

The Declining Equity Premium: What Role Does Macroeconomic Risk Play?*

Martin Lettau
NYU, CEPR, NBER

Sydney C. Ludvigson
NYU and NBER

Jessica A. Wachter
Wharton and NBER

Comments Welcome

First draft: April 12, 2003

This draft: April 24, 2006

*Lettau: Department of Finance, Stern School of Business, New York University, 44 West Fourth Street, New York, NY 10012-1126; Email: mlettau@stern.nyu.edu; Tel: (212) 998-0378; <http://www.stern.nyu.edu/~mlettau>. Ludvigson: Department of Economics, New York University, 269 Mercer Street, 7th Floor, New York, NY 10003; Email: sydney.ludvigson@nyu.edu; Tel: (212) 998-8927; <http://www.econ.nyu.edu/user/ludvigsons/>. Wachter: The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104; Tel: (215) 898-7634; Email: jwachter@wharton.upenn.edu; <http://finance.wharton.upenn.edu/~jwachter/>. Ludvigson acknowledges financial support from the Alfred P. Sloan Foundation, and the C.V. Starr Center at NYU. We thank Ravi Bansal, Jacob Boudoukh, John Campbell, John Cochrane, Diego Comin, George Constantinides, Darrel Duffie, Robert Engle, Raquel Fernandez, Mark Gertler, Robert Hall, Lars Peter Hansen, John Heaton, Timothy Johnson, Pat Kehoe, Anthony Lynch, Ellen McGrattan, Stijn van Nieuwerburgh, Lubos Pastor, B. Ravikumar, Tom Sargent, Karl Walentin, Robert Whitelaw, seminar participants at the 2003 CIREQ-CIRANO-MITACS Conference on Macroeconomics and Finance, the NBER Economic Fluctuations and Growth fall 2003 meeting, the 2004 SED annual meeting, the July 2005 FRB conference on Financial Market Risk-Premiums, the Bank of England, the Federal Reserve Bank of St. Louis, NYU, London Business School, Princeton, SUNY Stony Brook, the University of Illinois, and Wharton for helpful comments. Any errors or omissions are the responsibility of the authors.

The Declining Equity Premium: What Role Does Macroeconomic Risk Play?

Abstract

Aggregate stock prices, relative to virtually any indicator of fundamental value, soared to unprecedented levels in the 1990s. Even today, after the market declines since 2000, they remain well above historical norms. Why? We consider one particular explanation: a fall in *macroeconomic risk*, or the volatility of the aggregate economy. Empirically, we find a strong correlation between low frequency movements in macroeconomic volatility and low frequency movements in the stock market. To model this phenomenon, we estimate a two-state regime switching model for the volatility and mean of consumption growth, and find evidence of a shift to substantially lower consumption volatility at the beginning of the 1990s. We then use these estimates from post-war data to calibrate a rational asset pricing model with regime switches in both the mean and standard deviation of consumption growth. Plausible parameterizations of the model are found to account for a significant portion of the run-up in asset valuation ratios observed in the late 1990s.

JEL: G12

1 Introduction

It is difficult to imagine a single issue capable of eliciting near unanimous agreement among the many opposing cadres of economic thought. Yet if those who study financial markets are in accord on any one point, it is this: the close of the 20th century marked the culmination of the greatest surge in equity values ever recorded in U.S. history. Aggregate stock prices, relative to virtually any indicator of fundamental value, soared to unprecedented levels. At their peak, equity valuations were so extreme that even today, after the broad market declines since 2000, aggregate price-dividend and price-earnings ratios remain well above their historical norms (Figure 1). More formally, the recent run-up in stock prices relative to economic fundamentals is sufficiently extreme that econometric tests for structural change (discussed below) provide evidence of a break in the mean price-dividend ratio around the middle of the last decade.¹

How can such persistently high stock market valuations be justified? One possible explanation is that the equity premium has declined (e.g., Blanchard (1993); Jagannathan, McGrattan, and Scherbina (2000); Fama and French (2002)). Thus, stock prices are high because future returns on stocks are expected to be lower. These authors focus less on the question of why the equity premium has declined, but other researchers have pointed to reductions in the costs of stock market participation and diversification (Heaton and Lucas (1999); Siegel (1999); Calvet, Gonzalez-Eiras, and Sodini (2003)).

In this paper, we consider an alternative explanation for the declining equity premium and persistently high stock market valuations: a fall in *macroeconomic risk*, or the volatility of the aggregate economy. To understand intuitively why macroeconomic risk can affect asset prices, consider the following illustrative example. By the law of one price, there exists a stochastic discount factor, or pricing kernel, M_{t+1} , such that the following expression holds for any traded asset with gross return R_t at time t :

$$E_t [M_{t+1}R_{t+1}] = 1, \tag{1}$$

¹The full run-up in valuation ratios cannot be attributed to shifts in corporate payout policies that have led many firms to substitute share repurchases for cash dividends. Although the number of dividend paying firms has decreased in recent years, large firms with high earnings actually increased real cash dividend payouts over the same period; as a consequence, aggregate payout ratios exhibit no downward trend over the last two decades (DeAngelo, DeAngelo, and Skinner (2002); Fama and French (2001)). See also Campbell and Shiller (2003). This, along with the evidence that price-earnings ratios remain unusually high, means that changes in corporate payout policies cannot fully explain the sustained high levels of financial valuation ratios.

where E_t denotes the expectation operator conditional on information available at time t . Suppose the pricing kernel and returns are jointly lognormal. Then it follows from (1) that the Sharpe ratio, SR_t , may be written

$$SR_t \equiv \max_{\text{all assets}} \frac{E_t [R_{t+1} - R_{f,t+1}]}{\sigma_t(R_{t+1})} \approx \sigma_t(\log M_{t+1}),$$

where $R_{f,t+1}$ is a riskless return known at time t , and $\sigma_t(\cdot)$ denotes the standard deviation of the generic argument “ \cdot ,” conditional on time t information. Fixing $\sigma_t(R_{t+1})$, the equity premium, in the numerator of the Sharpe ratio, is approximately proportional to the conditional volatility of the log pricing kernel.² In many asset pricing models, the pricing kernel is equal to the intertemporal marginal rate of substitution in aggregate consumption, C_t . A classic specification assumes there is a representative agent who maximizes a time-separable power utility function given by $u(C_t) = C_t^{1-\gamma}/(1-\gamma)$, $\gamma > 0$. With this specification, the Sharpe ratio may be written, to a first order approximation, as

$$SR_t \approx \gamma \sigma_t(\Delta \log C_{t+1}).$$

Thus, macroeconomic risk plays a direct role in determining the equity premium: fixing $\sigma_t(R_{t+1})$, lower consumption volatility, $\sigma_t(\Delta \log C_{t+1})$, implies a lower equity premium and a lower Sharpe ratio. Of course, this stylized model has important limitations, but its very simplicity serves to illustrate the crucial point: macroeconomic risk plays an important role in determining asset values. Below, we investigate these issues using a more complete asset pricing model.

Why underscore macroeconomic risk? There is now broad consensus among macroeconomists of a widespread and persistent decline in the volatility of real macroeconomic activity over the last 15 years. Kim and Nelson (1999) and McConnell and Perez-Quiros (2000) were the first to formally identify structural change in the volatility of U.S. GDP growth, occurring sometime around the first quarter of 1984. Blanchard and Simon (2001), using a different set of econometric tools, also find a large decline in output volatility over the last 20 years. Following this work, Stock and Watson (2002) subject a large number of macroeconomic time series to an exhaustive battery of statistical tests for volatility change. They conclude that the decline in volatility has occurred broadly across sectors of the aggregate economy. It appears in employment growth, consumption growth, inflation and sectoral output growth, as well as in GDP growth. It is large and it is persistent. Reductions in

²Conditioning (1) on time 0 information, the same expression can be stated in terms of unconditional moments.

standard deviations are on the order of 60 to 70 percent relative to the 1970s and 1980s, and the marked change seems to be better described as a structural “break,” or regime shift, than a gradual, trending decline. The macroeconomic literature is currently involved in an active debate over the cause of this sustained volatility decline.³

The subject of this paper is not the cause of the volatility decline, but its possible consequences for the U.S. aggregate stock market. Indeed, it would be surprising if asset prices were not affected by this fundamental change in the structure of macroeconomic volatility. Consistent with this, we find that volatility in consumption is highly correlated with fluctuations in the aggregate dividend-price ratio over longer horizons. This phenomenon is not merely a feature of postwar U.S. data, but is also present in postwar international data for 10 countries, and in prewar U.S. data.

Our main investigation contains two parts. In the first part, we employ the same empirical techniques used in the macroeconomic literature to characterize the decline in volatility of various measures of aggregate consumer expenditure growth in U.S. data. In the second part, we investigate the behavior of the stock market in a theoretical asset pricing model when empirically plausible shifts in macroeconomic risk are introduced.

The empirical part of this paper follows much of the macroeconomic literature and characterizes the decline in volatility by estimating a regime switching model for the standard deviation and mean of consumption growth. The estimation produces evidence of a shift to substantially lower consumption volatility at the beginning of the 1990s.

The theoretical part of our study investigates whether an asset pricing model that incorporates empirically plausible shifts in both the mean and volatility of consumption growth can account for the sharp run-up in aggregate stock prices during the 1990s. Using the preference specification developed by Epstein and Zin (1989, 1991) and Weil (1989), we study two types of asset pricing models. Our primary focus is on a learning model with regime switches in both the mean and standard deviation of consumption growth, calibrated to match our estimates from post-war data. We assume that agents cannot observe the regime but must infer it from consumption data; this learning aspect is an important feature of the model, discussed further below. Feeding in the (estimated) historical posterior probabilities of being in low and high volatility and mean states, we find plausible parameterizations of the model that can account for an important fraction of the run-up in price-dividend ratios observed in the late 1990s. The model’s predicted valuation ratios move higher in the 1990s because the long-run equity premium declines, a direct consequence of the persistent decline

³See Stock and Watson (2002) for a survey of this debate in the literature.

in macroeconomic risk in the early part of the decade. A shift to a higher mean growth state also plays a role in generating the model’s predicted run-up in equity values, but is far less important than the sharp decline in volatility. Finally, although the volatility of consumption declines in the 1990s, the model predicts that the volatility of equilibrium stock returns does not—consistent with actual experience.

The learning model just described has a number of appealing properties that are especially useful for studying the questions raised in this paper: It can be calibrated to empirical estimates of how long regimes are expected to last, and it provides a means of explicitly modelling expectations about future changes in regime, as well as a means of studying interesting transitional dynamics of regime shifts. A drawback of this model, however, is that its learning aspect makes the numerical solution quite intensive, requiring that the relation between consumption and dividends be modeled in a particularly simple manner. Thus, as a robustness check to our main findings, we present additional results in which the learning aspect of the model is eliminated, but consumption and dividends follow cointegrated processes. As in the learning framework, we find plausible parameterizations of the model that can account for significant fractions of the run-up in valuation ratios in the 1990s.

The literature has offered other possible explanations for the persistently high stock market valuations observed in the 1990s. One is an increase in the expected long-run growth rates of corporate earnings or dividends. The plausibility of this explanation has been questioned by academic researchers who point out that neither recent experience nor historical data provide any basis for the hypothesis (Siegel (1999); Jagannathan, McGrattan, and Scherbina (2000); Fama and French (2002); Campbell and Shiller (2003)). Other hypotheses include behavioral stories of “irrational exuberance” (Shiller (2000)), higher intangible investment in the 1990s (Hall (2000)), changes in the effective tax rate on corporate distributions (McGrattan and Prescott (2002)), the attainment of peak saving years during the 1990s by the baby boom generation (Abel (2003)), and a redistribution of rents away from factors of production towards the owners of capital (Jovanovic and Rousseau (2003)). We view the story presented here as but one of several possible contributing factors to the stock market boom of the 1990s. But in this paper we leave aside these alternative explanations in order to isolate the influence of declining macroeconomic volatility on the low-frequency behavior of stock prices at the end of the 20th century.

Two further points about our theoretical framework bear emphasizing. First, we show that the fraction of the 1990s equity boom that can be rationalized by declining macroeconomic volatility depends on the perceived persistence of the volatility decline. If the decline is

expected to be very persistent, almost all of the boom in asset prices can be explained; if the decline is expected to be more transitory, less of the boom can be rationalized through this mechanism. The data provide some guidance, and we discuss this extensively below. Second, we stress that our concern in this paper is not the short to medium term movements in equity valuations that may be attributable to cyclical fluctuations in the conditional (point-in-time) expected stock market return.⁴ Instead, we are interested in the ultra low-frequency movements in valuation ratios corresponding to possible low-frequency movements in the equity premium, what Fama and French (2002) call the “unconditional” equity premium. Thus, the model we present below is designed to illustrate the possible impact of a regime shift in macroeconomic volatility, not to explain high or medium frequency fluctuations in valuation ratios. As such, the model we explore is not designed to explain the full run-up in the price-dividend ratio in the 1990s (and their subsequent decline), but rather that portion of the run-up that has been sustained, leaving the level of the price-dividend ratio persistently above its previous historical norm consistent with the type of structural change documented in Table 1 below.

A number of existing papers use theoretical and empirical techniques related to those employed here to investigate a range of asset pricing questions. One group of papers investigates asset pricing when there is a discrete-state Markov switching process in the conditional mean of the endowment process (Cecchetti, Lam, and Mark (1990); Kandel and Stambaugh (1991); Cecchetti, Lam, and Mark (1993); Abel (1994); Abel (1999); Wachter (2003)), or in technology shocks (Cagetti, Hansen, Sargent, and Williams (2002)). None of these studies investigate the impact of regime switches in the *volatility* of the endowment process, however, the focus of this paper. Veronesi (1999) studies an equilibrium model in which the drift in the endowment process follows a latent two-state regime switching process and finds that such a framework is better at explaining volatility clustering than a model without regime

⁴A large literature finds that excess stock returns on aggregate stock market indexes are forecastable, suggesting that the conditional expected excess stock return varies. Shiller (1981), Fama and French (1988), Campbell and Shiller (1989), Campbell (1991), and Hodrick (1992) find that the ratios of price to dividends or earnings have predictive power for excess returns. Harvey (1991) finds that similar financial ratios predict stock returns in many different countries. Lamont (1998) forecasts excess stock returns with the dividend-payout ratio. Campbell (1991) and Hodrick (1992) find that the relative T-bill rate (the 30-day T-bill rate minus its 12-month moving average) predicts returns, while Fama and French (1988) study the forecasting power of the term spread (the 10-year Treasury bond yield minus the one-year Treasury bond yield) and the default spread (the difference between the BAA and AAA corporate bond rates). Lettau and Ludvigson (2001) forecast returns with a proxy for the log consumption-wealth ratio.

changes. Whitelaw (2000) also investigates an equilibrium economy with regime-switching in the mean of the endowment process, and he allows for time-varying transition probabilities between regimes. He finds that such a model generates a complex nonlinear relation between expected returns and volatility in the stock market. In contrast to these studies, Bonomo and Garcia (1994, 1996) and Driffil and Sola (1998) allow for regime changes in the variance of macroeconomic fundamentals, but their sample ends in 1985 and therefore excludes the switch to a prolonged period of record-low macroeconomic volatility in the 1990s that is the focus of this study. Otrok, Ravikumar, and Whiteman (2002) study the temporal distribution of consumption variance and its implications for habit-based asset pricing models. The study here, by contrast, focuses on low-frequency shifts in the overall level of volatility, rather than on shifts in its temporal composition. Other papers linking consumption volatility to movements in equity valuation ratios include Bansal and Lundblad (2002) and Bansal, Khatchatrian, and Yaron (2005). Bansal and Lundblad argue that there has been a fall in the global equity risk-premium, and explore a model in which this decline is associated with a fall in the conditional volatility of the world market portfolio return. Because the conditional volatility of the world market portfolio is a magnified version of the conditional volatility of world cash flow growth in their model, they indirectly link the decline risk-premia to a decline in the volatility of underlying fundamentals. By contrast, Bansal, Khatchatrian and Yaron focus more directly on the volatility of underlying fundamentals and construct quarterly measures of volatility for aggregate consumption based on parametric models including GARCH, and from the residuals of an autoregressive specifications for consumption growth. They find that quarterly price-dividend ratios are predicted by these lagged volatility measures, with R-squared statistics as high as 25 percent. Our analysis differs from both of these papers in that our emphasis is on the ultra low frequency movements in consumption risk that have become the subject of a large and growing body of macroeconomic inquiry, rather than on the cyclical stock market implications of quarterly fluctuations in conditional (point-in-time) consumption volatility.

The rest of this paper is organized as follows. In the next section we present empirical results documenting regime changes in the mean and volatility of measured consumption growth. We then explore their statistical relation with movements in measures of the price-dividend ratio for the aggregate stock market. Next, we turn to an investigation of whether the observed behavior of the stock market at the end of the last century can be generated from rational, forward looking behavior, as a result of the decline in macroeconomic risk. Section 3 presents an asset pricing model that incorporates shifts in regime, and evaluates

how well it performs in explaining the run-up in stock prices during the 1990s. Additional results from a model without learning but with cointegrated consumption and dividends are also presented. Section 4 concludes.

2 Macroeconomic Volatility and Asset Prices: Empirical Linkages

In this section we document the decline in volatility for consumer expenditure growth. We investigate the volatility decline in total per capita personal consumer expenditures (PCE). The series is in 1996 chain-weighted dollars. As has been argued elsewhere (e.g., Cecchetti, Lam, and Mark (1990)), the equilibrium model studied below—in which consumption equals output—is somewhat ambiguous as to the appropriate time-series for calibrating the endowment process. We use the broad PCE measure of consumption to calibrate the model, since it exhibits lower volatility at the beginning of the 1990s, by which time the vast majority of other macroeconomic time-series also exhibited a volatility decline (Stock and Watson (2002)). This is important because it helps to insure that the timing of the decline in volatility that we feed into our theoretical model is not driven by measurement error in any one series. This is important because individual series will be an imperfect measure of the relevant theoretical concept provided by our model, and we are interested in when agents could have plausibly inferred that macroeconomic volatility reached a new, lower regime. The Appendix at the end of this paper gives a complete description of the data and our sources. Our data are quarterly and span the period 1952:1 to 2002:4. We focus our primary analysis on postwar data because prewar data on consumption and output are known to be measured with significantly greater error that exaggerates the size of cyclical fluctuations in the prewar period (Romer (1989)).

We begin by looking at simple measures of the historical volatility of this series. Figure 2 provides graphical evidence of the decline in volatility. The growth rates of this series is plotted over time along with (plus or minus) two standard deviation error bands in each estimated volatility “regime,” where a regime is defined by the estimated two-state Markov switching process described below. (For the purposes of this figure, a low volatility regime is defined to be a period during which the posterior probability of being in a low volatility state is greater than 50 percent.) All figures clearly show that volatility is lower in the 1990s than previously.

Another way to see the low frequency fluctuations in macroeconomic volatility is to look at volatility estimates for non-overlapping five-year periods. Figure 3 (top panel) plots the standard deviation of consumption growth for non-overlapping five-year periods. There is a significant decline in volatility in the five-year window beginning in 1992, relative to the immediately preceding five-year window. In particular, the series is about one-half as volatile in the 1990s as it is in the whole sample. To illustrate how these movements in volatility are related to the stock market, this panel also plots the mean value of the log dividend-price ratio in each five year period.⁵ Our measure of the log dividend-price ratio for the aggregate stock market is the corresponding series on the CRSP value-weighted stock market index. The bottom panel of Figure 3 plots the same, but with the log earnings-price ratio in place of the log dividend-price ratio. The data for the price-earnings ratio is taken from Robert Shiller's Yale web site.⁶ The figure shows how these low frequency shifts in macroeconomic volatility are related to low frequency movements in the stock market.

Figure 3 exhibits a striking correlation between the low frequency movements in macroeconomic risk and the stock market: both volatility and the log dividend-price ratio (denoted $d_t - p_t$) are high in the early 1950s, low in the 1960s, high again in the 1970s, and then begin falling to their present low values in the 1980s. The correlation between consumption volatility and $d_t - p_t$ presented in this figure is 72 percent. A similar picture holds for the price-earnings ratio (bottom panel).

The correlations between high asset valuations and low volatility are present in countries other than the U.S. Figure 4 plots the volatility estimates for non-overlapping five-year periods, along with the mean value of the log dividend-price ratio in each five year period, for ten countries: Australia, Canada, France, Germany, Italy, Japan, the Netherlands, Sweden, Switzerland, and the United Kingdom. The international data on quarterly consumption and dividend-price ratios are from Campbell (2003), and are typically available over a shorter time period than for U.S. data.⁷ Figure 4 uses the longest available sample for each country. The figure shows that international data also display a striking correlation between the low frequency movements in macroeconomic risk and the national stock market for the respective

⁵Replacing the mean with mid-point or end-points of $d_t - p_t$ in each five year period produces a similar picture.

⁶<http://aida.econ.yale.edu/~shiller/data.htm>

⁷The dataset uses Morgan Stanley Capital International stock market data covering the period since 1970. Data on consumption are from the International Financial Statistics of the International Monetary Fund. With the exception of a few countries, starting dates for consumption data in each country range from 1970, first quarter to 1982, second quarter.

country. For every country, the figure exhibits a strong positive correlation between low frequency movements in macroeconomic risk and stock market valuation ratios, similar to that obtained for the U.S. Virtually every country also experiences a significant decline in macroeconomic volatility and increase in equity valuation ratios in the last decade of the century relative to earlier decades. The one exception is Australia, which displays no visible trend in either macroeconomic volatility or the stock market. Hence even for this observation, the correlation between macroeconomic volatility and the stock market is remarkable. More generally, Figure 3 and 4 tell the same story: for the vast majority of countries, the 1990s were a period of record-low macroeconomic volatility and record-high asset prices.

Moving back to U.S. data, Figure 5 shows that the strong correlation between macroeconomic volatility and the stock market is also present in prewar data. Although consistently constructed consumption data going back to the 1800s are not available, we do have access to quarterly GDP data from the first quarter of 1877 to the third quarter of 2002. The data are taken from Ray Fair's web site,⁸ which provides an updated version of the GDP series constructed in Balke and Gordon (1989). Figure 5 plots estimates of the standard deviation of GDP growth for non-overlapping ten year periods along with the mean value of the log dividend-price ratio in each ten year period, for whole decades from 1880 to 2000. The absolute value of GDP volatility in pre-war data must be viewed with caution. As mentioned, we focus our primary analysis on postwar data because volatility in this period is overstated relative to the postwar period due to greater measurement error, and because consistent data collection methodologies were not in place until the postwar period. What Figure 5 does reveal, however, is that the strong correlation between macroeconomic volatility and the stock market is not merely a feature of postwar data or of a single episode in the 1990s. Rather, it is present in over a century of data spanning the period since 1880.

To characterize the decline in macroeconomic volatility more formally, the macroeconomic literature has generally taken two approaches: (i) tests for structural breaks in the variance at an unknown date, and (ii) estimates from a regime switching model.⁹ We follow both

⁸<http://fairmodel.econ.yale.edu/RAYFAIR/PDF/2002DTBL.HTM>

⁹As mentioned, other methods of modelling changes in volatility, such as GARCH, may be appropriate for describing high frequency, stationary fluctuations in variance, but are not appropriate for documenting prolonged periods of moderated volatility like that observed at the end of the last century. For example, results not reported show that GARCH models do not generate the observed magnitude of volatility decline during this period. Intuitively, the GARCH model does a reasonable job of modelling changes in volatility *within* regimes, once those have been identified by other procedures, but does not adequately capture movements in volatility across regimes. All observations in a GARCH procedure are treated as having been generated

of these approaches here. Table 1 provides the results of undertaking structural break tests for the volatility of each consumption measure described above, and for the mean of the price-dividend ratio on the CRSP value-weighted index.¹⁰ Notice that these tests test the hypothesis of a permanent shift in the volatility or mean of the series in question. The top panel of Table 1 shows the results of a test for the break in the variance of consumption growth using the Quandt (1960) likelihood ratio (QLR) statistic employed by Stock and Watson (2002).¹¹ The null hypothesis of no break is tested against the alternative of one. The null hypothesis of no break in the variance is rejected at the 1% significance level for consumption. The break date is estimated to be 1992:Q1, with 67% confidence intervals equal to 1991Q3-1994Q4.¹² Note that these tests, unlike estimates from the regime switching model discussed below, are ex post dating tests that use the whole sample and are therefore not appropriate for inferring the precise timing of when agents would most likely have assigned a high probability of being in a new, low volatility regime. Nevertheless, they provide evidence of a persistent shift down in macroeconomic volatility in our sample and give us a sense of when that break may have actually occurred.

The bottom panel of Table 1 presents results from considering a sup F type test (Bai and Perron (2003)) of no structural break versus one break in the mean of the price-dividend ratio.¹³ The sup F test statistic is highly significant (with a p -value less than 1%), implying structural change in the price-dividend ratio. The break date is estimated to be 1995:Q1, with a 90 percent confidence interval of 1994:Q1 to 1999:Q3. The mean price-dividend ratio before the break is estimated to be 28.22; after the break, the mean is estimated to be 66.69, from a single distribution with a stationary variance, rather than from a mixture of distinct distributions with constant variances.

¹⁰See Lettau and Van Nieuwerburgh (2005) for a recent study of the affects of structural breaks on the forecasting power of the price-dividend ratio for excess returns.

¹¹This test also allows for shifts in the conditional mean, by estimating an autoregression that allows for a break in the autoregressive parameters at an unknown date.

¹²As Stock and Watson point out, the break estimator has a non-normal, heavy-tailed distribution that renders 95% confidence intervals so wide as to be uninformative. Thus, we follow Stock and Watson (2002) and report the 67% confidence intervals for this test.

¹³The linear regression model has one break and two regimes:

$$y_t = z_t \tau_j + u_t \quad t = T_{j-1} + 1, \dots, T_j,$$

for $j = 1, 2$, where y_t denotes the price-dividend ratio here, z_t is a vector of ones and the convention $T_0 = 0$ and $T_{m+1} = T$ has been used. The procedure of Bai and Perron (2003) is robust potential serial correlation and heteroskedasticity both in constructing the confidence intervals for break dates, as well as in constructing critical values for the sup F statistic for the test of the null of no structural change.

an over two-fold increase. It is interesting that the break date is estimated to occur after the estimated break dates for consumption volatility, consistent with the learning model we present below.

Next, we follow Hamilton (1989) and much of the macroeconomic literature in using our postwar data set to estimate a regime-switching model based on a discrete-state Markov process.¹⁴ This approach has at least two advantages over the structural break approach for our application. First, the structural break approach assumes that regime shifts are literally permanent; by contrast, the regime switching model provides a quantitative estimate of how long changes in regime are expected to last, through estimates of transition probabilities. Second, unlike the structural break estimates, the regime switching model allows one to treat the underlying state as latent, and provides an estimate of the posterior probability of being in each state at each time t , formed using only observable data available at time t . The estimates from this regime-switching model will serve as a basis for calibrating the asset pricing model we explore in the next section.

Consider a time-series of observations on some variable X_t and let x_t denote $\log X_t$. A common empirical specification takes the form

$$\begin{aligned} \Delta c_t &= \mu(S_t) + \phi(\Delta c_{t-1} - \mu(S_{t-1})) + \epsilon_t \\ \epsilon_t &\sim N(0, \sigma^2(V_t)), \end{aligned} \tag{2}$$

where S_t and V_t are latent state variables for the states of mean and variance and Δc_t denotes the log difference of consumption. We assume that the probability of changing mean states is independent of the probability of changing volatility states, and vice versa. To model the volatility reduction we follow the approach taken in the macroeconomic literature (e.g., Kim and Nelson (1999), McConnell and Perez-Quiros (2000)), by allowing the mean and variance of each series to follow independent, two-state Markov switching processes. It follows that there are two mean states, $\mu_t \equiv \mu(S_t) \in \{\mu_l, \mu_h\}$ and two volatility states, $\sigma_t \equiv \sigma(V_t) \in \{\sigma_l, \sigma_h\}$, where l denotes the low state and h the high state.¹⁵ Note that

¹⁴We focus on the larger U.S. dataset for this procedure, as it is known to require a large number of data points to produce stable results.

¹⁵Although a greater number of states could be entertained in principle, there are important practical reasons for following the existing macro literature in a two-state process. On the empirical side, more regimes means more parameters and fewer observations within each regime, increasing the burden on a finite sample to deliver consistent parameter estimates. On the theory/implementation side, we use these empirical estimates to calibrate our regime switching model discussed below. The two-state model already takes several days to solve on a work-station computer; a three-state model would more than double the

independent regimes do not imply that the mean and volatility of consumption growth are themselves independent. Even with a single volatility regime, the volatility of consumption growth would be higher in recessions than in booms, because the probability of switching regimes is higher in the low mean state than in the high mean state. Note also that the posterior regime probabilities inferred by theoretical agents observing data, as well as by the econometrician, are not independent.

We denote the transition probabilities of the Markov chains

$$\begin{aligned} P(\mu_t = \mu_h | \mu_{t-1} = \mu_h) &= p_{hh}^\mu \\ P(\mu_t = \mu_l | \mu_{t-1} = \mu_l) &= p_{ll}^\mu \end{aligned}$$

and

$$\begin{aligned} P(\sigma_t = \sigma_h | \sigma_{t-1} = \sigma_h) &= p_{hh}^\sigma \\ P(\sigma_t = \sigma_l | \sigma_{t-1} = \sigma_l) &= p_{ll}^\sigma \end{aligned}$$

where the probabilities of transitioning between states are denoted $p_{hl}^\mu = 1 - p_{ll}^\mu$ and $p_{lh}^\mu = 1 - p_{hh}^\mu$ for the mean state, and $p_{hl}^\sigma = 1 - p_{ll}^\sigma$ and $p_{lh}^\sigma = 1 - p_{hh}^\sigma$ for the volatility state. Denote the transition probability matrices

$$\begin{aligned} \mathbf{P}^\mu &= \begin{bmatrix} p_{hh}^\mu & p_{hl}^\mu \\ p_{lh}^\mu & p_{ll}^\mu \end{bmatrix}, \\ \mathbf{P}^\sigma &= \begin{bmatrix} p_{hh}^\sigma & p_{hl}^\sigma \\ p_{lh}^\sigma & p_{ll}^\sigma \end{bmatrix}. \end{aligned}$$

The parameters $\Theta = \{\mu_h, \mu_l, \sigma_h, \sigma_l, \phi, \mathbf{P}^\mu, \mathbf{P}^\sigma\}$ are estimated using maximum likelihood, subject to the constraints $p_{ij}^k \geq 0$ for $i = l, h$, $j = l, h$ and $k = \{\mu, \sigma\}$.

Let lower case s_t represent a state variable that takes on one of $2^3 = 8$ different values representing the eight possible combinations for S_t , S_{t-1} and V_t . Equation (2) may be written as a function of the single state variable s_t .

Since the state variable, s_t , is latent, information about the unobserved regime must be inferred from observations on x_t . Such inference is provided by estimating the posterior probability of being in state s_t , conditional on estimates of the model parameters Θ and observations on Δx_t . Let $\mathbf{Y}_t = \{\Delta x_0, \Delta x_1, \dots, \Delta x_t\}$ denote observations in a sample of size T based on data available through time t . We call the posterior probability $P\{s_t = j | \mathbf{Y}_t; \hat{\Theta}\}$, where $\hat{\Theta}$ is the maximum likelihood estimate of Θ , the *state probability* for short.

state space and would be computationally infeasible.

The estimation results are reported in Table 2. The regime represented by $\mu(S_t) = \mu_h$ has average consumption growth equal to 0.611% per quarter, whereas the regime represented by $\mu(S_t) = \mu_l$, has an average growth rate of -0.469% per quarter. Thus, the high growth regime is an expansion state and the low growth regime a contraction state. These fluctuations in the conditional mean growth rate of consumption mirror cyclical variation in the macroeconomy.

The volatility estimates give a sense of the degree to which macroeconomic risk varies across regimes. For example, the high volatility regime represented by $\sigma(V_t) = \sigma_h$, has residual variance of 0.541 per quarter, whereas the low volatility regime represented by $\sigma(V_t) = \sigma_l$ has the much smaller residual variance of 0.151 per quarter. This corresponds to a 47 percent reduction in the standard deviation of consumer expenditure growth. The results for GDP growth (not reported) qualitatively similar.

How persistent are these regimes? The probability that high mean growth will be followed by another high mean growth state is 0.971, implying that the high mean state is expected to last on average about 33 quarters. The volatility states are more persistent than the mean states. The probability that a low volatility state will be followed by another low volatility state is 0.992, while the probability that a high volatility state will be followed by another high volatility state is 0.995. This implies that the low volatility state reached in the 1990s is expected to last about 125 quarters, over 30 years. In fact, a 95% confidence interval includes unity for these values, so we cannot rule out the possibility that the low macroeconomic volatility regime is an absorbing state, i.e., expected to last forever. This characterization is consistent with that in the macroeconomic literature, which has generally viewed the shift toward lower volatility as a very persistent, if not permanent, break.

Figure 6 shows time-series plots of the smoothed and unsmoothed posterior probabilities of being in a low volatility state, $P(\sigma_t = \sigma_l)$, along with the smoothed and unsmoothed probabilities of being in a high mean state, $P(\mu_t = \mu_h)$.¹⁶ consumption exhibits a sharp increase in the probability of being in a low volatility state at the beginning of the 1990s. The probability of being in a low volatility state switches from essentially zero, where it resided for most of the post-war period prior to 1991, to unity, where it remains for the rest of the decade. Thus, the series shows a marked decrease in volatility in the 1990s relative to previous periods.

¹⁶ $P(\sigma_t = \sigma_l)$ is calculated by summing the joint probabilities of all states s_t associated with being in a low volatility state. $P(\mu_t = \mu_h)$ is calculated by summing the joint probabilities of all states s_t associated with being in a high mean growth state.

3 An Asset Pricing Model With Shifts in Macroeconomic Risk

The results in the previous section show that the shift toward lower macroeconomic risk coincides with a sharp increase in the stock market in the 1990s. We now investigate whether such a relation can be generated in a model of rational, forward-looking agents. To do so, our primary analysis considers an asset pricing model augmented to account for regime switches in both the mean and standard deviation of consumption growth, with the shifts in regime calibrated to match our estimates from post-war data. We study the implications of this model first, which posits a particularly simple “levered” relationship between consumption and dividends. The final part of this section considers a model without learning but with cointegrated consumption and dividends.

Modelling such shifts as changes in regime is an appealing device for addressing the potential impact of declining macroeconomic risk on asset prices, for several reasons. First, the macroeconomic literature has characterized the moderation in volatility as a sharp break rather than a gradual downward trend, a phenomenon that is straightforward to capture in a regime-switching framework (e.g., McConnell and Perez-Quiros (2000); Stock and Watson (2002)). Second, changes in regime can be readily incorporated into a rational, forward-looking model of behavior without regarding them as purely forecastable, deterministic events, by explicitly modelling the underlying probability law governing the transition from one regime to another. The probability law can be readily calibrated from our previous estimates from post-war consumption data. Third, the regime switching model provides a way of modelling how beliefs about an unobserved state evolve over time, by incorporating Bayesian updating. Finally, notice that the regime switching framework encompasses a structural break model as a special case, since the model is free to estimate transition probabilities that are absorbing states.

Consider a representative agent who maximizes utility defined over aggregate consumption. To model utility, we use the more flexible version of the power utility model developed by Epstein and Zin (1989, 1991) and Weil (1989). Let C_t denote consumption and $R_{w,t}$ denote the simple gross return on the portfolio of all invested wealth. The Epstein-Zin-Weil objective function is defined recursively as

$$U_t = \left\{ (1 - \delta) C_t^{\frac{1-\gamma}{\alpha}} + \delta (E_t U_{t+1}^{1-\gamma})^{\frac{1}{\alpha}} \right\}^{\frac{\alpha}{1-\gamma}}, \quad (3)$$

where $\alpha \equiv (1 - \gamma) / (1 - 1/\psi)$, ψ is the intertemporal elasticity of substitution in consump-

tion (IES), and γ is the coefficient of relative risk aversion.

We consider a model of complete markets in which all wealth (including human capital) is tradeable. In this case, the aggregate wealth return $R_{w,t}$ can be interpreted as the gross return to an asset that represents a claim to aggregate consumption, C_t , and aggregate consumption is the dividend on the portfolio of all invested wealth. For our main learning model, we follow Campbell (1986) and Abel (1999), and assume that the dividend on equity, D_t , equals aggregate consumption raised to a power λ :

$$D_t = C_t^\lambda. \quad (4)$$

When $\lambda > 1$, dividends and the return to equity are more variable than consumption and the return to aggregate wealth, respectively. Abel (1999) shows that $\lambda > 1$ can be interpreted as a measure of leverage. We refer to the dividend claim interchangeably as the levered consumption claim. In what follows, we use lower case letters to denote log variables, e.g., $\log(C_t) \equiv c_t$.

The specification (4) implies that the decline in the standard deviation of consumption growth in the 1990s should be met with a proportional decline in the volatility of dividend growth, $\sigma(\Delta c_t) = \lambda \sigma(\Delta d_t)$. In fact, such a proportional decline is present in cash-flow data. The standard deviation of consumption growth declined of 43% from the period 1952:Q1 to 1989:Q4 to 1990:Q1 to 2002:Q4. In comparison, the standard deviation of Standard and Poor 500 dividend growth declined 58%,¹⁷ the standard deviation of NIPA dividends declined 42% and the standard deviation of NIPA Net Cash Flow declined 40%. We calibrate the model based on estimates of the consumption process, and model dividends as a scale transformation of consumption. This practice has an important advantage: we do not need to empirically model the short-run dynamics of cash-flows, which were especially affected in the 1990s by pronounced shifts in accounting practices, corporate payout policies, and in the accounting treatment of executive compensation.

To incorporate regime shifts in the mean and volatility of consumption growth, consider the following model for the first difference of log consumption:

$$\Delta c_t = \mu(s_t) + \sigma(s_t)\epsilon_t, \quad (5)$$

¹⁷The data for Standard and Poor dividend growth are monthly from Robert Shiller's Yale web site. These data are not appropriate for calibrating the *level* of dividend volatility because the monthly numbers are smoothed by interpolation from annual data. But they can be used to compare *changes* in volatility across subsamples of the data, as we do here.

where $\epsilon_t \sim N(0, 1)$ and s_t again represents a state variable that takes on one of N different values representing the possible combinations for the mean state S_t and the volatility state V_t . This model is the same as the empirical model (2), except that we do not allow for autocorrelation in the conditional mean process, $\mu(s_t)$. As the results in Table 2 suggest, the estimated autocorrelation coefficient for $\mu(s_t)$ is not large for either measure of consumption and is likely to be inflated by time-averaging of aggregate consumption data. The more parsimonious framework (5) is far more manageable, as it reduces the number of states over which the model must be solved numerically.

An important feature of our model is captured by the assumption that agents cannot observe the underlying state, but instead must infer it from observable consumption data. This learning aspect is important because it implies that agents only gradually discover over time very low frequency changes in volatility. As we shall see below, this assumption allows the model to deliver a sustained rise in equilibrium asset prices in response to a low frequency reduction in volatility, rather than implying an abrupt, one-time jump in the stock market.¹⁸

When agents cannot observe the underlying state, inferences about the underlying state are captured by the posterior probability of being in each state based on data available through date t , given knowledge of the population parameters. Define the $N \times 1$ vector $\hat{\xi}_{t+1|t}$ of unsmoothed posterior probabilities in the following manner, where its j th element is given by

$$\hat{\xi}_{t+1|t}(j) = P\{s_{t+1} = j \mid \mathbf{Y}_t; \Theta\}.$$

As before, \mathbf{Y}_t denotes a vector of all the data up to time t and Θ contains all the parameters of the model. Throughout it will be assumed that a representative agent knows Θ , which consequently will be dropped from conditioning statements unless essential for clarity.

Bayes' Law implies that the posterior probability $\hat{\xi}_{t+1|t}$ evolves according to

$$\hat{\xi}_{t+1|t} = \mathbf{P} \frac{(\hat{\xi}_{t|t-1} \odot \eta_t)}{\mathbf{1}'(\hat{\xi}_{t|t-1} \odot \eta_t)} \quad (6)$$

where \odot denotes element-by-element multiplication, $\mathbf{1}$ denotes an $(N \times 1)$ vector of ones, \mathbf{P}

¹⁸Our model should be contrasted with models in which there is learning, but a constant regime. In such models, the agent eventually learns the mean and volatility of consumption growth given enough data. By contrast, in our model the mean and volatility of consumption growth can each switch between two values with non-zero probability. In fact, the mean state switches relatively frequently given our empirical estimates. The agent's belief about what state she is in does not converge to zero or one because the probability of the state does not converge to zero or one.

is the $N \times N$ matrix of transition probabilities and

$$\eta_t = \begin{bmatrix} f(\Delta c_t | s_t = 1, \mathbf{Y}_{t-1}) \\ \vdots \\ f(\Delta c_t | s_t = N, \mathbf{Y}_{t-1}) \end{bmatrix}$$

is the vector of likelihood functions conditional on the state.¹⁹ Given the distributional assumptions stated in (5), the likelihood functions are based on the normal distribution.

As in the econometric model, we assume that the mean and variance of consumption growth follow two-state Markov switching processes. Thus, there are two possible values for the mean, μ , and two possible values for the variance, σ , of consumption growth, implying four possible combinations of the two. It follows that s_t takes on one of four different values representing the $2^2 = 4$ possible combinations for the mean state S_t , and the variance state V_t .

As above, let \mathbf{P}^σ be the 2×2 transition matrix for the variance and \mathbf{P}^μ be the 2×2 transition matrix for the means. Then the full 4×4 transition matrix is given by

$$\mathbf{P} = \begin{bmatrix} p_{hh}^\mu \mathbf{P}^\sigma & p_{hl}^\mu \mathbf{P}^\sigma \\ p_{lh}^\mu \mathbf{P}^\sigma & p_{ll}^\mu \mathbf{P}^\sigma \end{bmatrix}.$$

The elements of the four-state transition matrix (and the eight-state transition matrix in Section 2) can be calculated from the two-state transition matrices \mathbf{P}^μ and \mathbf{P}^σ . The theoretical model can therefore be calibrated to match our estimates of \mathbf{P} , $\hat{\xi}_{t+1|t}$ and Θ from the regime switching model for aggregate consumption data, and closed as an general equilibrium exchange economy in which a representative agent receives the endowment stream given by the consumption process (5).

3.1 Pricing the Consumption and Dividend Claims

This section discusses how we solve for the price of a consumption and dividend claim. The Appendix gives a detailed description of the solution procedure; here we give only a broad outline.

Let P_t^D denote the ex-dividend price of a claim to the dividend stream measured at the end of time t , and P_t^C denote the ex-dividend price of a share of a claim to the consumption stream. From the first-order condition for optimal consumption choice and the definition of

¹⁹See Hamilton (1994), Chapter 22.

returns

$$E_t [M_{t+1}R_{t+1}] = 1, \quad R_{t+1} = \frac{P_{t+1}^D + D_{t+1}}{P_t^D} \quad (7)$$

where M_{t+1} is the stochastic discount factor, given under Epstein-Zin-Weil utility as

$$M_{t+1} = \left(\delta \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{1}{\psi}} \right)^\alpha R_{w,t+1}^{\alpha-1}. \quad (8)$$

Again, $R_{w,t+1}$ is the simple gross return on the aggregate wealth portfolio, which pays a dividend equal to aggregate consumption, C_t . The return on a risk-free asset whose value is known with certainty at time t is given by

$$R_{t+1}^f \equiv (E_t [M_{t+1}])^{-1}.$$

In contrast to Cecchetti et. al (1990, 2000) and Bonomo and Garcia (1994, 1996), we assume that investors do not observe the state s_t directly, but must instead infer it from observable consumption data. Because innovations to consumption growth are i.i.d. conditional on regime, and because agents cannot observe the underlying state, the posterior probabilities $\hat{\xi}_{t+1|t}$ summarize the information upon which conditional expectations are based. The price-dividend ratio for either claim may be computed by summing the discounted value of future expected dividends across states, weighted by the posterior probabilities of being in each state. The price-dividend ratio for both the consumption and dividend claims are functions only of $\hat{\xi}_{t+1|t}$. The Appendix explains how we solve for these functional equations numerically on a grid of values for the state variables $\hat{\xi}_{t+1|t}$.

Given the price-dividend ratio as a function of the state, we calculate the model's predicted price-dividend ratio over time by feeding in our time-series estimates of $\hat{\xi}_{t+1|t}$ presented above.²⁰ We also compute an estimate of the L year equity premium (the difference between the equity return and the risk-free rate over an L -year period) as a function of time t information. For L large, this "long-run" equity premium is analogous to what Fama and French (2002) call the unconditional equity premium, as of time t . The Appendix provides details about how these quantities are computed.

²⁰The posterior probabilities from the empirical model (2) are eight rather than four-dimensional because the mean-state of consumption growth in the previous model is also an element of the state vector. The four dimensional vector $\hat{\xi}_{t+1|t}$ fed into the model is created by summing the appropriate elements of the eight dimensional vector estimated from data. For example, $P \{ \mu_{t+1}, \sigma_{t+1} = h \mid \mathbf{Y}_t; \Theta \}$ is created by summing $P \{ S_{t+1} = 1, V_{t+1} = 1, S_t = 1 \mid \mathbf{Y}_t; \Theta \}$ and $P \{ S_{t+1} = 1, V_{t+1} = 1, S_t = 2 \mid \mathbf{Y}_t; \Theta \}$.

3.2 Choosing Model Parameters

We calibrate the model above at a quarterly frequency. The rate of time-preference is set to $\delta = 0.9925$. The parameters of the consumption process, (5), are set to match the empirical estimates reported in Table 2. Other key parameters for our investigation are the leverage parameter, λ , the coefficient of relative risk aversion, γ , the IES, ψ , and the transition probabilities of staying in a high or low volatility state. We discuss these in turn.

To calibrate the transition probabilities, we use the estimates presented in Table 2. The probability of remaining in the same volatility state next period exceeds 0.99 regardless of whether the volatility state is high or low, and indeed a 95% confidence interval for these estimates includes unity. Thus, these estimates are statistically indistinguishable from those that would imply the low volatility regime reached in the 1990s is expected to continue indefinitely, and they coincide with evidence from the macroeconomic literature that the shift to lower macroeconomic volatility is well described as an extremely persistent, if not permanent, break (Kim and Nelson (1999), McConnell and Perez-Quiros (2000), Stock and Watson (2002)). Indeed, the reduction in volatility in the last decade has been dubbed “the great moderation,” by Stock and Watson (2002), consistent with a common perception that this is evidence not of a transitory decline in volatility, but as a structural change in the economy as a whole.

In order to capture a very persistent decline in macroeconomic volatility, we set $p_{hh}^\sigma = p_{ll}^\sigma = 0.9999$ for the baseline results, but we also examine the sensitivity of these results to alternative values in Table 3, discussed below. The transition probabilities for the mean state, p_{hh}^μ and p_{ll}^μ are set to their samples estimates for consumption.

To calibrate λ , we follow Abel (1999) and set the parameter to match the sample standard deviation of dividend growth relative to that of consumption growth, $\lambda = \frac{\sigma(\Delta \ln D_t)}{\sigma(\Delta \ln C_t)}$. As discussed, this specification has support in the data in that the volatility of dividend growth has decreased by about the same proportion as that of consumption growth. As reported in Lettau and Ludvigson (2005), the percent standard deviation of real, per capita dividend growth constructed from CRSP index returns is 12.2 at an annual rate in post war data, about 8 times as high as that of real, per capita consumption growth, equal to 1.52 percent.²¹

²¹Abel (1999) calculated a smaller value for λ (approximately 3), by calibrating his model to the 1889-1978 sample used in Mehra and Prescott (1985). The reason he obtained a smaller number is that this sample includes prewar consumption data, which is over three times as volatile, relative to dividends, as is postwar consumption data. But the greater volatility of prewar consumption data relative to postwar data has been attributed, not to greater volatility of economic fundamentals in the prewar period, but to

For our benchmark results we set this parameter slightly lower than 8, to $\lambda = 6$, but we also considered values for λ as low as 4 and as high as 7; the results of using some of these parameter permutations are presented below.

To study the financial effects of a secular decline in macroeconomic risk, it is essential that the model economy we expose to such a shift be consistent with the average *levels* of the stock market and the equity premium. Therefore, we calibrate the coefficient of relative risk aversion, γ , in order to insure that our model is able to roughly match the mean equity premium, level of dividend-price ratio, and risk-free rate in post-war data. To do so, the model presented above requires moderately high risk aversion, around $\gamma = 25$.²² We use this value for the baseline results reported in this paper.

Finally, to choose parameter values for the IES, ψ , we consider how macroeconomic volatility influences the behavior of the equilibrium price-dividend ratio in the model presented above. A change in macroeconomic volatility has three affects on the equilibrium price-dividend ratio. First, regardless of the IES, lower macroeconomic volatility reduces the long-run equity premium because it lowers consumption risk; this effect drives up the price-dividend ratio. Second, lower macroeconomic volatility reduces the precautionary motive for saving, increasing the desire to borrow and therefore the equilibrium risk-free rate; this effect drives down the price-dividend ratio. Third, because we model a regime shift in the volatility of log changes, lower macroeconomic volatility reduces the mean of dividend growth in levels, D_t/D_{t-1} , because of a Jensen's inequality effect; this drives down the price-dividend ratio. The magnitude of the latter two effects relative to the first depends on the value of ψ and γ . If $\gamma > 1$, the first effect will dominate the last two only if $\psi > 1$.

Empirical estimates of ψ using aggregate consumption data often suggest that the IES is

greater measurement error in the prewar consumption and output series (Romer (1989)). For this reason, we calibrate our model to postwar consumption data.

²²Introspection notwithstanding, there is no broadly-accepted body of empirical evidence that provides a direct measure of risk-aversion. Survey evidence based on thought experiments is controversial and difficult to interpret. For example, Kandel and Stambaugh (1991) argue that careful scrutiny of common thought experiments shows that high risk aversion is not as implausible as commonly presumed. Financial markets, by contrast, provide indirect evidence that risk aversion is unlikely to be low. As a result, high risk aversion is a common feature of leading asset pricing models (see, for example, Campbell and Cochrane (1999) and the models discussed in Campbell (1999a) and Cochrane (2005)). A separate objection to high risk aversion is that it delivers counterfactual implications for the mean and volatility of the risk-free rate in asset pricing models with power utility and time-separable preferences. But, like the Epstein-Zin-Weil preference specification employed here, leading asset pricing models that depart from power utility resolve these difficulties by delivering reasonable implications for the risk-free rate even with high risk aversion.

relatively small, and in many cases statistically indistinguishable from zero (e.g. Campbell and Mankiw (1989), Ludvigson (1999), Campbell (2003)). But there are several reasons to think that the IES may be larger than estimates from aggregate data suggest. First, other researchers have found higher values for ψ using cohort level data (Attanasio and Weber (1993), Beaudry and van Wincoop (1996)), or when the analysis is restricted to asset market participants using household-level data (Vissing-Jorgensen (2002)).²³ More recently, Vissing-Jorgensen and Attanasio (2003) estimate the IES using the same Epstein-Zin framework employed in this study and find that this parameter for stockholders is typically above 1 (depending on the specification), with the most common values ranging from 1.17 to 1.75.²⁴ Second, Bansal and Yaron (2004) suggest that estimates of ψ based on aggregate data will be biased down if the usual assumption that consumption growth and asset returns are homoskedastic is relaxed. Third, Guvenen (2003) points out that macroeconomic models with limited stock market participation imply that properties of aggregate variables directly linked to asset wealth are almost entirely determined by stockholders who have empirically higher values for ψ . For the results reported below, we set $\psi = 1.5$, in the mid-range of the estimates reported by Vissing-Jorgensen and Attanasio (2003).²⁵

3.3 Model Results

In this section we present results from solving the model numerically. We focus on how stock prices are influenced by the break in macroeconomic volatility documented in the empirical macroeconomic literature. To this end, we characterize the behavior of the equilibrium price-dividend ratio of a claim to the dividend stream by plotting the model's solution for this quantity as a function of the posterior probabilities, and by feeding the model the historical values of $\hat{\xi}_{t+1|t}$, estimated as discussed above.

²³Vissing-Jorgensen emphasizes that even though estimates of the IES for non-asset holders are lower than those of asset holders, the difference should not be interpreted as evidence of heterogeneity in the IES across households. The reason is that estimates of the IES are based on Euler equations. Since the Euler equation for a given asset return cannot be expected to hold for households who do not have a position in the asset, IES estimates for non-asset holders will be inconsistent estimates of the IES for those households, and may be substantially biased down.

²⁴Unlike the estimates reported in Vissing-Jorgensen and Attanasio (2003), which are typically greater than one, the estimates of the IES reported in Vissing-Jorgensen (2002) are close to but slightly less than one. Vissing-Jorgensen (2002) explains the reason for this discrepancy: Vissing-Jorgensen (2002) ignored taxes, which biases estimates of the IES down. Vissing-Jorgensen and Attanasio (2003) address the downward bias by assuming a marginal tax rate of 30 percent.

²⁵This value is also used in Bansal and Yaron (2004).

Figure 7 (left plot) presents the model's predicted log price-dividend ratio of the dividend claim as a function of the posterior probability of being in a low volatility state, $\sigma = \sigma_l$, when the mean state is high either with probability one or with probability zero. The price-dividend ratio increases with the posterior probability of being in a low volatility state, regardless of whether the mean state is high or low. The increasing function is not linear, but is instead a convex function of investor's posterior probability of being in the low volatility state.

The intuition for this convexity is similar to that given in Veronesi (1999) for an asset pricing model with regime shifts in the mean of the endowment process. Suppose the probability of being in a low consumption volatility state is initially zero. News that causes an increase in the posterior probability of being in a low volatility state has two effects on the price-dividend ratio. First, because investors believe that the probability of being in a low volatility state has risen, consumption risk is perceived to be lower, which works to decrease the equilibrium risk-premium and raise the price-dividend ratio. Second, because the probability of being in a low volatility state is farther from zero, investors are more uncertain about which volatility regime the economy is in, which works to lower the equilibrium price-dividend ratio. The two effects are offsetting. Consequently, as the posterior probability of being in a low volatility state increases from zero, the price-dividend ratio rises only modestly.

Conversely, suppose the probability of being in a low consumption volatility state is initially at unity. News that causes a decrease in this posterior probability again has two effects on the price-dividend ratio. First, consumption risk is perceived to be higher, which works to increase the equilibrium risk-premium and lower the price-dividend ratio. Second, because the probability of being in a low volatility state is now farther from unity, investors are more uncertain about which volatility regime the economy is in, which works to further lower the equilibrium price-dividend ratio. In this case, the two effects are reinforcing rather than offsetting. Consequently, as the posterior probability of being in a low volatility state declines from unity, the price-dividend ratio falls dramatically. This explains why the equilibrium price-dividend ratio is a convex function of the posterior probabilities. But the degree of convexity is affected by risk-aversion. The more risk-averse agents are, the higher the posterior probability of being in a low volatility state must be before it has a noticeable affect on the equilibrium price-dividend ratio.

Figure 7 (right plot) displays the log price-dividend ratio of the dividend claim as a function of the posterior probability of being in a high mean growth state, $\mu = \mu_h$, when

the volatility state is high either with probability one or with probability zero. The price-dividend ratio increases with the posterior probability of being in a high mean state. For reasons similar to those just given, the function is again convex in the investor’s posterior probability of being in the high mean state, but is substantially less convex than the function plotted against the low volatility probability. The effect of a change in mean probability on the price-dividend ratio is also much smaller than the effect of a change in the volatility probability on the price-dividend ratio. These differences appear to be attributable to the lower persistence of the mean regimes compared to the volatility regimes. For example, the probability that a low mean (contraction) state will be followed by another period of contraction is 0.8 for consumption growth, so that this regime will persist on average for only 5 quarters. The estimated high mean, or expansion, regime is more persistent, but is still only expected to last 33 quarters on average. By contrast, the volatility regimes we estimate are far more persistent, and we have calibrated them for this figure so that the shift to lower macroeconomic volatility in the early 1990s is expected to persist almost indefinitely, consistent with this characterization in the macro literature. Because of this persistence in regime, asset prices can rise dramatically as investors become increasingly certain that a low macroeconomic volatility state has been reached.

How well does this model capture the run-up in asset prices observed in the late 1990s? To address this question, we feed the model historical values of $\hat{\xi}_{t+1|t}$ for our post-war sample, 1951:Q4 to 2002:Q4. Figure 8 presents the actual log price-dividend ratio on the CRSP value-weighted index, along with the post-war history of the price-dividend ratio on the dividend claim implied by the model. The figure displays plots of the model’s prediction for $p_t - d_t$ using the estimated unsmoothed posterior probabilities, for two parameter configurations. In the top panel, benchmark parameter values (discussed above) are used; in the bottom panel, risk-aversion is increased slightly (to $\gamma = 27$) and the leverage parameter is decreased slightly ($\lambda = 5$). The bottom panel also raises the rate of time preference slightly (to $\delta = .996$) in order to keep the mean value of the price-dividend ratio at empirically plausible values and roughly consistent with a low risk-free rate. All other parameters in the second panel of Figure 8 are held at their benchmark values. Note that the “model” line in each graph is produced using only the posterior probabilities estimated from consumption data; no asset market data are used.

Using the historical values of the unsmoothed probabilities, the top panel of Figure 8 shows that the benchmark model provides a remarkable account of the longer-term tendencies in stock prices. In particular, it captures virtually all of the boom in equity values that began

in the early 1990s and continued through the end of the millennium. The bottom panel of Figure 8 shows the same result when risk-aversion is slightly higher but leverage is lower. This case captures less of the overall boom in asset prices, but still explains all of the run-up since 1990 that was sustained after the broad market declines since 2000. In fact, the model's predicted price-dividend ratio is almost identical to the actual price-dividend ratio reached at the end of 2002.

Notice that the increase in valuation ratios predicted by the model is not well described by a sudden jump upward, but instead occurs gradually over several years, consistent with the data. This is a result of the learning built in to the model by the assumption that agents cannot observe the underlying state directly. Thus, the model produces about the right average value for stock returns during the 1990s.

We noted above that the model considered here delivers reasonable implications for the risk-free rate of interest. If we feed the model historical values of $\hat{\xi}_{t+1|t}$ we may compute the post-war history of the risk-free rate predicted by the model. Using the baseline parameters discussed above and used to create the results in panel A of Figure 8, this rate has a mean of 3.3 percent and a standard deviation of 0.35 percent per annum, roughly in line with actual values for an estimate of the real rate of return on a short-term Treasury bill.²⁶ The mean of 3.3 percent is a slightly higher than that suggested by historical estimates, but this can be rectified by using the parameter configuration employed to obtain the results in panel B of Figure 9: for this case, the mean risk-free rate generated by the model is 1.59 percent per annum, the standard deviation is again 0.35 percent.

What drives up the price-dividend ratio in the 1990s in this model? Although the shift to a higher mean growth state during this period generates a small part of the increase, the vast majority of the boom is caused by a decline in the equity premium as a consequence of the shift to reduced macroeconomic volatility. This can be seen in Figure 7. The right panel shows that, fixing the volatility state, variation in the equilibrium price-dividend ratio across mean states is quite modest. For example, fixing the probability of being in a low volatility state at one, the log price-dividend ratio ranges between 3.24 (when the probability of being in a high mean state is zero), to 3.57 (when the probability of being in a high mean state is one). Thus, the maximum possible variation in $p_t - d_t$ across mean states is about

²⁶Empirical estimates of volatility of the risk-free rate are typically based on the annualized sample standard deviation of the ex post real return on US Treasury bills—about 2% per annum in postwar data. This figure likely overstates the true volatility of the ex ante real interest rate, however, since much of the volatility of these returns is due to unanticipated inflation.

10 percent. Fixing the probability of being in a low volatility state at zero, the maximum possible variation in $p_t - d_t$ across mean states is even smaller, about 8 percent. This variation should be contrasted with the results for variation across volatility states, shown in the left panel. Fixing the probability of being in a high mean state at one, the log price-dividend ratio ranges between 3.57 (when the probability of being in a low volatility state is zero), to 4.34 (when the probability of being in a low volatility state is one), a range of variation of over 22 percent. Fixing the probability of being in a high mean state at zero, the maximum possible variation in $p_t - d_t$ across volatility states is about 24 percent. In short, large swings in the price-dividend ratio in this model are generated not by shifts in the mean of the endowment process, but by changes in the posterior probability of being exposed to a less volatile endowment process.

To understand what happens to the equity premium in the model, Figure 9 plots the $L = 50$, and 100-year equity premia implied by the model, computed recursively from the one period equity premium. Given that we calibrate the model at quarterly frequency, the L -year equity premium is computed as the expectation as of year t of the annualized compound rate of return from investing in the dividend claim from years t to $t + L$, less the annualized compound return from investing in the risk-free rate over years t to $t + L$. This is simply the expected value of the sum of future one-quarter log excess returns on the dividend claim for L years, reported at an annual rate. (The Appendix provides details about how this quantity is computed numerically.) We use this quantity for L large to capture the model's predicted value for the "unconditional" equity premium as of time t . The figure displays, for $L = 50$ and $L = 100$, the predicted equity premium of the dividend claim as a function of the posterior probability of being in a low volatility state, $\sigma = \sigma_t$, when the mean state is high either with probability one or with probability zero. In all cases, the predicted equity premium declines as the probability of being in a low volatility state increases, and drops off sharply once that probability exceeds 90 percent.

Figure 10 plots the post-war history of the log annual (100-year) equity premium on the dividend claim implied by the model, feeding in this history of the posterior probabilities. The model equity premium is relatively flat for most of the post-war period, but begins to decline in the early 1990s. For the benchmark model, premium declines by a little over two percentage points from peak to trough. We should not be surprised that the percentage decline is not greater; even small changes in the equity premium can have a large impact on asset values if they are sufficiently persistent.

The level of the long-run equity premium predicted by this parameterization is a bit higher

than estimates based on the historical average return on stocks in excess of returns on money market instruments. For example, the average annual excess return on the Standard and Poor 500 Stock market index in excess of a short-term Treasury bill rate is about 8 percent in data from 1951 to 2000, whereas the model’s predicted premium declines from around 11.7 percent at an annual rate, to 9.7 percent at its trough. Nevertheless, several researchers have pointed out that the historical average excess return may not result in a good estimate of the long-run premium that investors actually expect to earn in the future. This is because such a computation misses any change in stock prices that result in an unexpected decline in the equity premium (Jagannathan, McGrattan, and Scherbina (2000); Fama and French (2002)). Jagannathan, McGrattan and Scherbina provide a computation of the expected stock return (by decade) that takes such changes into account, and find that the expected return for the aggregate stock market in postwar data up to 1990 is higher than that suggested by simple historical averages of returns, ranging from about 10 percent to 11.4 percent, depending on which measure of the aggregate stock market is studied. These estimates for the expected (long-run) equity return are closer to those predicted by the model explored here under the benchmark parameter values described above. The model’s predicted value of the long-run equity premium can be altered by changing the model parameterization; we discuss this further below.

We emphasize an additional aspect of this model: although the volatility of consumption declines in the 1990s, the volatility of equilibrium stock returns does not—consistent with actual experience. In fact, in the data, stock market volatility appears to be, if anything, slightly higher in the late 1990s than in much of the rest of the postwar sample.²⁷ Figure 11 plots the post-war history of the conditional quarterly standard deviation of the log stock market return implied by the model at benchmark parameter values, reported at an annual rate. The figure shows that stock market volatility in the model is no lower in the 1990s than previously in the sample, despite the lower macroeconomic volatility. This result is primarily attributable to the increased uncertainty about which volatility regime the economy was in during the transition from a high to low macroeconomic volatility state.

We now explore how the model’s predictions change when we depart from the benchmark parameter values for λ and, importantly, the key posterior probabilities p_{jj}^σ , which denote the agent’s inference that next period’s volatility state will be j given that this period’s volatility state is j . The results of permuting these parameters to other values within two-standard

²⁷Updated plots of volatility of aggregate stock market indexes are provided by G. William Schwert at his University of Rochester web site: <http://schwert.ssb.rochester.edu/volatility.htm>.

errors of the point estimate are summarized in Table 3, which exhibits the model’s predictions for the price-dividend ratio and the long-run (100-year) equity premium in 1990:Q1 (before the estimated volatility shift) and in 2002:Q4 (after the volatility shift). Also shown is the one-period risk-free rate in 1990:Q1 and 2002:Q4. The values in each time period are computed by feeding the model the historical values of the posterior probabilities for the relevant quarters in our sample. The first row of Table 3 presents the results from our benchmark parameter values, used to generate the results reported in Panel A of Figure 8. Subsequent rows show how those results are changed when we depart from the benchmark parameter configuration by assigning the values indicated in the first four columns of Table 3, for a subset of the parameters. For each permutation, all of the parameters not listed in columns of Table 3 are maintained at their benchmark values.

Several notable aspects of the model are exhibited in Table 3. First, observe that the price-dividend ratio in the data rises from about 30 to 58.44 over the period 1990:Q1 to 2002:Q4, an increase of 28.44 (Figure 1). Row 1 of Table 3 shows that, at benchmark parameter values, the model predicts an increase of 33 in the price-dividend ratio over this same period, larger than that observed. At benchmark parameter values, the model predicts a higher value for the price-dividend ratio at the end of our sample than what actually occurred (68.44 versus 58.44). Row 2 shows what happens when the transition probability of remaining in the same volatility state next period is lowered to the statistically indistinguishable level of $p_{hh}^\sigma = p_{ll}^\sigma = 0.999$, rather than the benchmark $p_{hh}^\sigma = p_{ll}^\sigma = 0.9999$; now the model predicts a run-up in stock prices over the period 1990:Q1 to 2002:Q4 equal to roughly 70 percent of that observed. For this parameterization, the equilibrium price-dividend ratio rises from 35 to 56. The third row shows what happens when we use the exact point estimates for these transition probabilities, which are also statistically indistinguishable from the benchmark values. In this case the model explains about a quarter of the total run-up, with the equilibrium price-dividend ratio rising from 37 to 44. The result is essentially the same if we set $p_{hh}^\sigma = p_{ll}^\sigma = 0.99$, slightly lower than the point estimates (row 4). These findings illustrate the importance of the perceived permanence of the volatility decline in determining the magnitude of the rise in the equilibrium price-dividend ratio. Even a modest decrease in macroeconomic volatility can cause a dramatic boom in stock prices when the decrease is perceived to be sufficiently permanent.

Table 3 also shows that when λ is lowered from 6 to 4, the model predicts a smaller fraction of the run-up from 1990:Q1 to 2002:Q2 in price-dividend ratios than does the benchmark case, but still captures about 44% of the increase. This case also delivers a lower long-run

equity premium: the predicted equity premium is 7.7% in 1990:Q1, and falls to 6.5% in 2002:Q4, as a result of the shift toward lower volatility.²⁸ The last row of Table 3 gives results for the parameter configuration used in panel B of Figure 8; this case captures about 68 percent of the run-up in asset values from 1990:Q1 to 2002:Q4.

As the analysis above demonstrates, the model explored here predicts a surge in the price-dividend ratio on the dividend claim in the 1990s. What about the price-dividend ratio of the consumption claim (the wealth-consumption ratio)? It turns out that this quantity is far less affected by the shift to lower consumption volatility. This occurs because the price of an unlevered consumption claim is much less sensitive to swings in consumption risk than is the price of a levered claim. Results (not shown) demonstrate that the wealth-consumption ratio in the benchmark model is hardly affected by the low frequency shift toward lower volatility. This prediction is consistent with empirical evidence in Lettau and Ludvigson (2004), which shows that—unlike the log price-dividend ratio—the log wealth-consumption ratio has been largely restored to its sample mean with the broad market declines since 2000. This feature of the model is also consistent with evidence presented in Lettau and Ludvigson (2005), which suggests the presence of a low frequency component in the sampling variation of the log dividend-price ratio that is not present in the wealth-consumption ratio.

Of course, these considerations also imply that some short-term fluctuations in asset market valuation ratios are not as well captured by the model studied here. The focus in this paper is on the low-frequency movements in the stock market and the unconditional equity premium, and in particular in understanding the boom in the 1990s, an episode that dominates the post-war sampling variation of the stock market. As such, the model we investigate is not designed to capture the higher frequency fluctuations observed in the log price-dividend ratio prior to 1990, or the degree to which the price-dividend ratio overshot its value at the end of our sample, displayed in Figure 8. Although both the mean and the general direction of these cyclical fluctuations are reasonably well captured, the magnitude of these cyclical fluctuations is not. For example, the model correctly predicts a decline in the price-dividend ratio during the early 1950s and throughout most of the 1970s, but it misses the magnitude of both declines. One framework that is better able to capture these shorter-term, cyclical fluctuations in equity values is the model explored by Campbell and

²⁸For all parameter-value combinations, price-dividend ratios rise over the period 1990:Q1 to 2002:Q4, not because the long-run risk-free rate falls, but because the long-run equity premium falls. In fact, results (not reported) show that the long-run risk-free rate actually rises modestly in each case, but not by enough to offset the decline in the equity premium and cause an increase in the total rate of return.

Cochrane (1999). Unlike the model explored here, this model has time-varying risk-aversion and therefore generates substantial business cycle movement in the equity risk-premium. A shortcoming of that model, however, lies with its inability to capture the extraordinary stock market boom in the 1990s and the low frequency movements in the price-dividend ratio that dominate its behavior at the close of the last century. It is precisely this behavior that our model is designed to capture.

More recently, researchers have focused on the empirical underpinnings of cyclical fluctuations in assets values. For example, Duffee (2005), suggests that the conditional correlation between stock returns and consumption growth fluctuates over time and reached a peak at the end of 2000. It is important to note that these findings—of interest in their own right—do not necessarily contradict the conclusions of this paper. The estimates of conditional moments reported in Duffee (2005) are based on a GARCH model for the conditional covariance. As mentioned, the GARCH specification does a reasonable job of modeling cyclical changes in volatility and covariance *within* regimes, once those have been identified by other procedures, but it is not well-suited for capturing prolonged periods of moderated volatility or covariance. Although we have not modeled it as such here, less persistent shifts in conditional correlations, as captured by GARCH models, could co-exist with rare, but extremely persistent shifts in regime for volatility. Since regime changes can be expected to have a much greater impact on asset markets than cyclical fluctuations in conditional moments, evidence of GARCH-type increases in the conditional covariance of consumption and returns is not necessarily inconsistent with a model in which high stock prices relative to measures of fundamental value are driven by a near-permanent fall in macroeconomic risk.

We close this section by discussing the importance of the expected persistence of the volatility moderation in generating a large boom in stock prices. As the results above demonstrate, if the volatility moderation is perceived to be very persistent—lasting many decades—a large fraction of the run-up in stock prices can be explained. If the volatility decline is expected to be more transitory, less of the run-up can be rationalized through this mechanism. Of course, this feature of our theoretical model is not unique to the volatility explanation explored here. Given the extraordinary behavior of equity valuation ratios in the 1990s, *any* rational explanation of the stock market during this period must rest on an extremely persistent change in some underlying fundamental. Fortunately, for macroeconomic volatility, we have independent evidence about this persistence, since the Hamilton procedure provides an empirical estimate of how long a regime is expected to last through estimates of the transition probabilities. Table 2 shows that the estimated probability of re-

maintaining in a low volatility state once it has been reached is in excess of 0.99 for consumption, implying that the regime will persist on average more than 40 years.

How likely is persistence of 80, 100 or even 1000 years? Figure 12 plots the log likelihood of our empirical model (2), as a function of p_{ll}^σ , the probability of remaining in a low volatility state next period given a low volatility state this period.²⁹ The likelihood has a clear peak at the point estimate, 0.992, but is almost as high at unity as it is at the point estimate. Thus values for p_{ll}^σ that imply the low volatility regime will persist indefinitely are just as empirically defensible, statistically, as those that suggest it will persist for 40 to 80 years. In the model, the difference between 40 years and indefinitely is not inconsequential for equilibrium asset prices, but even the low end of the empirically plausible range implies extreme persistence. This means that regardless of what value for p_{ll}^σ one favors, results in Table 3 suggest that the decline in volatility plays some role in the rise of equity values since 1990. One view of these theoretical results is that the stock market appears to be very informative about the expected persistence of the volatility moderation. These estimates are obtained without using any stock market data. Had we included data on the stock market in our estimation such estimates of the persistence would likely have been pushed to the very high end of the range obtained from pure macroeconomic data.

3.4 A Model With Cointegrated Consumption and Dividends

The learning model above imposes a particularly simple relationship between consumption and dividends. As a robustness check with our previous results, we solve a model where consumption and dividends are cointegrated. For simplicity, we abstract away from the learning aspect of our model. Unfortunately, it is difficult to introduce such a cointegrated process into the learning model above, since the state space becomes too large.³⁰

We assume the following error-correction representation for log consumption and log

²⁹We thank Lars Hansen for suggesting this plot.

³⁰An alternative specification that also implies a decline in dividend volatility is

$$\begin{aligned}\Delta d_{t+1} &= \mu_t + \lambda\sigma_t\varepsilon_t \\ \Delta c_{t+1} &= \mu_t + \sigma_t\varepsilon_t.\end{aligned}$$

We find this specification produces results that are qualitatively similar to those from the learning model described in section 3.1-3.3 above.

dividends:

$$\Delta c_{t+1} = \mu + \epsilon_{c,t+1} \tag{9}$$

$$\Delta d_{t+1} = \mu + k(c_t - d_t) + \lambda \epsilon_{c,t+1} + \epsilon_{d,t+1}, \tag{10}$$

where $\epsilon_{c,t+1}$ and $\epsilon_{d,t+1}$ are normal, mean-zero independent shocks, with $\text{Var}(\epsilon_{c,t+1}) = \sigma_c^2$ and $\text{Var}(\epsilon_{d,t+1}) = \sigma_d^2$. Log consumption and dividends are cointegrated, with cointegrating vector $(1, -1)'$. The leverage parameter λ in this model multiplies the consumption innovation $\epsilon_{c,t+1}$, rather than the overall level of consumption, allowing consumption and dividends to have the same long-run growth rate. The parameter k is an adjustment parameter that governs the persistence of deviations from the common long-run trend. The innovation $\epsilon_{d,t+1}$ is a purely idiosyncratic shock to dividends.

Since consumption and dividends are cointegrated but consumption follows a random walk, dividend growth must be predictable by last period's consumption-dividend ratio. This is so because in this specification consumption defines the trend in dividends; thus shocks to the consumption-dividend ratio must be accompanied by a subsequent adjustment in dividends in order to move these variables back toward their common long-run trend.

To study the affects of lower consumption volatility on asset prices in this model, we assume that the variance of consumption undergoes a perfectly anticipated one-off shift to a lower value. This amounts to comparing two steady states. The consumption volatilities in the two steady states are the estimated values in the high and low volatility regime reported in Table 2. In particular, we assume that σ_c^2 takes on the high volatility regime value in data up through 1990, and the low value in data after 1990. The value of μ is set to match the average growth of consumption in our sample, and the value of σ_d^2 is set to match the greater overall level of volatility of dividend growth as compared to consumption growth. Finally, we allow the adjustment parameter k to take on a number of values ranging from 0.005 to 0.20. These numbers are in line with those calibrated to match consumption and dividend data in Lettau and Wachter (2006). In particular, using data on consumption and dividends, Lettau and Wachter find that k is 0.09, about the median of the range considered here. The numerical solution procedure is described in the Appendix.

Table 4 exhibits the model's predictions for the price-dividend ratio in the two volatility regimes for a variety of parameter values of γ , λ and k . Other parameters are set as in the baseline calibration of the learning model above.

The first aspect of note in Table 4 is that risk aversion must be higher than in the learning model in order to match the overall level of the equity premium and price-dividend ratio

found in the data. This is attributable to the two factors: the presence of cointegration and the absence of persistent changes in consumption growth. Cointegration ties dividends to consumption in the long-run. This means that even though dividend growth is much more volatile than consumption growth in the short-run, it can be no more risky than consumption growth in the long-run (i.e., the annualized volatility of dividend growth must decline with the horizon over which it is measured to match the lower volatility of consumption growth). Persistent changes in cash-flow growth also make the dividend claim more risky. These two factors mean that we need to adjust risk aversion upward in this model (Table 4 reports results for $\gamma = 40$ or 60) in order to match the overall level of the equity premium and price-dividend ratio found in the data. With this adjustment, the model predicts a significant fraction of the run-up from 1990:Q1 to 2002:Q2 in price-dividend ratios for a range of values of k . Moreover, when k is sufficiently small, this model can explain more than 100% of the run-up. This is not surprising because when k is small, dividends deviate in a very persistent manner from the common trend defined by consumption, implying that the annualized volatility of dividend growth is greater than that of consumption growth at all but the very longest horizons. (In the limit when $k = 0$, there is no cointegration and the annualized volatility of dividend growth is greater than that of consumption growth at every horizon.) But even with k as high as 0.2—a value almost twice as high as that suggested by empirical estimates—the model predicts almost 80% of the run-up in valuation ratios as a result of the shift toward lower volatility. We conclude that our main findings are robust to modeling consumption and dividends as cointegrated processes.

4 Conclusions

This paper considers the low frequency behavior of post-war equity values relative to measures of fundamental value. Such longer-term movements are dominated by the stock market boom of the 1990s, an extraordinary episode in which price-dividend ratios on aggregate stock market indices increased three-fold over a period of five years. Indeed, Figure 1 shows this period to be the defining episode of postwar financial markets. As Campbell (1999b) notes, the relationship between stock prices and fundamentals in the 1990s appears to have changed. A growing body of literature is now working to understand this phenomenon, and explanations run the gamut from declining costs of equity market participation and diversification, to irrational exuberance, to changes in technology and demography.

In this paper, we consider a different explanation for why the relationship between stock

prices and fundamentals appears to have changed. We ask whether the phenomenal surge in asset values that dominated the close of the 20th century can be plausibly described as a rational response to macroeconomic factors, namely the sharp and sustained decline in macroeconomic risk. We find that, in large part, it can. There is a strong correlation between the low frequency movements in macroeconomic volatility and asset prices in post-war data, both in the US and internationally. We show that, when such a shift toward decreased consumption risk is perceived to be sufficiently persistent, an otherwise standard asset pricing model can explain a large fraction of the surge in equity valuation ratios observed in U.S. data in the 1990s. In the model economy, a boom in stock prices occurs because the decline in macroeconomic risk leads to a fall in expected future stock returns, or the equity risk-premium. An implication of these findings is that multiples of price to earnings or dividends may remain above previous historical norms into the indefinite future.

Of course, in the final analysis, the complexities of modern financial markets leave little doubt that several factors outside of our model are likely to have contributed to the surge in asset values relative to measures of fundamental value during the final part of the last century. Nevertheless, the analysis here suggests that the well documented break in macroeconomic volatility could be an important contributing factor.

5 Appendix

I. Data Description

The sources and description of each data series we use are listed below.

GDP

GDP is gross domestic product, measured in 1996 chain-weighted dollars. Our source is the U.S. Department of Commerce, Bureau of Economic Analysis.

CONSUMPTION

Consumption is measured as total personal consumption expenditures. The quarterly data are seasonally adjusted at annual rates, in billions of chain-weighted 1996 dollars. Our source is the U.S. Department of Commerce, Bureau of Economic Analysis.

POPULATION

A measure of population is created by dividing real total disposable income by real per capita disposable income. Consumption, wealth, labor income, and dividends are in per capita terms. Our source is the Bureau of Economic Analysis.

PRICE DEFLATOR

Real asset returns are deflated by the implicit chain-type price deflator (1996=100) given for the consumption measure described above. Our source is the U.S. Department of Commerce, Bureau of Economic Analysis.

PRICE-DIVIDEND RATIO

The price-dividend ratio is that of the CRSP value-weighted index, constructed as in Campbell (2003). Our source is the Center for Research in Security Prices, University of Chicago.

II. Pricing the Consumption and Dividend Claims, Computation of Unconditional Equity Premium: Learning Model

This appendix describes the algorithm used to solve for prices.

The first order conditions that the price-dividend ratio of a claim to the dividend stream satisfies

$$E_t \left[M_{t+1} \left(\frac{P_{t+1}^D}{D_{t+1}} (\hat{\xi}_{t+2|t+1}) + 1 \right) \frac{D_{t+1}}{D_t} \right] = \frac{P_t^D}{D_t} (\hat{\xi}_{t+1|t}), \quad (11)$$

and the price-consumption ratio for the consumption claim satisfies

$$E_t \left[M_{t+1} \left(\frac{P_{t+1}^C}{C_{t+1}} (\hat{\xi}_{t+2|t+1}) + 1 \right) \frac{C_{t+1}}{C_t} \right] = \frac{P_t^C}{C_t} (\hat{\xi}_{t+1|t}). \quad (12)$$

Notice that $\frac{P_t^C}{C_t}$ is the wealth-consumption ratio, where wealth here is measured on an ex-dividend basis. The posterior probabilities $\hat{\xi}_{t+1|t}$ are the only state variables in this framework, so the price-dividend ratio is a function only of $\hat{\xi}_{t+1|t}$. We substitute for M_{t+1} from (8) and solve these functional equations numerically on a grid of values for the state variables $\hat{\xi}_{t+1|t}$.

Equation (11) can be rewritten as:

$$E_t \left[M_{t+1} \left(\frac{P_{t+1}^D}{C_{t+1}^\lambda} + 1 \right) \left(\frac{C_{t+1}}{C_t} \right)^\lambda \right] = \frac{P_t^D}{C_t^\lambda}, \quad (13)$$

where M_{t+1} is given in (8). Applying the definition of returns with $\lambda = 1$, M_{t+1} can be rewritten as

$$M_{t+1} = \delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\alpha}{\psi} + \alpha - 1} \left(\frac{P_{t+1}^C}{C_{t+1}} + 1 \right)^{\alpha - 1} \left(\frac{P_t^C}{C_t} \right)^{1 - \alpha}. \quad (14)$$

Plugging (14) into (13) we obtain

$$E_t \left[\delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\alpha}{\psi} + \alpha - 1} \left(\frac{P_{t+1}^C}{C_{t+1}} + 1 \right)^{\alpha - 1} \left(\frac{P_t^C}{C_t} \right)^{1 - \alpha} \left(\frac{P_{t+1}^D}{C_{t+1}^\lambda} + 1 \right) \left(\frac{C_{t+1}}{C_t} \right)^\lambda \right] = \frac{P_t^D}{C_t^\lambda} \quad (15)$$

The price-dividend ratio on equity, $\frac{P_t^D}{C_t^\lambda}$, is defined recursively by (15).

We write the price-dividend ratio for a levered consumption claim as a function of the state vector $\hat{\xi}_{t+1|t}$:

$$\frac{P_t^D}{C_t^\lambda} = F_D(\hat{\xi}_{t+1|t}). \quad (16)$$

Similarly, the price-dividend ratio for the unlevered consumption claim can be written

$$\frac{P_t^C}{C_t} = F_C(\hat{\xi}_{t+1|t}), \quad (17)$$

for some function F_C . Notice that the price-dividend ratio in (17) is simply the wealth-consumption ratio, where wealth is defined to be ex-dividend wealth.

The wealth-consumption ratio is defined as the fixed point of (15) for $\lambda = 1$ and $P^D = P^C$ everywhere. Applying this case to (15) and substituting $\frac{P_t^C}{C_t} = F_C(\hat{\xi}_{t+1|t})$, the Euler equation for the consumption claim may be written

$$E_t \left[\delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\alpha}{\psi} + \alpha} \left(F_C(\hat{\xi}_{t+2|t+1}) + 1 \right)^\alpha \right] = \left(F_C(\hat{\xi}_{t+1|t}) \right)^\alpha.$$

From the definition of the conditional expectation, the left-hand-side of the expression above is given by

$$E_t \left[\delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\alpha}{\psi} + \alpha} \left(F_C(\hat{\xi}_{t+2|t+1}) + 1 \right)^\alpha \right] = \sum_{j=1}^N P \{s_{t+1} = j \mid \mathbf{Y}_t\} E \left[\delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)_j^{-\frac{\alpha}{\psi} + \alpha} \left(F_C(\hat{\xi}_{t+2|t+1}) + 1 \right)^\alpha \mid s_{t+1} = j, \mathbf{Y}_t \right], \quad (18)$$

where $\exp \left[\left(\frac{C_{t+1}}{C_t} \right)_j \right] \equiv \mu(s_j) + \sigma(s_j)\epsilon_{t+1}$ denotes consumption growth in state j . From the evolution equation (6) and the stochastic model (5), the distribution of $\hat{\xi}_{t+2|t+1}$ conditional on time- t data \mathbf{Y}_t and on the state s_{t+1} depends only on $\hat{\xi}_{t+1|t}$ and the state. Using the definition of $\hat{\xi}_{t+1|t}$, (18) can be written

$$E_t \left[\delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\alpha}{\psi} + \alpha} \left(F_C(\hat{\xi}_{t+2|t+1}) + 1 \right)^\alpha \right] = \sum_{j=1}^N \hat{\xi}_{t+1|t}(j) E \left[\delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)_j^{-\frac{\alpha}{\psi} + \alpha} \left(F_C(\hat{\xi}_{t+2|t+1}) + 1 \right)^\alpha \mid s_{t+1} = j, \hat{\xi}_{t+1|t} \right].$$

Thus, the wealth-consumption ratio, $F_C(\hat{\xi}_{t+1|t})$, is defined by the recursion:

$$\left(F_C(\hat{\xi}_{t+1|t}) \right)^\alpha = \sum_{j=1}^N \hat{\xi}_{t+1|t}(j) E \left[\delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)_j^{-\frac{\alpha}{\psi} + \alpha} \left(F_C(\hat{\xi}_{t+2|t+1}) + 1 \right)^\alpha \mid s_{t+1} = j, \hat{\xi}_{t+1|t} \right]. \quad (19)$$

It is straightforward to show that a similar recursion defines the price-dividend ratio of a levered equity claim, $F_D(\hat{\xi}_{t+1|t})$, by allowing λ to take on arbitrary values greater than unity. From (15), we have

$$\frac{P_t^D}{D_t} = F_C(\hat{\xi}_{t+1|t})^{1-\alpha} \delta^\alpha E_t \left[\left(\frac{C_{t+1}}{C_t} \right)_j^{-\frac{\alpha}{\psi} + \alpha - 1 + \lambda} \left(F_C(\hat{\xi}_{t+2|t+1}) + 1 \right)^{\alpha-1} \left(\frac{P_{t+1}^D}{D_{t+1}} + 1 \right) \right].$$

Substituting (16) into the above, we obtain

$$F_D(\hat{\xi}_{t+1|t}) = F_C(\hat{\xi}_{t+1|t})^{1-\alpha} \delta^\alpha \times \sum_{j=1}^N \hat{\xi}_{t+1|t}(j) E \left[\left(\frac{C_{t+1}}{C_t} \right)_j^{-\frac{\alpha}{\psi} + \alpha - 1 + \lambda} \left(F_C(\hat{\xi}_{t+2|t+1}) + 1 \right)^{\alpha-1} \left(F_D(\hat{\xi}_{t+2|t+1}) + 1 \right) \mid s_{t+1} = j, \hat{\xi}_{t+1|t} \right].$$

The expectation above is computed by numerical integration under the assumption that innovations to consumption growth are i.i.d., conditional on the state j .

Denote the log return of the dividend claim from t to $t+1$ as $\log(R_{D,t+1}) = r_{D,t+1}$. It is also possible to compute moments of log returns:

$$E_t[r_{D,t+1}] = E_t \left[\log \left(\frac{F_D(\hat{\xi}_{t+2|t+1}) + 1}{F_D(\hat{\xi}_{t+1|t})} \left(\frac{D_{t+1}}{D_t} \right) \right) \right] \quad (20)$$

$$\sigma_t^2[r_{D,t+1}] = E_t \left[\left(\log \left(\frac{F_D(\hat{\xi}_{t+2|t+1}) + 1}{F_D(\hat{\xi}_{t+1|t})} \left(\frac{D_{t+1}}{D_t} \right) \right) \right)^2 \right] - (E_t[r_{D,t+1}])^2$$

Equation (20) points the way to calculating the expected L period return. For example, suppose we were interested in the annualized compound rate of return from investing in the levered consumption claim for thirty years. Because the model is calibrated to t equals a quarter, we would compute

$$\frac{4}{120} E_t[r_{D,t+1} + r_{D,t+2} + \cdots + r_{D,t+120}] \quad (21)$$

To compute the thirty-year equity premium, we could subtract out the return from rolling over investments in the risk-free rate. Let $r_{t+1}^f = \log R_{t+1}^f$, then

$$\frac{4}{120} E_t[r_{t+1}^f + r_{t+2}^f + \cdots + r_{t+120}^f] \quad (22)$$

Note that r_{t+1}^f is known at time t and so could be brought outside the expectation. Subtracting (22) from (21) gives the thirty-year ahead risk premium.

The question is how to compute the elements in the sums of (21) and (22)? We show that these quantities can be computed recursively. Because the one-period ahead expected return and risk-free rate are functions of $\hat{\xi}_{t+1|t}$, we can write

$$G_1(\hat{\xi}_{t+1|t}) = E_t(r_{D,t+1})$$

Define

$$G_2(\hat{\xi}_{t+1|t}) = E_t(G_1(\hat{\xi}_{t+2|t+1}))$$

By the law of iterated expectations

$$G_2(\hat{\xi}_{t+1|t}) = E_t(G_1(\hat{\xi}_{t+2|t+1})) = E_t(E_{t+1}(r_{D,t+2})) = E_t(r_{D,t+2})$$

More generally, define G_m recursively as

$$G_m(\hat{\xi}_{t+1|t}) = E_t(G_{m-1}(\hat{\xi}_{t+2|t+1}))$$

Note that assuming $G_{m-1}(\hat{\xi}_{t+1|t}) = E_t[r_{t+m-1}]$ implies,

$$G_m(\hat{\xi}_{t+1|t}) = E_t(G_{m-1}(\hat{\xi}_{t+2|t+1})) = E_t(E_{t+1}(r_{D,t+m})) = E_t(r_{D,t+m}). \quad (23)$$

So by induction, (23) holds for all L .

Similarly define

$$H_1(\hat{\xi}_{t+1|t}) = r_{t+1}^f$$

and

$$H_m(\hat{\xi}_{t+1|t}) = E_t(H_{m-1}(\hat{\xi}_{t+2|t+1}))$$

If $H_{m-1}(\hat{\xi}_{t+1|t}) = E_t[r_{t+m-1}^f]$,

$$H_m(\hat{\xi}_{t+1|t}) = E_t(H_{m-1}(\hat{\xi}_{t+2|t+1})) = E_t(E_{t+1}(r_{t+m}^f)) = E_t(r_{t+m}^f). \quad (24)$$

So (24) holds for all L . Thus the L -period ahead expected return can be found by recursively calculating G_m for $m = 1, \dots, L$, summing up, and multiplying by $4/L$. The L -period ahead risk premium can be found by calculating H_m for $m = 1, \dots, L$, summing up the differences $G_m - H_m$, and multiplying by $4/L$.

III. Pricing the Consumption and Dividend Claims: Cointegration Model

To solve the cointegration model, rewrite the system (9)-(10) as

$$\begin{aligned} \Delta d_{t+1} &= \mu + z_t + \lambda \epsilon_{c,t+1} + \epsilon_{d,t+1} \\ z_{t+1} &= (1 - k)z_t + k(1 - \lambda)\epsilon_{c,t+1} - k\epsilon_{d,t+1}, \end{aligned}$$

The price-consumption ratio P_t/C_t solves the following:

$$E_t \left[\delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\alpha}{\psi} + \alpha} \left(\frac{P_{t+1}^C}{C_{t+1}} + 1 \right)^\alpha \right] = \left(\frac{P_t^C}{C_t} \right)^\alpha$$

This equation is solved by

$$\frac{P_t^C}{C_t} = \frac{R}{1 - R}, \quad (25)$$

where

$$R = \delta \exp \left\{ \left(1 - \frac{1}{\psi} \right) \mu + \frac{1}{2} \alpha \left(1 - \frac{1}{\psi} \right)^2 \sigma_c^2 \right\}.$$

The price-dividend ratio then solves the equation

$$E_t \left[\delta^\alpha \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\alpha}{\psi} + \alpha - 1} R^{1-\alpha} \left(\frac{D_{t+1}}{D_t} \right) \left(\frac{P_{t+1}^D}{D_{t+1}} + 1 \right) \right] \quad (26)$$

This follows from substituting (25) into the equation for the price-dividend ratio. Conjecture that the price-dividend ratio has a solution of the form

$$\frac{P_t^D}{D_t} = \sum_{n=1}^{\infty} \exp \{ A_n + B_n z_t \}. \quad (27)$$

Substituting (27) into (26) verifies the conjecture and implies that A_n and B_n satisfy

$$A_n = A_{n-1} + \alpha \log \delta + (1 - \alpha) \log R + \alpha \left(1 - \frac{1}{\psi} \right) \mu + \frac{1}{2} \left(\alpha \left(1 - \frac{1}{\psi} \right) + (\lambda - 1)(1 - B_{n-1}k) \right)^2 \sigma_c^2 + \frac{1}{2} (1 - B_{n-1}k)^2 \sigma_d^2 \quad (28)$$

$$B_n = (1 - k)B_{n-1} + 1, \quad (29)$$

with boundary conditions $A_0 = B_0 = 0$.

References

- ABEL, A. B. (1994): “Exact Solutions for Expected Rates of Return Under Markov Regime Switching: Implications for the Equity Premium,” *Journal of Money, Credit, and Banking*, 26, 345–361.
- (1999): “Risk Premia and Term Premia in General Equilibrium,” *Journal of Monetary Economics*, 43, 3–33.
- (2003): “The Effects of a Baby Boom on Stock Prices and Capital Accumulation in the Presence of Social Security,” *Econometrica*, 71(2), 551–578.
- ATTANASIO, O. P., AND G. WEBER (1993): “Consumption Growth, the Interest Rate and Aggregation,” *Review of Economic Studies*, 60(3), 631–649.
- BAI, J., AND P. PERRON (2003): “The Computation and Analysis of Multiple Structural Change Models,” *Journal of Applied Econometrics*, 18, 1–22.
- BALKE, N. S., AND R. J. GORDON (1989): “The Estimation of Prewar Gross National Product: Methodology and New Evidence,” *Journal of Political Economy*, 97(1), 38–92.
- BANSAL, R., V. KHATCHATRIAN, AND A. YARON (2005): “Interpretable Asset Markets?,” *European Economic Review* forthcoming.
- BANSAL, R., AND C. LUNDBLAD (2002): “Market Efficiency, Asset Returns, and the Size of The Risk Premium in Global Equity Markets,” *Journal of Econometrics*, 109, 195–237.
- BANSAL, R., AND A. YARON (2004): “Risks for the Long-Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 59(4), 1481–1509.
- BEAUDRY, P., AND E. VAN WINCOOP (1996): “The Intertemporal Elasticity of Substitution: An Exploration Using a U.S. Panel of States,” *Economica*, 63(251), 495–512.
- BLANCHARD, O. J. (1993): “Movements in the Equity Premium,” *Brookings Papers on Economic Activity*, 2, 75–138.
- BLANCHARD, O. J., AND J. SIMON (2001): “The Long and Large Decline in U.S. Output Volatility,” *Brookings Papers on Economic Activity*, 1, 135–174.
- BONOMO, M., AND R. GARCIA (1994): “Can a Well-Fitted Equilibrium Asset Pricing Model Produce Mean Reversion?,” *Journal of Applied Econometrics*, 9, 19:29.
- (1996): “Consumption and Equilibrium Asset Pricing: An Empirical Assessment,” *Journal of Empirical Finance*, 3, 239–265.

- CAGETTI, M., L. P. HANSEN, T. SARGENT, AND N. WILLIAMS (2002): “Robustness and Pricing with Uncertain Growth,” *Review of Financial Studies*, 15(2), 363–404.
- CALVET, L., M. GONZALEZ-EIRAS, AND P. SODINI (2003): “Financial Innovation, Market Participation and Asset Prices,” NBER Working Paper No. 9840.
- CAMPBELL, J. Y. (1986): “Bond and Stock Returns in a Simple Exchange Model,” *Quarterly Journal of Economics*, 101, 785–804.
- (1991): “A Variance Decomposition for Stock Returns,” *Economic Journal*, 101, 157–179.
- (1999a): “Asset Prices, Consumption and the Business Cycle,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, pp. 1231–1303. North-Holland, Amsterdam.
- (1999b): “Comment on John Heaton and Deborah Lucas ‘Stock Prices and Fundamentals’,” in *National Bureau of Economic Research Macroeconomics Annual, 1999*, ed. by B. S. Bernanke, and J. J. Rotemberg, pp. 253–261. MIT Press, Cambridge, MA.
- (2003): “Consumption-Based Asset Pricing,” in *Handbook of the Economics of Finance* forthcoming, ed. by G. Constantinides, M. Harris, and R. Stulz. North-Holland, Amsterdam.
- CAMPBELL, J. Y., AND J. H. COCHRANE (1999): “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107, 205–251.
- CAMPBELL, J. Y., AND G. MANKIW (1989): “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence,” in *NBER Macroeconomics Annual: 1989*, ed. by O. Blanchard, and S. Fischer, pp. 185–216. MIT Press, Cambridge, MA.
- CAMPBELL, J. Y., AND R. J. SHILLER (1989): “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *Review of Financial Studies*, 1(3), 195–228.
- (2003): “Valuation Ratios and the Long-Run Stock Market Outlook: An Update,” in *Advances in Behavioral Finance, Volume II* forthcoming, ed. by N. Barberis, and R. Thaler. Russell Sage Foundation, New York, NY.
- CECCHETTI, S. G., P. LAM, AND N. C. MARK (1990): “Mean Reversion in Equilibrium Asset Prices,” *American Economic Review*, 80(3), 398–418.
- (1993): “The Equity Premium and the Risk-Free Rate: Matching the Moments,” *Journal of Monetary Economics*, 31, 21–45.
- COCHRANE, J. H. (2005): *Asset Pricing*. Princeton University Press, Princeton, NJ.

- DEANGELO, H., L. DEANGELO, AND D. J. SKINNER (2002): “Are Dividends Disappearing? Dividend Concentration and the Consolidation of Earnings,” Unpublished paper, University of Michigan.
- DRIFFIL, J., AND M. SOLA (1998): “Comparing Predictive Accuracy,” *Journal of Monetary Economics*, 42, 357–373.
- DUFFEE, G. R. (2005): “Time Variation in the Covariances Between Stock Returns and Consumption Growth,” *Journal of Financial Economics*, forthcoming.
- EPSTEIN, L., AND S. ZIN (1989): “Substitution Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937–968.
- (1991): “Substitution, Risk Aversion, and the Temporal behavior of Consumption and Asset Returns: An Empirical Investigation,” *Journal of Political Economy*, 99, 555–576.
- FAMA, E. F., AND K. R. FRENCH (1988): “Dividend Yields and Expected Stock Returns,” *Journal of Financial Economics*, 22, 3–27.
- (2001): “Disappearing Dividends: Changing Firm Characteristics or Lower Propensity to Pay,” *Journal of Financial Economics*, 60(1), 3–43.
- (2002): “The Equity Premium,” *Journal of Finance*, 57(2), 637–659.
- GUVENEN, M. F. (2003): “Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective,” University of Rochester, RCER Working Paper, No 491.
- HALL, R. E. (2000): “E-Capital: The Link Between the Stock Market and the Labor Market in the 1990s,” *Brookings Papers on Economic Activity*, 2, 73–118.
- HAMILTON, J. D. (1989): “A New Approach to the Analysis of Nonstationary Returns and the Business Cycle,” *Econometrica*, 57, 357–384.
- (1994): *Time Series Analysis*. Princeton University Press, Princeton, NY.
- HARVEY, C. R. (1991): “The World Price of Covariance Risk,” *Journal of Finance*, 46, 111–117.
- HEATON, J., AND D. LUCAS (1999): “Stock Prices and Fundamentals,” in *NBER Macroeconomics Annual: 1999*, ed. by O. J. Blanchard, and S. Fischer. MIT Press, Cambridge, MA.
- HODRICK, R. (1992): “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement,” *Review of Financial Studies*, 5, 357–386.

- JAGANNATHAN, R., E. R. MCGRATTAN, AND A. SCHERBINA (2000): “The Declining U.S. Equity Premium,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 24(4), 3–19.
- JOVANOVIC, B., AND P. L. ROUSSEAU (2003): “Specific Capital and the Division of Rents,” Unpublished paper, New York University.
- KANDEL, S., AND R. F. STAMBAUGH (1991): “Asset Returns and Intertemporal Preferences,” *Journal of Monetary Economics*, 27, 39–71.
- KIM, C.-J., AND C. R. NELSON (1999): “Has the Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle,” *The Review of Economics and Statistics*, 81, 608–616.
- LAMONT, O. (1998): “Earnings and Expected Returns,” *Journal of Finance*, 53, 1563–87.
- LETTAU, M., AND S. C. LUDVIGSON (2001): “Consumption, Aggregate Wealth and Expected Stock Returns,” *Journal of Finance*, 56(3), 815–849.
- (2004): “Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption,” *American Economic Review*, 94(1), 276–299.
- (2005): “Expected Returns and Expected Dividend Growth,” *Journal of Financial Economics*, 76, 583–626.
- LETTAU, M., AND S. VAN NIEUWERBURGH (2005): “Reconciling the Return Predictability Evidence: In-Sample Forecasts, Out-of-Sample Forecasts, and Parameter Instability,” Unpublished paper, New York University.
- LETTAU, M., AND J. A. WACHTER (2006): “Why is Long-Horizon Equity Less Risky? A Duration Based Explanation of the Value Premium,” *Journal of Finance* forthcoming.
- LUDVIGSON, S. C. (1999): “Consumption and Credit: A Model of Time-Varying Liquidity Constraints,” *Review of Economics and Statistics*, 81, 434–447.
- MCCONNELL, M. M., AND G. PEREZ-QUIROS (2000): “Output Fluctuations in the United States: What Has Changed Since the early 1980s?,” *American Economic Review*, 90(5), 1464–1476.
- MCGRATTAN, E. R., AND E. C. PRESCOTT (2002): “Taxes, Regulations, and the Value of U.S. Corporations: A General Equilibrium Analysis,” Federal Reserve Bank of Minneapolis Staff Report No. 309.
- MEHRA, R., AND E. C. PRESCOTT (1985): “The Equity Premium Puzzle,” *Journal of Monetary Economics*, 15, 145–161.

- OTROK, C., B. RAVIKUMAR, AND C. H. WHITEMAN (2002): “Habit Formation: A Resolution of the Equity Premium Puzzle?,” *Journal of Monetary Economics*, 49, 1261–1288.
- QUANDT, R. E. (1960): “Tests of the Hypothesis that a Linear Regression Obeys Two Separate Regimes,” *Journal of the American Statistical Association*, 55, 324–330.
- ROMER, C. D. (1989): “The Pre-War Business Cycle Reconsidered: New Estimates of Gross National Product 1869-1908,” *Journal of Political Economy*, 97(1), 1–37.
- SHILLER, R. J. (1981): “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?,” *American Economic Review*, 71, 421–436.
- (2000): *Irrational Exuberance*. Princeton University Press, Princeton, NJ.
- SIEGEL, J. J. (1999): “The Shrinking Equity Premium,” *The Journal of Portfolio Management*, pp. 10–16.
- STOCK, J. H., AND M. W. WATSON (2002): “Has the Business Cycle Change and Why?,” in *NBER Macroeconomics Annual: 2002*, ed. by M. Gertler, and K. Rogoff. MIT Press, Cambridge, MA.
- VERONESI, P. (1999): “Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model,” *Review of Financial Studies*, 12(5), 975–1007.
- VISSING-JORGENSEN, A. (2002): “Limited Asset Market Participation and Intertemporal Substitution,” *Journal of Political Economy*, 110(4), 825–853.
- VISSING-JORGENSEN, A., AND O. P. ATTANASIO (2003): “Stock Market Participation, Intertemporal Substitution and Risk-Aversion,” *American Economic Review, AEA Papers and Proceedings*, 93(2), 383–391.
- WACHTER, J. (2003): “A Consumption Based Model of the Term Structure of Interest Rates,” Unpublished paper, New York University.
- WEIL, P. (1989): “The Equity Premium Puzzle and the Risk-Free Rate Puzzle,” *Journal of Monetary Economics*, 24(3), 401–421.
- WHITELAW, R. F. (2000): “Stock Market Risk and Return: An Equilibrium Approach,” *Review of Financial Studies*, 13(3), 521–547.

Table 1: Tests for Structural Breaks

Stock-Watson Test for Break in Variance					
	QLR statistic	p -value	Break Date	67% Confidence Interval	
Δc	14.34	0.0034	1992Q1	1991Q3	1994Q3
Bai-Perron Test for Break in Mean					
	$supF$ Test	p -value	Break Date	90% Confidence Interval	
$p - d$	33.85	< 0.01	1995Q1	1994Q1	1999Q3

Notes: This table reports results from structural break tests. The Quandt Likelihood Ratio test is described in detail in Appendix 1 of Stock and Watson (2002). The bottom panel reports Bai and Perron's (2003) $supF$ test statistic for a break in the mean of the log CRSP-VW price-dividend ratio. Both tests test the null hypothesis of no structural break against the alternative of a single structural break. The data are quarterly and span the period from 1952 to 2002.

Table 2: A Markov-Switching Model

x_t	μ_h	μ_l	σ_h^2	σ_l^2	ϕ	p_{hh}^μ	p_{ll}^μ	p_{hh}^σ	p_{ll}^σ
Δc	0.611 (0.061)	-0.469 (0.439)	0.541 (0.080)	0.151 (0.048)	0.084 (0.088)	0.971 (0.020)	0.772 (0.154)	0.995 (0.006)	0.992 (0.012)

Notes: This table reports the maximum likelihood estimates of the model

$$\Delta x_t = \mu(S_t) + \phi(\Delta x_{t-1} - \mu(S_{t-1})) + \epsilon_t$$

$$\epsilon_t \sim N(0, \sigma^2(V_t)).$$

We allow for two mean states and two volatility states. μ_h denotes the growth rate in the high mean state, while μ_l denotes the growth rate in the low mean state. σ_h^2 denotes the variance of the shock in the high volatility state and σ_l^2 denotes the variance of the shock in the low volatility state. S_t and V_t are latent variables that are assumed to follow independent Markov chains. The probabilities of transiting to next period's state j given today's state i are p_{ij}^μ and p_{ij}^σ , respectively. Standard errors are in parentheses. Standard errors are in parentheses. The data are quarterly and span the period from the first quarter of 1952 to the fourth quarter of 2002.

Table 3: Model Implications in 1990Q1 and 2002Q4

row	γ	λ	δ	$p_{j,j}^\sigma$	P/D_{90}	P/D_{02}	% of boom	$r_{90}^p(100)$	$r_{02}^p(100)$
1	25	6	0.9925	0.9999	35.06	68.44	119%	11.32	9.71
2	25	6	0.9925	0.9990	35.44	55.61	71%	11.26	10.17
3	25	6	0.9925	est	36.93	44.46	27%	11.14	10.83
4	25	6	0.9925	0.9900	39.39	45.91	25%	10.91	10.69
5	25	4	0.9925	0.9999	30.97	43.42	44%	7.69	6.54
6	27	5	0.9960	0.9999	32.71	51.02	68%	11.05	9.75


Notes: This table reports the model implications for asset prices using the estimated state probabilities in 1990:Q1 and 2002:Q4. λ is the leverage factor, δ is the discount rate and $p_{j,j}^\sigma$ is the probability that next period is a volatility state j given that today's state is volatility state j , $j \in \{l, h\}$. The elasticity of intertemporal substitution is set to 1.5 for all cases. P/D , $r^p(100)$, $r^f(100)$ and $r^f(1)$ are the price-dividend ratio and the 100-year risk premium, respectively. The entry for $p_{j,j}^\sigma$ labelled "est" show the results when the point estimates for all transition probabilities are used. All returns are annualized in percent. The variables with subscript "90" ("02") report the model's predictions using historical state probabilities in 1990:Q1(2002:Q4). The columns denoted "% of boom" reports the change of the P/D ratio from 1990:Q1 to 2002:Q4 in the model relative to the change of the CRSP-VW P/D ratio.

Table 4: Model with Cointegration

row	γ	λ	k	P/D -high vol.	P/D -low vol.	% of boom
1	40	6	0.005	20.01	49.27	149%
2	40	6	0.05	43.05	80.97	112%
3	40	6	0.20	75.47	95.65	77%
4	60	4	0.005	20.62	50.58	151%
5	60	4	0.05	42.96	80.70	110%
6	60	4	0.20	71.82	94.12	79%


Notes: This table reports price-dividend ratios for steady states with high and low consumption volatility for the model with cointegrated consumption and dividend processes. Consumption growth is assumed to be a random walk $\Delta c_{t+1} = \mu + \epsilon_{c,t+1}$ while dividend growth is forecastable by the log consumption-dividend ratio: $\Delta d_{t+1} = \mu + k(c_t - d_t) + \lambda \epsilon_{c,t+1} + \epsilon_{d,t+1}$. Consumption volatilities in the two steady states are the estimated values reported in Table 2. The volatility of ϵ_d is chosen to match the volatility of dividend growth to that in the data. The mean of growth is 2% (annualized). The values of the error correction coefficient k in the fourth column are annualized. The preference parameters δ and ψ are set to 0.996 and 1.5, respectively. The columns denoted “% of boom” reports the change of the P/D ratio from 1990:Q1 to 2002:Q4 in the model relative to the change of the CRSP-VW P/D ratio.

Figure 1: Price Ratios





Notes: This figure plots the CRSP-VW price-dividend ratio in the top panel, the S&P500 price-earnings ratio where earnings are a one-year lagged moving average in the middle panel and the S&P500 price-earnings ratio where earnings are a 10-year lagged moving average in the bottom panel. The earnings-ratios are from Robert Shiller's webpage. The quarterly data start in the first quarter of 1952 and include the most recent data available (the fourth quarter of 2002 for the price-dividend ratio and the third quarter of 2003 for the earnings ratios).

Figure 2: Growth Rates




Notes: This figure shows the growth rates of personal consumption expenditures. The lines in the plot correspond to the volatility regimes estimated from the Hamilton regime switching model. The data are quarterly and span the period from the first quarter of 1952 to the fourth quarter of 2002.

Figure 3: 5-Year Volatility Estimates and log Price Ratios




Notes: This figure plots the standard deviation of consumption growth and the average CRSP-VW log dividend-price ratio in 5-year windows. All series are demeaned and divided by their standard deviation. The data are quarterly and span the period from the first quarter of 1952 to the fourth quarter of 2002.

Figure 4: 5-Year Volatility Estimates and the log D/P Ratio – International Evidence



Notes: This figure plots the standard deviation of consumption growth and the average log dividend-price ratio in 5-year windows for ten countries. The plots in each each panel use the longest available data in each country. The data are from Campbell (2003).


Figure 5: GDP volatility and the D/P Ratio - Pre-war Evidence



Notes: This figure plots the standard deviations of GDP growth and the mean D/P ratio by decade starting in 1880 until 2000. Both series are demeaned and divided by their standard deviation. The GDP data are from Ray Fair's website (<http://fairmodel.econ.yale.edu/RAYFAIR/PDF/2002DTBL.HTM>) based on Balke and Gordon (1989). The dividend yield data is from Robert Shiller's website (http://aida.econ.yale.edu/~shiller/data/ie_data.htm).


Figure 6: State Probabilities

Consumption Growth



Notes: This figure plots the time series of estimated state probabilities. $P(\text{low variance})$ is the unconditional probability of being in a low consumption volatility state next period, calculated by summing the probability of being in a low volatility state and high mean state, and the probability of being in a low volatility state and low mean state. $P(\text{high mean})$ is calculated analogously. The data are quarterly and span the period from the first quarter of 1952 to the fourth quarter of 2002.


Figure 7: The Price-Dividend Ratio




Notes: The price-dividend ratio as a function of the probability that consumption volatility is low (left panel) and the probability that consumption mean is high (right panel). In the left panel, the probability that the consumption mean is high is set to be zero (solid line) or one (dashed line). In the right panel, the probability that consumption volatility is low is set to be zero (solid line) and one (dashed line). The probability of a change in the consumption volatility state is assumed to be .0001; otherwise the parameters of the endowment process are set equal to their maximum likelihood estimates. The rate of time preference $\delta = .9925$, the elasticity of intertemporal substitution, $\psi = 1.5$, risk aversion $\gamma = 25$ and leverage $\lambda = 6$.

Figure 8: Time Series of the P/D Ratio

Panel A: $\gamma = 25, \delta = .9925, \lambda = 6$




Panel B: $\gamma = 27, \delta = .996, \lambda = 5$




Notes: Time series of the log price-dividend ratio from the data and implied by the model. The probability of a change in the consumption volatility state is assumed to be .0001; otherwise the parameters of the endowment process are set equal to their maximum likelihood estimates. The elasticity of intertemporal substitution is $\psi = 1.5$. The data are quarterly and span the period from the first quarter of 1952 to the fourth quarter of 2002.

Figure 9: The Long-run Equity Premium




Notes: The conditional long-run equity premium as a function of the probability that consumption volatility is low. The probability that the consumption mean is high is set to be zero (solid line) or one (dashed line). The means is reported in annual terms. The probability of a change in the consumption volatility state is assumed to be .0001; otherwise the parameters of the endowment process are set equal to their maximum likelihood estimates. The rate of time preference $\delta = .9925$, the elasticity of intertemporal substitution, $\psi = 1.5$, risk aversion $\gamma = 25$ and leverage $\lambda = 6$.

Figure 10: Time Series of the Long-run Expected Return and Equity Premium




Notes: Time series of the 100-year expected equity return and equity premium implied by the model. The probability of a change in the consumption volatility state is assumed to be .0001; otherwise the parameters of the endowment process are set equal to their maximum likelihood estimates. The rate of time preference $\delta = .9925$, the elasticity of intertemporal substitution, $\psi = 1.5$, risk aversion $\gamma = 25$ and leverage $\lambda = 6$. The data are quarterly and span the period from the first quarter of 1952 to the fourth quarter of 2002.

Figure 11: Time series of volatility



Notes: Time series of conditional volatility of one-quarter ahead equity returns. Volatility is annualized, i.e. $2\sigma_t$. The probability of a change in the consumption volatility state is assumed to be .0001; otherwise the parameters of the endowment process are set equal to their maximum likelihood estimates. The rate of time preference $\delta = .9925$, the elasticity of intertemporal substitution, $\psi = 1.5$, risk aversion $\gamma = 25$ and leverage $\lambda = 6$. The data are quarterly and span the period from the first quarter of 1952 to the fourth quarter of 2002.

Figure 12: The Likelihood Function



Notes: This figure shows the log-likelihood function of the Hamilton regime switching model. The probabilities of remaining in the high (low) volatility state given that today's volatility state is high (low) are set to the same value, $P_sig(ii)$. The figure plots the log-likelihood as a function of $P_sig(ii)$. All other parameters are set the optimized values reported in Table 3. The data are quarterly and span the period from the first quarter of 1952 to the fourth quarter of 2002.