (1-\theta)E_t r_{a,t+1}$$

$$- \frac{1}{2}\text{var}_t\left[\frac{\theta}{\psi}g_{t+1} + (1-\theta)r_{a,t+1}\right], \tag{A25}$$

subtract $(1-\theta)r_{f,t}$ from both sides and divide by θ, where it is assumed that $\theta \neq 0$. It then follows that

$$r_{f,t} = -\log(\delta) + \frac{1}{\psi}E_t[g_{t+1}] + \frac{(1-\theta)}{\theta}E_t[r_{a,t+1} - r_{f,t}]$$

$$- \frac{1}{2\theta}\text{var}_t\left[\frac{\theta}{\psi}g_{t+1} + (1-\theta)r_{a,t+1}\right]. \tag{A26}$$

Further, to solve the above expression, note that $\text{var}_t[\frac{\theta}{\psi}g_{t+1} + (1-\theta)r_{a,t+1}] \equiv \text{var}_t(m_{t+1})$, and therefore,

$$\text{var}_t(m_{t+1}) = (\lambda_{m,\eta}^2 + \lambda_{m,e}^2)\sigma_t^2 + \lambda_{m,w}^2\sigma_w^2. \tag{A27}$$

The unconditional mean of $r_{f,t}$ is derived by substituting the expression for the risk premium for $r_{a,t+1}$ given in (A11) and (A27) into (A26). This substitution yields

$$E(r_{f,t}) = -\log(\delta) + \frac{1}{\psi}E(g) + \frac{(1-\theta)}{\theta}E[r_{a,t+1} - r_{f,t}]$$

$$- \frac{1}{2\theta}\big[(\lambda_{m,\eta}^2 + \lambda_{m,e}^2)E[\sigma_t^2] + \lambda_{m,w}^2\sigma_w^2\big], \qquad \text{(A28)}$$

since $E[\sigma_t^2] = \text{var}(\eta)$.

The unconditional variance of $r_{f,t}$ is:

$$\text{var}(r_{f,t}) = \left(\frac{1}{\psi}\right)^2 \text{var}(x_t) + \left\{\frac{1-\theta}{\theta}Q_1 - Q_2\frac{1}{2\theta}\right\}^2 \text{var}(\sigma_t^2), \qquad \text{(A29)}$$

where $Q_2 = (\lambda_{m,\eta}^2 + \lambda_{m,e}^2)$, and $Q_1 = (-\lambda_{m,\eta} + (1-\theta)B^2 - 0.5(1+B^2))$, where B is defined above. Note that Q_1 determines the time-varying portion of the risk premium on $r_{a,t+1}$. For all practical purposes, the variance of the risk-free rate is driven by the first term.

REFERENCES

Abel, Andrew B., 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review* 80, 38–42.

Abel, Andrew B., 1999, Risk premia and term premia in general equilibrium, *Journal of Monetary Economics* 43, 3–33.

Anderson, Evan, Lars P. Hansen, and Thomas Sargent, 2002, A quartet of semi-groups for model specification, detection, robustness, and the price of risk, Unpublished manuscript, University of Chicago.

Ang, Andrew, and Geert Bekaert, 2001, Stock return predictability: Is it there? Unpublished manuscript, Columbia University.

Ansley, Craig F., and Paul Newbold, 1980, Finite sample properties of estimators for autoregressive moving average models, *Journal of Econometrics* 13, 159–184.

Attanasio, P. Orazio, and Guglielmo Weber, 1989, Intertemporal substitution, risk aversion and the Euler equation for consumption, *Economic Journal* 99, 59–73.

Bansal, Ravi, and Wilbur J. Coleman, II, 1997, A monetary explanation of the equity premium, term premium and the risk-free rate puzzles, *Journal of Political Economy* 104, 1135–1171.

Bansal, Ravi, Varoujan Khatchatrian, and Amir Yaron, 2002, Interpretable asset markets? NBER Working paper 9383.

Bansal, Ravi, and Christian Lundblad, 2002, Market efficiency, asset returns, and the size of the risk premium in global equity markets, *Journal of Econometrics* 109, 195–237.

Bansal, Ravi, and Amir Yaron, 2000, Risks for the long run: A potential resolution of asset pricing puzzles, NBER Working paper 8059.

Barberis, Nicholas, Ming Huang, and Tano Santos, 2001, Prospect theory and asset pricing, *Quarterly Journal of Economics* 116, 1–53.

Barsky, Robert, and Bradford J. DeLong, 1993, Why does the stock market fluctuate? *Quarterly Journal of Economics* 108, 291–312.

Bollerslev, Timothy, Robert F. Engle, and Jeffrey Wooldridge, 1988, A capital asset pricing model with time-varying covariances, *Journal of Political Economy* 96, 116–131.

Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.

Campbell, John Y., 1999, Asset prices, consumption and the business cycle, in John B. Taylor, and Michael Woodford, eds.: *Handbook of Macroeconomics*, vol. 1 (Elsevier Science, North-Holland, Amsterdam).

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Campbell, John Y., and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility is stock returns, *Journal of Financial Economics* 31, 281–318.

Campbell, John Y., and Robert J. Shiller, 1988, The dividend–price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–227.

Cecchetti, Stephen G., Pok-Sang Lam, and Nelson C. Mark, 1990, Mean reversion in equilibrium asset prices, *American Economic Review* 80, 398–419.

Cecchetti, Stephen G., Pok-Sang Lam, and Nelson C. Mark, 1993, The equity premium and the risk free rate: Matching the moments, *Journal of Monetary Economics* 31, 21–46.

Chapman, David, 2002, Does intrinsic habit formation actually resolve the equity premium puzzle? *Review of Economic Dynamics* 5, 618–645.

Cochrane, John H., 1988, How big is the random walk in GNP? *Journal of Political Economy* 96, 893–920.

Cochrane, John H., 1992, Explaining the variance of price dividend ratios, *Review of Financial Studies* 5, 243–280.

Constantinides, George, 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519–543.

Constantinides, George, and Darrell Duffie, 1996, Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219–240.

Drost, Feike C., and Theo E. Nijman, 1993, Temporal aggregation of GARCH processes, *Econometrica* 61, 909–927.

Epstein, Lawrence, and Stanley Zin, 1989, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Evans, Martin, 1998, Real rates, expected inflation and inflation risk premia, *Journal of Finance* 53, 187–218.

Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.

French, Kenneth R., William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

Glosten, Lawrence, Ravi Jagannathan, and David Runkle, 1993, On the relation between expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.

Goyal, Amit, and Ivo Welch, 1999, The myth of predictability: Does the dividend yield forecast the equity premium? Unpublished manuscript, UCLA.

Guvenen, Fatih, 2001, Mismeasurement of the elasticity of intertemporal substitution: The role of limited stock market participation, Unpublished manuscript, University of Rochester.

Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 339–357.

Hansen, Lars P., and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.

Hansen, Lars P., Thomas Sargent, and Thomas Tallarini, 1999, Robust permanent income and pricing, *Review of Economic Studies* 66, 873–907.

Hansen, Lars P., and Kenneth Singleton, 1982, Generalized instrumental variables estimation of nonlinear rational expectations models, *Econometrica* 50, 1269–1286.

Heaton, John, 1995, An empirical investigation of asset pricing with temporally dependent preference specifications, *Econometrica* 63, 681–717.

Heaton, John, and Deborah Lucas, 1996, Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 443–487.

Hodrick, Robert J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–386.

Hurwicz, Leonid, 1950, Least square bias in time series, in T. Koopmans, ed.: *Statistical Inference in Dynamic Economic Models* (Wiley, New York).

Jagannathan, Ravi, Ellen R. McGrattan, and Anna Scherbina, 2000, The declining U.S. equity premium, *Federal Reserve Bank of Minneapolis, Quarterly Review* 24, 3–19.

Judd, Kenneth, 1998, *Numerical Methods in Economics* (MIT Press, Cambridge, MA).

Kandel, Shmuel, and Robert F. Stambaugh, 1991, Asset returns and intertemporal preferences, *Journal of Monetary Economics* 27, 39–71.

LeRoy, Stephen, and Richard Porter, 1981, The present value relation: Tests based on implied variance bounds, *Econometrica* 49, 555–574.

Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Nelson, Daniel, 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347–370.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Shephard, Neil G., and Andrew C. Harvey, 1990, On the probability of estimating a deterministic component in the local level model, *Journal of Time Series Analysis* 11, 339–347.

Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71, 421–436.

Vissing-Jorgensen, Annette, 2002, Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* 110, 825–853.

Wachter, Jessica, 2002, Habit formation and returns on stocks and bonds, Unpublished manuscript, NYU.

Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 401–421.

Whitelaw, Robert F., 1994, Time variations and covariations in the expectation and volatility of stock market returns, *Journal of Finance* 49, 515–541.

Whitelaw, Robert F., 2000, Stock market risk and return: An equilibrium approach, *Review of Financial Studies* 13, 521–547.

Wilcox, David, 1992, The construction of the U.S. consumption data: Some facts and their implications for empirical work, *American Economic Review* 82, 922–941.