Great Moderation or Great Mistake: Can overconfidence in low macro-risk explain the boom in asset prices?

Tobias Broer, IIES, Stockholm University*
Afroditi Kero, European University Institute

June 2011
Preliminary. Please Do Not Quote.

Abstract

The fall in US macroeconomic volatility from the mid-1980s coincided with a strong rise in asset prices. Recently, this rise, and the crash that followed, have been attributed to overconfidence in a benign macroeconomic environment of low volatility. This paper introduces learning about the persistence of volatility regimes in a standard asset pricing model. It shows that the fall in US macroeconomic volatility since the mid-1980s only leads to a relatively small increase in asset prices when investors have full information about the highly persistent, but not permanent, nature of low volatility regimes. When investors infer the persistence of low volatility from empirical evidence, however, the model can deliver a strong rise in asset prices by up to 45%. Moreover, depending on the learning scheme, the end of the low volatility period leads to a strong and sudden crash in prices.

JEL classification: E, G.
Keywords: Macroeconomic Risk, Asset Prices

*Contact: tobias.broer@iies.su.se, afroditi.kero@eui.eu. Broer gratefully acknowledges financial support from the European Research Council under Advanced Grant 230574.

"*From the Great Moderation to the Great Conflagration: The decline in volatility led the financial institutions to underestimate the amount of risk they faced, thus essentially (though unintentionally) reintroducing a large measure of volatility into the market.*"

Thomas F. Cooley, Forbes.com, 11 December 2008

"*The stress-tests required by the authorities over the past few years were too heavily influenced by behavior during the Golden Decade. [...] The sample in question was, with hindsight, most unusual from a macroeconomic perspective. The distribution of outcomes for both macroeconomic and financial variables during the Golden Decade differed very materially from historical distributions.*"

Andrew Haldane, Bank of England, 13 February 2009

"*But what matters is how market participants responded to these benign conditions. They are faced with what is, in essence, a complex signal-extraction problem. But whereas many such problems in economics involve learning about first moments of a distribution, this involves making inferences about higher moments. The longer such a period of low volatility lasts, the more reasonable it is to assume that it is permanent. But as tail events are necessarily rarely observed, there is always going to be a danger of underestimating tail risks.*"

Charles Bean , European Economic Association, 25 August 2009

1 Introduction

The fall in macroeconomic volatility in the United States and other countries from the mid-1980s, later coined the "Great Moderation", coincided with a strong rise in asset prices. After the economic crisis that started in 2007, both policy-makers and academics attributed part of this rise, and the subsequent fall in prices, to overconfidence in the benign macroeconomic environment of the "golden decade" (Haldane et al 2009). According to this argument, in their attempt to infer the distribution of future shocks on the basis of observed data, investors overestimated the persistence of a low volatility environment, thus bidding up the price of assets beyond their fundamental value. This paper introduces learning about the persistence of volatility regimes in a standard asset pricing model. It shows that the fall in US macroeconomic volatility since the mid-1980s only leads to a relatively small increase in asset prices when investors have full information about the highly persistent, but not permanent, nature of low volatility regimes. When investors optimally infer the persistence of low volatility from empirical evidence using Bayes' rule, however, the model can deliver a much stronger rise in asset prices similar to that observed in the data. Moreover, depending on the learning scheme, the end of the low volatility period leads to a strong and sudden crash in prices.

Previous studies have found that a fall in macroeconomic volatility of the magnitude observed in the United States between the late 1980s and early 1990s would have to be, essentially, permanent to explain a significant proportion of the subsequent boom in equity prices (Lettau et al 2007). However, while some authors have attributed the great moderation to stuctural changes in developed economies that are indeed very persistent, or potentially permanent, such as central bank independence, the increase in world trade, or the development of new financial products to diversify risk, others have pointed to its transitory origins, such as an unusually long period of small exogenous shocks ("good luck") that hit western economies during this period (see section 2 for more detail). Moreover, similar uncertainty about the origins and persistence of the Great Moderation can be found in analysis by market participants. After the economic crisis that started in 2007, both policymakers and academics have attributed the boom in asset prices and their subsequent crash to overconfidence of investors in a benign macroeconomic environment of low volatility (Cooley 2008, Haldane 2009, Bean 2009). For example, Haldane (2009) argues that data availability was such that the high volatility period preceding the Great Moderation was often neglected in the estimation of quantitative asset pricing models. Similarly, Bean (2009) attributes part of the boom and bust in asset prices to rising investor confidence that the low volatility environment would be permanent.

This paper looks at the behaviour of asset prices in an environment where investors have to infer the persistence of changes in macro-volatility from the data. Specifically, we interpret the economic experience of the US economy after the second World-War as consisting of realisations of high and low volatility regimes, whose transition probabilities are unknown to investors. This allows us to analyse the behaviour of asset prices in a general equilibrium where investors use optimal bayesian learning rules to infer the persistence of periods of low macro-volatility. Specifically, we study an economy where investors simply update their priors about transition probabilities in line with observed realisations of high and low volatility regimes according to Bayes' rule (Cogley and Sargent (2008)). The model delivers a boom and bust in asset prices much stronger than in the absence of uncertainty about transition probabilities, and explains about ~~XX~~ percent of the boom in US asset prices between the early 1980s and their fall that started in 2007. As a robustness exercise, we look at alternative learning schemes. First, we analyse an alternative Bayesian learning rule based on two popular hypotheses that explained the great moderation either by an unusually long sequence of small shocks ("good luck") or by permanent structural change ("good policy"). Under this alternative learning scheme we assume that transition probabilities during normal times are known, but that there is a small ex ante-probability that a low-volatility regime turns out to be permanent. This scheme leads to asset price dynamics that are qualitatively similar, but even stronger in magnitude, compared to our benchmark learning scheme. Finally, we also look at non-optimal, "adaptive" learning schemes, where investors use simple statistical rules to update their inference about volatility on the basis of observed data. This ad hoc learning results in strong overvaluation of assets, relative to the prices implied by full information about data generating process, but does not yield a strong crash after the end of the Great Moderation (which we identify with the beginning of the economic crisis in 2007).

This paper is most related to the literatures on asset pricing with time-varying volatility, and with learning about features of the economic environment. After earlier papers on the effect of changes in economic volatility for asset prices in stationary environments (Bonomo and Garcia (1994, 1996) and Drifil and Sola (1998)), more recently Bansal and Lundblad (2002)), Lettau et al (2008) ask whether a persistent change to a low macro-volatility regime can help explain the boom in US asset prices of the 1990s and early 2000s. They find that the low volatiilty environment would have to be, essentially, permanent to explain the data.[1] Most papers look at environments where agents learn about the mean growth rate of output or consumption. For example, Cogley and Sargent (2008) assume that after the Great Depression, investors had pessimistic priors about the probability of transitions from a high to a low-growth state. Using a learning mechanism that is identical to one of those analysed in our study, they show how this may explain a sustained

[1]Lower macro-volatility is only one item on a long list of potential reasons behind the asset price boom of the 1990s and 2000s. Others are a lower equity premium (Blanchard (1993), Jagannathan, McGrattan, and Scherbina (2000), Fama and French (2002)), higher long-run growth (Jagannathan, McGrattan, and Scherbina (2000), Fama and French (2002), Campbell and Shiller (2003), although Siegel (1999) finds no evidence for this), stronger intangible investment in the 1990s (Hall (2000)) saving during the 1990s by the baby boom generation (Abel (2003)), redistribution of rents towards owners of capital (Jovanovic and Rousseau (2003)) or reduced costs of stock market participation and diversification (Heaton and Lucas (1999), Siegel (1999), Calvet, Gonzalez-Eiras, and Sodini (2003)).

fall over time from an initially high equity premium, as learning leads to rising confidence in high growth. More recently, Adam and Marcet (2010), show how learning about an unknown process for cum-dividend equity returns introduces a self-referential element in equity prices that leads to persistent bubbles and occasional crashes. There has also been a growing number of contributions that study learning about risk. Branch and Evans (2010) employ self-referential adaptive learning about asset prices and return volatility in order to explain high frequency booms and busts in asset prices. Weitzman (2007) adopts a consumption-based asset-pricing model and replaces rational expectations with Bayesian learning about consumption growth rate volatility, which allows him to solve a number of asset pricing puzzles.

Most relevant for this paper are two studies that link the asset price boom and bust of 1990s and 2000s to learning about regime changes in key parameters of the economic environment. ~~Mendoza et al (2010)~~ study a partial equilibrium model where investors face an exogenous leverage constraint that follows a two-state markov process with unknown transition probabilities. Assuming Bayesian learning as in Cogley and Sargent (2008), the authors show that with little prior information, the observation of a string of high leverage periods can lead to overoptimism about their persistence, and thus a boom in asset prices, leverage and consumption which crashes abruptly once the economy switches back to a tighter constraint. While one of our learning mechanisms also follows Cogley and Sargent (2008), we analyse, in general equilibrium, exogenous changes in macro-volatility, rather than in regimes of financial regulation. This focus is similar to Lettau et al. (2008), who also study the asset price effect of changes in macro volatility-regimes under limited information about the environment. Particularly, while knowing all parameters of the environment, including the persistence of volatility regimes, agents in their model ignore whether the economy is currently in a high or low volatility regime.[2] Rather than incorporating learning explicitly, they then calculate asset prices given the sequence of posterior state probabilities implied by a regime-switching model estimated on post-war consumption data for the US. Our work differs to theirs in several ways: first, based on our reading of the academic literature and the business press (see section 2), we assume agents were sure that the US economy had experienced a change in aggregate volatility with the Great Moderation, but were uncertain about its persistence. Second, we explicitly look at different optimal and ad hoc learning mechanisms to ask how their

[2]Lettau et al (2008) also have two states of different mean growth, leaving four states of the economy in total.

implied asset price behaviour relates to that of an environment of full information about transition probabilities, and US data. Importantly, in Lettau et al (2008) asset prices are, essentially, weighted averages of full information prices. The model-implied prices are therefore always lower than those that would prevail in the most benign low-volatility regime with full information. In our model, on the other hand, agents may overestimate the persistence of the Great Moderation, leading to significant overvaluation of asset prices relative to full information.

The rest of the paper is organised as follows. To motivate our approach in more detail, section II reviews the main empirical facts on the Great Moderation as well as the debate about its causes among academics and market participants. Section III presents the model. Section IV gives the main results and section V shows how robust these are to changes in the underlying assumptions.

2 Motivation: The Great Moderation, its uncertain cause and persistence, and the boom in asset prices

2.1 Asset Prices and the Great Moderation: Stylized Facts

Figure 1 and figure 2 present the time series of real GDP and consumption growth rates and their corresponding volatilities (computed as the standard deviation over 10-quarter rolling windows). Both series exhibit a significant and abrupt fall in volatility, which persisted until the beginning of the current crisis. The timing of the drop, however, differs: while GDP volatility declined around the middle of the 1980s, the fall occured somewhat later, at the beginning of the 1990s, for consumption growth.

Enter Figure 1 about here

Using quarterly data from 1952$Q2$ to 2010$Q2$, table 1 and 2 quantify this decline in volatility for different subperiods.[3] The end dates of the first subperiod are 1984$Q1$ for GDP (see McConnell and Perez-Quiros (2000)) and 1992$Q1$ for consumption (Lettau et al. (2008)), while the second ends with the start of the financial crisis in 2007. Whereas there

[3] See the Data Appendix for a more detailed description of the data series.

Moments of GDP growth

Date	*Mean*	*StDev*
$1952Q2 : 1983Q4$	0.53%	1.1%
$1984Q1 : 2006Q4$	0.51%	0.51%
$2007Q1 : 2010Q2$	−0.16%	0.90%

Table 1: The table reports sample estimates for the mean and the standard deviation of real GDP growth rate. Output is defined in real per-capita terms. The GDP and the population data are taken from Bureau of Economic Analysis. The data are quarterly and span the period $1952Q2 - 2010Q2$. The estimates are in percent.

Moments of Consumption Growth

Date	*Mean*	*StDev*
$1952Q2 : 1991Q4$	0.57%	0.82%
$1992Q1 : 2006Q4$	0.61%	0.36%
$2007Q1 : 2010Q2$	−0.19%	0.50%

Table 2: The table reports sample estimates for the mean and the standard deviation of real consumption growth rate. Consumption is defined in real per-capita terms. The consumption and the population data are taken from BEA. The data are quarterly and span the period $1952Q2 - 2010Q2$. The estimates are in percent.

is almost no change in mean growth across the first two subperiods, there is a significant fall in volatility of more than 50 percent for both aggregate output and consumption growth. In the third sub-sample that covers the recent crisis, we observe a sharp decrease in mean growth for both GDP and consumption and a strong rise in volatility.

Figure 2 shows how the decline in macroeconomic volatility coincided with a strong rise in asset prices and a fall in the US price-dividend ratio for the $S\&P$ 500. Importantly, this fall was much less abrupt than the decline in volatility itself. Again, table 3 quantifies this effect for 3 subperiods, choosing $1995Q1$ as the start of the second subperiod (Lettau et al (2008)). The price-dividend ratio more than doubled across the first two periods, but fell back to levels seen in the 1960s and seventies with the start of the recent crisis.

Enter Figure 2 about here.

US Equity Prices

Date	*Mean* $\frac{p}{d}$	*Mean* $\frac{p}{e}$
$1952Q2 : 1994Q4$	27.49	15.54
$1995Q1 : 2006Q4$	62.25	30.03
$2007Q1 : 2010Q2$	46.48	21.18

Table 3: The table reports sample estimates for the mean of the price-dividend and the price-earning ratio based on the $S\&P$ 500. Consumption is defined in real per-capita terms. The data are taken from the Robert Shiller's homepage. The data on prices are monthly instead the data on dividends and on the price-earning ratio are quarterly. We calculate quarterly estimates for the prices by taking quarterly averages over the monthly data. The data span the period $1952Q2 - 2010Q2$. The estimates are in percent.

2.2 Uncertainty about Origin and Persistence of the Great Moderation

By the second half of the 1990s, both the academic literature (Kim and Nelson (1999), McConnell and Perez-Quiros (1997, 2000)) and the business press had noticed a break in the volatility properties of US output growth around the middle of the preceding decade. Somewhat later, a similar decline in volatility was documented for a broader set of US macro-economic variables (Blanchard and Simon (2001), Stock and Watson (2002)), as well as for other industrial countries (Stock and Watson 2003). However, although the Great Moderation itself had become a stylised fact, there was no consensus about its causes. While some authors explained the phenomenon by changes in the structure of industrial economies, such as financial innovation (Dynan et al 2006), improved inventory management, or financial and trade liberalisation (see Wachter (2006) for a brief summary), the two perhaps most prominent hypotheses competed under the heading of "Good Policy or Good Luck?". Specifically, following the seminal article by Stock et al (2003), several studies[4] used time-varying VAR models to find that a string of unusually small shocks, rather than changes in their transmission to main macroeconomic variables or in the conduct of monetary policy, were at the root of the decline in macro-volatility. Against this, both academics (Benati et al 2008) and policymakers (Tucker 2005, Bernanke 2004) argued that reduced-form models were likely to mistake effects of

[4]Giorgio E. Primiceri (2005), Christopher Sims and Tao Zha (2006), and Luca Gambetti, Evi Pappa, and Fabio Canova (2008)

improved monetary policy, such as more stable but unobserved inflation expectations, for changes in the variance-covariance-properties of exogenous economic shocks. For example, Bernanke (2004) argued that "some of the benefits of improved monetary policy may easily be confused with changes in the underlying environment". Importantly, the lack of consensus about the causes of the observed fall in macro-volatility left it unclear wether the phenomenon was likely to be permanent, as suggested by structural change or possibly improved policy environments, or transitory, in line with the "good luck" hypothesis.

How did market participants perceive the Great Moderation and its effect on prices? Investment analysts explicitly attributed part of the observed fall in the equity risk premium since the late 1980s to the decline in macro-volatility. For example, Goldman Sachs research (2002) noted that an estimated 8 percentage point fall in the risk premium since the 1970s was "underpinned by dramatic improvements in the economic environment. Inflation fell sharply, and the volatility of GDP growth, inflation and interest rates all declined significantly." (p. 2). But while investors acknowledged the effect of the Great Moderation on asset prices, they were also aware of the uncertain persistence of this low-volatility environment, and thus of the decline in equity premia. For example, regarding risk premia in fixed income securities, Unicredit analysts (2006) argued that "the ongoing deterioration in surprise risk should be seen as one of the arguments behind the declining risk premium. Whether this is due to a more effective central bank policy, a major improvement in the forecast ability of economic observers around the globe, sheer luck or maybe a mix of all three factors can't finally be answered." (p. 10). Researcher at JP Morgan (2005), on the other hand, attribute most of the fall in volatility to a changed orientation of policymakers towards a "Stability Culture", which, however, is not certain to persist.

We draw three conclusions from this evidence: first, the fall of macro-volatility since the mid-1980s was accepted as a stylised fact, and widely seen as a contributing factor to higher asset prices during the 1990s and 2000s. Second, as Lettau et al (2008) have shown, standard asset pricing models predict significantly higher asset prices only when a fall in volatility is permanent, or extremely persistent. Finally, during the Great Moderation it was exactly this persistence that investors were uncertain about. This paper therefore puts learning about the persistence of volatility changes at the center of its analysis. Particularly, it asks two questions: first, can rising confidence in the persistence of low volatility explain the strong and gradual rise in asset prices during the Great Moderation. And second, how much of this rise is due to overconfidence in the great moderation, equivalent to an overvaluation of assets relative to the prices implied by the true value of persistence.

3 The model

This section presents a standard general equilibrium asset pricing model and adds learning about the persistence of volatility regimes.

3.1 Preferences

We consider an endowment economy with an infinitely-lived representative agent who solves the following problem

$$\max_{C_t, S_t} U_t \tag{1}$$

$$s.t.\ S_t P_t + C_t = S_{t-1} P_t + D_t \tag{2}$$

$$S_{-1}\ given \tag{3}$$

where U_t denotes an expected utility index at time t, C_t denotes consumption, S_t are the agent's stockholdings, P_t is the stock price and D_t are dividends. Preferences U_t are as in Epstein and Zin (1989, 1991) or Weil (1989)

$$U(C_t) = [(1-\beta)C_t^{\frac{1-\gamma}{\alpha}} + \beta(E_t U_{t+1}^{1-\gamma})^{\frac{1}{\alpha}}]^{\frac{\alpha}{1-\gamma}}$$

where E_t is the mathematical expectation with respect to the agents subjective probability distribution conditional on period t information, $\alpha = \frac{1-\gamma}{1-\frac{1}{\psi}}$, γ is the coefficient of relative risk aversion, and ψ the intertemporal elasticity of substitution.
The first order condition associated to this problem is

$$P_t = E_t^s[M_{t+1}(P_{t+1} + D_{t+1})] \tag{4}$$

where M_{t+1} is the stochastic discount factor, which, with Epstein-Zin preferences, equals

$$M_{t+1} = (\beta(\frac{C_{t+1}}{C_t})^{-\frac{1}{\psi}})^{\alpha} R_{w,t+1}^{\alpha-1}$$

Here, $R_{w,t+1}^{\alpha-1}$ is the return on the aggregate wealth portfolio of the representative agent, whose returns equal aggregate consumption.

3.2 The Processes for Consumption and Dividend Growth

We choose a simple and transparent way to model an economy that goes through periods of low and high macro-volatility by assuming that the log consumption follows an exogenous random walk with drift

$$g_t = \Delta \ln C_t = \bar{g} + \varepsilon_t$$

where $\bar{g}$ is constant mean consumption growth.[5] Shocks ε_t are independently normally distributed, and their variance follows a two-state markov process

$$\varepsilon_t \sim N(0, \sigma_t^2) \; \sigma_t^2 \in \{\sigma_l^2, \sigma_h^2\}$$

The transition probabilities for the markov process are

$$\begin{aligned} \Pr(\sigma_{t+1}^2 &= \sigma_l^2 \mid \sigma_t^2 = \sigma_l^2) = F_{ll} \\ \Pr(\sigma_{t+1}^2 &= \sigma_h^2 \mid \sigma_t^2 = \sigma_h^2) = F_{hh} \end{aligned}$$

which yields the transition probability matrix as

$$\mathbf{F} = \begin{bmatrix} F_{ll} & 1 - F_{ll} \\ 1 - F_{hh} & F_{hh} \end{bmatrix}$$

Following Mehra and Prescott (1985), and in line with the endowment nature of the economy, it is common to assume that dividend flows equal consumption flows. To capture the higher empirical volatility of dividends, we follow Campbell (1986), Abel (1999), or Lettau et al (2008), and use a generalised version of the standard model where shocks to dividend growth are a multiple of those to consumption

$$\Delta lnD_t = \overline{g} + \lambda\varepsilon_t \quad \lambda \geq 1$$

Dividends thus follow the same volatility pattern as consumption, but are on average more volatile.

[5] Previous studies have looked at time-variation in $\bar{g}$. Here, we assume $\bar{g}$ to be constant over time, and concentrate instead on changes over time in the variance of shocks ε_t.

3.3 Full Information Price-Dividend Ratios

We can use the first-order condition for share holdings to express the price-dividend ratio $p_t = \frac{P_t^D}{D_t}$ as

$$p_t = (\frac{P_t^C}{C_t})^{1-a} E_t[\beta^a (\frac{C_{t+1}}{C_t})^{-\frac{\alpha}{\psi}+\alpha-1} (\frac{P_{t+1}^C}{C_{t+1}} + 1)^{\alpha-1} (p_{t+1}+1)\frac{D_{t+1}}{D_t}] \tag{5}$$

where $\frac{P_t^C}{C_t}$ is the price-consumption ratio, with P_t^C the price of a claim to aggregate consumption and $\frac{P_t^C}{C_t}$ equals $\frac{P_t^D}{D_t}$ whenever $\lambda = 1$. When the agent knows the true structure of uncertainty, given the random walk nature of consumption and dividends, the price-dividend and price-consumption ratios are functions only of the volatility state, $p_t = p(\sigma_t^2)$, and thus non-random conditional on σ_t^2. We can thus simplify (5) by taking expectations across realisations of log-normal consumption and dividend growth conditional on σ_{t+1}^2, which gives a recursive expression for the price-dividend and ratios.[6]

Note that in the special case when $\lambda = 1$ both consumption and dividend growth follow the same the log-normal distribution. With $\psi = \frac{1}{\gamma}$ (CRRA preferences), this yields an analytical solution to the vector of price-dividend ratios p as

$$p = \beta F(1+p)e^{(-\gamma+1)\bar{g}}(e^{(-\gamma+1)^2 2\sigma_i^2}] \tag{10}$$

$$= \mathbb{F}\beta F e^{(-\gamma+1)\bar{g}}(e^{(-\gamma+1)^2 2\sigma_i^2}] \tag{11}$$

where $\mathbb{F} = [I - \beta F[\begin{matrix} \sigma_l^2 & 0 \\ 0 & \sigma_h^2 \end{matrix}]]^{-1}$.

[6]Specifically, $p(\sigma_t^2)$ is defined by

$$p(\sigma_t^2 = \sigma_i^2) = \rho_i^{1-\alpha}\beta^\alpha e^{(-\frac{\alpha}{\psi}+\alpha)\bar{g}} \tag{6}$$

$$\left(F_{ii} e^{\frac{(-\frac{\alpha}{\psi}+\alpha-1+\lambda)^2}{2}\sigma_i^2}(1+\rho_i)^{\alpha-1}(1+p_i) + F_{ij} e^{\frac{(-\frac{\alpha}{\psi}+\alpha-1+\lambda)^2}{2}\sigma_j^2}(1+\rho_j)^{\alpha}(1+p_j)\right) \tag{7}$$

where $\rho = \frac{P_t^C}{C_t}$ follows

$$\rho^\alpha(\sigma_t^2 = \sigma_i^2) = \rho_i^{1-\alpha}\beta^\alpha e^{(-\frac{\alpha}{\psi}+\alpha)\bar{g}} \tag{8}$$

$$\left(F_{ii} e^{\frac{(-\frac{\alpha}{\psi}+\alpha)^2}{2}\sigma_i^2}(1+\rho_i)^{\alpha-1} + F_{ij} e^{\frac{(-\frac{\alpha}{\psi}+\alpha)^2}{2}\sigma_j^2}(1+\rho_j)^{\alpha}\right) \tag{9}$$

.

3.4 Learning and Subjective Beliefs

To study whether a long spell of σ_l can lead to a boom in asset prices via overconfidence in the persistence of a low-volatility environment, we assume that the representative agent does not know the full probabilistic structure of the economy. Specifically, the agent knows that dividend growth is log-normal with mean $\bar{g}$ but ignores some information about its variance. Particular, she uses observations on realised dividends to infer the transition probabilities between high and low volatility states. A later sensitivity analysis studies how the results change when assuming other Bayesian or non-Bayesian learning schemes. In this section, we look at an environment where the agent learns about the transition probabilities between volatility states F_{hh} and F_{ll}. The model is thus very similar to Cogley and Sargent (2008), with the difference that the agent learns about transitions between volatility states, rather than mean-growth states. More specifically, the agent knows the structure of the model and all parameter values except the true transition probabilities F_{hh} and F_{ll}. Every period, she observes a dividend realization and the distribution that this specific realization is drawn from, parameterised by σ_t^2. The agent thus forms a best guess about F_{hh} and F_{ll} on the basis of the history of volatility-states $\Sigma^t = \{\sigma_t^2, \sigma_{t-1}^2, ..., \sigma_2^2, \sigma_1^2\}$.

We assume that the agent has independent beta-binomial prior distributions about F_{hh} and F_{ll}

$$f_0(F_{hh}, F_{ll}) \propto f_0(F_{hh}) f_0(F_{ll})$$

with

$$\begin{aligned} f(F_{hh}) &= f(F_{hh} \mid \Sigma^0) = beta(n_0^{hh}, n_0^{hl}) \propto F_{hh}^{n_0^{hh}-1}(1 - F_{hh})^{n_0^{hl}} \\ f(F_{ll}) &= f(F_{ll} \mid \Sigma^0) = beta(n_0^{ll}, n_0^{lh}) \propto F_{ll}^{n_0^{ll}-1}(1 - F_{ll})^{n_0^{lh}}. \end{aligned}$$

where Σ^0 denotes a prior belief about frequencies n_0^{ij} of transitions from state i to state j.

The agent updates this prior on the basis of the likelihood function for the history of volatility states Σ^t conditional on F_{hh} and F_{ll}, which is the product of two independent binomial density functions, thus

$$L(\Sigma^t \mid F_{hh}, F_{ll}) \propto L(\Sigma^t \mid F_{hh}) L(\Sigma^t \mid F_{ll})$$

where

$$L(\Sigma^t \quad | \quad F_{hh}) = binomial(F_{hh}, F_{hl}) \propto F_{hh}^{n_t^{hh} - n_0^{hh}} (1 - F_{hh})^{n_t^{hl} - n_0^{hl}}$$
$$L(\Sigma^t \quad | \quad F_{ll}) = binomial(F_{ll}, F_{lh}) \propto F_{ll}^{n_t^{ll} - n_0^{ll}} (1 - F_{ll})^{n_t^{lh} - n_0^{lh}}$$

Here, n_t^{ij} is a "counter", that equals the number of transitions from state i to state j up to time t plus the prior frequencies n_0^{ij}. The posterior kernel is the product of the beta prior and the binomial likelihood function,

$$f(F_{hh}, F_{ll} \quad | \quad \Sigma^t) \propto \underbrace{L(\sigma_t^2 \mid F_{hh}, F_{ll})}_{Likelihood} \cdot \underbrace{f(F_{hh}, F_{ll} \mid \Sigma^{t-1})}_{Prior}$$

which after normalizing by $M(\Sigma^t) = \int \int F_{hh}^{n_t^{hh}-1} (1 - F_{hh})^{n_t^{hl}-1} F_{hh}^{n_t^{hh}-1} (1 - F_{hh})^{n_t^{hl}-1} dF_{hh} F_{ll}$ yields the posterior density function as the product of independent *Beta* distributions

$$f(F_{hh} \quad | \quad \Sigma^t) = beta(n_t^{hh}, n_t^{hl}) \propto F_{hh}^{n_t^{hh}-1} (1 - F_{hh})^{n_t^{hl}-1}$$
$$f(F_{ll} \quad | \quad \Sigma^t) = beta(n_t^{ll}, n_t^{lh}) \propto F_{ll}^{n_t^{ll}-1} (1 - F_{ll})^{n_t^{lh}-1}$$

Note that in this context the counters n_t^{ij} are sufficient statistics for the posterior.

Let $p(\sigma_t^2, F)$ denote the price-dividend ratio when the transition probability matrix is F. Following Cogley and Sargent (2008), p_t^{BL}, the vector of price-dividend ratios under Bayesian learning about transition probabtilities can then be written as

$$p_t^{BL} = \int p(\sigma_t^2, F) f(F, \Sigma^t) dF \tag{12}$$

where $f(F, \Sigma^t)$ is the posterior distribution of F[7]. Note that for given F_{hh}, F_{ll}, $p(\sigma_t^2, F)$ is described by the same pair of equations as under full information ((9), (7)). And the law of iterated expectations implies that we can compute $p(\sigma_t^2, F)$ as a fixed point of these two equations. p_t^{BL} can then easily be calculated by numerical integration across the independent beta posteriors for F_{hh}, F_{ll}.

[7]For a derivation of equation (1) see Appendix B.

4 Quantitative Results for the Benchmark Economy

4.1 The exercise

This section presents the results of numerical simulations to answer the two main questions of this paper: Can learning about the persistence of the great moderation explain the observed boom and bust in US asset prices? And can overconfidence in this persistence lead to an overvaluation of assets, and a larger fall in prices at the end of the low-volatility period, relative to the case of full information. To answer these questions we analyse a scenario that is similar to the economic experience of the US after World-War II. Particularly, we interpret this experience as a long realisation of high volatility followed by the Great Moderation that ends with the recent crisis. Our data generation process thus consists of three sequences of shocks corresponding to three subperiods of different consumption growth volatility σ_t^2. Specifically, our analysis starts with a high volatility regime in 1952Q3. Since in our highly stylised model there is no distinction between consumption and GDP, we use a start date for the Great Moderation at the beginning of 1984, as suggested by the fall in GDP volatility, but also look at later dates as suggested by the consumption growth series. In line with the observed rise in volatility in figure 1, we locate the end of the Great Moderation at the beginning of 2007, the start year of the crisis. To compute the fall in asset prices around this end of the Great Moderation we also make the stronger assumption that the economy returned to the high volatility environment observed before the Great Moderation. This assumption is largely heuristical. It allows us to isolate the crash in asset prices implied by the disappearing overconfidence in the Great Moderation from other factors that this paper abstracts from.

4.2 Parameter choice

4.2.1 Preferences

As Bansal and Yaron (2004) have shown, for a rise in consumption volatility to increase asset prices with Epstein-Zin preferences, the intertemporal elasticity of substitution ψ has to be greater than unity. We thus follow Lettau et al (2008) and set $\psi = 1.5$. For our statements about the size of boom and bust to be interesting, the model has to deliver a level of asset prices that is approximately equal to the data in the period before the

Great Moderation. Rather than changing parameters across different learning rules to target asset prices exactly, however, we choose $\beta = 0.99325$ to target an interest rate of 2 percent p.a. (which varies very little across models), and set $\gamma = 30$ which yields equity prices that are, on average across models, close to US data, but not exactly equal to it for any particular model.

4.2.2 The Process for Consumption and Dividends

The consumption process in this model is characterised by three parameters: constant mean growth $\overline{g}$, and the variances in the two subperiods σ_h^2, σ_l^2, which we estimate directly from quarterly data on US personal real per capita consumption expenditure, using the subperiods from table 1. This yields mean growth 0.6 percent per quarter and variances of 0.82 percent and 0.37 percent respectively.

To capture the underlying uncertainty about the persistence of the Great Moderation, we choose a particularly simple and transparent calibration of the transition matrix $\mathbf{F}$, by choosing F_{ll}, F_{hh} such that the expected durations of high and low volatility regimes equal the subperiods identified from US data. Specifically, we set $F_{ii} = 1 - \frac{1}{T_i}$, where T_l, T_h are the durations of the Great Moderation and the high-volatility period preceding it, which yields transition matrices equal to

$$\mathbf{F}^{cons} = \begin{bmatrix} 0.989 & 1-0.989 \\ 1-0.992 & 0.992 \end{bmatrix}$$

It is interesting to note that these transition probabilities are extremely close to those in Lettau et al (2008), based on an estimated markov process on the same data.[8]

Unless otherwise mentioned, we set $\lambda = 4.5$ as suggested by Lettau et al (2008) on the basis of the relative volatility of US consumption and dividends. Table 4 summarises the parameters of preferences and the endowment process for the baseline model.

[8] Their point estimates are

$$\mathbf{F} = \begin{bmatrix} 0.991 & 1-0.991 \\ 1-0.994 & 0.994 \end{bmatrix}$$

Their process, however, is more complex, as they also include uncertainty about mean growth.

Parameter Values for the Baseline Model		
Preferences		
β	0.99325	Discount Factor
γ	30	Relative Risk Aversion
ψ	1.5	Intertemporal Elasticity of Substitution
Endowment Process		
$\bar{g}$	0.0059	Mean of Consumption Growth
σ_l	0.0037	Low Volatility of Consumption Growth
σ_h	0.0082	High Volatility of Consumption Growth
λ	4.5	Leverage

Table 4: Parameter values in the benchmark model.

Prior for the benchmark model with Learning		
n_0^{ll}	3.6	Prior frequency for low-to-low transitions
n_0^{lh}	0.4	Prior frequency for low-to-high transitions
n_0^{hh}	36	Prior frequency for high-to-high transitions
n_0^{hl}	4	Prior frequency for high-to-low transitions

Table 5: Learning parameters in the benchmark model

4.2.3 Learning Parameters

In our benchmark Bayesian learning scheme, the only additional free parameters are the prior probabilities. We use relative frequencies that imply high but not extreme persistence of both regimes equal to $f_{ii} = 0.9$. In line with the relative lack of prior knowledge about the Great Moderation, however, in our benchmark case we choose weak priors by giving agents relatively little information about the transitions in a low-volatility regime (equivalent to 1 year of data). In the high volatility regime, with its longer history, we choose 10 years. In another scenario, we look at stronger priors (equivalent to 5 and 20 years of data, respectively).

Table 4 summarises the parameter values.

Asset prices - Benchmark Model

	Boom	Overvaluation	Bust
Full Info	15%	0	15%
Learning	44%	28%	34%

Table 6: "Boom" denotes the increase in prices until the end of the Great Moderation relative to the high-volatility regime preceding it (computed over windows of 5 years). "Overvaluation" is the overvaluation at the end of the Great Moderation relative to the prices under full information. And "Bust" is the fall in prices in the first period after the Great Moderation.

4.3 Asset Price Dynamics: Mean and Variance Effects of Learning and Uncertainty about Transition Probabilities

Figure 4 presents the time path of the PD ratio in US data in the upper panel. The bottom panel depicts both the PD ratio with learning (solid lines) and those under full information about the the data generating process (dashed lines). The first thing to note is that with and without learning, the model delivers realistic levels of asset prices before the 1980s. Together with a model-implied interest rate of 2%, which is independent of regime and the information structure, the benchmark calibration thus delivers a realistic equity premium. Without learning, however, there is no sustained asset price boom during the Great Moderation. Rather, the model with full information delivers a small jump in prices of around 15% at the beginning of the Great Moderation. With learning, on the other hand, asset prices during the Great Moderation continue to rise after an initial jump. While the model is not able to replicate the hump-shape in PD ratios during the Great Moderation, nor the more than two-fold rise in prices between the early 1980s and 2007, it can explain a strong sustained boom in prices until the early 2000s of around 45%, and a subsequent bust of similar magnitude. The results are summarised in table 6.

The model with learning thus differs from that with full information both in terms of the average level of asset prices and in terms of its dynamics. What is behind this? The slight fall in prices before the great moderation, and the concave shape after the jump that follows the fall in volatility in the 1980s, are very intuitive: as agents observe how the economy remains in the same regime of low or high volatility, they increase their persistence estimates from their relatively loose and moderate prior centered around 0.9.

This leads to a fall in prices in the high volatility regime, as agents infer a falling likelihood of moving to a more benign macroeconomic climate of lower volatility. Similar reasoning explains the rise in prices during low volatility. Finally, as information accumulates, the marginal contribution of an additional observation to persistence estimates falls, explaining the concave shape of prices during the Great Moderation. However, with a moderate prior of $F_{ll}^0 = 0.9$, prices at the beginning of the great moderation should, in principle be lower, as agents anticipate a move back to the high volatility environment of the 1960s and 70s with higher probability than under full information. Instead, the model with learning predicts asset prices to be around slightly higher than with full information at the beginning of the Great Moderation. Why is this? The answer follows from the fact that, in addition to the 'mean effect' of a less persistent average prior, there is an additional effect of uncertainty about transition probabilities on price levels, which results from the variance of posterior persistence estimates around their expected value. Particularly, PD ratios turn out to be a strongly convex function of persistence at high values of f_{ll}. This implies a Jensen's inequality effect, whereby PD ratios with uncertainty about persistence parameters are higher than without. And the difference is larger the looser are priors, and the higher is therefore the variance of the posterior estimates.

Figure 5 looks at the simplified case of symmetric transition probabilities ($f_{ll} = f_{hh}$), to illustrate both effects. The solid lines depict PD ratios in the absence of persistence. As persistence rises, high-volatility prices fall, as agents are less willing to pay for assets whose payoffs they anticipate to remain volatile with a larger probability. Interestingly, low-volatility prices initially fall, but rise strongly for high values of persistence above 0.99. This non-linearity of the certainty price leads to higher level of prices as priors about persistence become looser. Figure 5 demonstrates this by showing the PD ratios when decreasing the size of the parameters in the beta distribution n^{ii}, n^{ij} while keeping their ratio constant.

Figure 5, however, does not explain why PD ratios are a convex function of persistence in the first place. Figure 6 gives a partial answer by plotting the diagonal and off-diagonal elements of the present discount value matrix $V = \sum_{i=0}^{\inf} \beta^i F^i = [I - \beta * P]^{-1}$. As the figure shows, for other than very high persistence, the geometrically declining probability of remaining in the same state for $1, 2, ..., n$ periods leads to mixing coefficients in V that are close to $\frac{1}{2}$, and thus asset prices that differ little between regimes. It is thus the geometric nature of present discounted probabilities that leads to the highly non-linear relationship between asset prices and persistence in figure 5.

Note how uncertainty about the value of regime persistence has a fundamentally different effect from that of uncertainty about dividend realisations in this model. The latter reduces prices, as agents are risk averse. The former, however, increases prices, as agents anticipate with positive probability persistence to be at levels above its expected value, where prices rise much more strongly than below.

4.3.1 Sensitivity

In the benchmark parameterisation, we assumed an ad hoc prior of moderately persistent regimes ($f_{hh} = f_{ll} = 0.9$). Also, priors were relatively weak, corresponding to 1 and 5 year(s) of data in the low and high volatility regime respectively. In this section, we look at stronger priors and priors equal to the true transition probabilities.

Panel 1 and 2 of figure 7 show how a stronger ad hoc prior reduces the initial jump in the price-dividend ratio at the beginning of the Great Moderation, which is exactly the Jensen's inequality effect described in the previous section. Also, with a stronger prior, the marginal contribution of observed transitions to the posterior is reduced, which lowers the increase in prices during the course of the Great Moderation. The convexity effect is further illustrated by comparing panels 3,4 and 5 of figure 7, which show results for a correct prior about transition probabilities and strengthening information (corresponding to 1 and 5 years, 5 and 10 years, and 50 and 100 years of prior observations on low and high volatility regimes respectively). The results show how, with a true prior, the initial jump in prices is higher than with the less persistent ad hoc prior. More importantly, the jump is much stronger than under full information when the prior is weak, but close to the full information case when it is strong.

The assumption of high risk-aversion was made to target price-dividend ratios that are close to those observed in the period before the Great Moderation. The relative volatility of log-dividend growth $\lambda = 4.5$ is equal to the benchmark value in Lettau et al (2008), whose choice is based on estimates of the relative volatility in post-war US data. Table 8 shows how the results for Bayesian learning change with lower risk aversion of $\gamma = 20$ and dividend volatility of $\lambda = 2.5$. Lower risk aversion reduces the impact of the fall in volatility during the Great Moderation on full information prices by about a third. Similar reductions in the effect can be observed for the Bayesian learning schemes. With lower dividend volatility, the reduction is even stronger, to about half the benchmark value.

5 Asset Prices under Alternative Learning Schemes

5.1 "Good policy or good luck": Learning when low-volatility is suspected to be permanent

In this section, we propose an alternative learning scheme that tries to explicitly capture the uncertainty about whether the Great Moderation was permanent in nature or not. For this, we assume, as before, that the agent observes the current variances of shocks σ_t^2. Also, and contrary to the previous section, we assume she knows the transition probabilities between high and low volatility regimes during normal times. However, whenever the agent observes a move to low-volatility, she attaches a small prior probability $\widehat{p}$ to this change being permanent. She then updates this prior probability according to the likelihood of the observed sequence of regimes in normal times relative to that after a permanent change to low volatility. More specifically, the likelihood of observing a sequence of σ_l^N of N low-variance periods when transitions probabilities are given by $\mathbf{F}$ is simply

$$L(\sigma_l^N|\sigma_t^2 = \sigma_l^2, \mathbf{F})) = P(\sigma_{t+1}^2 = \sigma_l, \sigma_{t+2}^2 = \sigma_l, ..., \sigma_{t+N}^2 = \sigma_l|\sigma_t^2 = \sigma_l^2, \mathbb{F}) = F_{ll}^N$$

where $P(A|B, F)$ denotes the probability of event A conditional on event B and transition matrix F. The posterior probability of a permanent shift having occured in period t, denoted $P(F = \mathbf{1}|\sigma_l^N)$, for $\mathbf{1}$ the identity matrix, is thus

$$\begin{aligned} P(F = \mathbf{1}|\sigma_l^N) &= \frac{P(F = \mathbf{1} \wedge \sigma_l^N)}{P(F = \mathbf{1} \wedge \sigma_l^N) + P(F = \mathbf{F} \wedge \sigma_l^N)} \\ &= \frac{\widehat{p}}{\widehat{p} + F_{ll}^N(1 - \widehat{p})} \end{aligned} \tag{13}$$

In our analysis we focus on a scenario similar to the experience of the US economy after World War II, which we interpret as a long realisation of high volatility followed by the Great Moderation and a move back to higher volatility with the recent crisis. Accordingly, we assume that a move to permanently low volatility can only happen once. It is thus immediate that equity prices in any high volatility period after the start of the Great Moderation are equal to the full information price p_h. Similar to the full information case, the vector of price dividend ratios under Bayesian learning about a 'permanent vs. transitory' Great Moderation, denoted p_t^{PT} is then described by equations similar to (9), (7). With $\lambda = 1$, this yields

$$p_{it}^{PT} = \beta e^{\left(-\frac{a}{\psi}+a\right)\bar{g}} \left(\mathbf{P}_{ii,t} e^{\frac{\left(-\frac{\alpha}{\psi}+\alpha\right)^2}{2}\sigma_i^2} (1+p_{i,t+1}^{PT})^a + \mathbf{P}_{ij,t} e^{\frac{\left(-\frac{\alpha}{\psi}+\alpha\right)^2}{2}\sigma_j^2} (1+p_{j,t+1}^{PT})^a \right)^{1/a} \tag{14}$$

where again $i, j \in \{h, l\}$ and $\mathbf{P}_{ij,t}$ is the probability of moving from regime i to regime j given the period t posterior probability of the change to low volatility being permanent in equation (13). Note, however, that price-dividend ratios under this learning schemes are not simply fixed points to equation (14). Rather, the representative agent anticipates that, should low volatility persist next period, the probability of a permanent change increases, and so does the price-dividend ratio. We thus have to compute the whole path of price-dividend ratios jointly.[9]

To implement this model quantitatively, we choose the same moderately persistent prior about transition probabilities in normal times as in the previous section ($F_{ll} = F_{hh} = 0.9$). And we set the conditional probability of the Great Moderation being permanent conditional on an observed change from high to low volatility to 2 percent, which yields an unconditional probability 0.2%. Figure 8 shows the time path that results from this learning scheme, compared to US data. The less persistent prior probabilities lead to a PD ratio that is higher until the early 1980s, but also a smaller initial rise when the agent observes a move to low volatility in 1985. As the economy persists in low volatility, however, the rising posterior probability of a permanent moderation in macro-volatility leads to an S-shaped increase in prices. The magnitude of the boom is with 37% similar to our benchmark learning scheme. The observed end of the Great Moderation comes with a strong bust in asset prices of 49%, as agents update the probability of being in a permanently more benign macroeconomic environment to zero.

[9]This is simplified by the fact that price-dividend ratios are easily calculated for permanent regimes with $\mathbb{F} = \mathbf{1}$. Also, it is simple to calculate the path of posterior probabilities $P(F = \mathbf{1}|\sigma_l^N)$ as N rises, and thus the transition probabilities $\mathbf{P}_{ij,t}$ increases. Once $P(F = \mathbf{1}|\sigma_l^N)$ is close enough to 1, say after $\overline{N}$ low volatility periods, we know that $p_{l,t+\overline{N}}^{PT}$ equals the price-dividend ratio under permanently low volatlity $p(F = \mathbf{1})$, which we can easily calculate as a fixed point to (14) for $F = \mathbf{1}$. This allows us to calculate the sequence of price-dividend ratios at for periods $s = \overline{N} - 1, \overline{N} - 2, ...t$ by backward induction.

Asset prices - Learning about transitory vs. Permanent GM

	Boom	Overvaluation	Bust
No Learning	2%	0	1%
Learning	37%	41%	43%

Table 7: "Boom" denotes the increase in prices until the end of the Great Moderation relative to the high-volatility regime preceding it (computed over windows of 5 years). "Overvaluation" is the overvaluation at the end of the Great Moderation relative to the prices without learning. And "Bust" is the fall in prices in the first period after the Great Moderation.

5.2 Ad hoc learning

It has been argued for example by Haldane (2009), that overconfidence in a low volatility environment may arise when agents base their inferences about the future predominantly on recent observations of small shocks. This over-reliance on the recent past is opposed to the optimal nature of learning schemes so far, but in line with a large number of studies where agents follow ad hoc learning rules that map observations into estimates of parameters of interest (see for example Evans and Honkapohja (1999)). To see whether non-optimal learning rules can deliver a boom and bust in asset prices similar to those observed in US data, we assume that the representative agent knows the mean dividend growth $\bar{g}$ and observes the history of shocks $\Omega_t = \{\varepsilon_s\}_{s=0}^{t}$. But she ignores, or chooses to ignore, in her estimate about future macro-volatility the two-stage nature of the data generating markov process. Rather she uses simple ad hoc rules that map observed histories into estimates $\widehat{\sigma_{t+1}^2}$ of the variance of future shocks σ^2

$$\widehat{\sigma_{t+1}^2} = G(\Omega_t)$$

where $G : R^t \longrightarrow R^+$. Specifically, we consider 3 simple mappings G

$$G^{OLS} = \frac{1}{N}\sum_{s=0}^{t}(\varepsilon_s)^2 \tag{15}$$

$$G^{CG} = \xi(\varepsilon_t)^2 + (1-\xi)G_{t-1}^{CG} = \sum_{s=0}^{t}\xi(1-\xi)^{t-s}\varepsilon_t^2,\ 0 < \xi < 1 \tag{16}$$

$$G^{CW} = \frac{1}{n}\sum_{s=t-n}^{t}(\varepsilon_s)^2 \tag{17}$$

Thus, under G^{OLS} agents simply compute their best guess of the future variance as an average over the entire history of shocks. G^{CG} describes a simple "constant-gain" learning rule: the agent computes the variance as a weighted average of his best guess in the previous period and the squared shock today. Relative to G^{OLS}, this overweighs more recent observations, as the weight on more distant observations decays geometrically at rate $1 - \xi$. Finally, G^{CW} uses windows of n most recent observations to compute the variance.

Figure 9 presents the time path of asset prices under the three ad-hoc learning rules, together with full-information prices in the case of high persistence. With OLS learning, the fall in the variance estimate for consumption growth is relatively slow. Moreover, since each estimate weighs all past periods equally, the variance estimate remains an average across high and low volatility periods, resulting in a relatively small rise in prices. The price nevertheless rises above that under full information, which is, again, due to a Jensen's inequality effect.

With constant gain learning, the contribution of past periods to the variance estimate falls geometrically over time. This implies that the estimate of consumption variability during the great moderation falls faster, and further, than with OLS learning. The boom in prices is thus steeper and stronger, amounting to around 70 % at the end of the Great Moderation, far above that implied by full information. When agents compute their estimate of the consumption growth variance as an average across a window of constant length, asset prices reach an even higher plateau as under constant gain learning, although their path is slightly more convex, as at the beinning of a new regime estimates adjust more slowly.

Under all three ad-hoc learning rules, the fall in prices at the end of the Great Moderation is relatively slow: only as information about a change in volatility accumulates, adjust agents their estimates. Contrary to their Bayesian counterparts, recursive, ad hoc learning rules are thus not able to deliver sudden crashes in prices.[10] Table 7 summarises the results for the benchmark case.

[10]Since the true persistence of regimes has no effect under ad hoc learning, we omit the results for low persistence.

Asset prices - Ad hoc learning

	Boom	Overvaluation	Bust
Full Info	17%	0	17%
OLS	22%	3%	0%
Constant Gain	72%	46%	2%
Constant Window	83%	54%	1%

Table 8: "Boom" denotes the relative increase in prices until the end of the Great Moderation relative to the high-volatility regime preceding it (computed over windows of 5 years). "Overvaluation" is the overvaluation at the end of the Great Moderation relative to the prices under full information. And "Bust" is the fall in prices in the first period after the Great Moderation.

5.3 Sensitivity

5.3.1 Alternative prior probabilities of a permanent Great Moderation

When learning about the permanent vs. transitory character of the Great Moderation, in the benchmark version of the model agents believed the observed move to low volatility to be permanent with a conditional prior probability of $P(F_{ll} = 1) = 2\%$. Here we look at how the results change when reducing this prior probability to 1 and 0.1 percent. As suggested by intuition, the rise in price-dividend ratios during the Great Moderation is slower when agents attach a lower prior probability to it being permanent. As figure 9 shows, with $P(F_{ll} = 1) = 1\%$, the difference is small. But when $P(F_{ll} = 1) = 0.1\%$, the Great Moderation comes to an end before the posterior converges to 1, such that the rise in prices is "cut off", and both the overall boom and the fall in asset prices lower.

5.4 Different assumptions for ad hoc learning rules

Figure 10 shows the time path of price-dividend ratios under our ad hoc learning rules with a smaller gain parameter $\xi = 0.01$ and a longer window for the estimation of volatilities of $w^n = 30$ years. As anticipated, the boom is now less steep, as individuals update their volatility estimates less quickly as information about a new low-volatility environment accumulates. The size of the boom in the case of constant gain is reduced from more then 70 to around 35 percent. That with constant estimation windows from 83 to 47 percent.

Asset prices - Lower Risk Aversion and Dividend Volatility

	Boom	Overvaluation	Bust
Full Info (benchmark)	17%	0	17%
Full Info ($\gamma = 20$)	14%	0	14%
Full Info ($\lambda = 2.5$)	8%	0	8%
Trans Probabilities (benchmark)	46%	24%	39%
Trans Probabilities ($\gamma = 20$)	40%	17%	28%
Trans Probabilities ($\lambda = 2.5$)	24%	11%	18%
Perm vs Trans (benchmark)	48%	47%	50%
Perm vs Trans ($\gamma = 20$)	33%	32%	34%
Perm vs Trans ($\lambda = 2.5$)	25%	26%	27%

Table 9: "Boom" denotes the relative increase in prices until the end of the Great Moderation relative to the high-volatility regime preceding it (computed over windows of 5 years). "Overvaluation" is the overvaluation at the end of the Great Moderation relative to the prices under full information. And "Bust" is the fall in prices in the first period after the Great Moderation.

5.5 Lower Risk Aversion and Dividend Volatility

Lower risk aversion has a similar effect on asset price dynamic as under the benchmark learning scheme. For example, with learning about a permanent vs. transitory Great Moderation the predicted boom falls from around 50 percent in the benchmark case to 33 percent with lower risk aversion, and to 25 percent with lower dividend volatility. The changes with ad-hoc learning are similar, and therefore omitted.

5.6 Alternative Timing of the Great Moderation

All results so far were based on a start of the Great Moderation in 1984, in line with the estimates for GDP. An later start date, for example in the early 1990s, as suggested by the fall in consumption volatility, has very little impact on the sizes of boom and bust in cases where learning is relatively fast, and price-dividend ratios therefore converge to a constant value during the early or middle years of the Great Moderation. When learning is slow, however, for example in the ad-hoc cases with low gain or long windows, or when the prior probability of a permanent move to low-volatility is low, both the size of the boom and the crash are smaller, as the increase in price-dividend ratios in the figures becomes “interrupted” by the end of the Great Moderation.

6 Conclusion

From a review of both academic and investment research we conclude that, first, the “Great Moderation” in macro-volatility was perceived as an important factor behind the asset price boom of the 1990s and 2000s, and, second, that academics and investors alike were uncertain about the origins and persistence of the new low-volatility environment. Using different learning mechanisms, we modelled this uncertainty explicitly in a general equilibrium asset pricing model with time-varying volatility. The results confirmed the intuition of policymakers (Bean 2009, Haldane 2009) that overconfidence in a benign macroeconomic environment may have led to an overvaluation of assets beyond their fundamental value. Particularly, we find that Bayesian learning can lead to an asset price boom of around 50 percent at the end of the Great Moderation, as agents had become increasingly confident in its persistence. The end of the low-volatility period, which we identified with the beginning of the recent crisis, in turn leads to a strong crash in prices. Moreover, both boom and crash are an order of magnitude larger than in a model with full information about the data generating process. More ad hoc, or statistical, learning rules predict an even stronger boom in prices, but cannot replicate the fast crash at the end of the Great Moderation period.

Future research could extend this study in several directions, by, for example, including time variation in the mean growth of the economy, or by looking at an alternative scenario where agents form expectations about future prices directly, rather than the distribution of dividends as in the model analysed here. Adam and Marcet (2010) show how this can lead so self-fulfilling bubbles and crashes in asset prices, as a rise in prices is sustained by generating expectations of rises in the future. When learning about volatility, this self-referential mechanism is less clear, as higher expected volatility primarily feeds into the level of prices, and not into their second moment. An in-depth analysis of this issue should be done in future work.

References

[1] Abel, A. B. (2003): "The Effects of a Baby Boom on Stock Prices and Capital Accumulation in the Presence of Social Security," *Econometrica,* vol.71, issue 2, pages 551-578.

[2] Adam, K. and A. Marcet (2010): "Booms and Bust in Asset Prices," Department of Economics, London School of Economics.

[3] Bansal, R., V. Khatchatrian and A. Yaron (2005): "Interpretable asset markets?," *European Economic Review,* vol. 49, issue 3, pages 531-560.

[4] Bansal, R. and C. Lundblad (2002): "Market Efficiency, Asset Returns, and the Size of the Risk Premium in Global Equity Markets," *Journal of Econometrics*, vol. 109, issue 2, pages 195-237.

[5] Bansal, R. and A. Yaron (2004): "Risks For The Long Run: A Potential Resolution Of Asset Pricing Puzzles," *Journal of Finance*, vol. 59, pages 1481-1509.

[6] Bean, C. (2009): "The Great Moderation, the Great Panic and the Great Contraction," *BIS Review,* 101/2009.

[7] Benati, L. and P. Surico (2009): "VAR Analysis and the Great Moderation," *American Economic Review*, American Economic Association, vol. 99(4), pages 1636-52.

[8] Bernanke, B. S. (2004): "The Great Moderation," Remarks by Governor Ben S. Bernanke at the Meetings of the Eastern Economic Association, Washington, DC.

[9] Blanchard, O. J. (1993): "Movements in the Equity Premium," *Brookings Papers on Economic Activity*, vol. 1993, No. 2, pages 75-138.

[10] Blanchard, O. J. and J. Simon (2001): "The Long and Large Decline in U.S. Output Volatility," *Brookings Papers on Economic Activity*, vol. 2001, No. 1, pages 135-164.

[11] Bonomo, M. and R. Garcia (1994): "Disappointment Aversion as a Solution to the Equity Premium and the Risk-Free Rate Puzzles," CIRANO Working Papers 94s-14.

[12] Bonomo, M and R. Garcia (1996). "Consumption and Equilibrium Asset Pricing: An Empirical Assessment". Journal of Empirical Finance vol. 3, pages 23965.

[13] Branch, W. A. and G. W. Evans (2010): "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," Centre for Dynamic Macroeconomic Analysis Working Paper Series, CDMA10/10.

[14] Canova, F., L. Gambetti and E. Pappa (2007): "The Structural Dynamics of Output Growth and Inflation: Some International Evidence," *Economic Journal*, Royal Economic Society, vol. 117(519), pages C167-C191, 03.

[15] Calvet, L. E., M. Gonzalez-Eiras, and P. Sodini (2004): "Financial Innovation, Market Participation, And Asset Prices," *Journal of Financial and Quantitative Analysis*, vol 39, pages 431-459.

[16] Campbell, S. D. (2005): "Stock Market Volatility and the Great Moderation," Finance and Economics Discussion Series, Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board, Washington, D.C.

[17] Campbell, J. Y. and R. J. Shiller (2004): "The MD Interview: Robert J. Shiller," by John Y. Campbell, *Macroeconomic Dynamics*, vol 8, pages 649–683.

[18] Cogley, T. and T. J. Sargent (2008): "The market price of risk and the equity premium: A legacy of the Great Moderation?," *Journal of Monetary Economics, vol.* 55, pages 454-476.

[19] Cooley, T. F. (2008): "How We Got Here: From the Great Moderation to the Great Conflagration," Forbes.com.

[20] Drifil, J. and M. Sola (1998): "Intrinsic Bubbles and Regime Switching," *Journal of Monetary Economics*, vol. 42, pages 357-373.

[21] Dynan, K. E., D. W. Elmendorf and D. E. Sichel (2006) "Can financial innovation help to explain the reduced volatility of economic activity?," *Journal of Monetary Economics*, Elsevier, vol. 53(1), pages 123-150.

[22] Evans, G. and S. Honkapohja (1999): "Learning Dynamics', in: Handbook of Macroeconomics, vol. 1, Part 1, pages 449-542.

[23] Fama, E. F. and K. R. French (2002): "The Equity Premium," *Journal of Finance*, vol. 57(2), pages 637-659.

[24] Fogli, A. and Perri F. (2009): "Macroeconomic Volatility and External Imbalances," University of Minnesota.

[25] Goldman Sachs (2002): "The Equity Risk Premium: It's Lower Than You Think", CEO confidential, Issue 2002/14 November 2002.

[26] Haldane, A. G. (2009): "Why the Banks Failed the Stress Test," Bank of England.

[27] Hall, R. E. (2000): "E-Capital: The Link Between the Stock Market and the Labor Market in the 1990s," *Brookings Papers on Economic Activity*, vol 2000(2), pages 73–118.

[28] Heaton, J. and D. Lucas (2000): "Stock Prices and Fundamentals," NBER Chapters, in: *NBER Macroeconomics Annual* 1999, vol. 14, pages 213-264.

[29] Jagannathan, R., E. R. McGrattan and A. Scherbina (2000): "The declining U.S. equity premium," *Quarterly Review*, Federal Reserve Bank of Minneapolis, vol. 24, No. 4, pages 3–19.

[30] Jovanovic, B. and P. L. Rousseau (2003): "Two Technological Revolutions," *Journal of the European Economic Association*, MIT Press, vol. 1(2-3), pages 419-428.

[31] Global Market Strategy, J.P. Morgan Securities Ltd.: "Volatility, Leverage and Returns", London October 19, 2005.

[32] Kim, C-J. and C. R. Nelson (1999): "Has The U.S. Economy Become More Stable? A Bayesian Approach Based On A Markov-Switching Model Of The Business Cycle," *The Review of Economics and Statistics*, MIT Press, vol. 81(4), pages 608-616.

[33] Lettau, M., S. C., Ludvigson and J. A. Wachter (2007): "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?", *The Review of Financial Studies,* vol. 21, n 4, pages 1653-1687.

[34] McConnell, M and G. Perez-Quiros (1997): "Output Fluctuations in the United States: What has changed since the early 1980s?," Research Paper 9735, Federal Reserve Bank of New York.

[35] McConnell, M and G. Perez-Quiros (2000): "Output Fluctuations in the United States: What has changed since the early 1980s?," *American Economic Review,* vol. 90(5), pages 1464-1476.

[36] Mehra, R., and E. C. Prescott (1985): "The Equity Premium Puzzle". Journal of Monetary Economics vol. 15. pages 14561.

[37] Boz, Emine and Enrique G. Mendoza (2010). "Financial Innovation, the Discovery of Risk, and the U.S. Credit Crisis". NBER Working Papers 16020, National Bureau of Economic Research.

[38] Primiceri, G. (2005): "Why Inflation Rose and Fell: Policymakers' Beliefs and US Postwar Stabilization Policy," NBER Working Papers 11147, National Bureau of Economic Research.

[39] Shiller, R. J. (2000): Irrational Exuberance, Princeton University Press.

[40] Siegel, J. J. (1999): "The Shrinking Equity Premium," *Journal of Portfolio Management,* vol. 26, pages 10-17.

[41] Sims, C. A. and T. Zha, 2006. "Were There Regime Switches in U.S. Monetary Policy?," *American Economic Review*, American Economic Association, vol. 96(1), pages 54-81.

[42] Stock, J. H. and M. W. Watson (2002): "Has the Business Cycle Changed and Why?," NBER Macroeconomics Annual 2002, vol.17 (2003), MIT Press.

[43] Stock, J. H. and M. W. Watson (2005): "Understanding Changes In International Business Cycle Dynamics," *Journal of the European Economic Association*, MIT Press, vol. 3(5), pages 968-1006, 09.

[44] Tucker, P. (2005): "Monetary Policy, Stability and Structural Change," for the Confederation of British Industry in Guildford, vol 45, no 2, pages. 247-255.

[45] UniCredit Global Economics and FI/FX Research Economics and Commodity Research (2006): "Friday Notes", November 10 2006.

[46] Veronesi, P. (1999): "Stock Market Overreaction to Bad News in Good Times: A Rational Expectation Equilibrium Model," *Review of Financial Studies*, vol. 12(5), pages 975-1007.

[47] Jessica A. Wachter (2006), Comment on: "Can financial innovation help to explain the reduced volatility of economic activity?', Journal of Monetary Economics 53 (2006) 151-154

[48] Weitzman, M. L. (2007): "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, vol. 97(4), pages 1102–1130.

7 Appendix

7.1 Data Appendix

Consumption is quantified as the *Total Real Personal Consumption Expenditures* measured in quantity index [index numbers, 2005 = 100]. The data are quarterly, seasonally adjusted and their source is the National Economic Accounts of the Bureau of Economic Analysis (BEA).

GDP is quantified as the *Real Gross Domestic Product*, measured in 2005-chained dollars. The data are quarterly, seasonally adjusted and their source is the National Economic Accounts of the Bureau of Economic Analysis (BEA).

Population is quantified as the *Midperiod Population* of each quarter. The data source is the National Economic Accounts of the Bureau of Economic Analysis (BEA).

Asset Price is quantified as the average *S&P 500 Stock Price Inde*x of each quarter. The data source is the Robert Shiller's homepage. The original data are monthly averages of daily closing prices.

Dividend is quantified as the original quarterly *Dividend Payment* reported in the Robert Shiller's homepage.

Price-Earning Ratio is quantified as the *Cyclically Adjusted Price Earnings Ratio* (P/E10), known also as the *CAPE*. The data source is the Robert Shiller's homepage. The price-earning-ratio series used by us contains only the original quarterly earning data.

7.2 Appendix B

7.2.1 Numerical solving for the price when agents use Bayesian leaning

The Bayesian agent enters each period with a prior. He observes the realization of the exogenous process and he updates the counters

$$\begin{aligned} n_{t+1}^{ij} &= n_t^{ij} + 1 \; if \; s_{t+1} = j \; and \; s_t = i \\ n_{t+1}^{ij} &= n_t^{ij} \qquad if \qquad otherwise. \end{aligned}$$

The posterior density function is

$$f(F_{hh}, F_{ll} \mid \Sigma^t) = beta(n_t^{hh}, n_t^{hl}) * beta(n_t^{ll}, n_t^{lh}). \tag{18}$$

We would like to calculate

$$p_t = \int p(S_t, F) f(F \mid \Sigma^t) dF$$

which can be also expressed as

$$\int p(S_t, F) f(F \mid \Sigma^t) dF = E_{\Sigma^t}[p(F)]. \tag{19}$$

Therefore equation (2) can be approximated as

$$E_{\Sigma^t}[p(F)] \approx \frac{\sum_{i=1}^{n} p(S_t, F_i)}{n} \tag{20}$$

In order to compute equation (3) at each time t we generate a sample of n transition probability matrixes, F, as random observations from equation (1). Therefore in each period the price function can be numerically approximated by the sample average, so

$$p_t \approx \frac{\sum_{i=1}^{n} p(S_t, F_i)}{n}$$

8 Figures

Figure 1: GDP growth

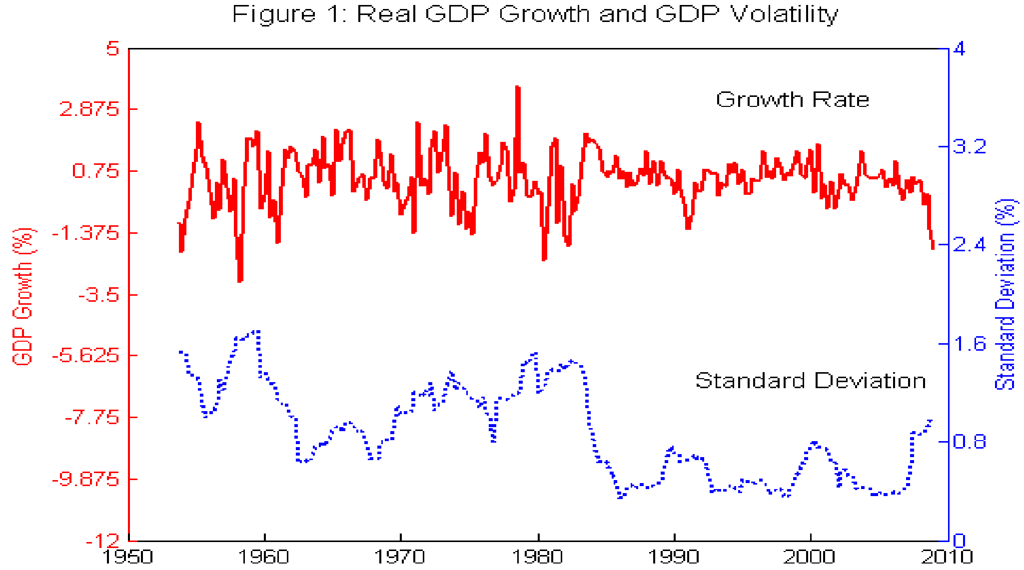

The figure plots the growth rate of the real GDP and its standard deviation estimated in 10-quarter rolling windows. Output is defined in per-capita terms, calculated as ratio of the real gross domestic product, measured in 2005-chained dollars, over total population. The data are quarterly and span the period 1952$Q2-2010Q2$. The data source is BEA. The estimates are in percent.

Figure 2: Consumption growth

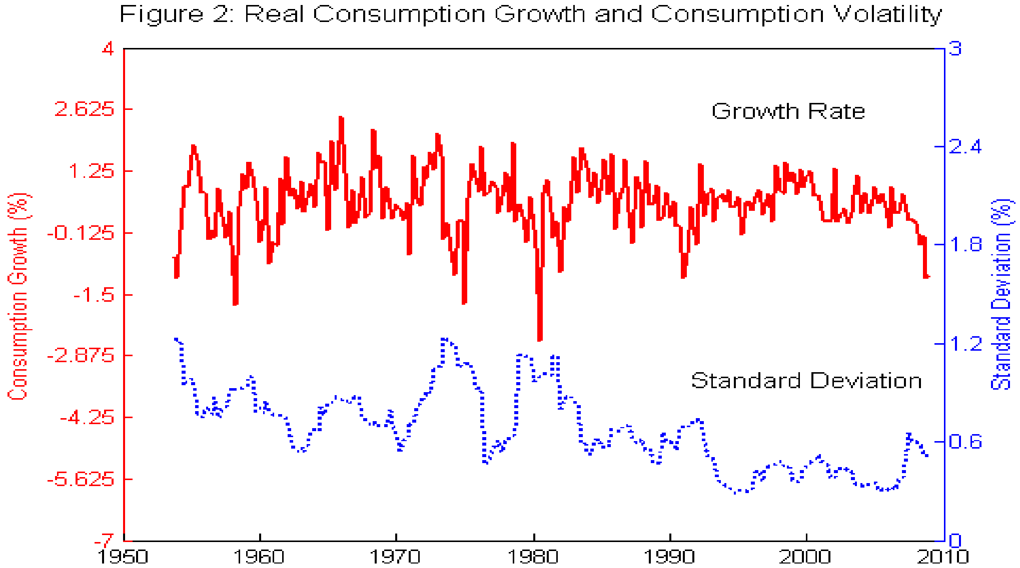

The figure plots the growth rate of the real consumption and its standard deviation estimated in 10-quarter rolling windows. Consumption is defined in per-capita terms, calculated as ratio of the total real personal consumption expenditures, measured in quantity index [index numbers, 2005 = 100], over total population. The data are quarterly and span the period $1952Q2 - 2010Q2$. The data source is BEA. The estimates are in percent.

Figure 3: Price-Dividend Ratio

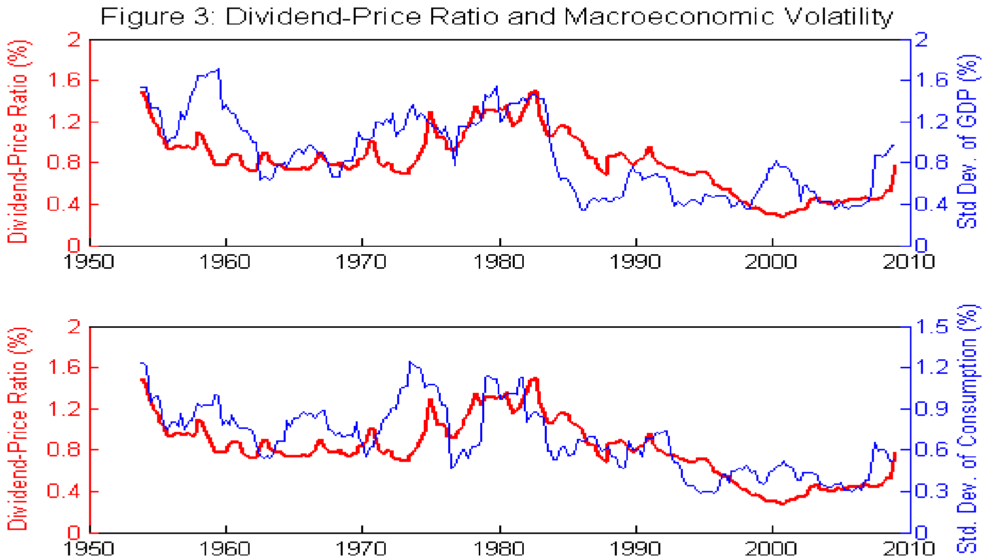

The figure plots the dividend price ratio against the standard deviation of real GDP growth rate (first subplot) and the standard deviation of the real consumption growth rate (second subplot), estimated in 10-quarter rolling windows. Output is defined in per-capita terms, calculated as ratio of the real gross domestic product, measured in 2005-chained dollars, over total population. Consumption is defined in per-capita terms, calculated as ratio of the total real personal consumption expenditures, measured in quantity index [index numbers, 2005 = 100], over total population. The financial data are taken from the Robert Shiller's homepage and the rest of the data are taken from BEA. The estimates are in percent.

Figure 4: Dividend Ratios: Benchmark Model

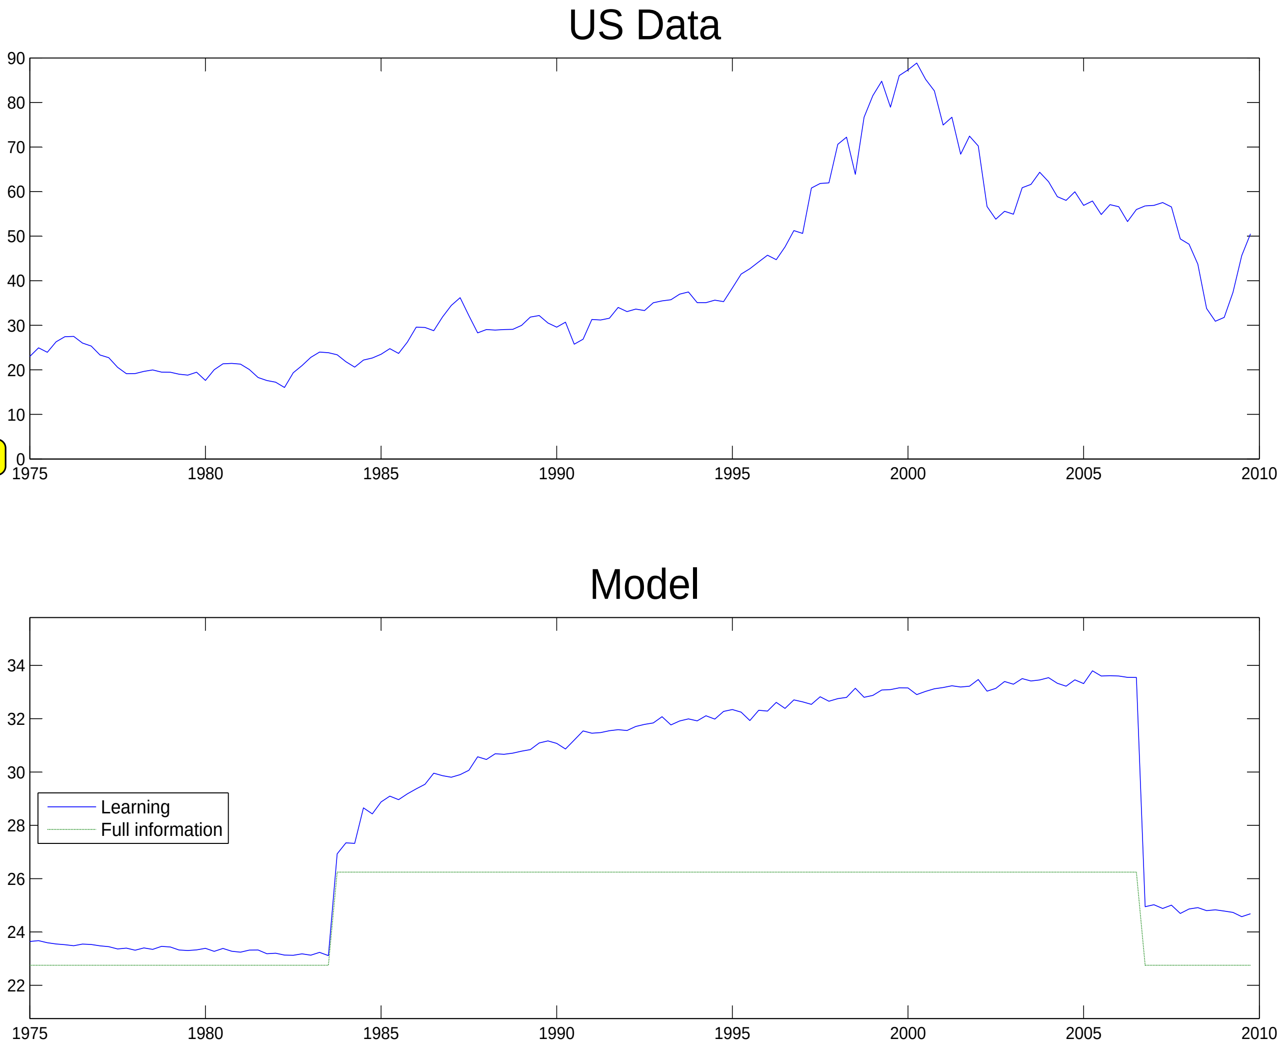

The figure plots the price dividend ratio in US data (upper Panel), and under learning about transition probabilities (lower panel), for the benchmark calibration of the model.

Figure 5: Dividend Ratios as a Function of Persistence and Prior Tightness

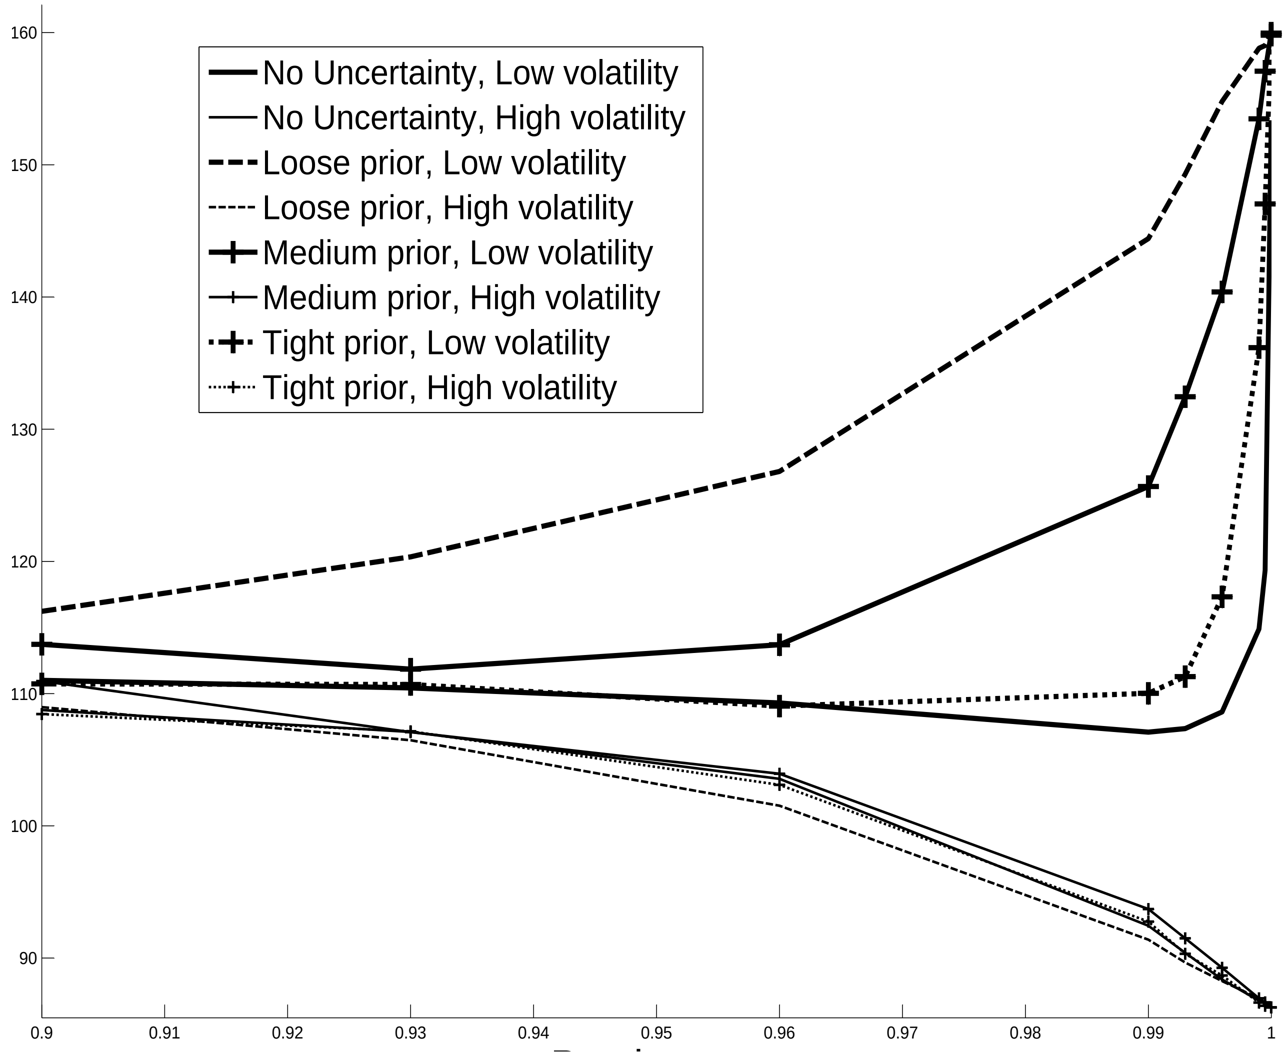

Using the simplified case of symmetric transitions ($f_{ll} = f_{hh}$), the figure plots the price dividend ratio as a function of persistence for different values of the tightness of priors for the benchmark calibration of the model.

Figure 6: Non-Linear Asset Price-Persistence Relation

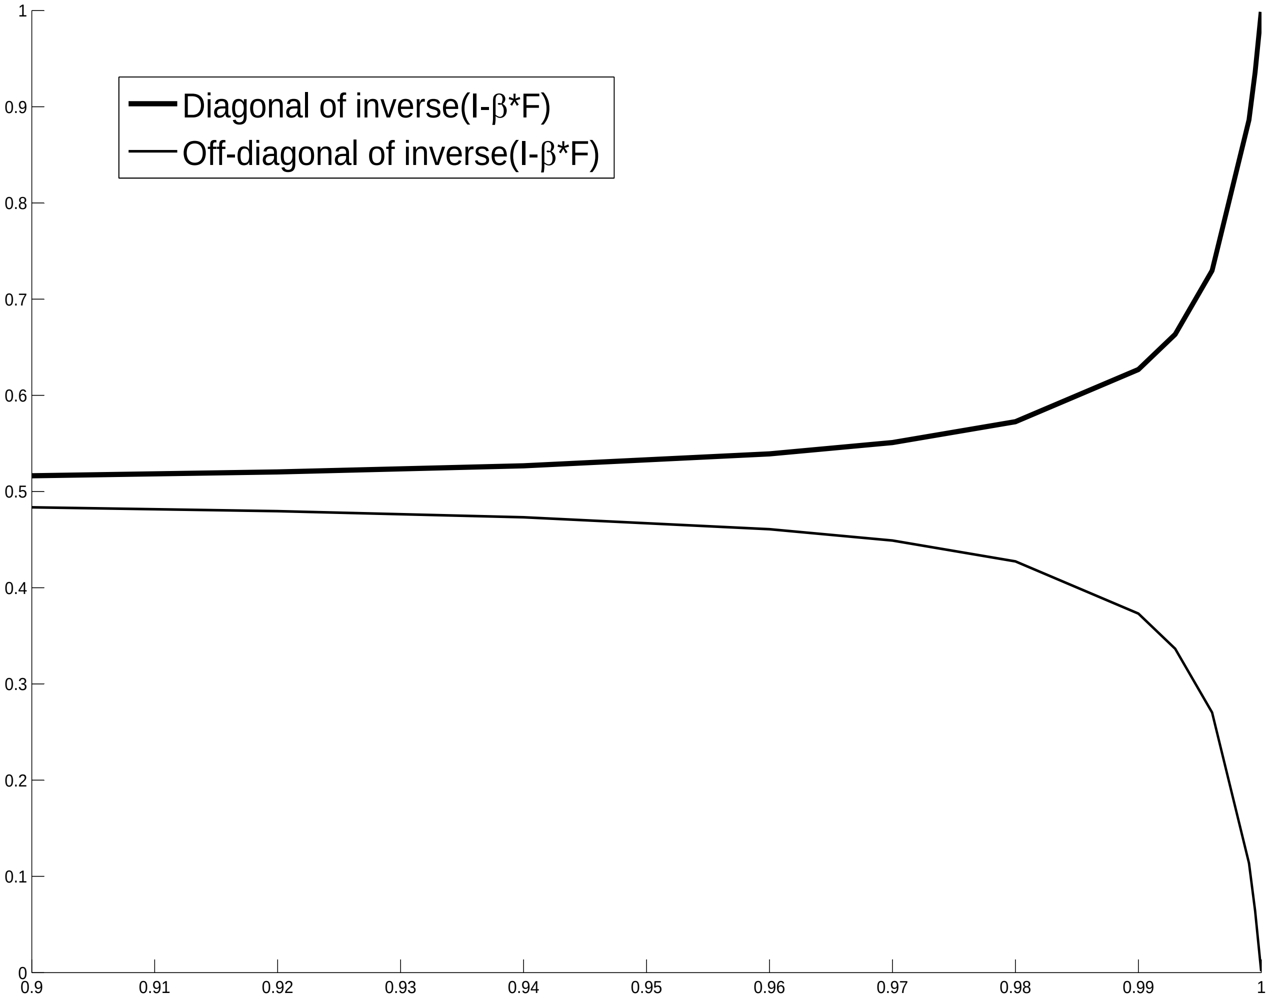

The figure depicts the diagonal and non-diagonal elements of the present discounted value matris $inv(I - \beta * P)$.

Figure 7: Learning about Transition Probabilities with Different Priors, $F = \mathbf{F}^{hp}$

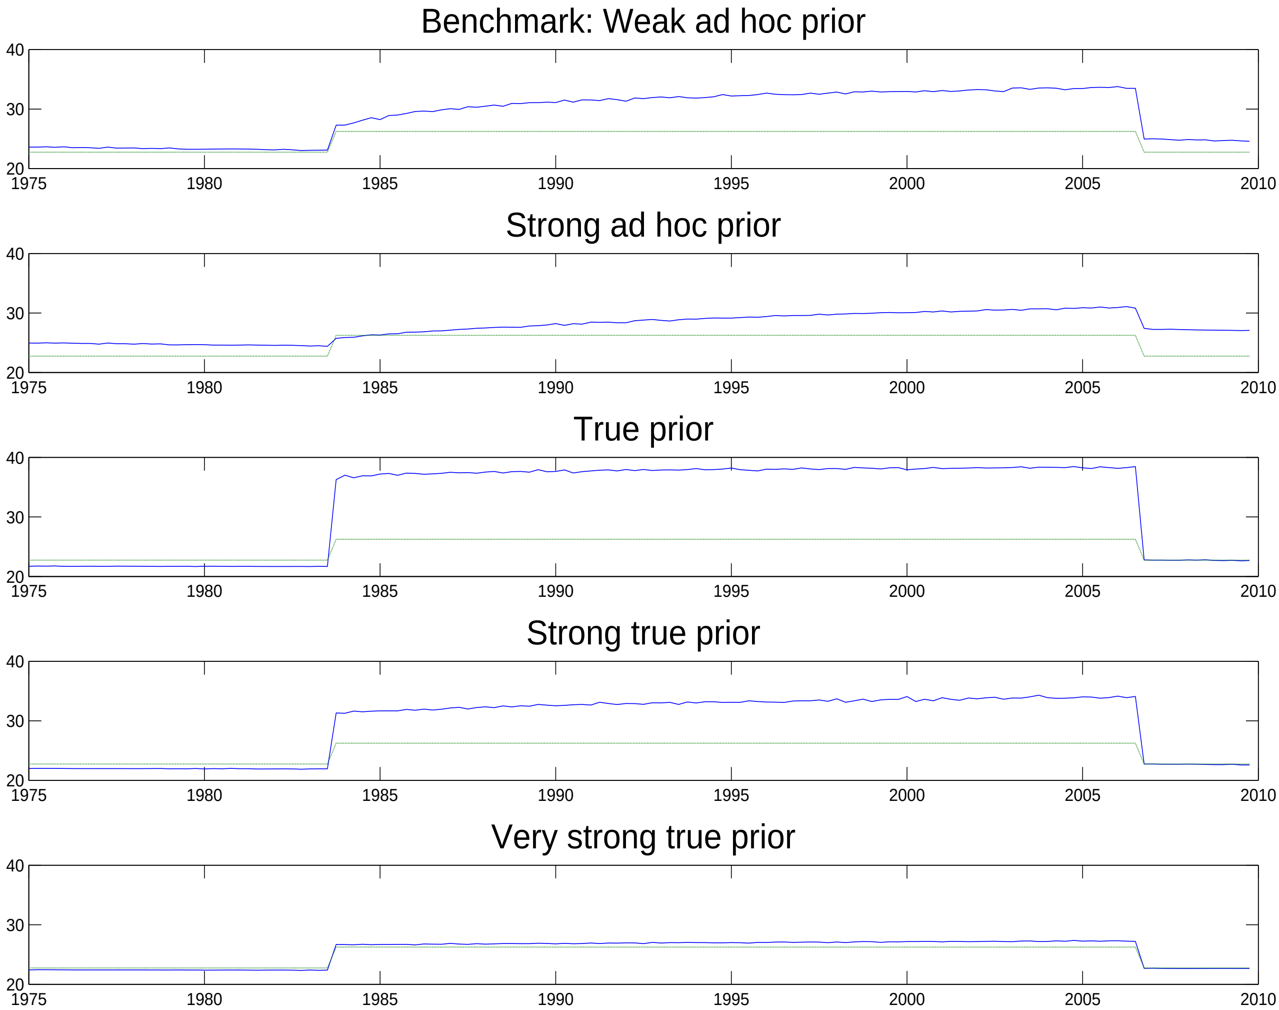

The figure shows the time-path of dividends for the benchmark model with weak ad hoc priors (panel 1), for stronger priors (corresponding to 5 and 10 years of prior observations, panel 2), and for a correct prior about $F = \mathbf{F}^{hp}$ with strengthening information (panel 3 to 5, corresponding to 1 and 5 years, 5 and 20 years, and 50 and 100 years of prior observations respectively)

Figure 8: Dividend Ratios - Learning about a Permanent vs. Transitory Great Moderation

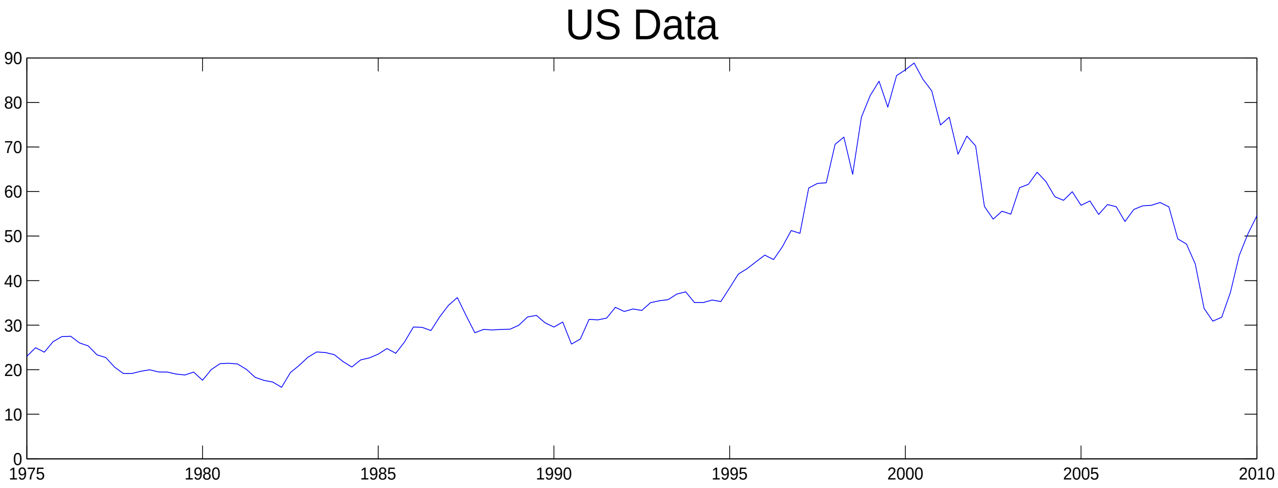

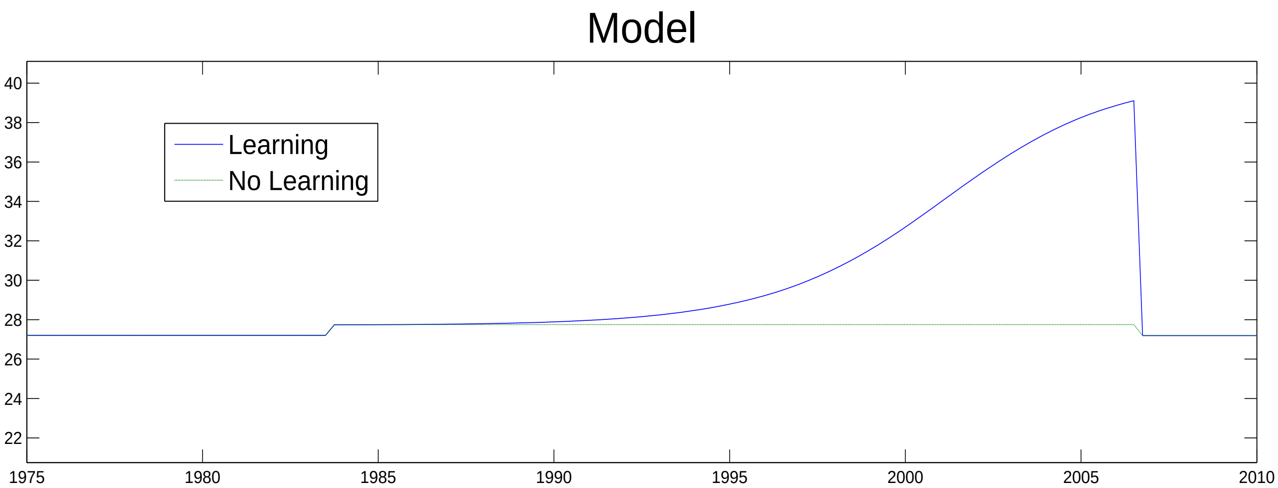

The figure plots the price dividend ratio in US data (upper Panel), and under learning about a permanent vs. transitory Great Moderation (lower panel), for the benchmark calibration of the model.

Figure 9: Dividend Ratios - Ad Hoc Learning

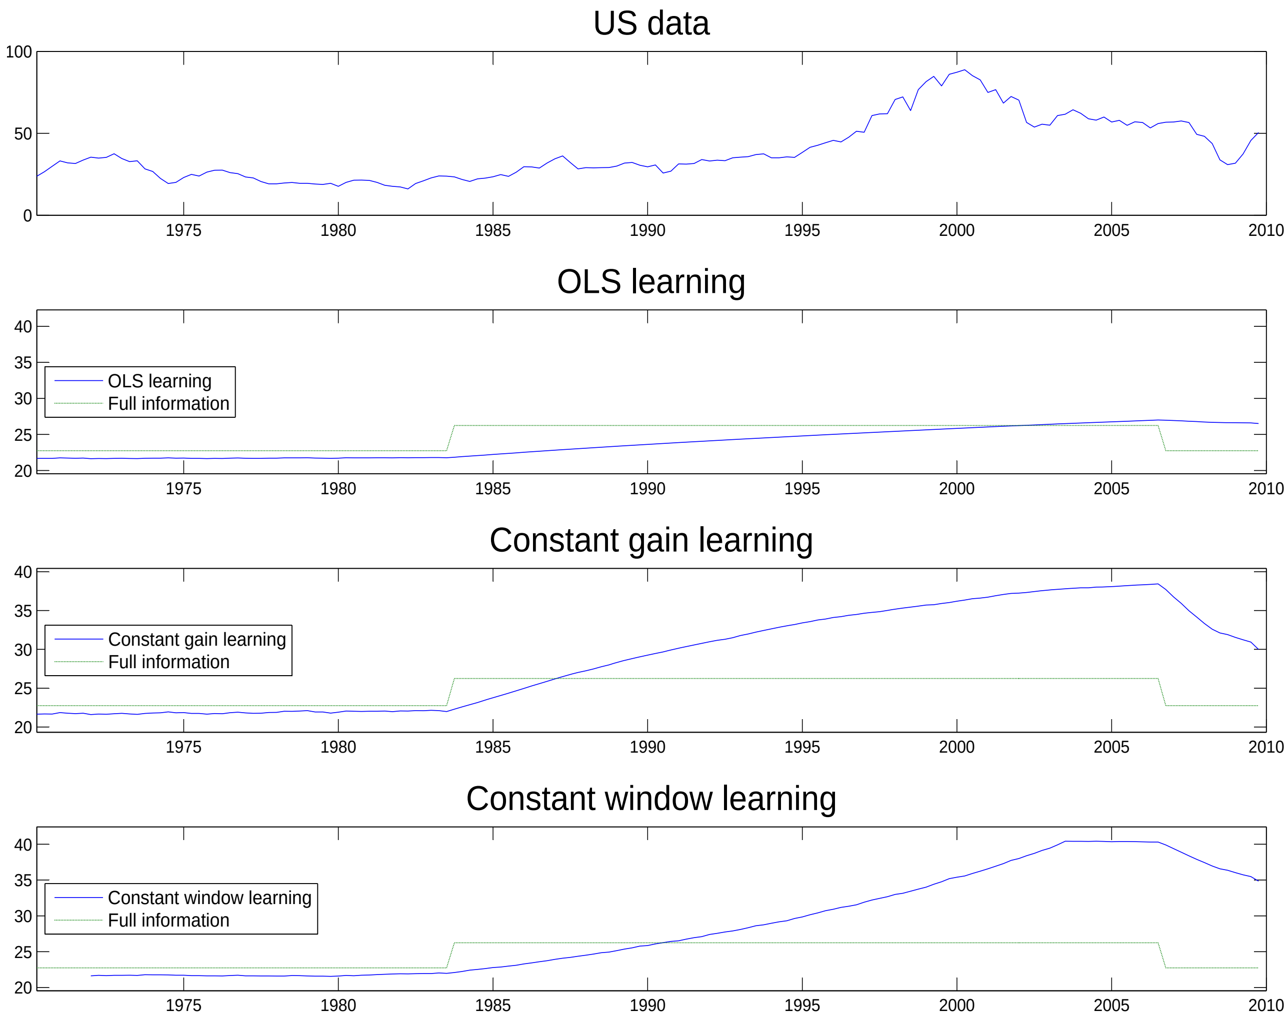

The figure plots the price dividend ratio in US data (upper Panel), and under three ad hoc learning rules: OLS (second panel), constant gain (third panel), and constant window (bottom panel), for the benchmark calibration of the model. The full information prices correspond to the case of high persistence.

Figure 10: Learning about permanent vs. transitory GM with Different Priors

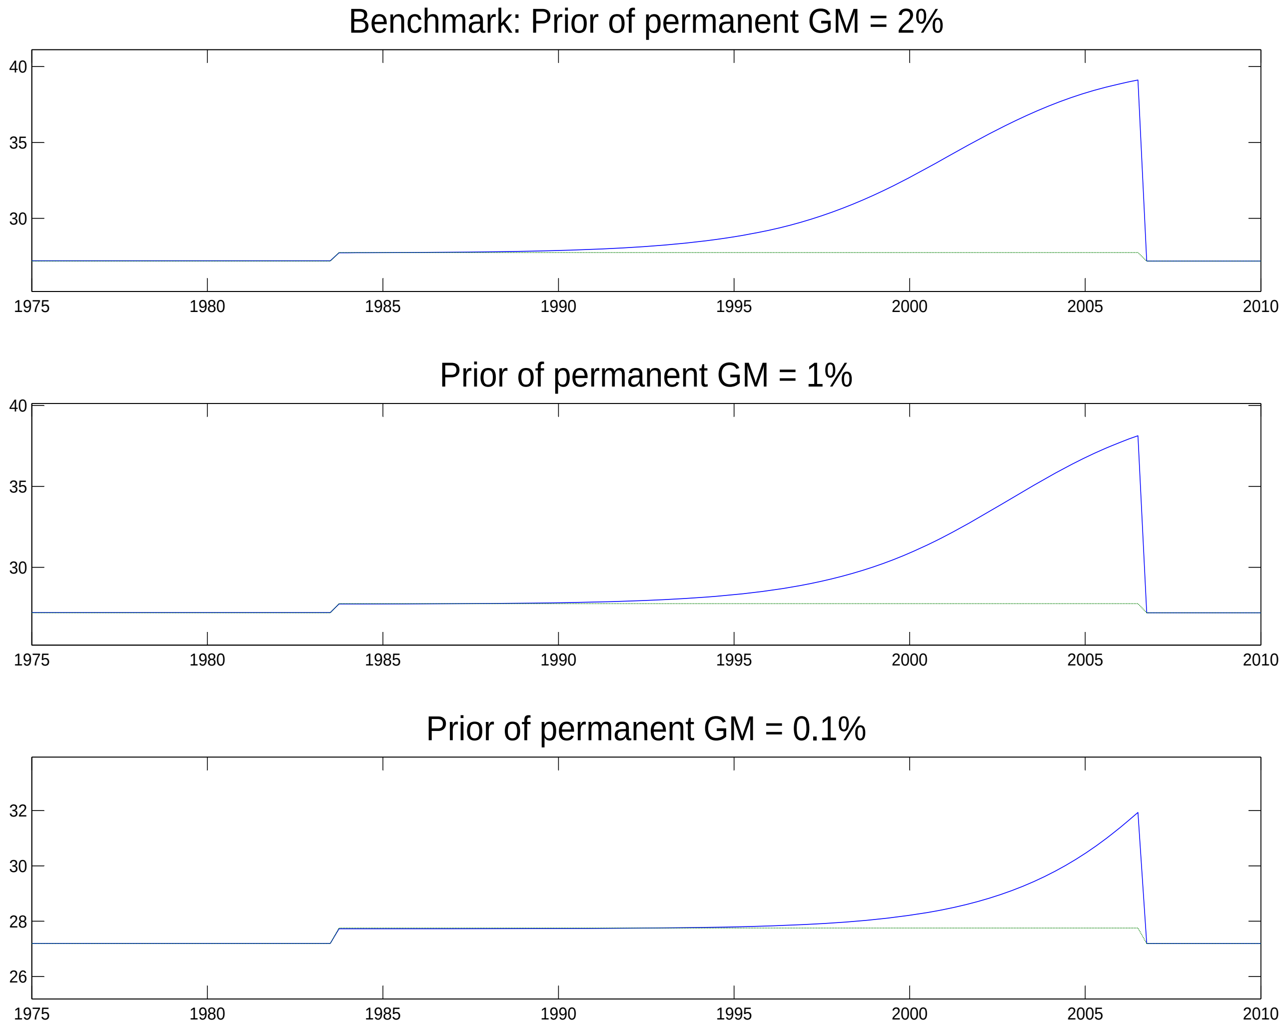

The figure shows the time-path of dividends with learning about a permanent vs. transitory Great Moderation with different prior probabilities.

Figure 11: Ad-hoc learning with lower gain and longer windows

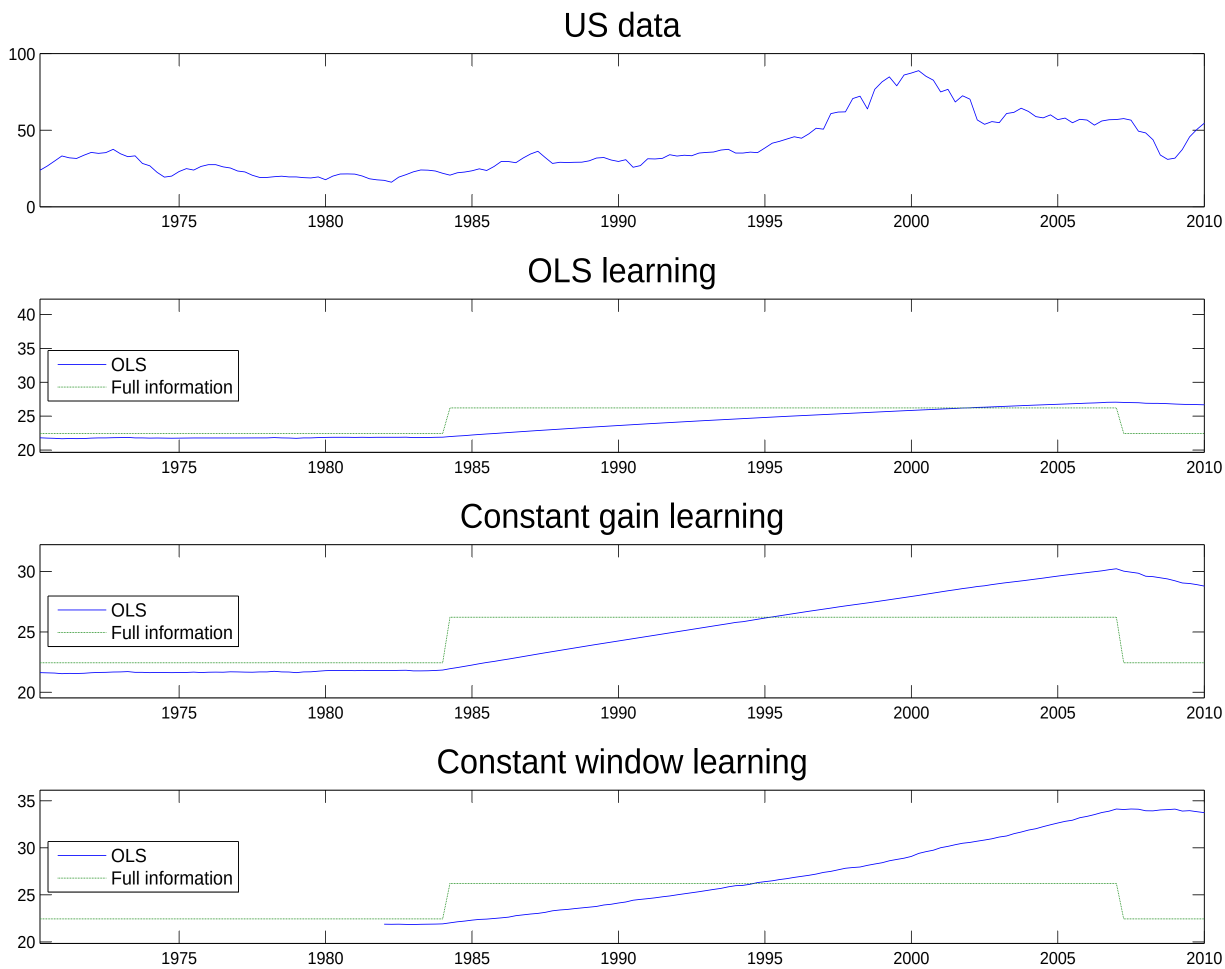

The figure shows the time-path of dividends in the data, and in the model with ad-hoc learning for $\xi = 0.01$ and $w^n = 30$ years.