

ECON 2021 - FINANCIAL ECONOMICS I

Lecture 9 – Recursive Preferences & Long-Run Risk

November 12, 2018

MOTIVATION

- Observed risk premia and return volatility require volatile SDFs. In traditional models, $SDF = M_t \sim e^{-\delta t} U'(C_t) = e^{-\delta t} C_t^{-\gamma}$.
- When discussing the CCAPM we showed that a volatile SDF could be generated by making C_t volatile. This required a particular sort of idiosyncratic consumption/labor income risk. Last week we generated a volatile SDF by introducing state variables that index an agent's **beliefs**, which fluctuate due to learning.
- Today we explore another possibility. We abstract from incomplete markets and learning, and instead relax the assumption of time-separability, so marginal utility, $U'(C)$, depends on things besides current consumption. This allows $U'(C)$ to be much more volatile than C . With additive preferences, the only way to achieve this is to make risk aversion large.
- There are 2 basic strategies for doing this:
 - 1 **Habit Persistence**. MU depends on past consumption.
 - 2 **Recursive Preferences**. MU depends on expected future consumption.
- Today we study the 2nd strategy, recursive preferences.

DRAWBACKS OF TIME-ADDITIVE PREFS

- The assumption of time-additive preferences has no compelling theoretical justification. Its widespread use is based **solely** on mathematical convenience.
- Unfortunately, this convenience comes at a cost. Time-additive/Expected Utility preferences impose several restrictions:

- 1 It implies that the Elasticity of Intertemporal Substitution (EIS) is the reciprocal of the Coefficient of Relative Risk Aversion (RRA). Consider the following 2-period example

$$U(C_0, C_1) = \frac{1}{1-\gamma} C_0^{1-\gamma} + \delta E_0 \left[\frac{1}{1-\gamma} \tilde{C}_1^{1-\gamma} \right]$$

We already know $RRA = -CU''/U' = \gamma$. The EIS is just the percentage change in consumption growth in response to a 1 percent change in the (riskless) rate of return:

$$EIS = \frac{d \ln(C_1/C_0)}{d \ln(1+r)} = \frac{d \ln(C_1/C_0)}{d \ln(U'(C_0)/U'(C_1))} = \frac{1}{\gamma}$$

This is an arbitrary restriction. Why should the willingness to substitute across **dates** be related to the willingness to substitute across **states**? This restriction is the fundamental cause of the Equity Premium/Risk-Free Rate Puzzle. We need a high γ to explain the mean return on risky assets, but then this implies a very low EIS, which generates a very high risk-free rate.

- 2 Time-additive/Expected Utility preferences imply an indifference to the timing of the resolution of uncertainty. Consider the following 2 lotteries
- (I) At time 0 a single coin is tossed. If it's Heads, then $C_t = H$ for all t . If it's Tails, then $C_t = T$ for all t .
 - (II) At each date t a coin will be tossed. If it's Heads then $C_t = H$. If Tails, then $C_t = T$.

An agent with time-additive/expected-utility preferences will be indifferent to these 2 lotteries, since the expected utility of both is

$$\sum_{t=1}^{\infty} \delta^t \left[\frac{1}{2} U(H) + \frac{1}{2} U(T) \right]$$

Is this reasonable? Maybe. But maybe not. In the first lottery, uncertainty is resolved earlier. A basic axiom of expected utility theory is that agents reduce **compound lotteries**. This has always been a questionable assumption, but it is especially doubtful in a dynamic setting. A preference for early/late resolution reflects an unwillingness to reduce (intertemporal) compound lotteries. The recursive preferences we study impose expected utility w.r.t. to atemporal lotteries, but relax the expected utility axioms w.r.t. intertemporal lotteries.

- 3 As extended by Savage (1954) to the case of subjective probabilities, expected utility admits no distinction between risk and uncertainty. Many have argued (e.g., Knight & Keynes) that this distinction is especially important in financial markets. Next week we shall see that recursive preferences can be reinterpreted in a way that allows this distinction to be operationalized.

RECURSIVE PREFERENCES


- Despite these drawbacks, time-additive CRRA utility does have a couple of desirable features, which we want to retain:
 - ① **Dynamic Consistency**: This allows us to use DP to characterize optimal behavior.
 - ② **Scale Invariance**: CRRA utility has the attractive property that interest rates and risk premia are stationary in the presence of growth.
- Kreps & Porteus (1978) show that dynamic consistency can be preserved while relaxing time-additivity. They define current utility recursively, using two distinct functions:

$$u_t = W(C_t, \mathcal{R}_t(u_{t+1}))$$

The **Certainty Equivalent** operator, $\mathcal{R}_t(u_{t+1})$, translates random future utility into consumption units. The **Time Aggregator**, $W(\cdot)$, combines current consumption and the certainty equivalent of future utility into a measure of current utility. The Certainty Equivalent operator captures risk aversion, while the Time Aggregator captures intertemporal substitution.

- KP show that attitudes toward the timing of the resolution of uncertainty are captured by the curvature of W w.r.t. to its second argument. If it is convex ($W_{22} > 0$) then the agent prefers early resolution of uncertainty. If it is concave, then the agent prefers late resolution. This follows from Jensen's inequality.
- Of course, additive preferences are recursive too, and are a special case, i.e, when the aggregator and certainty equivalent are linear, $W(x, y) = U(x) + \delta y$ and $\mathcal{R}_t(u_{t+1}) = E_t(u_{t+1})$

VISUAL SUMMARY



Let $\psi = \rho^{-1}$ denote the EIS. The above figure shows the sense in which time additive/power utility is special, i.e., it is a 1-dimensional manifold in a 2-dimensional parameter space, defined by the hyperbola $\gamma\psi = 1$. EZ preferences allow us to consider separately the special cases of myopic (horizon independent) portfolio choice ($\gamma = 1$) and a constant Consumption/Wealth ratio ($\psi = 1$). With power utility, if you want one you must accept the other, which implies log utility. Note that along the power utility manifold $\gamma\psi = 1$ the agent is indifferent to the resolution of uncertainty. Above it he prefers early resolution.

THE SDF WITH EZ PREFERENCES

- Given our interest in asset pricing, we want to know what EZ preferences imply about the SDF.
- EZ show that it is a **combination** of the CCAPM SDF, based on consumption growth, and the classical Sharpe-Lintner SDF, based on the return on the market portfolio.
- With EZ prefs we can write the Euler equation as

$$W_1(t) = W_2(t)E_t[W_1(t+1)R_{t+1}]$$

where subscripts denote partial derivatives of the time aggregator.

- In words, giving up a unit of consumption today costs you $W_1(t)$ utils. At an optimum, this should equal the expected utility value of the future payoff, $E_t[W_1(t+1)R_{t+1}]$, expressed in units of time- t utility, $W_2(t)E_t[W_1(t+1)R_{t+1}]$.
- Hence we can write

$$\frac{M_{t+1}}{M_t} = W_2(t) \frac{W_1(t+1)}{W_1(t)} = \frac{\frac{\partial U_t}{\partial U_{t+1}} \cdot \frac{\partial U_{t+1}}{\partial C_{t+1}}}{\frac{\partial U_t}{\partial C_t}}$$

- From the EZ aggregator we get

$$\frac{\partial U_t}{\partial C_t} = (1 - \delta) C_t^{-\rho} U_t^\rho$$

$$\frac{\partial U_t}{\partial U_{t+1}} = \delta U_t^\rho [\mathcal{R}_t(U_{t+1})]^{-\rho} U_{t+1}^{-\gamma} [\mathcal{R}_t(U_{t+1})]^\gamma$$

- Hence, we can write the EZ SDF as follows

$$\frac{M_{t+1}}{M_t} = \frac{(1 - \delta) C_{t+1}^{-\rho} U_{t+1}^\rho \delta U_t^\rho [\mathcal{R}_t(U_{t+1})]^{-\rho} U_{t+1}^{-\gamma} [\mathcal{R}_t(U_{t+1})]^\gamma}{(1 - \delta) C_t^{-\rho} U_t^\rho}$$

$$= \delta \left(\frac{C_{t+1}}{C_t} \right)^{-\rho} \left[\frac{U_{t+1}}{\mathcal{R}_t(U_{t+1})} \right]^{\rho - \gamma} \quad \text{EZ SDF}$$

- As usual, the SDF increases in response to contemporaneous reductions in consumption growth. However, notice that when $\gamma > \rho$ it also increases in response to negative innovations in anticipated **future** utility. This injects an additional source of volatility into the SDF, which helps the model fit the data.
- Also note that when $\gamma = \rho$ the SDF collapses to the standard time-additive SDF.

- Remember that a key advantage of the complete markets Lucas/Breeden CCAPM is that it is based on **observable** data (albeit of dubious quality). In contrast, the ICAPM model of Merton (1973) required observable proxies for the marginal value of wealth.
- In a sense, the EZ model combines the CCAPM and ICAPM models. When $\rho \neq \gamma$ we must find observable proxies for innovations in continuation utility, $\mathbf{u}_{t+1}/\mathcal{R}_t(\mathbf{u}_{t+1})$.

There are 2 approaches:

- 1 Relate $\mathbf{u}_{t+1}/\mathcal{R}_t(\mathbf{u}_{t+1})$ to the return on the 'market portfolio', defined as an asset that pays aggregate consumption as its dividend. This is the approach followed by EZ.
 - 2 Use a homoskedastic/log-linear approximation to express the return on the market portfolio in terms of revisions in expected future aggregate consumption growth. This has been the approach followed more recently in the long-run risks literature.
- For both strategies we need to relate returns on the market portfolio to continuation utility. To do this, note that the EZ recursion is homogeneous of degree 1 in current consumption and continuation utility. Hence, if we let $MC_t \equiv \partial \mathbf{u}_t / \partial C_t$ and $MU_{t+1} = \partial \mathbf{u}_t / \partial \mathbf{u}_{t+1}$ we can use Euler's Theorem to write:

$$\mathbf{u}_t = MC_t \cdot C_t + E_t[MU_{t+1} \cdot \mathbf{u}_{t+1}]$$

which we can rewrite as

$$\begin{aligned} \frac{\mathbf{u}_t}{MC_t} &= C_t + E_t \left[\frac{MU_{t+1} MC_{t+1}}{MC_t} \cdot \frac{\mathbf{u}_{t+1}}{MC_{t+1}} \right] \\ &= C_t + E_t \left[\frac{M_{t+1}}{M_t} \cdot \frac{\mathbf{u}_{t+1}}{MC_{t+1}} \right] \end{aligned}$$

- Now notice that \mathbf{u}_t/MC_t can be interpreted as an agent's wealth. To see this, remember from our previous discussion of the Martingale Method that optimally managed wealth can be viewed as an asset that pays consumption as its dividend. If we let $J_t = \text{wealth} = \mathbf{u}_t/MC_t$, that's exactly what the above equation says,

$$J_t = C_t + E_t \left[\frac{M_{t+1}}{M_t} \cdot J_{t+1} \right]$$

where (as always) we use the SDF to discount future payoffs.

- By definition, the return on this asset is

$$R_{w,t+1} = \frac{J_{t+1}}{J_t - C_t}$$

Luckily, we've already computed MC_t , so we can easily compute $J_t = \mathbf{u}_t/MC_t$

$$J_t = \frac{1}{1 - \delta} C_t^\rho (\mathbf{u}_t)^{1-\rho}$$

Therefore, using the EZ recursion, we can write

$$J_t - C_t = \left(\frac{C_t}{1 - \delta} \right) \left[\left(\frac{\mathbf{u}_t}{C_t} \right)^{1-\rho} - (1 - \delta) \right] = \left(\frac{C_t}{1 - \delta} \right) \delta \left[\frac{\mathcal{R}_t(\mathbf{u}_{t+1})}{C_t} \right]^{1-\rho}$$

- Hence, we get the following expression for the return on the market portfolio

$$R_{w,t+1} = \frac{1}{\delta} \left(\frac{C_{t+1}}{C_t} \right)^\rho \left[\frac{U_{t+1}}{\mathcal{R}_t(U_{t+1})} \right]^{1-\rho}$$

- This means that if we can observe the return on the market portfolio, we can infer the innovation in continuation utility (assuming we can also observe consumption growth). In particular, we can use this expression to substitute out for $U_{t+1}/\mathcal{R}_t(U_{t+1})$ in our previous expression for the SDF

$$\frac{M_{t+1}}{M_t} = \delta^\theta \left(\frac{C_{t+1}}{C_t} \right)^{-\rho\theta} R_{w,t+1}^{\theta-1} \quad \theta = \frac{1-\gamma}{1-\rho}$$

- Notice when $\rho = \gamma$, then $\theta = 1$, and we get back the time-additive SDF.
- Also note that when $\gamma > \rho$ and $\rho < 1$ we get $\theta < 1$, and so declines in the market portfolio return produce increases in the SDF. This increases the risk premium, and therefore this has been the parameterization most often used in empirical work. (Although this parameterization mutes role of consumption growth, this is offset by the fact that observed proxies for $R_{w,t+1}$ are considerably more volatile than aggregate consumption growth).

- Of course, the major drawback of this approach is that we **don't** observe $R_{w,t+1}$. EZ used the return on an equity index as a proxy, but this is a rather narrow definition (e.g., in principle, R_w should include the returns to housing and human capital).
- In response, in recent years empirical work has been based on homoskedastic log-linear approximations, which allow observed consumption growth to replace the unobserved market portfolio return. To see this, first let $\tilde{x}_{t+1} = x_{t+1} - E_t x_{t+1}$ denote the innovation in some variable x_{t+1} . Taking logs of the above expression for the SDF,

$$\tilde{m}_{t+1} = -\theta \rho \tilde{c}_{t+1} - (1 - \theta) \tilde{r}_{w,t+1}$$

- Iterating forward the standard Campbell-Shiller linearization of the return innovation

$$\begin{aligned} \tilde{r}_{w,t+1} &= \Delta E_{t+1} \sum_{j=0}^{\infty} \beta^j \Delta c_{t+1+j} - \Delta E_{t+1} \sum_{j=1}^{\infty} \beta^j r_{w,t+1+j} \\ &= \tilde{c}_{t+1} + (1 - \rho) \Delta E_{t+1} \sum_{j=1}^{\infty} \beta^j \Delta c_{t+1+j} \end{aligned}$$

where $\Delta E_{t+1} = E_{t+1} - E_t$ represents innovations in expectations, and $\beta \equiv [1 + \exp(\overline{c - w})]^{-1}$ is a linearization constant. The second line uses the fact that the risk premium is constant with homoskedasticity, so that $r_{w,t} = \text{constant} + \rho E_t \Delta c_{t+1}$

- If we sub this into the above expression for \tilde{m}_{t+1} we get

$$\tilde{m}_{t+1} = -\gamma \tilde{c}_{t+1} - (\gamma - \rho) \Delta E_{t+1} \sum_{j=1}^{\infty} \beta^j \Delta c_{t+1+j}$$

EZ PREFERENCES IN CONTINUOUS-TIME

- As usual, moving to the continuous-time limit confers analytical advantages. The continuous-time version of EZ preferences is known as **Stochastic Differential Utility** (Duffie & Epstein (1992)).
- Since the Certainty Equivalent, $\mathcal{R}_t(\mathbf{U}_{t+1}) = [E_t(\mathbf{U}_{t+1})^{1-\gamma}]^{1/(1-\gamma)}$, is nonlinear, Ito's lemma implies that in the continuous-time limit a variance term will appear. DE show that this variance term can be avoided by using the ordinally equivalent utility index $\mathbf{V}_t = \frac{1}{1-\gamma} \mathbf{U}_t^{1-\gamma}$. The resulting aggregator is then called a **Normalized Aggregator**. (Next time, when discussing robust control, the unnormalized aggregator will appear).
- In terms of \mathbf{V}_t the EZ recursion takes the form

$$\mathbf{V}_t = \frac{1}{1-\gamma} \left[(1 - e^{-\delta}) C_t^{1-\rho} + e^{-\delta} ((1-\gamma) E_t \mathbf{V}_{t+1})^{\frac{1-\rho}{1-\gamma}} \right]^{\frac{1-\gamma}{1-\rho}}$$

- Instead of a unit time step, consider an arbitrarily small time step, ε . Solving for $E_t[\mathbf{V}_{t+\varepsilon}]$ gives

$$E_t[\mathbf{V}_{t+\varepsilon}] = \frac{1}{1-\gamma} \left[e^{\delta\varepsilon} ((1-\gamma) \mathbf{V}_t)^{\frac{1-\rho}{1-\gamma}} - (e^{\delta\varepsilon} - 1) C_t^{1-\rho} \right]^{\frac{1-\gamma}{1-\rho}}$$

- Now compute the drift of V by taking limits,

$$\begin{aligned}
 \lim_{\varepsilon \rightarrow 0} \frac{E_t[V_{t+\varepsilon}] - V_t}{\varepsilon} &= \lim_{\varepsilon \rightarrow 0} \frac{\left[e^{\delta\varepsilon} ((1-\gamma)V_t)^{\frac{1-\rho}{1-\gamma}} - (e^{\delta\varepsilon} - 1) C_t^{1-\rho} \right]^{\frac{1-\gamma}{1-\rho}} - (1-\gamma)V_t}{\varepsilon(1-\gamma)} \\
 &= \frac{\delta}{1-\rho} [(1-\gamma)V_t]^{\frac{\rho-\gamma}{1-\gamma}} \left[((1-\gamma)V_t)^{\frac{1-\rho}{1-\gamma}} - C_t^{1-\rho} \right] \\
 &\equiv -f(C_t, V_t)
 \end{aligned}$$

where the second equality follows from L'Hospital's rule.

- The function $f(C_t, V_t)$ is the continuous-time (normalized) aggregator. Assuming the transversality condition $\lim_{T \rightarrow \infty} E_t[V_T] = 0$ holds, we can integrate dV forward to get the following expression for the value function,

$$V_t = E_t \int_t^{\infty} f(C_s, V_s) ds$$

- Again, in the special case $\rho = \gamma$, the nonlinearity disappears and the aggregator becomes $f(C, V) = \delta \left[\frac{1}{1-\gamma} C^{1-\gamma} - V \right]$, which gives us the usual time separable preferences

$$V_t = \delta E_t \int_t^{\infty} \frac{1}{1-\gamma} C_s^{1-\gamma} e^{-\delta(s-t)} ds$$

- Because SDU preferences are recursive, we can use DP as usual to characterize optimal consumption/portfolio decisions. In fact, the only change we need to make to our previous analysis of the Merton problem is to replace $U(C)$ with $f(C, V)$ in the HJB equation! [Exercise: Show that if μ , r , and σ are constant, recursive preferences still produce the myopic optimal portfolio, $\alpha = (\mu - r)/\gamma\sigma^2$. Hence, when the investment opportunity set is constant, portfolio choice is independent of the EIS. Is this still true in general equilibrium? How would the EIS influence r ? How would it change μ ?]
- Although the previous analysis implicitly assumed $\rho \neq 1$ and $\gamma \neq 1$, by subtracting appropriate scaled constants from the utility function, the analysis can be extended to these cases as well. One that is often used in practice is the $EIS = 1$ case, pertaining to $\rho = 1$. One can readily verify that in this case the (normalized) aggregator becomes

$$f(C, V) = \delta(1 - \gamma)V \left[\log(C) - \frac{1}{1 - \gamma} \log((1 - \gamma)V) \right]$$

In this case the consumption/wealth ratio is constant ($C = \delta W$), for **all** values of the risk aversion coefficient, γ . Hence, we can use γ to influence the portfolio/risk premium, without disturbing the savings rate (which remains constant when $\rho = 1$).

- For asset pricing we need to compute the SDF. Although we could take continuous-time limits of the previously derived discrete-time SDF, it is easier and more illuminating to consider a classic perturbation/variational argument.

- As always, we can think of the equilibrium price of an asset as the change in the expected present value of utility that buying or selling the asset enables

$$P_t = E_t[(\beta U'(C_{t+1})/U'(C_t)) \cdot X_{t+1}]$$

To extend this intuition to continuous time with nonseparable preferences, consider the following perturbed value function

$$V_t^\eta = E_t \left[\int_t^\infty f(C_s + \eta X_s, V_s^\eta) ds \right]$$

where η is the perturbation parameter.

- Next, define the 'utility gradient' as the following limit

$$\begin{aligned} dV_t &= \lim_{\eta \rightarrow 0} \frac{V_t^\eta - V_t}{\eta} = \lim_{\eta \rightarrow 0} E_t \left[\int_t^\infty \frac{1}{\eta} [f(C_s + \eta X_s, V_s^\eta) - f(C_s, V_s)] ds \right] \\ &= E_t \left[\int_t^\infty [f_c(C_s, V_s) X_s + f_v(C_s, V_s) dV_s] ds \right] \end{aligned}$$

where subscripts denote partial derivatives.

- Solving this equation gives us

$$dV_t = E_t \left[\int_t^\infty \exp \left(\int_t^s f_v(C_u, V_u) du \right) f_c(C_s, V_s) X_s ds \right]$$

Note that this is just an asset pricing equation, with X_s as the payoff stream and the following expression as the SDF

$$M_t = \exp \left[\int_0^t f_v(C_s, V_s) ds \right] f_c(C_t, V_t)$$

- In the special time-additive case, where $\rho = \gamma$, we know $f(C, V) = U(C) - \delta V$, so that $f_v = -\delta$ and $f_c = U'(C)$. Thus, we get back our old friend, $M_t = e^{-\delta t} U'(C)$. Note that in the more general nonlinear case, the partial derivative of the aggregator w.r.t. continuation utility plays the role of the rate of time preference. This goes back all the way to Koopmans (1960).
- As always, pricing is determined by the drift and volatility of M_t , $dM/M = -r_t dt - \kappa_t dB$, where r_t is the riskless rate and κ_t is the price of risk. From above,

$$d \log M_t = f_v(C_t, V_t) dt + d \log f_c(C_t, V_t)$$

so that the diffusion coefficient of dM/M , and hence the price of risk, is determined by the diffusion coefficient of $d \log f_c(C_t, V_t)$. Notice that this will now depend on the dynamics of continuation utility, V , and not just consumption. This is where the long-run risk literature comes in.

LONG-RUN RISKS

- In a very influential paper, Bansal & Yaron (*JF*, 2004) combined EZ preferences with an environment featuring 'long-run risk', i.e., persistent changes in consumption **growth** and stochastic volatility. They discovered that their model could match many features of observed asset market data.
- The innovation in their paper was the **interaction** between nonseparable preferences and long-run risk. Previous researchers had applied EZ preferences, but only in i.i.d environments. The only real benefit was to mitigate the Risk-Free Rate Puzzle, by separating risk aversion from intertemporal substitution. Very high risk aversion was still needed to explain risky asset returns. Likewise, previous researchers had studied environments with changing consumption growth, but always with time-additive CRRA preferences, where such growth rate risk is unpriced.
- To see why it is the interaction that is crucial, let's return to our previous expression for the log-linearized SDF

$$\tilde{m}_{t+1} = -\gamma \tilde{c}_{t+1} - (\gamma - \rho) \Delta E_{t+1} \sum_{j=1}^{\infty} \beta^j \Delta c_{t+1+j}$$

Notice that revisions in expectations about future consumption growth are irrelevant with separable preferences ($\gamma = \rho$). At the same time, notice that even if $\gamma \neq \rho$, the implications of the model collapse to those of the additive/CRRA model if consumption growth is i.i.d., since then beliefs about future consumption growth never change.

- One way to interpret this interaction is in terms of the timing of the resolution of uncertainty. Note that if revisions about consumption growth are to make assets riskier, we must have $\gamma > \rho$, or in terms of the IES $\psi = \rho^{-1}$, $\psi\gamma > 1$. From our earlier discussion we know that this parameter configuration signals a preference for the early resolution of uncertainty. Growth rate uncertainty implies risk is resolved slowly, and EZ agents with $\gamma > \rho$ don't like this, and so require a return premium for bearing such risk.
- The following four log-linearized equations are the heart of the BY model:

$$\begin{aligned} \Delta c_{t+1} &= \mu_c + x_t + \sigma_t \varepsilon_{c,t+1} \\ x_{t+1} &= \rho x_t + \phi_x \sigma_t \varepsilon_{x,t+1} \\ \sigma_{t+1}^2 &= \bar{\sigma}^2 + \nu(\sigma_t^2 - \bar{\sigma}^2) + \sigma_w w_{t+1} \\ \Delta d_{t+1} &= \mu_d + \phi x_t + \phi \sigma_t u_{t+1} + \pi \sigma_t \varepsilon_{d,t+1} \end{aligned}$$

The important variable here is x_t , which captures stochastic consumption (and dividend) growth. To both fit the data and explain asset prices, these growth rate fluctuations must be **persistent**. BY set $\rho = .98$ per month, implying a half-life of expected consumption growth of 2-3 years.

- In order to capture the cyclical nature of risk premia and price/dividend ratios, BY find that it is important to allow stochastic volatility. As with the conditional mean, these volatility fluctuations must be persistent. BY set $\nu = .99$.

RESULTS

- Of course, the other two important parameters are γ and $\rho^{-1} = EIS$. BY find that in order to match the risk-free rate and the cyclical covariance between volatility and asset prices, the EIS must exceed unity. They set it to $EIS = 1.5$. Given this, they see how large γ must be to approximate the observed equity premium.
- The model is solved both numerically and using log-linear approximations. The following table displays the key results, based on data from the USA for the period 1929-1998.

Variable	Data	$\gamma = 10/EIS = 1.5$	$\gamma = 7.5/EIS = 1.5$
$E(r_m - r_f)$	6.33	6.84	4.01
$E(r_f)$	0.86	0.93	1.44
$\sigma(r_m)$	19.42	18.65	17.81
$\sigma(r_f)$	0.97	0.57	0.44
$E[\exp(p - d)]$	26.56	19.98	25.02
$\sigma(p - d)$	0.29	0.21	0.18
$\beta(3)/R^2$	-0.37/0.19	-0.47/0.10	—

- BY distinguish between consumption and dividends (implicitly allowing for labor income). The leverage ratio parameter is set to $\phi = 3$. r_m refers to the return on the 'market portfolio', defined as a claim to the dividends on the CRSP value-weighted market portfolio.
- The bottom row displays the regression coefficient and R^2 from a regression of 3-year excess returns on the price/dividend ratio.

CAVEATS

- At the time of publication, BY's results were the most positive yet obtained in the macro-finance literature. They have been replicated many times.
- Still, many have raised questions about the plausibility of the underlying assumptions (Beeler & Campbell (2012) provide a detailed discussion):
 - 1 Evidence in favor of both stochastic growth and stochastic volatility in aggregate consumption data is weak. Although the data don't strongly reject the BY specification, nor do they strongly support it.
 - 2 The assumption $EIS > 1$ is crucial to the results, but most other empirical work (especially with micro data) finds that $EIS < 1$.
 - 3 In the model, long-term bonds are a good hedge against changes in expected future consumption growth, since a decrease in expected consumption growth lowers interest rates and raises bond prices. This produces a downward-sloping yield curve, which is inconsistent with the data.
 - 4 The model predicts that price/dividend ratios should not only forecast returns, but should also forecast volatility. There is little evidence for this in the data.