

INTERPRETING AGGREGATE STOCK MARKET BEHAVIOR: HOW FAR CAN THE STANDARD MODEL GO?*

MASSIMILIANO DE SANTIS[†]

ABSTRACT. We add incomplete markets, time variation in expected dividend growth, and fluctuating uncertainty to an otherwise standard economy to explain the major features of aggregate U.S. stock market behavior. The economy matches the high equity premium and low risk free rate with reasonable levels of relative risk aversion, ten and lower. It also replicates a volatile, persistent and mean reverting price-dividend ratio, time varying risk premia, predictable returns, and the low correlation between consumption growth and excess returns. The price-dividend ratio predicted by the model using U.S. historical consumption data (1891-2001) has a correlation of 72% with the S&P 500 price-dividend ratio. The main contribution of our model is that risk premia come from standard expected utility theory. Agents fear stocks primarily because they do poorly in recessions. A small, time varying component in expected dividend growth amplifies and gives persistence to downward price movements, as in the data, which makes it harder for agents to smooth consumption.

*I am grateful to John Campbell, George Constantinides, Timothy Cogley, Martine Quinzii, Aaron Smith, Jason Lepore, seminar participants at Dartmouth College, The University of Toronto, Fordham University, The European Central Bank, Tulane University, and Clemson University for many helpful comments and discussions. I especially thank Louis Makowski, who helped me a lot from the very inception of this project, giving generously of his time.

[†]Department of economics, Dartmouth College, 6106 Rockefeller Hall, Hanover, NH 03755. E-mail: massimiliano.desantis@dartmouth.edu.

1. INTRODUCTION

A number of empirical observations challenge our economic understanding of the relationship between risk and returns. The standard consumption-based CAPM fails to replicate U.S. data in several dimensions. First, post-war realized excess returns on the S&P 500 over six month commercial paper average more than 7%, a level too high for standard consumption models. Second, the volatility of real stock returns is too high in relation to the volatility of short term interest rates. Third, consumption growth and real stock returns have very low correlation in U.S. data. Fourth, excess returns on U.S. stocks relative to short term bonds are forecastable.

Understanding these puzzles is important for both finance and macroeconomics. From a finance view, the observations above imply that we do not have a successful theory and measure of the sources of risk that drive expected returns, needed to correctly price assets. In macroeconomics, identification of variables that affect the level and movement of risk premia is important for understanding business cycles. Standard macroeconomic models predict that people are almost indifferent about business cycles (Lucas, 1987). Asset prices reveal that they are not indifferent; the risk-return tradeoff in the stock market implies that agents require a substantial premium to shoulder an additional unit of income volatility.

Our objective is to gain a better understanding of major movements in aggregate stock prices. In contrast to previous literature, we explain the aforementioned asset pricing anomalies using standard expected utility theory and CRRA preferences with reasonable levels of risk aversion. We do so by specifying an economy in which markets are incomplete, there is a persistent component in shocks to dividend growth, and agents face uncertainty of varying degree over the business cycle.

Early work on the equity premium puzzle (Mehra and Prescott, 1985) focused on matching a high equity premium with a low risk free rate (also see Reitz 1988, Weil 1989a, among others). More recent models have been able to match other features of the data, linking rational expectations and fundamentals to the return predictability and excess volatility literatures. Among the successful models are Campbell and Cochrane's (1999) model with external habit; Barberis, Huang and Santos' (2001) model with elements of behavioral finance; Bansal and Yaron's (2004) and Lettau, Ludvigson and Wachter's (2003) models with Epstein-Zin-Weil preferences.¹

All successful models have so far relied on substantial modifications of standard preferences. By adding parameters and/or state variables to the utility function, these generalizations allow for extra degrees of freedom in matching relevant dimensions of the data. Epstein-Zin-Weil

¹Surveys of the literature are found in Kocherlakota (1996), Campbell (2003), and Cochrane (2001).

preferences allow agents to express a preference for flexibility in the resolution of uncertainty. This implies that marginal rates of substitution across states may differ from marginal rates of substitution over time. Such preferences can explain asset pricing features only by assuming high risk aversion and/or high intertemporal elasticity of substitution.

Alternatively, habit formation models postulate that agents care not only about the level of their consumption, but also the distance between the level and some reference point. In some cases, as in Campbell and Cochrane (1999), the reliance on habit is more extreme: utility does not depend on the level of consumption, only on the relative distance to a habit level. Thus the interpretation of all stock market features depends on habit formation. Habit models match asset pricing features such as the equity premium only by assuming extremely high levels of risk aversion (50 to 100).

Here we propose an economy with heterogeneous agents in which recessions (expansions) are the driving factor for the joint behavior of the equity premium, the risk free rate, the Sharpe ratio, and the price-dividend ratio. Preferences are entirely standard. The model is based on the idea that during recessions, the degree of uncertainty about future economic prospects is greater. This has two manifestations: a negative persistent shock to dividend growth, and an increase in uncertainty (how much the agents know) about the distribution of idiosyncratic consumption shocks.

During uncertain times, the current state of the economy is less informative about future prospects. As a result, investment decisions are harder to make, some investment projects may look less profitable, and some irreversible investment decisions may be delayed. This lowers expected future dividend growth. Waiting for uncertainty to clear and the process of learning about available technologies and/or skills takes some time, which is reflected in a persistent effect of the recessionary shock.

In our economy, agents receive idiosyncratic income shocks, so no agent knows how much of the aggregate growth he/she will share. This reflects the fact that people have jobs and cannot fully insure their labor income. Agents know that nature will draw individual shocks from a distribution and hand them out to the population. But we assume that agents do not know the distribution of shocks nature will pick. So they not only face “risk,” but also “uncertainty,” in the sense of Hart (1942), and also Jones and Ostroy (1984). If nature draws its shocks from a fixed, commonly known distribution, individuals face risk. If even the distribution is unknown (although individuals have a probability distribution over the family of distributions), individuals are said to face not just risk, but also uncertainty. We will assume that recessions are times in which individuals face higher uncertainty, in the sense that they know less about the distribution nature will pick. Greater uncertainty makes labor income

more volatile and insurance more important (e.g. precautionary savings). In our model stocks are feared because they are a bad insurance tool against negative consumption shocks, which occur at times of higher individual uncertainty.

The main ingredients of the economy are the following:

- standard CRRA time-separable preferences,
- heterogeneous agents subject to idiosyncratic income shocks and incomplete markets as in Constantinides and Duffie (1996),
- a small persistent component in shocks to dividend growth as in Barsky and De Long (1993).

The absence of complete markets implies that individual consumption is more variable than per capita consumption, even if individuals are identical *ex ante*. This is important to reconcile the low variance of per capita consumption growth (3.6% in Mehra and Prescott's annual data) with the higher cross-sectional volatility observed in panel data (greater than 10% in Carroll, 1992). With a greater probability of a very low realization of consumption, risk averse individuals are willing to save at low interest rates (the precautionary saving motive). Uncertainty about the idiosyncratic shock fluctuates over the business cycle, and so will the precautionary motive, being stronger in recessions than in expansions. This fluctuating uncertainty is important to match the countercyclical variation in the price of risk, that is, the fact that people require higher returns to invest in the stock market during recessions than they do during expansions, as emphasized in Campbell and Cochrane (1999).

The small persistent component in shocks to dividend growth is meant to capture time-varying economic prospects. Barsky and DeLong (1993) show that a small persistent component in shocks to dividend growth leads to extremely volatile stock prices. Here we show that, in a general equilibrium model, such a persistent component is important for explaining the high equity premium, the persistence of the price-dividend ratio, and the predictability of excess returns.

In the next section, we relate the model to existing literature, discuss how it goes beyond it, and how it improves upon it. In sections 3 and 4 we detail the model. To gain intuition, we begin in section 3 with a simplified model that can be solved analytically. This model matches the historical equity premium, low risk free rate, and procyclical behavior of the price-dividend ratio. Section 4 details a more general model that, in addition to these three features, generates a countercyclical Sharpe ratio and predictable excess returns. In section 5, we present simulation results based on this second model. Section 6 concludes.

2. RELATIONSHIP TO EXISTING LITERATURE

Our theory touches on several strands of the vast literature that seeks to document and resolve asset pricing puzzles. Surveys are found in Kocherlakota (1996), Campbell (2003), and Cochrane (2001). The objective here is to relate our model to both previous work that directly motivated our maintained assumptions and also to alternative theories that are able to generate asset pricing features similar to ours.

In our treatment of incomplete markets and idiosyncratic shocks, we follow Constantinides and Duffie (1996). Their objective is to study the theoretical possibility that any combination of equity premium and risk free rate can be generated by an economy with persistent and uninsurable individual income shocks, even in economies with relatively low per-capita consumption risk and low risk aversion. All that is required in their economy is cross-sectional variation in consumption growth, and a cross-sectional variation that is higher in periods of low stock market returns (usually recessions). Because their focus is theoretical in nature, they do not test their economy empirically. Recent work has studied the relevance of cross-sectional variation in consumption for explaining the equity premium and the risk free rate. Evidence is mixed. Using CEX data, Cogley (2002) finds that cross-sectional variation generates premia of 2% or less for preferences with a low degree of risk aversion. Brav, Constantinides and Geczy (2002) find that a pricing equation that takes individual risk into account is not rejected in the CEX data.

In this paper we go beyond the work of Constantinides and Duffie and propose a specification of the idiosyncratic shock distribution that has not been considered in the literature. We also show that this has important asset pricing implications that go beyond matching the equity premium. The process is important for generating a low risk free rate and excess return predictability. Furthermore, our results are based on a conservative specification of the idiosyncratic shock process, in the sense of being consistent with the findings of Cogley (2002).

Our model borrows the idea of a persistent component in shocks to dividend growth from Barsky and De Long (1993). In their paper, they are concerned with long run swings in the U.S. stock market and with the evidence of excess volatility documented in Shiller (1981), and in LeRoy and Porter (1981). They show that if expected dividend growth is time varying and persistent, prices respond more than proportionately to long run movement in dividends. They also show that the time varying component of dividend growth need not be detectable in the dividend data for it to have large effects on stock prices. Here we incorporate their idea in an asset pricing model and show its importance for the model's ability to replicate the high equity premium and the volatile, persistent price-dividend ratio.

The model of Bansal and Yaron (2004) also uses Barsky and De Long’s idea, but it does so in a representative agent economy in which agents have preferences of the Epstein-Zin-Weil type. Epstein-Zin-Weil preferences include standard preferences as a special case, provided the relative risk aversion parameter equals the inverse of the inter-temporal rate of substitution. The parameter specifications used in Bansal and Yaron makes their preferences far from standard. They assume that both relative risk aversion and intertemporal rate of substitution are greater than one. In their model, the low risk free rate, countercyclical Sharpe ratio, and excess return predictability features depend crucially on an elasticity of substitution greater than one. While the authors rightly point out the difficulties with estimating this elasticity, most empirical evidence indicates a value less than unity (Motohiro 2004, Campbell 2003, Vissing-Jorgensen 2002, Hall 1988). Evidence in Hall (1988), Campbell (2003), and Motohiro (2004) even casts doubts on a positive inter-temporal elasticity of substitution. By contrast, we assume a low intertemporal elasticity of substitution, 0.1 to 0.2, consistent with the findings in Motohiro (2004). In terms of interpretation of the various market features, our model is very different, but both models have in common the fact that risk premia are time-varying because of changing risk in the economy. In our model, risk premia are time varying because of changing individual risk, whereas in Bansal and Yaron it is because of changing per-capita risk.²

By contrast, in Campbell and Cochrane (1999), henceforth CC, the factor that drives risk premia is changing risk aversion. This is because, as consumption moves closer to “habit,” the curvature of their utility function (the second derivative with respect to consumption) increases. This makes marginal rates of substitution more sensitive to consumption variation, an important feature to explain the equity premium. In both CC and our model, a strong precautionary saving motive in the model generates a low risk free rate. But the source of the need for insurance is different in the two models. In ours it is individual risk, and the uncertainty about the distribution of individual risk, that generate a strong precautionary saving motive. In CC, the source is per capita risk, and the need for keeping consumption above habit.

Any model that wants to generate time variation in the price of risk and predictable excess returns needs to have a stochastic discount factor that is heteroskedastic. In CC this is achieved by postulating that the sensitivity of the marginal rates of substitution to consumption shocks changes over the business cycle, and so heteroskedasticity is part of the definition of the

²Another, more technical difficulty is also worth pointing out. Solving and empirically testing the Epstein-Zin-Weil model requires the assumption that the stock market index is a good proxy for the wealth portfolio. Mehra and Prescott (2003) feel that this is not a satisfactory assumption as it overstates the correlation between wealth portfolio and asset returns.

habit process. In Bansal and Yaron (2004), heteroskedasticity is embedded in the per-capita consumption growth process, and it is interpreted as changing risk; time variation in risk is assumed independent of the business cycle. In our model, it is fluctuating uncertainty about the idiosyncratic shock that gives rise to a heteroskedastic stochastic discount factor.³

In Mehra and Prescott's seminal paper, the economy is based on three assumptions: individuals maximize expected discounted value of a stream of utilities generated by a *power utility function*; markets are *complete*; asset trading is *costless*. Taking their economy as a point of departure, our model relaxes the assumption of completeness, postulates a different process for dividend growth, and adds fluctuating uncertainty about individual uninsurable shocks.

Even successful models like CC have difficulty matching the price dividend ratio in the 1990's. In confronting the model's predictions (when fed actual consumption shocks) with the historical behavior of the price-dividend ratio in U.S. data, it helps to take stock repurchases into account. Grullon and Michaely (2002) report evidence for the period 1972-2000 that repurchases have become an important source of payout for U.S. corporations, and firms finance their share repurchases with funds that would otherwise be used to increase dividends. Boudoukh, Michaely, Richardson, and Roberts (2003) construct a payout yield that adjusts the dividend yield for repurchases. They find that the adjustment can explain the lack of predictive power of the dividend yield over the 1990's, and they suggest that asset pricing models that relate cash flow to asset pricing (like ours) should take this into account. We do so and find that model predictions are much improved over the period, confirming their intuition.

3. EQUILIBRIUM AND SOLUTION FOR A SIMPLIFIED ECONOMY

Consider an exchange economy with a single non-durable consumption good and two traded assets, a risk-free discount bond and a risky equity. Bonds are issued at time $t-1$, matured at t , and each bond has a par value of one. We assume the bond is in zero net supply. The risky equity (whose net supply we normalize to be one) pays dividend D_t and has ex-dividend price P_t . Each consumer i is endowed with labor income I_t^i and consumes C_t^i at time t . Aggregate labor income is I_t , and aggregate consumption is $C_t = I_t + D_t$. It is assumed that $I_t + D_t > 0$ for all times t . There is an infinite set of distinct consumers denoted by A . The increasing sequence of information sets $\{\mathcal{F}_t : t = 0, 1, 2, \dots\}$ available to each consumer includes the equity's dividend history, the history of equity and bond prices, and the disaggregated labor income history $\{I_s^i : i \in A, 0 \leq s \leq t\}$. At time t , consumer i holds a portfolio of shares of

³In our setting, the assumption that recession times are times of greater uncertainty is enough to generate heteroskedasticity, as will be clear in section 4.

the risky asset θ_t^i and of the bond b_t^i . The time t budget constraint is:

$$(3.1) \quad C_t^i + \theta_t^i P_t + b_t^i \leq I_t^i + \theta_{t-1}^i (P_t + D_t) + b_{t-1}^i R_t^f,$$

where R_t^f denotes the return on a bond issued at $t - 1$. Consumers have homogeneous preferences represented by a time-separable von Neumann-Morgenstern utility function with constant relative risk aversion coefficient γ and a constant subjective discount factor δ . At time 0, each consumer maximizes

$$(3.2) \quad E \left[\frac{\sum_{t=0}^{\infty} \delta^t (C_t^i)^{1-\gamma}}{1-\gamma} \mid \mathcal{F}_0 \right]$$

subject to the sequence of budget constraints (3.1) by choosing a sequence $(\theta^i, b^i, C^i) \equiv (\theta_t^i, b_t^i, C_t^i)$, $t = 0, 1, 2, \dots$

An equilibrium is a security price and bond return process (P, R^f) , and strategies $\{(\theta^i, b^i, C^i) : i \in A\}$ for the consumers such that

- (i) (θ^i, b^i, C^i) maximizes (3.2) subject to (3.1)
- (ii) markets clear, i.e., $\sum_{i \in A} \theta_t^i = 1$ and $\sum_{i \in A} b_t^i = 0$ for all t .

Market clearing implies that $\sum_{i \in A} C_t^i = C_t \equiv I_t + D_t$ for all t .

An equilibrium price process for the risky asset will satisfy the following condition for all $i \in A$:

$$(3.3) \quad P_t = E \left[\delta \left(\frac{C_{t+1}^i}{C_t^i} \right)^{-\gamma} (P_{t+1} + D_{t+1}) \mid \mathcal{F}_t \right],$$

where the expectation is taken conditional on \mathcal{F}_t . Constantinides and Duffie (1996) show existence of an equilibrium for a carefully chosen uninsurable labor income process. They show that the equilibrium price process satisfies:

$$(3.4) \quad P_t = E \left[\delta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \exp \left[\frac{\gamma(\gamma+1)}{2} y_{t+1}^2 \right] (P_{t+1} + D_{t+1}) \mid \phi_t \right],$$

where the variable y_{t+1}^2 represents the variance of the cross-sectional distribution of individual consumption growth relative to aggregate growth, i.e.

$$y_{t+1}^2 = \text{Var} \left(\log \left(\frac{C_{t+1}^i / C_{t+1}}{C_t^i / C_t} \right) \right).$$

So P_{t+1} depends only on aggregate quantities. The information set ϕ_t differs from \mathcal{F}_t in that it does not include the disaggregated labor income histories, i.e., it is a subfield of \mathcal{F}_t and is interpreted as the information set observed by the econometrician. But the expectation in

(3.4) is the same if we condition on \mathcal{F}_t , because the extra information contained in \mathcal{F}_t is not relevant in calculating the expected value.⁴

The economy so far follows Constantinides and Duffie. We now extend beyond their model and describe a simple process for aggregate consumption growth, dividend growth, and y_t^2 that leads to an analytical solution of the equilibrium and can be used to determine the level of the risk free rate, the equity premium, and the procyclical variation in the price-dividend ratio.

Let $g_{t+1} = \Delta c_{t+1}$ and $g_{t+1}^d = \Delta d_{t+1}$, where lower case letters denote logs of the uppercase counterparts. The process for g_t, g_t^d, y_t is as follows:

$$(3.5) \quad \begin{aligned} g_{t+1} &= \mu + \sigma \eta_{t+1} \\ g_{t+1}^d &= \mu^d + \phi x_t + \sigma \varphi_d u_{t+1} \\ x_{t+1} &= \rho x_t + \sigma \varphi_e \eta_{t+1} \\ y_{t+1}^2 &= \bar{y}^2 - \varphi_y \eta_{t+1} \end{aligned}$$

with the two shocks, η_{t+1} and u_{t+1} having correlation $\rho_{\eta,u}$. The shocks are assumed to be i.i.d. with normal distribution $N(0,1)$. Hence σ is the standard deviation (volatility) of consumption growth. The volatility of dividend growth is $\sigma \varphi_d$, where $\varphi_d > 1$, reflecting the fact that dividends are more volatile than consumption in the data. The mean of dividend growth includes a small (relative to μ^d) time varying and persistent component x_t . Hence φ_e will be very small, to ensure that x_t is small.

Agents do not know the variance of the individual shocks nature will pick, which is reflected in the fact that y_{t+1}^2 is a random variable with a positive standard deviation φ_y . Hence, using Hart's terminology, agents face uncertainty about the variance of their incomes.⁵ In this simple model, y_{t+1}^2 is assumed to be i.i.d. with mean value \bar{y}^2 and a constant variance; hence uncertainty is constant through time (not linked to recessions).

A recession in the model is viewed as a large negative value of the aggregate shock η_t . Negative aggregate shocks have two effects: a negative, persistent shock to dividend growth through x_t , and a higher y_t^2 , the variance of individual shocks. The interpretation of the x_t component is that times of negative aggregate shocks are times of higher uncertainty about future economic prospects. As a result investment decisions are harder to make, some investment projects may look less profitable, and some irreversible investment decisions may be delayed. This lowers future expected dividend growth. Waiting for uncertainty to clear and

⁴We detail the specification of the income shocks, which contain a martingale component as in Constantinides and Duffie (1996), in appendix C. Interestingly, Meghir and Pistaferri (2004), and Deaton and Paxson (1994) provide empirical evidence that supports the presence of a martingale component in household earnings.

⁵Individual shocks come from a normal distribution with unknown variance y_{t+1}^2 , so the equation for y_{t+1}^2 defines a probability distribution on probability distributions.

the process of learning about profitable investments takes some time, which is reflected in a persistent effect of the negative aggregate shock.

Log-consumption growth in this economy is a pure i.i.d. process and it does not contain the persistent component x_t . This may trouble the careful reader because persistence in the dividend process could lead to persistence in consumption growth. Here are two ways to think about the i.i.d. process to alleviate this possible fear. First, the consumption we are calibrating in this model is the consumption of non-durables and services. Since aggregate shocks are felt disproportionately in the durable goods sectors, we can think of the x_t shock to affect the consumption of durable goods in the same way it does for the dividend process, but not the non-durables and services. A second way to think about it is to think that households have access to a linear saving technology. With a constant return technology, under log-normality assumptions about the stochastic income process, log of consumption would follow a random walk with drift. Without log-normality assumptions, consumption is still very close to a random walk with drift as shown in Robert Hall's 1978 paper.⁶

The parameter ρ in x_t determines the persistence of the expected growth rate. For the time varying component to be economically relevant, the persistence parameter has to be quite high, close to one. Barsky and De Long (1993), in a partial equilibrium model, show that a small and highly persistent process (they choose ρ to be one) plus a volatile i.i.d. process, as in our definition of g_{t+1}^d , cannot be distinguished from an i.i.d. process. Further, Shephard and Harvey (1990) show that standard identification techniques would favor the i.i.d. process. In other words, it is difficult to distinguish between these two types of processes in a finite sample. But their economic implications are very different.

We can now solve for the price-dividend ratio, the risk free rate, and the equity premium of this economy. We can solve for the price-dividend ratio by using the log-linear approximation of returns in Campbell and Shiller (1988):

$$(3.6) \quad \ln(R_{t+1}) \equiv r_{t+1} = \kappa_0 + \kappa_1 v_{t+1} - v_t + g_{t+1}^d$$

where v_t is the log of the price-dividend ratio, κ_0 and κ_1 constants of approximation. In Appendix A we show that, by using the approximation in the Euler equation (3.7), the log of the price-dividend ratio v_t is an affine function of the state x_t , i.e, in this economy,

$$v_t = A_0 + A_1 x_t, \quad \text{with}$$

$$A_1 = \frac{\phi}{1 - \kappa_1 \rho} > 0.$$

⁶See Hall's Corollary 5. It can also be shown that with CRRA preferences and a linear saving technology, up to a second order approximation of the intertemporal MRS, log of consumption follows a random walk with drift, which is our assumption.

If there is a recessionary shock ($\eta < 0$), time varying expected growth x_t falls and so does the price of the equity relative to dividends. The approximating constant κ_1 is close to one, and so is ρ , which implies that even a small change to the conditional mean of dividend growth can have important implications for asset prices. This “Barsky and De Long effect” can explain both the procyclical behavior of the price-dividend ratio and the high volatility of aggregate stock prices.

We can use the Euler equation (3.4) to find the risk free rate and the expected return on the risky asset by using our normality assumptions. First re-write (3.4) as

$$(3.7) \quad E_t \left[\exp \left\{ \ln \delta - \gamma g_{t+1} + \frac{\gamma(\gamma+1)}{2} y_{t+1}^2 + r_{t+1} \right\} \right] = 1.$$

Then use the fact that if X is normal, $Ee^X = e^{EX+0.5\text{Var}(X)}$, and take logs of both sides to obtain

$$(3.8) \quad \ln \delta - \gamma E_t g_{t+1} + \alpha E_t y_{t+1}^2 + E_t r_{t+1} + \frac{1}{2} \text{Var}_t(-\gamma g_{t+1} + \alpha y_{t+1}^2 + r_{t+1}) = 0,$$

where $\alpha = 0.5\gamma(\gamma+1)$.⁷ The Euler condition (3.8) is valid for all asset, so we can substitute r_{t+1}^f for r_{t+1} . Therefore the risk free rate,

$$(3.9) \quad r_{t+1}^f = -\ln \delta + \gamma \mu - \alpha \bar{y}^2 - \frac{1}{2}(\sigma\gamma + \alpha\varphi_y)^2.$$

By contrast, in Lucas’ economy $r_{t+1}^f = -\ln \delta + \gamma \mu - \frac{\gamma^2 \sigma^2}{2}$. Since σ is only 3.6% in Mehra and Prescott’s data, σ^2 is very small. So, even with high risk aversion, the precautionary saving term $\frac{\gamma^2 \sigma^2}{2}$ is second order, implying an unrealistically large risk free rate when γ is large (the risk free rate puzzle). With individual uninsurable shocks to income, the precautionary saving motive is not second order. For example, a relative risk aversion coefficient of 10 and a cross sectional standard deviation \bar{y} of 5% leads to a risk free-rate of less than 2%, even if $\varphi_y = 0$.

We can use the Euler equation again (this time with r_{t+1} instead of r_{t+1}^f) to find the expected return on the risky asset, which will now contain the covariances implied by the variance term in (3.8):

$$\begin{aligned} \text{Var}_t(-\gamma g_{t+1} + \alpha y_{t+1}^2 + r_{t+1}) &= (\gamma\sigma + \alpha\varphi_y)^2 + \text{Var}_t(r_{t+1}) + \\ &\quad - 2\gamma \text{Cov}_t(g_{t+1}, r_{t+1}) + 2\alpha \text{Cov}_t(y_{t+1}^2, r_{t+1}). \end{aligned}$$

This implies the risk premium on equity is

$$(3.10) \quad E_t(r_{t+1} - r_{t+1}^f) = \gamma \text{Cov}_t(g_{t+1}, r_{t+1}) - \alpha \text{Cov}_t(y_{t+1}^2, r_{t+1}) - 0.5 \text{Var}_t(r_{t+1}).$$

⁷Notice that normality of X is justified here. For the risky asset, it follows from (3.6) and the solution for v_t that log returns r_{t+1} are approximately normally distributed. For the risk free rate, since it is known at time t , normality follows from our assumption about g and \bar{y}^2 .

The first term is the covariance found in the standard model. The second covariance contributes to a positive risk premium if it is negative, i.e. returns are low in times of larger idiosyncratic shocks. This is the case in our model. The third term is a Jensen's inequality term arising from the fact that we are describing expectations of log returns. In effect, this term converts the expected excess returns from a geometric average to an arithmetic average.

We can substitute the log-linear approximation for r_{t+1} in the equity premium expression (3.10) using our solution for v_t , and calculate the premium for this economy:

$$(3.11) \quad E_t(r_{t+1} - r_{t+1}^f) = \gamma\sigma^2\varphi_d\rho_{\eta,u} + \gamma\sigma^2\kappa_1A_1\varphi_e \\ + \alpha\varphi_y\kappa_1A_1\sigma\varphi_e + \alpha\varphi_y\varphi_d\rho_{\eta,u} - 0.5\text{Var}_t(r_{t+1}).$$

The first two terms come from $\text{Cov}(g_{t+1}, r_{t+1})$. The following two terms come from $\text{Cov}(y_{t+1}^2, r_{t+1})$ through the recessionary shock η_t , which links idiosyncratic shocks to the persistent component of dividends. Since the small persistent component of shocks to dividend growth causes large swings in equity prices, the risk effect is also magnified. Notice in fact that the premium depends on A_1 , which is large for ρ sufficiently close to one. For the parameterization that we will use later, the premium in this economy is about 6%. The first term, $\gamma\sigma^2\varphi_d\rho_{\eta,u}$, accounts for a premium between 0.7 and 1.8%, depending on the value of $\rho_{\eta,u}$ used. The second term, $\gamma\sigma^2\kappa_1A_1\varphi_e$, is the most important, adding about 3.2% to the premium. This is the covariance between the permanent shock and aggregate consumption growth; permanent negative shocks occur in recessions. The term $\alpha\varphi_y\kappa_1A_1\sigma\varphi_e$ adds about 1% to the premium as does the last term, $\alpha\varphi_y\varphi_d\rho_{\eta,u}$. These last two terms are the covariances between the shocks to dividend growth (permanent and temporary) and individual risk. The Barsky and De Long effect accounts for more than half of the premium, thus it is the most important source of risk. Idiosyncratic risk on the other hand is important for a low risk free rate (solving the risk free rate puzzle).

4. A MORE GENERAL MODEL AND HOW IT WORKS

4.1. The Stochastic Process.

Notice that the equity premium is constant in the economy above, and so is the variance of returns on the risky asset. The predictability literature indicates countercyclical variation in the expected excess return (Campbell and Shiller 1988a, 1988b). Estimates of the conditional variance of returns also change over time, but they do not move one for one with the conditional mean of excess returns. This implies time varying Sharpe ratios. The Sharpe ratio on the risky asset satisfies the following inequality:

$$\frac{E_t(R_{t+1} - R_{t+1}^f)}{\sigma_t(R_{t+1})} \leq \frac{\sigma_t(M_{t+1})}{E_t(M_{t+1})},$$

where M_{t+1} is the stochastic discount factor. Since $E_t M_{t+1}$ is almost constant in the data, a model that explains countercyclical Sharpe ratios should have a stochastic discount factor that is heteroskedastic, with greater variance in bad times. We now specify such a model.

One would expect that as the economy goes into a recession, uncertainty rises, and with it the desire for income insurance. According to this view of the economy, equities should be less valuable relative to bonds in recessions than in expansions. Recessions are times of high uncertainty, which increases the volatility of individual consumption and lowers the return on stocks. There is more than one way to specify a process that generates the desired time variation in the Sharpe ratio. We specify a model for changing uncertainty so that the risk free rate in the economy stays constant; in this we follow Campbell and Cochrane (1999). The dividend growth and consumption growth processes are specified as

$$(4.1) \quad \begin{aligned} g_{t+1} &= \mu + \sigma \eta_{t+1} \\ g_{t+1}^d &= \mu^d + \phi x_t + \sigma \varphi_d u_{t+1} \\ x_{t+1} &= \rho x_t + \sigma \varphi_e \eta_{t+1}. \end{aligned}$$

The model for y_{t+1}^2 is the following MA(1) with heteroskedastic error term,

$$(4.2) \quad y_{t+1}^2 = \bar{y}^2 - \sigma \lambda_t \eta_{t+1} + \theta \sigma \lambda_{t-1} \eta_t.$$

The term $\sigma \lambda_t$ is the conditional standard deviation of y_{t+1}^2 .

In terms of uncertainty, agents do not know the variance of the individual shocks nature will pick, which is again reflected in the fact that y_{t+1}^2 is a random variable with a positive standard deviation. As in the model of Section 3, agents face uncertainty about the variance of their incomes. But the fact that the standard deviation λ_t is not constant implies that uncertainty is time varying, i.e, it depends on the state of the economy at time t . In particular, recessions will be periods of higher uncertainty.

We use the equations of the maximal Sharpe ratio and the risk free rate for this economy to specify λ_t and the sign of the constant θ . The time varying conditional variance will generate a heteroskedastic stochastic discount factor. The time varying conditional mean, the MA(1) component $\sigma \lambda_{t-1} \eta_t$, affects variation in the risk free rate. Using the normality assumption, we can get the risk free rate from

$$E_t m_{t+1} + E_t r_{t+1}^f + 0.5 \text{Var}(m_{t+1}) = 0.$$

Recall that $m_{t+1} = \ln \delta - \gamma g_{t+1} + \alpha y_{t+1}^2$, hence:

$$\begin{aligned} E_t m_{t+1} &= \ln \delta - \gamma \mu + \alpha E_t [y_{t+1}^2], \\ \text{Var}_t(m_{t+1}) &= \sigma^2 (\gamma + \alpha \lambda_t)^2. \end{aligned}$$

So that

$$(4.3) \quad r_{t+1}^f = -\ln \delta + \gamma\mu - \alpha E_t[y_{t+1}^2] - \frac{1}{2}\sigma^2(\gamma + \alpha\lambda_t)^2,$$

with $E_t[y_{t+1}^2] = \bar{y}^2 + \theta\sigma\lambda_{t-1}\eta_t$.

Let's consider the Sharpe ratio next. In a Gaussian economy, the maximal Sharpe-ratio (from the Hansen-Jagannathan bound) is

$$\max_{\{\text{all assets}\}} \frac{E_t(R_{t+1} - R_{t+1}^f)}{\sigma_t(R_{t+1})} = \frac{\sigma_t(M_{t+1})}{E_t(M_{t+1})} = (e^{\sigma_t^2} - 1)^{1/2} \simeq \sigma_t,$$

where $\sigma_t = \sqrt{\text{Var}_t(m_{t+1})}$. So the maximal Sharpe ratio for this economy is approximately

$$(4.4) \quad \frac{\sigma_t(M_{t+1})}{E_t(M_{t+1})} = \sigma(\gamma + \alpha\lambda_t).$$

A countercyclical Sharpe ratio implies that the last term in the risk free rate equation (4.3) is higher in recessions (in absolute value), which means the risk free rate will be lower as the precautionary savings motive rises. In the data, the risk free rate is procyclical, but the standard deviation is only about 1.5%, very low compared to the standard deviation of the S&P 500, which is about 18%. This means that the effect of greater precautionary saving is in large part offset by an intertemporal substitution motive that makes agents want to borrow against future growth during recessions. Given the low variation in interest rate data, we decided to follow Campbell and Cochrane (1999) and match a constant interest rate, which means that in our model intertemporal substitution completely offsets the larger precautionary saving motive in recessions. This amounts to choosing θ to be positive in the MA(1) process for y_{t+1}^2 . Say at time $t - 1$ the system is at steady state. Then $E_{t-1}y_t^2$ is the steady state constant \bar{y}^2 . If there is a recession at t , λ_t is high, which tends to increase the precautionary saving motive, thus lowering interest rates. With a positive θ , $E_t y_{t+1}^2$ is lower than \bar{y}^2 . This gives the agents an incentive to borrow to get out of the recession, thus increasing the risk free rate. This is the inter-temporal substitution effect. We will assume that the two motives exactly counterbalance. This implies that our model will generate an equity premium without a term premium, as in Campbell and Cochrane. Many models with changing risk, like ours, imply high variation in R_t^f and a high term premium (Jermann 1998, Boldrin, Christiano, and Fisher 1997), which are counterfactual.

These considerations lead to the following definition of λ_t :

$$(4.5) \quad \lambda_t = \frac{1}{\sigma\alpha} \sqrt{2[c - \alpha(\theta\lambda_{t-1}\sigma\eta_t)]} - \frac{\gamma}{\alpha}.$$

Notice that λ_t is greater in recessions (when η_{t+1} is negative) than in expansions (when η_{t+1} is positive). The constant c is calibrated so that λ_t is always defined by (4.5) in our simulations,

and it yields a risk free rate of 1.5% in (4.3).⁸ Appendix A details how the model is solved numerically and simulated.

It is important to notice that the time variation in λ_t is needed to match the Sharpe ratio, (i.e., the price of risk), not the size of the equity premium and the level of the risk free rate. The latter features can be matched even with the simpler process (3.5).

By way of comparison, consider the maximal Sharpe ratio generated by the economy of Campbell and Cochrane (their equation (7)):

$$\max_{\{\text{all assets}\}} \frac{E_t(R_{t+1} - R_{t+1}^f)}{\sigma_t(R_{t+1})} = \gamma\sigma(1 + \lambda(s_t)).$$

The variable s_t is the log of what they call the surplus consumption ratio, i.e. consumption above habit divided by consumption. The variable s_t is low in bad times. The function $\lambda(s_t)$ is the conditional standard deviation of s_t , and it is high in bad times. The process s_t in CC is a persistent AR(1). The analog of $\lambda(s_t)$ in CC is our λ_t . But of course the interpretation is completely different because ours is not a habit model. We believe that recessions are times of high uncertainty, and this has two effects in a world with incomplete markets. First, it makes people more subject to idiosyncratic shocks because of greater uncertainty (in the sense of Hart), and second, it makes investment decisions harder and therefore lowers growth expectations (of dividends). This implies that stocks are not a good hedge against individual consumption volatility, they perform badly when consumption is low and when uncertainty is high.

4.2. Parameter Specification.

Parameters are presented in Table 1. The mean and the standard deviation of consumption growth, μ and σ , are the same values used by Bansal and Yaron (2004). They match the BEA data on real per capita consumption of non-durables and services for the period 1929-1998. The risk free rate r^f that we decided to match is somewhat higher than the one used by Campbell and Cochrane, who use 0.94%, but it is lower than the average annualized log-rate on six month commercial paper of 2.4% in Shiller's dataset, which covers the sample 1871-2002.⁹ Siegel (1999) presents evidence that a value of 1% underestimate the return on treasury

⁸To derive (4.5), rewrite (4.3) as

$$\underbrace{r^f + \ln \delta - \gamma\mu + \alpha\bar{y}^2}_{-c} = -\alpha\theta\sigma\lambda_{t-1}\eta_t - \frac{1}{2}\sigma^2(\gamma + \alpha\lambda_t)^2$$

and solve for λ_t .

⁹See Robert Shiller's website. As a technical aside, the use of a higher rate in our simulations helps ensure that draws from the conditional distribution of y^2 (which is a normal variate) are not negative.

bills. There is evidence that the real rate both during the nineteenth century and after 1982 has been substantially higher than 1%.

The coefficients ϕ , φ_e , and ϕ_d in x_t , the persistent component in the dividend process, are chosen so that the Barksy and De Long effect powerfully affects price movements, and at the same time the implied dividend process is consistent with the data. Table 2 presents evidence that our dividend process has time series features similar to the observed data. The persistence parameter $\rho = .94$ is less than unity, so the dividend growth process is stationary.

TABLE 1. Parameter Choices

Parameter		Value
Mean consumption growth (%)	μ	1.89
Standard deviation of consumption growth (%)	σ	2.9
Log risk-free rate (%)	r^f	1.5
Mean dividend growth (%)	μ^d	1.3
Persistence shock coefficient	ϕ	3.2
Volatility of dividend growth	φ_d	3.2
Persistence Parameter in x_t	ρ	.94
Volatility of persistence shock (%)	φ_e	9
Mean idiosyncratic shock (%)	\bar{y}^2	6.1 ²
Correlation between g_t and g^d	$\rho_{\eta,u}$.3
Subjective discount factor	δ	.91
Relative risk aversion	γ	10
MA(1) coefficient in y_t^2	θ	0.45
Constant in λ_t	c	0.268

Annual values. $\alpha = \frac{1}{2}\gamma(\gamma + 1) = 55$.

The average standard deviation of the idiosyncratic shock is 6.1%. This is well below the cross-sectional variation reported in Carrol (1992) using PSID data and in Cogley (2002) using CEX data. We choose 0.3 as the correlation between consumption growth and dividend growth. There are different estimates of this correlation. Campbell and Cochrane use 0.2; Bansal and Yaron use 0.55. We found that values in this range do not alter the results significantly. The subjective discount factor is chosen to match the risk free rate of 1.5%. The risk aversion parameter we use, $\gamma = 10$, is admittedly high, although some estimates do exceed 10 (see Parker and Julliard (2005) for example). A high value of the risk aversion parameter helps to get a high equity premium and, in this model, to get a low risk free rate. Notice that

high risk aversion does not result in a high and volatile risk-free rate in our model. We chose 10 because it is the highest value in the range considered plausible by Mehra and Prescott (1985), and at the same time it is a lower bound among the models that succeed in matching similar dimensions of the data. Section 5.4 provides further discussion on risk aversion. For the MA parameter θ , since it affects the stationary distribution of y_t^2 , we pick it so that we can control the tails of the stationary distribution and prevent y_t^2 from being negative, while allowing for variation in the Sharpe ratio. To give time variation to the Sharpe ratio, we need the conditional variance of y_t^2 to be time varying. We keep the variance of y_t^2 small to avoid negative values and implausibly high positive ones. The unconditional distribution of y_{t+1}^2 is summarized on the right hand panel of Table 2. We argue below that the unconditional distribution of y_{t+1}^2 lies on the conservative side of the evidence found from panel data.

4.3. Implied Processes.

Statistics for the simulated dividend process are reported in Table 2. It's clear that the time series properties of our generated process are similar to the ones found in the data. The autocorrelation functions in the data are more complex and show some negative signs. But these are not significant, and could be matched if we gave a weak lag structure to u_t in the dividend process rather than a simple i.i.d. Complex time series dynamics are just not detected in the data for the dividend process. The autoregression coefficient $\rho_{AR(1)}$ of the simulated series is even lower than in the data, hence our parameterization for dividends does not impose an unrealistic persistence (at least as measured by a standard AR(1) model). Also, based on the data, we would not be able to reject the AR(1) parameters implied by the simulated series.¹⁰

The right panel of Table 2 shows some summary statistics for the stationary distribution of y_t^2 . The mean of the generated process is 6.2%. We report a 95% confidence interval for y (not y^2), which represents the standard deviation of the idiosyncratic shock, to get an idea of the implications of the generated distribution in terms of cross-sectional inequality. With 95% probability, the cross-sectional standard deviation of individual shocks is between 1.1% and 8.7%. These are definitely not extreme values considering panel data evidence on cross-sectional variation in income and consumption (see for example Cogley, 2002 or Carrol, 1992). Carrol, 1992 uses a value of 10%, after adjusting for measurement error, in his study of precautionary savings. Using CEX data, Cogley finds values for the cross-sectional variation on the order of 35-50%! He also warns us of the high measurement error in the data. A value of 6.2% corresponds to assuming that only 2.5% of the overall variation found in the data is true cross-sectional variation in consumption growth, while 97.5% is due to measurement error in the CEX data. Cogley's conclusion that cross-sectional variation can account for at most

¹⁰The AR(1) is a common specification for the law of motion of dividend growth.

TABLE 2. Implied g^d and y^2 Processes

Statistic	Model	Data		y^2 process		
		Estimate	2×S.E.			
μ^d	1.3%	1.0%			\bar{y}^2	$(6.2\%)^2$
$\sigma(g^d)$	9.5%	11.6%			$\sigma(y^2)$.00186
AC(1)	.08	.13	.19		CI_y^U	8.7%
AC(5)	.07	-.02	.19		CI_y^L	1.1%
AC(10)	.04	.15	.2		$c_{AR(1)}$.005
$c_{AR(1)}$.012	.009	.022		$\rho_{AR(1)}$	-.25
$\rho_{AR(1)}$.085	.123	.188		$\rho(y^2, g)$	-.69
$\rho(g, g^d)$.30	.28				

Notes: μ^d and $\sigma(g^d)$ are the mean and standard deviation of dividend growth. $AC(j)$ is the j -th autocorrelation. $c_{AR(1)}$ and $\rho_{AR(1)}$ are the constant and autoregression coefficients in the AR(1) model. $\rho(g, g^d)$ is the correlation between per-capita consumption growth and dividend growth. \bar{y}^2 is the average cross sectional variance in the model, $\sigma(y^2)$ is the unconditional standard deviation of y_t^2 , and $\rho(y^2, g)$ is the correlation between the cross sectional variance of consumption growth (y^2) and per-capita growth g . The data used is the annual dataset 1890-2001 on consumption and the S&P 500 dividend downloadable from Robert Shiller's website.

2% of the equity premium comes out of our model too. In the model of section 3 we find that our parameterization implies a premium of about 1% coming from uninsurable consumption shocks. In other words, the amount of cross-sectional variation imposed in the model is not the main risk factor here.

5. SIMULATION RESULTS

We simulate a history of 100,000 draws from our economy and calculate sample statistics that we then compare to the post war sample and long sample moments from Robert Shiller's data on the S&P 500 (1871-2001) and consumption (1889-2001).¹¹ Table 3 summarizes the comparison. We then look at the time series properties of the simulated price-dividend ratio in Table 5 and at the long horizon predictability of returns in Table 6.

5.1. Matching Moments.

We match the low risk free rate by choice of parameters. Our parameterization also implies a high equity premium, which in this economy is 6%. This is the level that Mehra and Prescott tried to match. Notice that there is some in-sample variation in excess returns (5.7% and

¹¹The data can be downloaded at <http://www.econ.yale.edu/~shiller/data.htm>

TABLE 3. Moments of Simulated and Historical Data.

Statistic	Model	Long Sample*	Postwar Sample*
$\exp(E[p - d])$	20.92	22.7	28.0
$\sigma(p - d)$.24	.35	.40
$\sigma_a(p - d)$.24	.26	.29
$\sigma_{93}(p - d)$.24	.26	.26
$E_t[R - R^f]$	6.0%	5.7%	7.2%
$\sigma(R - R^f)$	18%	18%	15.0%
Sharpe Ratio	.33	.32	.46
$\rho(g, r)$.57	.003	.02
$\rho(g, p - d)$.34	.16	.20

Notes: The Long Sample is 1871-2001, and the Postwar Sample is 1950-2001. $\exp(E[p - d])$ is the average price dividend ratio. $\sigma(p - d)$ is the standard deviation of the log price-dividend ratio. $\sigma_a(p - d)$ is the standard deviation of the log dividend-price ratio adjusted for repurchases as in Appendix B and $\sigma_{93}(p - d)$ is the standard deviation using data up to 1993. $R - R^f$ is the excess return on the stock market relative to the risk free interest rate. $\rho(g, r)$ is the correlation between consumption growth and log returns, and $\rho(g, p - d)$ is the correlation between consumption growth and the log price-dividend ratio.

7.2% in the two samples); our value of 6% indicates that the model can explain most of the equity premium. The model also generates a standard deviation that is consistent with the data, 18%, and therefore a Sharpe ratio consistent with the data.

The level of the price-dividend ratio is 20.9, close to the long sample value, but lower than the postwar average. The standard deviation of the price-dividend ratio is also lower than the level observed, 24% versus 35% and 40% in the two samples. Competing explanations, like Campbell and Cochrane (1999) and Bansal and Yaron (2004), generate a volatility of 26% and 21% respectively. It is of interest to notice that the higher observed variation is entirely due to the period 1994-2001, even in the century long sample. The value of the standard deviation for both long and post-war samples up to 1993 is only 26%.

There is empirical evidence that the price-dividend ratio shifted to a new regime over the nineties (see Lettau, Ludvigson and Watcher, 2003; Boudoukh et al., 2003), which could explain a higher mean and a higher variance in the 1990's. In particular, Boudoukh et al. study the role of stock repurchases on the price-dividend ratio. They show that a constructed payout ratio (i.e., price divided by dividend plus repurchases) has the same time series properties as the price-dividend ratio prior to the 1990's, but much different during that decade. The

values under $\sigma_a(p - d)$ in the table refer to the standard deviation of the price-dividend ratio adjusted to take stock repurchases into account. They are 26% and 29% respectively for the two samples. We discuss repurchases and the adjustment made to the price-dividend ratio in section 5.3 and Appendix B. This seems to explain almost entirely both the higher volatility and the higher mean in the late nineties.

The model has a consumption growth-excess return correlation of 57%. Though high, this is far lower than the almost perfect correlation that Mehra and Prescott’s economy or most consumption based asset pricing models imply, and without accounting for time aggregation in our simulations. Campbell and Cochrane (1999) simulate their model at monthly frequencies and then time aggregate the resulting variables to compare with annual data. By doing so, they reduce their correlation from 79% to 40%. In our model, consumption growth is conditionally strongly correlated with returns because recessionary shocks affect prices through the persistent component of dividends. Unconditionally, though, the correlation decreases because price-dividend ratios are very persistent. This is reflected in a correlation between the price-dividend ratio and consumption growth of 34%, which is mostly due to the assumed correlation of 30% between dividend growth and consumption growth. Time aggregation would further reduce the conditional correlation between consumption growth and price-dividend ratio, which in turns would degrade the unconditional correlation between consumption growth and returns. So we do not feel that the value of 57% is too high.

TABLE 4. Comparison of Models

Statistic	Benchmark	Model 1*	$\phi = 0$
$\exp(E[p - d])$	20.92	25.67	39.16
$\sigma(p - d)$.24	.24	.04
$E_t[R - R^f]$	6.0%	5.1%	2.8%
$\sigma(R - R^f)$	18%	18.16%	9.0%
Sharpe Ratio	.33	.28	.29
$\rho(g, r)$.57	.57	.21
$\rho(g, p - d)$.34	.34	-.07

Notes: The Benchmark model is the model in Table 3. Model 1 is the model of Section 3, without time-varying uncertainty. The last column, under the heading $\phi = 0$ shows results from a model in which the “Barsky De Long” component is shut down.

Table 4 compares our benchmark simulation with the model of section 3 (Model 1) and a simulation in which the Barsky and De Long component to dividend growth is shut down, i.e., $\phi = 0$.

As the table shows, the Barsky and De Long component is the one that generates the premium and the volatility of the price-dividend ratio. Model 1 matches the moments almost as well as our benchmark simulation, while removing the Barsky De Long component destroys the model's ability to match the data. Time-varying uncertainty is important in replicating the observed predictability of excess returns, as shown below.

5.2. Autocorrelations and Return Predictability.

The next two tables present autocorrelations of simulated variables and long horizon regressions. The price-dividend ratio is very persistent and slowly mean reverting in the data. Persistence and the source of persistence is at the heart of the return predictability and excess volatility literature.

The simulated price-dividend ratio has a first autocorrelation of .91, very close to the data estimate of .84-.86. The .91 is a little lower than the given persistence to the Barsky-De Long component in dividend growth, which is .94, and it's most likely due to the slight negative persistence implied by the variance of idiosyncratic shocks y_{t+1}^2 . Persistence declines at longer lags, as in the data. The two lower panels of Table 5 show autocorrelations of log excess returns and the partial sums of the autocorrelations. The model excess returns have very low, though negative, autocorrelation values, which imply univariate mean reversion. The mean reversion seems to be much stronger in the data, though the data is not very informative about the sign and magnitude of individual coefficients. Our model cannot generate complex dynamics, as all the stochastic processes are basically first order autoregressive. In order to match autocorrelation coefficients in the data, one should impose a more complex lag structure to the model.

The next table, Table 6, shows results from long horizon regressions, i.e. we regress cumulative log excess returns, $r_{t,t+1}^e + r_{t,t+2}^e \cdots + r_{t,t+j}^e$, on $p_t - d_t$, the log of the price-dividend ratio at time t . We do this for $j = 1, \dots, 5$. Both the R^2 from the regressions and the size of the coefficients are very close to the ones observed in the data. The sign is the right one, an increase in the price-dividend ratio forecasts lower future excess returns, as in Campbell and Shiller (1988), and the R^2 increases with the horizon. Recall that the first model (section 3) implies constant expected excess returns, hence it could not generate this kind of predictability. Both coefficients and R^2 are zero in that model.

As Table 2 shows, we gave little time variation to the standard deviation of y_t^2 , λ_t , in order to constrain the process for y_t^2 within plausible bounds. Table 6 shows that this time variation is enough to reproduce the return predictability in the data.

5.3. Model Implications from Historical Consumption Data.

We take the consumption series from Shiller's data (1890-2001) and use the parameterization

TABLE 5. Autocorrelations of Simulated and Historical data

Variable and Source	Lag (Years)				
	1	2	3	5	7
<i>p - d</i>					
Model	.91	.83	.78	.68	.59
Long Sample	.84	.70	.60	.42	.35
Post War	.86	.69	.56	.30	.13
<i>r - r^f</i>					
Model	-.01	-.02	.01	.00	.02
Long Sample	.08	-.19	.08	-.17	.13
Post War	.05	-.16	.17	-.02	.17
$\sum_{i=1}^j \rho(r_t^e, r_{t-i}^e)$					
Model	-.01	-.04	-.03	-.02	.01
Long Sample	.08	-.11	-.02	-.29	-.12
Post War	.05	-.12	.05	.27	.36

Notes: $r - r^f \equiv r^e$ is log excess return, and $\rho(r_t^e, r_{t-j}^e)$ is the j -th autocorrelation of log excess returns.

TABLE 6. Long-Horizon Regressions

Horizon	Model		Long Sample		Post War	
	$10 \times \beta$	R^2	$10 \times \beta$	R^2	$10 \times \beta$	R^2
1	-1.21	.12	-3.4	.08	-5.6	.24
2	-1.83	.20	-4.8	.13	-8.3	.36
3	-1.95	.26	-6.9	.20	-11.1	.49
4	-2.13	.37	-10.4	.31	-19.2	.57
5	-2.04	.44	-13.5	.44	-26.6	.64

Notes: The dependent variable is $r_{t,t+1}^e + r_{t,t+2}^e \cdots + r_{t,t+j}^e$, where j is the horizon. The independent variable is $p_t - d_t$.

in (4.1) to calculate the series of η_t . We then feed the η_t shocks to the model, which produces a series of price-dividend ratios (so no stock market data is used in this simulation). We adjust model predictions for the period 1972-2000 to take into account that repurchases have become an important source of payout for U.S. corporations using the data of Table I in Grullon and Michaely (2002). This makes model predictions comparable with the observed price-dividend ratio. Details of the adjustment are in Appendix B. Figure 1 presents a comparison between

data on the S&P 500 (from Shiller's dataset), and the price generated by the model using consumption data. Figure 2 shows the behavior of the implied price-dividend ratio versus the real S&P 500 price series.

Clearly, given the simple nature of the model, we do not expect the implied price to match the data point by point, but the impression from the figures is that the model does a satisfactory job at catching some of the main events in the stock market. The correlation coefficient between the simulated price-dividend series and the observed price-dividend ratio is 72%.¹²

Successive declines in consumption drive down future expected growth. As a consequence, prices fall. The model thus accounts for the decline in prices relative to dividends in the sharp recession of 1908 and in the post WWI period. The model also captures the rise and then decline of the stock market in the early 20th century, 1901-1918, the roaring 20's and the Great Depression. It predicts an even larger drop in 1932 than actually happened, so extreme was the drop in consumption growth. Then the model tracks the recovery during WWII, and the consumption and stock market boom of the sixties, though with a lag of about 5 years. It also tracks the poor performance of the 70's and the recovery of 1982-87.

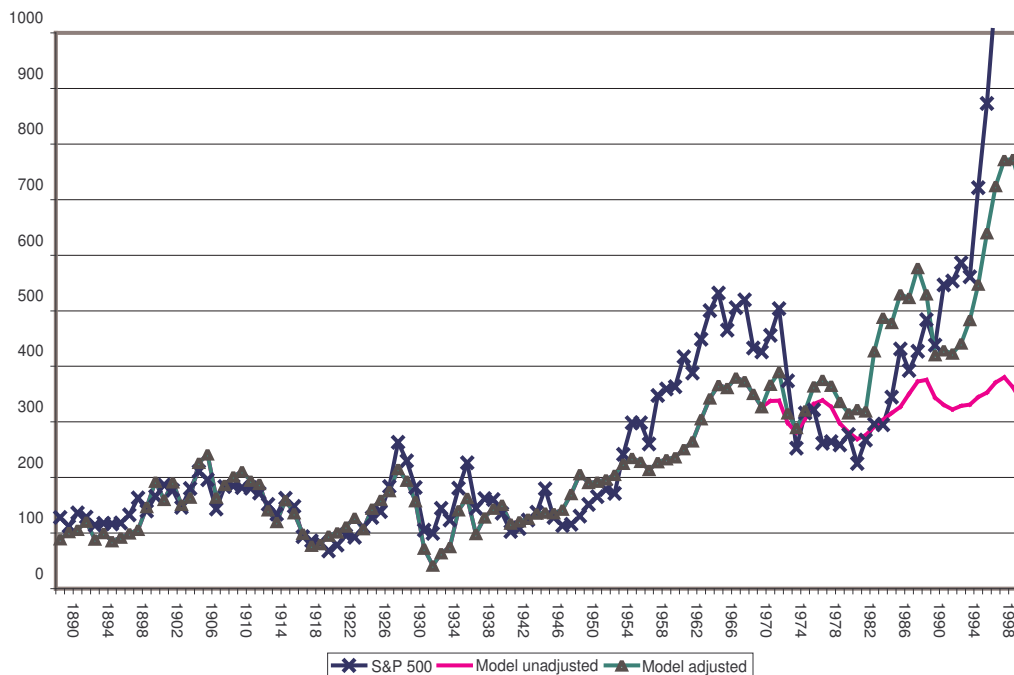
The worst performance of the unadjusted series (i.e., without adjusting for repurchases) is from the second half of the eighties and during the nineties when prices hit unprecedented levels at high speed. The unadjusted series moves with the adjusted series, but it does not vary enough. This is a common feature of models that try to explain stock market behavior using consumption data (see Campbell and Cochrane, 1999). Consumption growth was very smooth over the nineties, there was not a series of large positive shocks that is needed to move prices upward relative to dividends in the model. More generally, the smoothness of consumption applies to the entire second half of the figure, say after 1950. Consequently, model predictions after 1950 are much smoother relative to the S&P movements, than before 1950.

Figure 3 plots the consumption growth series used in our simulation, downloaded from Robert Shiller's web site. The standard deviation of the series is about 3%, but notice the huge difference in swings before and after 1950.

We find that we can reconcile part of the smoothness of consumption with greater swings in the price-dividend ratio if stock repurchases are taken into account. Our theoretical model's predicted price-dividend ratio should really be thought of as the predicted total payout rate, that is price divided by dividends plus repurchases. During most of our long sample, which runs from 1871 to 2001, repurchases were not significant. Hence the model's predicted payout ratio and predicted price-dividend ratio would approximately coincide. But in the 1980's

¹²The unadjusted series and the adjusted one have similar correlation coefficients with the S&P data up to 1990, 47% and 44% respectively. The coefficient of the adjusted series increases during the run-up in prices of the 1990's, when repurchases were highest.

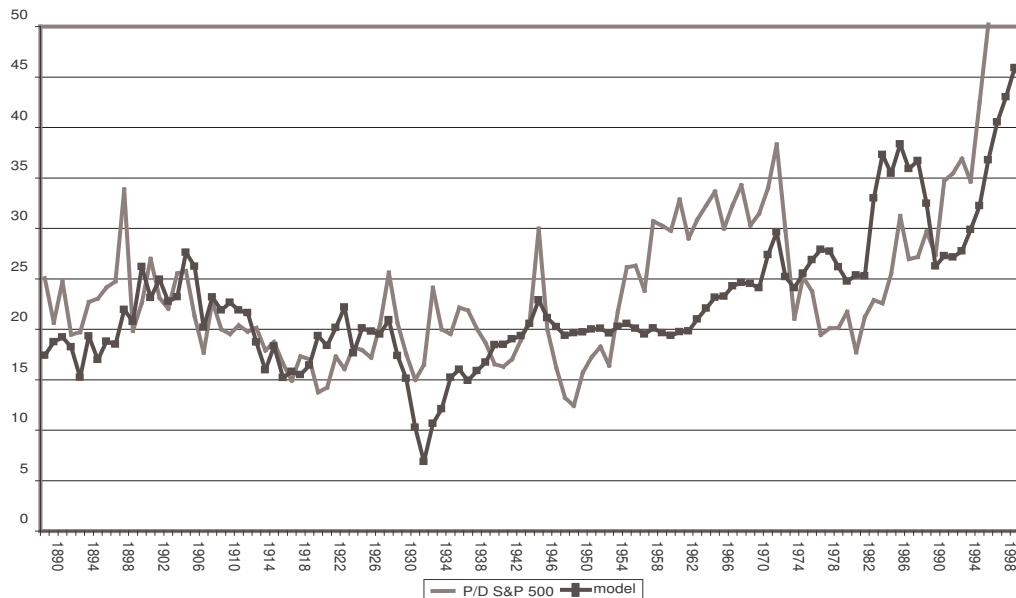
FIGURE 1. Historical S&P 500 Real price and Model Predictions



and 1990's repurchases became a significant fraction of total payout ratio to shareholders as documented in Grullon and Michaely (2002). While between 1972 and 1983 repurchases amounted to an average of 10.9% of dividend payments, between 1984 and 2000 repurchases were 57.7% of dividends, reaching a maximum of 113.11% in 2000. Most importantly, firms finance their share repurchases with funds that would otherwise be used to increase dividends. This means that the price-dividend ratio in the data is much larger than the true payout ratio towards the end of the sample, i.e., there was a regime shift. To take account of this shift, we adjust the model's predicted price-dividend ratio series as suggested by Boudoukh et al. using Grullon and Michaely's Table 1. Details are in Appendix B. The adjustment can account for most the movements in the eighties and nineties while it does not help at the end of sample, say 1997-2001.

As yet, it is not clear how and whether we can rationalize the level of prices in the late 1990's. Possible explanations include compositional shifts of consumption towards durable goods, or an increase in the consumption of stock market investors that is not captured in the data, due perhaps to demographic effects of the baby boom generation entering peak saving years. Eugene White (2004) argues that there is fair evidence of a bubble in the late 1990s. As Campbell (1999) points out:

FIGURE 2. Historical Price-Dividend Ratio and Model Predictions



The recent run-up in stock prices is so extreme relative to fundamental determinants such as corporate earnings, stock-market participation, and macroeconomic performance that it will be very hard to explain using a model fit to earlier historical data. (p.261)

While we leave open the possibility of a bubble in the late 1990s, we believe the model (and therefore fundamentals and rational expectations) can explain major movements in prices over the century-long sample.

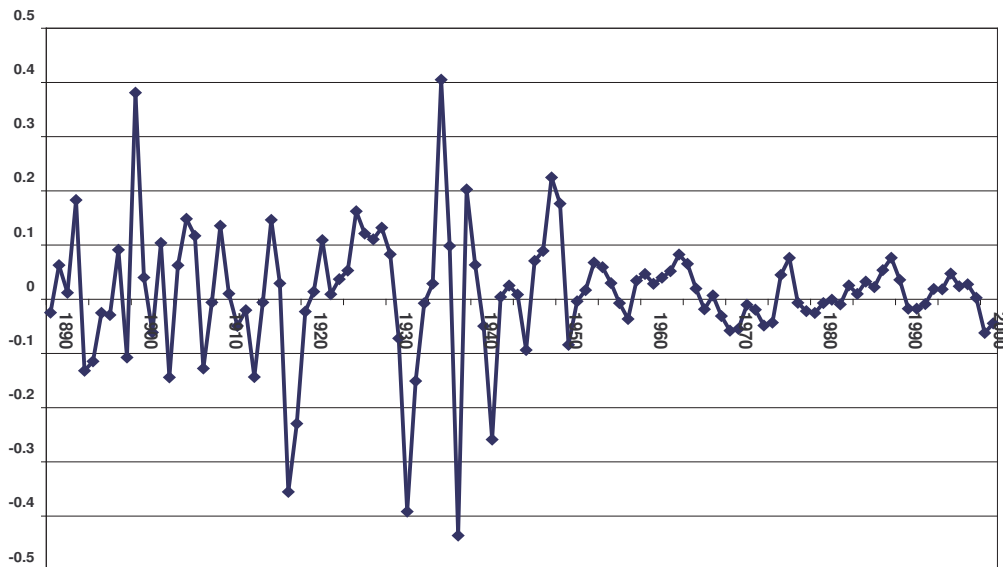
5.4. Risk Aversion.

The risk aversion coefficient used is 10, which is the upper bound of plausible parameters according to Mehra and Prescott (1985). The general view among economists is that 10 is too high (see Kocherlakota, 1996).

There is a sense in which a value of 10 can be justified for this economy and it comes from comparing our pricing equation with the one from a representative agent model. We can think of our economy as isomorphic to a representative agent economy in which the utility and the MRS's depend on y_t^2 . The representative agent will then exhibit high risk aversion, as he does in Campbell and Cochrane where risk aversion is as high as 80 at steady state.

Compared to 80, our economy shows that adding a simple form of heterogeneity to income shocks goes a long way toward reducing the level of risk aversion of individual agents. It is

FIGURE 3. Consumption Growth

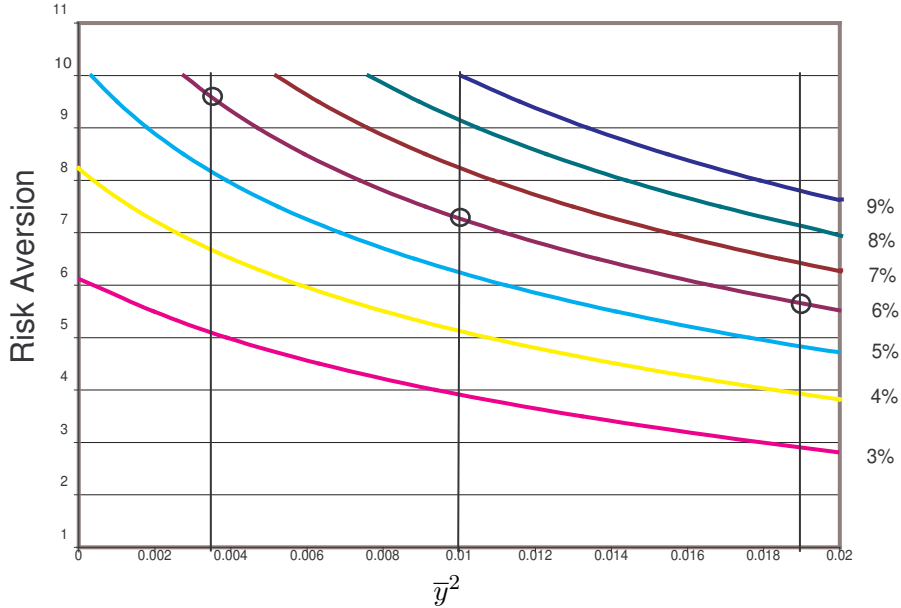


plausible to think that allowing for more general forms of heterogeneity could reduce further the level of risk aversion. It is unlikely though that such a model would yield a pricing function that depends only on aggregate variables.

Also notice that we have used a conservative specification for y_t^2 . If one finds higher values of cross-sectional variation plausible, equation (3.11) offers a trade off between the level of risk aversion and the level of uncertainty about idiosyncratic shocks (φ_y in (3.11) and λ_t in the model of section 4) for a given level of the premium. Figure 4 shows contour levels in $\gamma - \bar{y}^2$ space, i.e., for any fixed level of the equity premium, each contour line shows the combinations of risk aversion and \bar{y}^2 that generate it. These are calculated from our benchmark distribution for y_t^2 by increasing the mean, \bar{y}^2 , and the standard deviation, φ_y in (3.11), keeping the coefficient of variation constant.

For example, consider the vertical line through $\bar{y}^2 = 0.01$. This value for \bar{y}^2 corresponds to a mean value for y_t of 10%, and φ_y is calculated so that with 99% probability y_t is between 0 and 20%. When $\bar{y}^2 = 0.01$ and risk aversion $\gamma = 10$, the equity premium will equal about 9%, since the 9%-contour line crosses the vertical line at about $\gamma = 10$. Even if risk aversion falls to $\gamma = 5$, we would get a premium only slightly lower than 5% when $\bar{y}^2 = 0.01$. The vertical line through $\bar{y}^2 = 0.0038$ corresponds to our benchmark case, hence when $\gamma = 10$ the

FIGURE 4. Tradeoff Between Risk aversion and idiosyncratic risk



The circles correspond to the three alternative pairing of risk aversion and idiosyncratic risk that give a 6% equity premium, as discussed in the body of the paper. The circle on the upper left corresponds to the benchmark parameterization.

premium equals approximately 6%.¹³ Finally, the vertical line through $\bar{y}^2 = 0.019$ corresponds to assuming that with 95% probability, the cross sectional variance is less than 15% (close to Cogley's 2002 value), which implies a standard deviation of 38%. In other words, we make the 38% standard deviation a very unlikely event, a $2\text{-}\sigma$ event. In this case, when risk aversion γ lies between 5 and 6, the risk premium still lies between 5 and 6%.

Table 7 presents two simulations suggested by Figure 4. The first is with $\gamma = 7$ and a value of $\bar{y}^2 = .01$, and the second is with $\gamma = 5$ and $\bar{y}^2 = .019$. Figure 4 shows that these combinations should give us an equity premium close to 6%. This is confirmed in the table. The model can reproduce a high premium with a risk aversion as low as 5, while still matching other relevant moments.¹⁴

There are other possible modifications of the model that would go in the same direction as higher cross-sectional variance, and could allow a lower risk aversion for a given level of the premium. Consider the effect of a probability mass of consumption close to zero, as

¹³The contour lines are computed using the approximate solution (3.11) so the vertical line does not exactly crosses at 6% as we would expect from the benchmark simulations.

¹⁴Parameters used are the same as in table 1, with the exception of $\phi = 3.5$, $\varphi_d = 2.8$, $\rho = .925$, $\varphi_e = 10\%$, $\delta = .8$. These values slightly improve the fit of the simulation and do not alter the implied stochastic process of dividend growth sensibly.

TABLE 7. Lower Risk Aversion and Greater Heterogeneity

Statistic	Benchmark $\gamma = 10$	$\gamma = 7$	$\gamma = 5$
$\exp(E[p - d])$	20.92	22.5	24.8
$\sigma(p - d)$.24	.24	.28
$E_t[R - R^f]$	6.0%	5.6%	5.4%
$\sigma(R - R^f)$	18%	16.4%	17.8%
Sharpe Ratio	.33	.34	.30
$\rho(g, r)$.57	.79	.82
$\rho(g, p - d)$.34	.37	.36

Notes: We pair $\gamma = 10, 7, 5$ with $\bar{y}^2 = 0.00372, 0.01, 0.0184$ respectively. That is, higher levels of risk aversion are paired with lower levels of idiosyncratic risk. These levels of variance correspond to standard deviations of 6.1%, 10%, and 13.6% respectively.

argued in Carrol (1992), or a possible multi-dimensionality of individual shocks whose relative importance varies over the business cycle, as in Krebs (2003).

6. CONCLUSIONS

This paper tries to assess how far a model with standard preferences can go in interpreting aggregate stock market behavior. The model is not simple, in the sense that we have to add incomplete markets, a persistent component to dividend growth, and time varying uncertainty about the distribution of idiosyncratic shocks. Our position is that this model is as standard as the literature has gone in explaining the features that we try to replicate. Because preferences are standard, our explanation of price movements does not rely on psychologies that are hard to interpret and may not hold in the aggregate. We do not need liquidity constraints, or constraints on the supply of bonds. Our model is based on the idea that stocks are a bad insurance instrument relative to short term bonds; hence they perform badly in recessions, when the need to insure against consumption shocks is greatest. So, the consumption smoothing motive is the main factor driving major stylized facts about the stock market. To our knowledge, this is the first paper in the literature that matches relevant feature of the data with standard preferences.

Our model is close to Campbell and Cochrane's in the sense that the *MRS*'s have similar dynamics, as can be seen from the expression for the maximal Sharpe ratio. But the models are quite different. Asset pricing implications such as the equity premium, low risk free rate, and persistence of the price-dividend ratio in our model do not depend on the behavior imposed on y_t^2 . Our interpretation of the time-varying risk premium is also different from Campbell

and Cochrane. In our model the risk premium is time varying because of changing risk as opposed to changing risk aversion. In this sense, our interpretation is closer to Bansal and Yaron (2004). We also specify the dividend process as in Bansal and Yaron (2004). But otherwise the working of the model is very different.

As an example, the risk-free rate in Bansal and Yaron is low and has low variance because they assume the intertemporal elasticity of substitution is very high, 1.5.

There is a long tradition in economics, starting perhaps with Bernoulli who used log utility, to interpret risk aversion as decreasing marginal utility. Under this interpretation, the risk aversion parameter is necessarily the inverse of the intertemporal elasticity of substitution. If risk aversion is 10, then the intertemporal elasticity of substitution should be 0.1, a value consistent with Hall's 1988 findings. A value of 1.5 is then far from the standard psychological point of view.

How well our model performs depends on how much of the data one believes should be explained by an asset pricing model without frictions, and only using aggregate quantities. We claim that our model goes as far as one can go. We provide explanations for a low risk free rate, a high equity premium, high price volatility, persistence in the price-dividend process, return predictability, and low correlation between consumption-growth and excess-returns. These are the major features of aggregate stock market behavior, and the ones that have been most studied and we know most about. We are able to explain these features with a relatively low risk aversion coefficient. We believe 10 is high, indeed it is at the upper bound of the range that Mehra and Prescott proposed, but it is not extreme, and we show that lower values are possible if one is willing to increase the assumed cross-sectional variation in consumption growth. Bansal and Yaron use 10. Thus our value is the lowest we know of, for a model that can match the same features as we do. Campbell and Cochrane's model implies a steady state value of 80 for risk aversion, and it remains as high as 60 even at the lower bound.

Finally, we provide some evidence for Barsky and De Long's idea that the dividend process is much more important for price movements than simple autoregressive specifications would suggest. While we do not have production and investment in our economy, one possible interpretation is in terms of irreversible choice under uncertainty, which also explains why recessions are felt disproportionately in the durable goods sector. The persistent component in the expected dividend process may simply represent this fact. After all, shares represent ownership of firms, so it's hard to believe that dividend prospects are not related to investment opportunities.

APPENDIX A. SOLVING FOR THE EQUILIBRIUM PRICE-DIVIDEND RATIO

A.1. **Model 1.** Notice that the stochastic process in (3.5) is linear with constant variances and covariances. This implies that the variance term in the Euler equation (3.8)

$$\ln \delta - \gamma E_t g_{t+1} + \alpha E_t y_{t+1}^2 + E_t r_{t+1} + \frac{1}{2} \text{Var}_t(-\gamma g_{t+1} + \alpha y_{t+1}^2 + r_{t+1}) = 0$$

is constant. We can substitute for r_{t+1} using the log linear approximation (3.6). Rearranging terms, this gives

$$\ln \delta - \gamma \mu + \alpha \bar{y}^2 + \frac{c}{2} = -E_t[\kappa_0 + \kappa_1 v_{t+1} - v_t + \mu^d + \phi x_t],$$

where c is the constant variance. Solving for v_t

$$(A.1) \quad v_t = B + \kappa_1 E_t v_{t+1} + \phi x_t,$$

with $B = \ln \delta - \gamma \mu + \alpha \bar{y}^2 + c/2 + \mu^d + \kappa_0$. The equation (A.1) can be solved by recursive substitution using the fact $E_t x_{t+j} = \rho^j$ to yield

$$v_t = \underbrace{\frac{B}{1 - \kappa_1}}_{A_0} + \underbrace{\frac{\phi}{1 - \kappa_1 \rho}}_{A_1} x_t.$$

The variance

$$\text{Var}_t(-\gamma g_{t+1} + \alpha y_{t+1}^2 + r_{t+1}) = \text{Var}_t(-\gamma g_{t+1} + \alpha \bar{y}^2) + \text{Var}_t(r_{t+1}) + \text{Cov}_t(r_{t+1}, -\gamma g_{t+1} + \bar{y}_{t+1}^2)$$

The first variance equals $(\gamma \sigma + \alpha \varphi_y)^2$, while the second is

$$\begin{aligned} \text{Var}_t(r_{t+1}) &= \text{Var}(r_{t+1} - E_t r_{t+1}) = \text{Var}_t(\kappa_1(v_{t+1} - E_t v_{t+1}) + g_{t+1}^d - E_t g_{t+1}^d) \\ &= \text{Var}_t(\kappa_1 A_1 \varphi_e \sigma \eta_{t+1} + \sigma \varphi_d u_{t+1}) \\ &= (\kappa_1 A_1 \varphi_e \sigma)^2 + \sigma^2 \varphi_d^2 + 2\kappa_1 A_1 \varphi_e \sigma^2 \varphi_d \rho_{\eta,u}. \end{aligned}$$

So both variances are constant. Lastly, the covariance term is equal to the covariance between the innovations to the stochastic discount factor and returns, i.e.

$$\begin{aligned} \text{Cov}_t(-\gamma g_{t+1} + \bar{y}_{t+1}^2, r_{t+1}) &= \text{Cov}_t(-(\gamma \sigma + \alpha \varphi_y) \eta_{t+1}, \kappa_1 A_1 \varphi_e \sigma \eta_{t+1} + \sigma \varphi_d u_{t+1}) \\ &= -(\gamma \sigma + \alpha \varphi_y)(\kappa_1 A_1 \varphi_e \sigma + \sigma \varphi_d \rho_{\eta,u}), \end{aligned}$$

which is the equation of the constant premium in this economy.

A.2. **Model 2.** First we get rid of the shock u_t , since it does not determine state variables and it can be integrated out. Then we use quadrature-based rules to approximate the functional equation implied by the first order conditions. Notice now that the process (4.1)-(4.5) is a function of two shocks (u_{t+1}, η_{t+1}) , and two state variables $(x_t, \lambda_{t-1} \eta_t)$, or equivalently $(x_t, E_t y_{t+1}^2)$. In fact, it is enough to know x_t and $E_t y_{t+1}^2$ to know the distribution of the process (4.1)-(4.2). We have:

$$(A.2) \quad \begin{aligned} y_{t+1}^2 &= E_t y_{t+1}^2 - \sigma \lambda (E_t y_{t+1}^2) \eta_{t+1}, \\ E_{t+1} y_{t+2}^2 &= \bar{y}^2 + \theta \sigma \lambda (E_t y_{t+1}^2) \eta_{t+1}, \\ \lambda (E_t y_{t+1}^2) &= \frac{1}{\sigma \alpha} \sqrt{2[c - \alpha E_t y_{t+1}^2]}, \end{aligned}$$

so we can define the state vector $s_t = (s_{1t}, s_{2t}) \equiv (x_t, E_t y_{t+1}^2)$.

Consider the Euler equation for the price-dividend ratio from (3.4) and denote by s_t the vector of state variables. Then

$$(A.3) \quad \frac{P_t}{D_t}(s_t) = E \left[\underbrace{\delta \exp(-\gamma g_{t+1} + \alpha_t y_{t+1}^2) \left(1 + \frac{P_{t+1}}{D_{t+1}}\right) \exp(g_{t+1}^d)}_{h(s_{t+1})} \middle| s_t \right] \\ = E \left[E \left[h(s_{t+1}) \exp(g_{t+1}^d) \middle| \eta_{t+1}, s_t \right] \middle| s_t \right]$$

by the law of iterated expectations. Given η_{t+1} and s_t , s_{t+1} is measurable, hence we have:

$$\frac{P_t}{D_t}(s_t) = E \left[h(s_{t+1}) E \left[\exp(g_{t+1}^d) \middle| \eta_{t+1}, s_t \right] \middle| s_t \right].$$

We can solve the integral $E \left[\exp(g_{t+1}^d) \middle| \eta_{t+1}, s_t \right]$ using the normality assumption:

$$E \left[\exp(g_{t+1}^d) \middle| \eta_{t+1}, s_t \right] = \exp \left\{ E[g_{t+1}^d \middle| \eta_{t+1}, s_t] + \frac{1}{2} \text{Var}(g_{t+1}^d \middle| \eta_{t+1}, s_t) \right\}$$

with

$$E[g_{t+1}^d \middle| \eta_{t+1}, s_t] = \mu^d + \phi x_t + \sigma \varphi_d \rho_{\eta, u} \eta_{t+1} \\ \text{Var}(g_{t+1}^d \middle| \eta_{t+1}, s_t) = \sigma^2 \varphi_d^2 (1 - \rho_{\eta, u}^2)$$

given that u_t and η_t are jointly normal. This yields

$$\frac{P_t}{D_t}(s_t) = \underbrace{\exp \left\{ \mu^d + \phi x_t + \frac{1}{2} \sigma^2 \varphi_d^2 (1 - \rho_{\eta, u}^2) \right\}}_{c(s_t)} E \left[h(s_{t+1}) \exp(\sigma \varphi_d \rho_{\eta, u} \eta_{t+1}) \middle| s_t \right],$$

so we can write

$$(A.4) \quad \frac{P}{D}(s_t) = c(s_t) E \left[\delta \exp \left\{ -\gamma g_{t+1} + \alpha y_{t+1}^2 + \sigma \varphi_d \rho_{\eta, u} \eta_{t+1} \right\} \left(1 + \frac{P}{D}(s_{t+1}) \right) \right] (s_t),$$

which means that we need to integrate only over one dimension, η_{t+1} .

Denote $P/D(s_t)$ by $v(s)$, dropping the time subscript for convenience. We can write:

$$(A.5) \quad v(s) = c(s) \int K(s, s') (1 + v(s')) f(s'|s) ds'$$

$$(A.6) \quad = \int \psi(s, s') (1 + v(s')) f(s'|s) ds'$$

making the appropriate substitutions for $K(\cdot)$ and $\psi(\cdot)$. Define $\lambda(s, s') \equiv \psi(s, s')(1 + v(s'))$ and

$$I[\lambda](s) = \int \lambda(s, s') f(s'|s) ds' = \int \lambda(s, s') \frac{f(s'|s)}{\omega(s')} \omega(s') ds'$$

where $\omega(s')$ is a strictly positive weighting function. The integral can be approximated by the quadrature rule for $\omega(\cdot)$. Let s'_k and w_k , $k = 1, 2, \dots, N$, denote the abscissa and weights for an N point quadrature rule for the density $\omega(s')$. The approximation based on this rule to $I[\lambda](s)$ is

$$(A.7) \quad I_N[\lambda](s) = \sum_{k=1}^N \lambda(s'_k, s) \pi_k(s),$$

where

$$(A.8) \quad \pi_k(s) = \frac{f(s'_k|s)}{N(s)\omega(s'_k)} w_k$$

and

$$N(s) = \sum_{i=1}^N \frac{f(s'_i|s)}{\omega(s'_k)} w_i$$

so that the weights π_k sum up to unity. Estimating (A.6) at points $s_k = s'_k$ for $k = 1, \dots, N$ using $I_N[\lambda](s)$ gives

$$v_j = \sum_{k=1}^N (1 + v_k) \psi_{j,k} \pi_{j,k}, \quad j = 1, \dots, N$$

where $\psi_{j,k} = \psi(s'_j, s'_k)$, $\pi_{j,k}(s'_j)$ and $v_j = v(s'_j)$.

The $\{v_{j=1}^N\}$ are the solutions to the asset pricing equations if one views the law of motion of the state vector as a discrete Markov chain with range $\{s'_k\}$ and transition probabilities $\pi_{j,k} = P(s' = s'_k | s = s'_j)$. We choose the weight function $\omega(\cdot)$ to be the distribution of $(x_{t+1}, E_{t+1}y_{t+2}^2)$ conditional on the steady state values of the variables, $(x_0, y_0^2) = (0, \bar{y}^2)$, so $\omega(s') = f(s'|s_0)$.

APPENDIX B. ADJUSTING FOR REPURCHASES

We solve for the price-dividend ratio as detailed in Appendix A. Denote by P_t/D_t the observed S&P 500 price-dividend ratio, and by v_t^* the model's payout ratio. We want to compute a price-dividend ratio using model predictions so that it is comparable to P_t/D_t . Denote this model-generated price-dividend ratio as v_t , so that we can write

$$v_t = a_t v_t^*,$$

for some value of a_t . If there were no repurchases, a_t would be one all the time. But with repurchases, $a_t > 1$, i.e., the price-dividend ratio should be greater than the payout ratio. The figures discussed in section 2 assume that $a_t = 1$ for the sample 1891-1971, while for the sample 1972-2001 we calculate

$$a_t = \frac{R_t^p}{D_t} + 1,$$

where R_t^p is expenditure on repurchase of common stock. Notice that the total payout ratio is $\frac{R_t^p + D_t}{P_t}$. If R_t^p is zero, $a_t = 1$. We compute R_t^p/D_t for the period 1973-2001 using data from the column denoted $\sum_i REPO / \sum_i DIV$ in Table I of Grullon and Michaely. The assumption is that their sample is representative enough so that the same a_t is applicable to the S&P 500. The figures in the text report v so calculated.

APPENDIX C. IDIOSYNCRATIC SHOCKS AND EQUILIBRIUM

Individual consumption depends both on labor income and the return from a portfolio of assets. Labor income (I_t^i in (3.1)) is defined by

$$I_t^i = \delta_t^i C_t - D_t,$$

where δ_t^i is the individual shock to labor income. Since aggregate consumption satisfies $C_t = I_t + D_t$, an infinite number of agents is needed so that a law of large numbers can be

applied to yield $\sum_i \delta_t^i = 1$ at each point in time. For idiosyncratic shocks to be relevant, they have to be non-stationary. A common feature of earlier models with uninsurable income, like Lucas (1994), Telmer (1993) is that the time series of the ratio of each consumer's labor income to aggregate labor income I_t^i/I_t is stationary. With low persistence, consumers are able to come close to the complete-market rule of complete risk sharing. Meghir and Pistaferri (2004) provide evidence of a significant martingale component in households' earning processes using data from the PSID. Further, they show that the variance of the idiosyncratic shock is related to the business cycle. In Constantinides and Duffie, the process δ_t^i is the following martingale:

$$\delta_t^i = \exp \left\{ \sum_{s=1}^t \left(\eta_s^i y_s - \frac{y_s^2}{2} \right) \right\}$$

y_t is the cross sectional standard deviation of consumption growth, and it depends on aggregates at t . The aggregates are determined first, then the shocks η_t^i are handed out. η_t^i is assumed to be standard normal $N(0, 1)$. Recall that for η normal $E[\exp(\eta k - (k^2/2))] = 1$, which implies that δ^i is a geometric martingale. Further, we have

$$(C.1) \quad \frac{\delta_{t+1}^i}{\delta_t^i} = \exp\{\eta_{t+1}^i y_{t+1} - \frac{1}{2} y_{t+1}^2\}.$$

Constantinides and Duffie (1996) prove that there exist a unique equilibrium with no trade in this economy.¹⁵ Under no-trade and incomplete markets, any agent's marginal rate of substitution is a valid stochastic discount factor, as is any weighed average of the agents' marginal rates of substitution. The cross sectional average of marginal rates of substitution is (given any y_{t+1}^2)

$$\begin{aligned} E \left[\delta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \exp \left[-\gamma \left(\eta_{t+1}^i y_{t+1} - \frac{y_{t+1}^2}{2} \right) \right] \mid \mathcal{F}_t \cup C_{t+1} \cup y_{t+1}^2 \right] \\ = \delta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \exp \left\{ \frac{\gamma(\gamma+1)}{2} y_{t+1}^2 \right\}, \end{aligned}$$

where the expectation is taken over the cross sectional distribution at time $t+1$. Substituting this into the expression for P_t yields

$$P_t = E \left[\delta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \exp \left[\frac{\gamma(\gamma+1)}{2} y_{t+1}^2 \right] (P_{t+1} + D_{t+1}) \mid \mathcal{F}_t \right],$$

which is equation (3.4) in the body of the paper.

REFERENCES

- [1] Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59:1491–1509, 2004.
- [2] N. Barberis, M. Huang, and T. Santos. Prospect theory and asset prices. *Quarterly Journal of Economics*, 116:1–53, 2001.
- [3] Robert B. Barsky and Bradford J. De Long. Why does the stock market fluctuate? *Quarterly Journal of Economics*, 108:291–311, 1993.

¹⁵See the appendix in their paper for a formal proof.

- [4] Michele Boldrin, Lawrence Christiano, and Jonas Fisher. Habit persistence and asset returns in an exchange economy. *Macroeconomic Dynamics*, 1(2):312–32, 1997.
- [5] J. Boudoukh, Roni Michaely, Matthew Richardson, and Michael Roberts. On the importance of measuring payout yield: Implications for empirical asset pricing. manuscript, 26 pages, 2003.
- [6] Alon Brav, George Constantinides, and Christopher Geczy. Asset pricing with heterogeneous consumers and limited participation: Empirical evidence. *Journal of Political Economy*, 110:793–824, 2002.
- [7] J. H. Campbell and R.J. Shiller. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1:195–227, 1988.
- [8] J. H. Campbell and R.J. Shiller. Stock prices, earnings, and expected dividends. *Journal of Finance*, 43:661–676, 1988.
- [9] John Campbell. Comment. In Ben Bernanke and Julio Rotemberg, editors, *NBER Macroeconomics Annual*, pages 253–262. Cambridge: MIT Press, 1999.
- [10] John Y. Campbell. Consumption-based asset pricing. In George Constantinides, Milton Harris, and Rene Stulz, editors, *Handbook of the Economics of Finance*. Norton, 2003.
- [11] John Y. Campbell and John H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107:205–251, 1999.
- [12] Christopher D. Carroll. The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity*, (2):61–156, 1992.
- [13] John Cochrane. *Asset Pricing*. Princeton University Press, Princeton, New Jersey, 2001.
- [14] Timothy Cogley. Idiosyncratic risk and the equity premium: Evidence from the consumer expenditure survey. *Journal of Monetary Economics*, 49:309–334, 2002.
- [15] G. Constantinides and D. Duffie. Asset pricing with heterogeneous consumers. *Journal of Political Economy*, 104:219–40, 1996.
- [16] Angus Deaton and Christina Paxson. Intertemporal choice and inequality. *The Journal of Political Economy*, 102(3):437–467, 1994.
- [17] G. Grullon and R. Michaely. Dividends, share repurchases, and the substitution hypothesis. *Journal of Finance*, 57:1649–1684, 2002.
- [18] Robert Hall. Stochastic implications of the life cycle permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86:971–987, Dec 1978.
- [19] Robert Hall. Intertemporal substitution in consumption. *Journal of Political Economy*, 96:339–357, 1988.
- [20] Albert G. Hart. Risk, uncertainty and the unprofitability of compounding probabilities. In O. Lange, F. McIntyre, and T.O. Yntema, editors, *Studies in Mathematical Economics and Econometrics*, pages 110–118. 1942.
- [21] Urban Jermann. Asset pricing in production economies. *Journal of Monetary Economics*, 42:167–247, 1998.
- [22] Robert Jones and Joseph Ostroy. Flexibility and uncertainty. *Review of Economic Studies*, 51(1):13–32, 1984.
- [23] Narayana R. Kocherlakota. The equity premium: It’s still a puzzle. *Journal of Economic Literature*, 34:42–71, 1996.
- [24] Tom Krebs. Welfare cost of business cycles when market are incomplete. Brown University, manuscript, 28 pages, 2004.
- [25] Stephen Le Roy and Richard Porter. The present value relation: Tests based on variance bounds. *Econometrica*, 49:555–557, 1981.
- [26] Martin Lettau, Sidney C. Ludvigson, and Jessica A. Wachter. The declining equity premium: What role does macroeconomic risk play? manuscript, Dec 2003.
- [27] Deborah J. Lucas. Asset pricing with undiversifiable risk and short sales constraints: Deepening the equity premium puzzle. *Journal of Monetary Economics*, 34:325–341, 1994.
- [28] Robert E. Jr. Lucas. *Models of Business Cycles*. Blackwell, New York, 1987.
- [29] Costas Meghir and Luigi Pistaferri. Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32, 2004.
- [30] R. Mehra and E. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15:145–162, 1985.
- [31] Rajnish Mehra and Edward Prescott. The equity premium in retrospect. NBER Working Paper No. W9525, March 2003.

- [32] Yogo Motohiro. Estimating the elasticity of intertemporal substitution when instruments are weak. *The Review of Economics and Statistics*, 86(3):797–810, August 2004.
- [33] Jonathan Parker and Christian Julliard. Consumption risk and the cross section of stock returns. *Journal of Political Economy*, 113(1):185–222, February 2005.
- [34] Tom Rietz. The equity risk premium: A solution? *Journal of Monetary Economics*, 21:117–132, 1988.
- [35] N. G. Shepard and A. Harvey. On the probability of estimating a deterministic component in the local level model. *Journal of Time Series Analysis*, 11:339–347, 1990.
- [36] Robert Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71:421–436, 1981.
- [37] Jeremy J. Siegel. The shrinking equity premium. *Journal of Portfolio Management*, 26:10–17, 1999.
- [38] Chris Telmer. Asset-pricing puzzles and incomplete markets. *Journal of Finance*, 48:1803–32, December 1993.
- [39] Annette Vissing-Jørgensen. Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy*, 110(4):825–853, August 2002.
- [40] Philippe Weil. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24:401–421, 1989a.
- [41] Eugene N. White. Bubbles and busts: the 1990s in the mirror of the 1920s. Duke and the University of North Carolina Conference: Understanding the 1990s, the Long-Run Perspective, March 2004.