UNIL | Université de Lausanne
Département de finance

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

FINANCE RESEARCH SEMINAR SUPPORTED BY UNIGESTION

"Learning About Consumption Dynamics"

Prof. Lars Lochstoer

Columbia University, Graduate School of Business

Abstract

This paper studies the asset pricing implications of Bayesian learning about the parameters, states, and models determining aggregate consumption dynamics. Our approach is empirical and focuses on the quantitative implications of learning in real-time using post World War II consumption data. We characterize this learning process and find that revisions in beliefs stemming from parameter and model uncertainty are significantly related to realized aggregate equity returns. This evidence is novel, providing strong support for a learning-based story. Further, we show that beliefs regarding the conditional moments of consumption growth are strongly time-varying and exhibit business cycle and/or long-run fluctuations. Much of the long-run behavior is unanticipated ex ante. We embed these subjective beliefs in a general equilibrium model to investigate further asset pricing implications. We find that learning significantly improves the model's ability to fit standard asset pricing moments, relative to benchmark model with fixed parameters.

Friday, April 19th, 2013, 10.30-12.00

Room 126, 1st floor of the Extranef building at the University of Lausanne

Learning about Consumption Dynamics

Michael Johannes, Lars Lochstoer, and Yiqun Mou*
Columbia Business School

April 18, 2013

Abstract

This paper studies the asset pricing implications of Bayesian learning about the parameters, states, and models determining aggregate consumption dynamics. Our approach is empirical and focuses on the quantitative implications of learning in real-time using post World War II consumption data. We characterize this learning process and find that revisions in beliefs stemming from parameter and model uncertainty are significantly related to realized aggregate equity returns. This evidence is novel, providing strong support for a learning-based story. Further, we show that beliefs regarding the conditional moments of consumption growth are strongly time-varying and exhibit business cycle and/or long-run fluctuations. Much of the long-run behavior is unanticipated ex ante. We embed these subjective beliefs in a general equilibrium model to investigate further asset pricing implications. We find that learning significantly improves the model's ability to fit standard asset pricing moments, relative to benchmark model with fixed parameters.

*All authors are at Columbia Business School, Department of Finance and Economics. We would like to thank Francisco Barillas (CEPR Barcelona discussant), Pierre Collin-Dufresne, Kent Daniel, Lars Hansen (AFA discussant), Lubos Pastor (NBER discussant), Tano Santos, and seminar participants at AFA Denver 2011, CEPR conference in Barcelona May 2011, Columbia University, NBER Asset Pricing meeting in Chicago 2011, Stanford, UT Austin, and the University of Wisconsin (Madison) for helpful comments. Corresponding author: Lars Lochstoer, 405B Uris Hall, 3022 Broadway, New York, NY 10027. Email: LL2609@columbia.edu

1 Introduction

This paper studies the asset pricing implications of learning about aggregate consumption dynamics. We are motivated by practical difficulties generated from the use of complicated consumption-based asset pricing models with many difficult-to-estimate parameters and latent states. For example, parameters or states controlling long-run consumption growth are at once extremely important for asset pricing and particularly difficult to estimate. Thus, we are interested in studying an economic agent who is burdened with some of the same econometric problems faced by researchers, a problem suggested by Hansen (2007).[1]

A large existing literature studies asset pricing implications of statistical learning – the process of updating beliefs about uncertain parameters, state variables, or even model specifications. Pastor and Veronesi (2009) provide a recent survey. In theory, learning can generate a wide range of implications relating to stock valuation, levels and variation in expected returns and volatility, and time series predictability, with many of the results focussed on the implications of learning about dividend dynamics.

Our analysis differs from existing work along three key dimensions. First, we focus on the empirical implications of simultaneously learning about parameters, state variables, and even model specifications. Most existing work focuses on learning a single parameter or state variable. Learning about multiple unknowns is more difficult as additional unknowns often confounds inference, slowing the learning process. Second, we focus on the specific implications of real-time learning about consumption dynamics from macroeconomic data during the U.S. post World War II experience. Thus, we are not expressly interested in general asset pricing implications of learning in repeated sampling settings, but rather the *specific* implications generated by the historical macroeconomic shocks realized in the United States over the last 65 years. Third, we use a new and stringent test of learning that relates updates in investor beliefs to contemporaneous, realized equity returns.

In studying the implications of learning, we focus on the following types of questions. Could an agent who updates his beliefs rationally detect non-i.i.d. consumption growth dynamics in real time? How rapidly does the agent learn about parameters and models? Are the revisions in beliefs about consumption moments correlated with asset returns, as a

[1] Hansen (2007) states: *"In actual decision making, we may be required to learn about moving targets, to make parametric inferences, to compare model performance, or to gauge the importance of long-run components of uncertainty. As the statistical problem that agents confront in our model is made complex, rational expectations' presumed confidence in their knowledge of the probability specification becomes more tenuous. This leads me to ask: (a) how can we burden the investors with some of the specification problems that challenge the econometrician, and (b) when would doing so have important quantitative implications" (p.2).*

learning story would require? Is there evidence that learning effects can help us understand standard asset pricing puzzles, such as the high equity premium, return volatility, and degree of return predictability?

One of the key implications of learning is that the agent's beliefs are nonstationary. For example, the agent may gradually learn that one model fits the data better than an alternative model or that a parameter value is higher or lower than previously thought, both of which generate nonstationarity in beliefs. The easiest way to see this is to note that the posterior mean of a parameter, $\mathbb{E}[\theta|y^t]$, where y^t is data up to time t, is trivially a martingale. Thus revisions in beliefs represent permanent, nonstationary shocks, that can have important asset pricing implications. For instance, nonstationary dynamics can generate a quantitatively important wedge between *ex post* outcomes and *ex ante* beliefs, providing an alternative explanation for standard asset pricing quantities such as the observed equity premium or excess return predictability.[2]

We study learning in the context of three standard Markov switching models of consumption growth: unrestricted 2- and 3-state models and a restricted 2-state model that generates i.i.d. consumption growth. The hidden states capture business cycle fluctuations and can be labeled as expansion and recession in 2-state models, with an additional 'depression' state in 3-state models.[3] Our key assumption is that the agent views the parameters, states, and even models as unknowns, using Bayes rule to update beliefs using consumption data, as well as additional macroeconomic data such as GDP growth in extensions.

To highlight the effects of parameter and model learning, we contrast the asset pricing results from the models with parameter and model uncertainty to the standard full-information implementation of the underlying models—that is, to the 2- and 3-state models where the state of the Markov chain is hidden, but where parameters are assumed known and set equal to the post-WW2 sample estimates. For the agent who is unsure about the parameters and even the model specification, the priors at the beginning of the post-WW2 sample are 'estimated' using all available U.S. consumption (and later, also GDP) data, starting from 1889, up until that time, in combination with quite flat initial priors in 1888. This is a common

[2] See also Cogley and Sargent (2008), Timmermann (1993), and Lewellen and Shanken (2002).

[3] Markov switching models for consumption or dividends are a benchmark specification in the literature, see, e.g., Mehra and Prescott (1985), Rietz (1988), Cecchetti, Lam, and Mark (1990, 1993), Whitelaw (2000), Cagetti, Hansen, Sargent, and Williams (2002), Barro (2006), Barro and Ursua (2008), Chen (2008), Bhamra, Kuehn, and Strebulaev (2008), Barro, Nakamura, Steinsson and Ursua (2009), Backus, Chernov, and Martin (2009), and Gabaix (2009). Rietz (1988) and, more recently, Barro (2006, 2009) argue that consumption disaster risk can help explain some of the standard macro-finance asset pricing puzzles.

approach to generate 'objective' priors.[4]

Our first results characterize the beliefs about parameters, states, models, and future consumption dynamics (e.g., moments) through the sample. The perceived dynamic behavior of aggregate consumption is at the heart of consumption-based asset pricing as it, jointly with preferences, determines the dynamic properties of the pricing kernel. In terms of beliefs, we compute at each point in time the posterior distribution of parameters, states, and models. As new data arrives, we update beliefs using Bayes rule. In addition to usual summaries of parameters and states, we also compute model probabilities and perform 'model monitoring' in real time as new data arrives. We find that the posterior probability of the i.i.d. model falls dramatically over time, provided the prior weight is less than one. Thus our agent is able to learn in real-time that consumption growth is not i.i.d., but has persistent components.[5] The agent believes that expected consumption growth is low in recessions and high in expansions, with the opposite pattern for consumption growth volatility. The 2-state model quickly emerges as the most likely, but the 3-state model with a depression state has $1-2\%$ probability at the end of the sample. At the onset of the financial crisis in 2008, the probability of the disaster model increases.[6]

There is significant learning about the expansion state parameters, slower learning about the recession state, and almost no learning about the disaster state, as it is rarely, if ever, visited. Thus, there is an observed differential in the speed of learning. Standard large sample theory implies that all parameters converge at the same rate, but the realized convergence rate depends on the actual observed sample path. We document a strong confounding effect of jointly learning about the hidden state of the economy and the parameters of the model, which slows down parameter learning considerably and leads to longer-lasting asset pricing implications of parameter learning. There is also strong evidence for nonstationary time-variation in the conditional means and variances of consumption growth, as well as measures of non-normality such as skewness and kurtosis. The agent's perception of the long-run mean (volatility) of consumption growth generally increases (decreases) over the sample.

[4] We do account for measurement error, which likely increased reported macroeconomic volatility during the pre-war period, as argued in Romer (1989). Malmendier and Nagel (2011) present evidence that the experience of the Great Depression affected investors' subsequent beliefs about risk and return, broadly consistent with the our prior calibration approach.

[5] This result is robust to persistence induced by time-aggregation of the consumption data (see Working (1960)).

[6] The posterior probability of the three-state model would change dramatically, if visited. For example, if a -3% quarterly consumption growth shock were realized at the end of the sample, the posterior probability of the three-state model would increase to almost 50%.

The perceived persistence of recessions (expansions) decreases (increases).[7] As the agent's beliefs about these parameters and moments change, asset prices and risk premia will also change.

The first formal test of the importance of learning regresses contemporaneous excess stock market returns on revisions in beliefs about expected consumption growth. This test, which to our knowledge is new to the literature, is a fundamental implication of any learning-based explanation: for learning to matter, unexpected revisions in beliefs about expected consumption growth should be reflected in the unexpected aggregate equity returns.[8]. We find strong statistical evidence that this relationship is positive. To disentangle parameter from state learning, we include revisions in beliefs generated by the fixed parameter prior as a control. Revisions in beliefs obtained from the model with parameter and model learning remain statistically significant, but revisions in beliefs generated by models with known parameters are statistically insignificant.

These results imply that learning about parameters and models is a statistically significant determinant of asset returns in our sample, confirming our main hypothesis. This result is strengthened if the agent learns from both consumption and GDP growth. It is important to note that our agent only learns in real-time and from macroeconomic fundamentals, as no asset price data (such as the dividend-price ratio) is used when forming beliefs. Since the revisions in beliefs obtained from the models with fixed parameters are statistically insignificant, the evidence questions the standard full-information, rational expectations implementation of the standard consumption-based model, at least for the models of consumption dynamics that we consider.[9]

As mentioned earlier, parameter and model learning generate nonstationary dynamics and permanent shocks that could have important implications. To investigate these implications, we consider a formal asset pricing exercise assuming Epstein-Zin preferences. Because the specific time-path of beliefs is important, the usual calibration and simulation approach used in the literature is not applicable, and we consider the following alternative pricing

[7] All of the results described in the current and previous paragraphs are robust to learning from additional GDP growth data.

[8] The sign of the effect would in a model depend on the elasticity of intertemporal substitution, and also on the other moments that change at the same time (volatility, skewness, kurtosis, etc.). In the model section, we show that this positive relation is consistent with a model with an elasticity of intertemporal substitution greater than 1.

[9] Parameter and model learning, on the one hand, and state learning on the other hand are distinct in our setting because the former generates a non-stationary path of beliefs, while the latter, after an initial burn-in period, is stationary.

procedure. At time t, given beliefs over parameters, models, and states, our agent prices a levered claim to a future consumption stream, computing quantities such as *ex-ante* expected returns and dividend-price ratios.[10] Then, at time $t+1$, our agent updates beliefs using new macro realizations at time $t+1$, recomputes prices, expected returns and dividend-price ratios. From this time series of prices, we compute realized equity returns, volatilities, etc. Thus, we feed historically realized macroeconomic data into the model and analyze the asset pricing implications for various models and prior specifications. This process is required when the time path matters and was previously used in, for example, Campbell and Cochrane (1999), where habit is a function of past consumption growth. We use standard preference parameters taken from Bansal and Yaron (2004).

Solving the full pricing problem with priced parameter uncertainty is computationally prohibitive, as the dimensionality of the problem is too large.[11] To price assets in a tractable way, while still incorporating learning, we follow Piazzesi and Schneider (2010) and Cogley and Sargent (2009) and use a version of Kreps' (1994) anticipated utility in our main asset pricing analysis. This implies that our agent prices claims at each point in time using current posterior means for the parameters and model probabilities, assuming those values will persist into the indefinite future. We do account for state uncertainty in this pricing exercise.

This pricing experiment provides additional evidence, along multiple dimensions, for the importance of learning. Focussing on the 3-state model, we first note that the model with parameters fixed at the full-sample values has a difficult time with standard asset pricing moments: the realized equity premium and Sharpe ratio are less than half the values observed in the data. The volatility of the price-dividend ratio is eighty percent less than the observed value. Parameter learning uniformly improves *all* of these statistics, bringing them close to observed values. It is important to note that this is not a calibration exercise – we did not choose the structural parameters to generate these returns.

The increase in the realized equity premium and return volatility is due to unexpected revisions in beliefs resulting from the parameter and model learning. In particular, the sample risk premium in the model with learning about both parameters and models is 3.7%

[10] We do price a levered consumption claim and introduce idiosyncratic noise to break the perfect relationship between consumption and dividend growth. The dividends are calibrated to match the volatility of dividend growth and the correlation between dividend and consumption growth.

[11] As an example, for the 3-state model there are twelve parameters, each with two hyperparameters characterizing the posteriors. This implies that we would have to have to solve numerically for prices on a very high dimensional grid, which is infeasible. There are additional difficult technical issues associated with priced parameter uncertainty, as noted by Geweke (2001) and Weitzman (2007).

(learning from consumption data only) and 5.2% learning from both consumption and GDP data), close to the 4.7% observed over the sample. In both cases, the sample average excess returns is about twice that of the sample average *ex ante* returns. Thus, the specific time path of beliefs about parameters and models realized over the post-WW2 period throws a large wedge between ex ante expectations and ex post outcomes and therefore matters quantitatively for standard asset pricing sample statistics. The high *ex post* average excess returns are due to revisions in beliefs in the direction that, overall, the economy is less risky and has higher growth than thought at the beginning of the sample. This also implies, looking forward, that the perceived equity premium is much smaller than the realized equity premium over the post World War II period. These points are consistent with the results in Cogley and Sargent (2008).[12]

In terms of predictability, the returns generated by learning over time closely match the data. In forecasting excess market returns with the lagged log dividend-price ratio, the generated regression coefficients and R^2's are increasing with the forecasting horizon and similar to those found in the data. The fixed parameters case, however, does not deliver significant *ex post* predictability, although the *ex ante* risk premium is in fact time-varying in these models as well because the time-variation in the risk premium assuming fixed parameters is too small relative to the volatility of realized returns to result in significant t-statistics. The intuition for why *in-sample* predictability occurs when agents are uncertain about parameters and models is the same as in Timmermann (1993) and Lewellen and Shanken (2002) – unexpected updates in growth and discount rates impact the dividend-price ratio and returns in opposite directions leading to the observed positive in-sample relation. Thus, in-sample predictability can be expected with parameter and model learning. The quantitatively large degree of in-sample relative to out-of-sample predictability we find is consistent with the literature.[13]

We also note that the model exhibits volatile long maturity risk-free yields, consistent

[12] Cogley and Sargent (2008) assume negatively biased beliefs about the consumption dynamics to highlight the same mechanism and also consider the role of robustness. In their model, the subjective probability of recessions is higher than the 'objective' estimate from the data. The results we present here are consistent with their conclusions, but our models are estimated from fundamentals in real-time, which allows for an out-of-sample examination of the time-series of revisions in beliefs. Further, we allow for learning over different models of the data generating process, as well as *all* the parameters of each model.

[13] For example, Fama and French (1988) document a high degree of in-sample predictability of excess (long-horizon) stock market returns using the price-dividend ratio as the predictive variable. On the other hand, Goyal and Welch (2008) and Ang and Bekaert (2007) document poor *out-of-sample* performance of these regressions in the data, and the historical and look-ahead prior learning models presented here are consistent with this evidence.

with the data. Learning about fixed quantities such as models or parameters generate permanent shocks that affect agents' expectations of the long-run (infinite-horizon) distribution of consumption growth. This is different from existing asset pricing models where only stationary variables affect marginal utility growth (see, e.g., Bansal and Yaron (2004), and Wachter's (2005) extension of Campbell and Cochrane (1999) model, as well as our fixed parameters benchmark model). In these models, long-run (infinite-horizon) risk-free yields are constant as the transitory shocks to marginal utility growth die out in the long run. This is additional evidence supporting a learning-based explanation relative to the fixed parameters alternative.

The permanent nature of updates in beliefs about the consumption dynamics is a source of subjective long-run consumption risks, as shown in Collin-Dufresne, Johannes, and Lochstoer (2013). In particular, if the agent considers parameter uncertainty in a fully rational setting, as opposed to the anticipated utility setting, shocks to beliefs resulting from parameter updates are priced. If the agent is aware beliefs will change in the future, the asset price response to shocks to beliefs may be less pronounced, questioning the robustness of our results. To investigate this hypothesis, we solve for the fully rational pricing problem, the very high-dimensional case of the 2-state model where all parameters are unknown. Consistent with the analysis in Collin-Dufresne, Johannes, and Lochstoer (2013), the ex ante risk premium increases strongly due to the presence of long-run consumption risks arising from belief updates. Overall, however, the conclusions of our analysis remain the same. The time-path of beliefs derived from the learning problem strongly increases the covariance between stock prices in the data with those implied by the model, and the ex post sample risk premium is almost twice the average ex ante risk premium.

In conclusion, our results strongly support the importance of parameter and model learning for understanding the joint behavior of consumption and asset prices in the U.S. post World War II sample. First, parameter and model learning leads to a time path of belief revisions that are correlated with realized equity returns, controlling for realized consumption growth. Second, the time series of beliefs help explain the time-series of the price level of the market (the time-series of the price-dividend ratio) in a general equilibrium model. Third, beliefs display strong nonstationarity over time, driving a large wedge between *ex ante* beliefs and *ex post* realizations. Fourth, permanent shocks to beliefs generate permanent shocks to marginal utility growth. These features help explain common asset pricing puzzles such as excess return volatility, the high sample equity premium, the high degree of in-sample return predictability, and the high volatility of long-run yields, all relative to a fixed parameter alternative. The results are generated by real-time learning from consumption (and GDP growth), using standard preference parameters without directly calibrating to asset returns. In this sense the results are entirely "out-of-sample."

2 The Environment

2.1 Model

We follow a large literature and assume an exogenous Markov or regime switching process for aggregate, real, per capita consumption growth dynamics. Log consumption growth, Δc_t, evolves via:

$$\Delta c_t = \mu_{s_t} + \sigma_{s_t}\varepsilon_t, \tag{1}$$

where ε_t are i.i.d. standard normal shocks, $s_t \in \{1, ..., N\}$ is a discretely-valued Markov state variable, and $\left(\mu_{s_t}, \sigma^2_{s_t}\right)$ are the Markov state-dependent mean and variance of consumption growth. The Markov chain evolves via a $N \times N$ transition matrix Π with elements π_{ij} such that $\text{Prob}[s_t = j | s_{t-1} = i] = \pi_{ij}$, with the restriction that $\sum_{j=1}^{N} \pi_{ij} = 1$. The fixed parameters of the N-state model contain the means and variances in each state, $\left\{\mu_n, \sigma_n^2\right\}_{n=1}^{N}$ as well as the elements of the transition matrix. The transition matrix controls the persistence of the Markov state.

Markov switching models are flexible and tractable and have been widely used since Mehra and Prescott (1985) and Rietz (1988). By varying the number, persistence, and distribution of the states, the model can generate a wide range of economically interesting and statistically flexible distributions. Although the ε_t's are i.i.d. normal and the distribution of consumption growth, conditional on s_t and parameters, is normally distributed, the distribution of future consumption growth is neither i.i.d. nor normal due to the shifting Markov state. This time-variation induces very flexible marginal and predictive distributions for consumption growth. These models are also tractable, as it is possible to compute likelihood functions and filtering distributions, given parameters.

We consider two and three state models and also consider a restricted version of the two state model generating i.i.d consumption growth by imposing the restriction $\pi_{11} = \pi_{21}$ and $\pi_{22} = \pi_{12} = 1 - \pi_{11}$. Under this assumption, consumption growth is an i.i.d. mixture of two normal distributions, essentially a discrete-time version of Merton's (1976) mixture model. The general two and 3-state models have 6 and 12 parameters, respectively. The i.i.d. two state model has 5 parameters ($\mu_1, \mu_2, \sigma_1, \sigma_2$ and π_{11}).

It is common in these models to provide business cycle labels to the states. In a 2-state model, we interpret the two states as 'recession' and 'expansion,' while the three state model additionally allows for a 'disaster' state.[14] Although rare event models have been used for understanding equity valuation since Rietz (1988), there has been a recent resurgence in research using these models (see, e.g., Barro (2006, 2009), Barro and Ursua (2008), Barro, Nakamura, Steinsson and Ursua (2009), Backus, Chernov, and Martin (2009), and Gabaix (2009)).

2.2 Information and learning

To operationalize the model, additional assumptions are required regarding the economic agent's information set. Since we want to model learning similar to that faced by the econometrician, we assume agents observe aggregate consumption growth, but are uncertain about the Markov state, the parameters, and the total number of Markov states. We label these unknowns as state, parameter, and model uncertainty, respectively. We assume agents are Bayesian, which means they update initial beliefs via Bayes' rule as data arrives. Later in the paper, we develop an extension to this model where agents can also learn from a vector of additional macro variables and consider the case of additional learning from GDP growth data.

The learning problem is as follows. We consider $k = 1, ..., K$ models, $\{\mathcal{M}_k\}_{k=1}^K$, and in model $\mathcal{M}_k$, the state variables and parameters are denoted as s_t and θ, respectively.[15] The distribution $p(\theta, s_t, \mathcal{M}_k|y^t)$ summarizes beliefs after observing data $y^t = (y_1, ...y_t)$. To understand the components of the learning problem, we can decompose the posterior as:

$$p\left(\theta, s_t, \mathcal{M}_k|y^t\right) = p\left(\theta, s_t|\mathcal{M}_k, y^t\right) p\left(\mathcal{M}_k|y^t\right). \tag{2}$$

$p(\theta, s_t|\mathcal{M}_k, y^t)$ solves the parameter and state "estimation" problem conditional on a model and $p(\mathcal{M}_k|y^t)$ provides model probabilities. It is important to note that this is a non-trivial,

[14] We do not consider, for instance, 1- or 4-state models as the Likelihood ratios of these relative to the 2- or 3-state model show that the 2- and 3-state models better describe the data. As we will show, however, there is some time-variation in whether a 2- or 3-state model matches the data better, which is one of the reasons we entertain both of these as alternative models.

[15] This is a notational abuse. In general, the state and dimension of the parameter vector should depend on the model, thus we should superscript the parameters and states by 'k', θ^k and s_t^k. For notational simplicity, we drop the model dependence and denote the parameters and states as θ and s_t, respectively.

high-dimensional learning problem, as posterior beliefs depend in a complicated manner on past data and can vary substantially over time. The dimensionality of the posterior can be high, in our case more than 10 dimensions.

One of our primary goals is to characterize and understand the asset pricing implications of the transient process of learning about the parameters, states, and models.[16] Learning generates a form of nonstationarity, since parameter estimates and model probabilities are changing through the sample. When pricing assets, this can lead to large differences between *ex ante* beliefs and *ex post* outcomes, as shown in Cogley and Sargent (2008). Given this nonstationarity, we are concerned with understanding the implications of learning based on the specific experience of the U.S. post-war economy.[17]

To operationalize the learning problem, we need to specify the prior distribution, the data the agent uses to update beliefs, and develop an econometric method for sampling from the posterior distribution. In terms of data, we in a benchmark case assume that agents learn only from observing past and current consumption growth, a common assumption in the learning literature (see, e.g., Cogley and Sargent (2008) and Hansen and Sargent (2009)). The primary data used is the 'standard' data set consisting of real, per capita quarterly consumption growth observations obtained from the Bureau of Economic Analysis (the National Income and Product Account tables) from 1947:$Q1$ until 2009:$Q1$.

2.3 Initial beliefs

The learning process begins with initial beliefs or the prior distribution. In terms of functional forms, we assume proper, conjugate prior distributions (Raiffa and Schlaifer (1956)). One alternative would be flat or 'uninformative' priors, but this is not possible in Markov switching models, as this creates identification issues (the label switching problem) and causes problems sampling from the posterior.[18] Conjugate priors imply that the functional form

[16] These type of problems received quite a bit of theoretical attention early in the rational expectations paradigm - see for example Bray and Savin (1986) for a discussion of model specification and convergence to rational expectations equilibria by learning from observed outcomes.

[17] This is different from the standard practice of looking at population or average small-sample unconditional asset price and consumption growth moments from a model calibrated to the U.S. postwar data – we are looking at a single outcome corresponding to the U.S. post-war economy.

[18] The label switching problem refers to the fact that the likelihood function is invariant to a relabeling of the components. For example, in a two-state model, it is possible to swap the definitions of the first and second states and the associated parameters without changing the value of the likelihood. The solution is to impose parameter constraints in optimization for MLE or to use informative prior distributions for Bayesian approaches. These constraints/information often take the form of an ordering of the means or variances of the parameters. For example in a two state model, it is common to impose that $\mu_1 < \mu_2$ and/or $\sigma_1 < \sigma_2$ to

of beliefs is the same before and after sampling, are analytically tractable for econometric implementation, and are flexible enough to express a wide range initial beliefs.

For the mean and variance parameters in each state, (μ_i, σ_i^2), the conjugate prior is $p(\mu_i|\sigma_i^2)p(\sigma_i^2) \sim \mathcal{NIG}(a_i, A_i, b_i, B_i)$, where $\mathcal{NIG}$ is the normal/inverse gamma distribution. The transition probabilities are assumed to follow a Beta distribution in 2-state specification and its generalization, the Dirichlet distribution, in models with three states. Calibration of the hyperparameters completes the specification.

We endow our agent with economically motivated initial beliefs to study how learning proceeds from various starting points. We consider an 'objective' approach that uses a training sample to calibrate the prior distribution. Training samples are the most common way of generating objective prior distributions (see, e.g., O'Hagan (1994)). In this case, an initial data set is used to provide information on the location and scale of the parameters. In our application, we use the annual consumption data from Shiller from 1889 until 1946. Given the prior generated from the training sample, learning proceeds on the second data set – in our case, the post World War II sample.[19] To alleviate issues with prior misspecification due to the change in the quality of the macro data when going from annual to quarterly data, which only are available from 1947, we also use a 10-year burn-in period in the post-WW2 sample, where beliefs are updated from data from 1947 through 1956. Thus, in terms of the pricing exercises, we consider as priors the beliefs this agent holds in $1957Q1$.

We contrast the results from the cases with parameter and model uncertainty to models with 'fixed parameter' priors. This is a point-mass prior located at the maximum likelihood estimates of the parameters of a given model estimated on the pricing sample – i.e., over the post-WW2 period. In this case, the agent only learns about the latent Markov state. This way of calibrating an exchange economy model mimics the typical rational expectations approach. Contrasting this model to the full learning model allows us to identify the differential roles of state and parameter learning.

[Table 1 about here]

We will discuss the evolution of beliefs in the post-WW2 sample in detail in what follows. Still, it is useful to preempt these results and discuss briefly some general properties of the

breaks the symmetry of the likelihood function.

[19]Romer (1989) presents evidence that a substantial fraction of the volatility of macro variables such as consumption growth pre-WW2 is due to measurement error. To alleviate this concern, we set the prior mean over the variance parameters to a quarter of the value estimated over the training sample. See the Online Appendix for further details.

prior beliefs at this stage. Table 1 shows the mean and dispersion of beliefs about each of the parameters in each of the three models that will later be used as the priors for the pricing exercise.[20] The 2-state model with persistent states have parameters that correspond to business cycle fluctuations in consumption dynamics, with slightly negative consumption growth in recessions that are persistent but shorter-lived than the high growth expansions. The 3-state model has in addition a rare and quite short-lived 'Depression' state that has a mean expected quarterly growth rate of -1.8%. The bad state in the i.i.d. 2-state model has a bad state that is in between the recession and the depression states of the 3-state model. The prior standard deviation of, say, the mean in the good state in the 3-state model is 0.17% whereas the prior standard deviation of the mean in the Depression state is 0.47%. This reflects the fact that historically there are more observations drawn from the good state than from recession state than from the depression state. Thus, the information in the training sample leads to priors that embody reasonable properties given the pre-WW2 data.

3 The time-series of subjective beliefs

This section characterizes the learning process. We first discuss state, parameter, and model learning and their implications for the time series of conditional consumption moments, as perceived by the Bayesian agent. Next, we empirically investigate how revisions in the agent's beliefs are related to stock market returns. We also consider the case of learning from GDP data, in addition to consumption data. In the following section, we embed these beliefs in a general equilibrium model and discuss the asset pricing implications in more detail.

3.1 State, parameter, and model learning

Conditional on a model specification, our agent learns about the Markov state and the parameters, with revisions in beliefs generated by a combination of data, model specification, and initial beliefs. To start, consider the agent's beliefs about the current state of the economy, s_t, where state 1 is an 'expansion' state, state 2 the 'recession' state and, if a

[20] Further details of how we arrive at the priors, the specific prior parameters chosen, as well as a description of the econometric technique we apply to solve this high-dimensional learning problem (particle filtering) are given in the Online Appendix.

3-state model, state 3 the 'depression' state. Estimates are given by

$$E\left[s_t|\mathcal{M}_k, y^t\right] = \int s_t p\left(\theta, s_t|\mathcal{M}_k, y^t\right) d\theta ds_t.$$

Note that these are marginal *mean* state beliefs, as parameter uncertainty is integrated out. Although s_t is discrete, the mean estimates need not be integer valued. Figure 1 displays the posterior state beliefs over time, for each model and for different priors.

There are a number of notable features of these beliefs. NBER recessions (shaded yellow) and expansions are strongly related to the beliefs about the current state, especially in the general 2- and 3-state models. The only exceptions are the recessions in the late 1960s and 2001, which were not associated with substantial consumption declines. Comparing the panels, the models generate strong differences is persistence of the states. The i.i.d. model identifies recessions as a one-off negative shock, but since shocks are i.i.d., the agent does not forecast that the recession state will persist with high likelihood. In contrast, the 2- and 3-state models clearly show the persistence of the recession states. Depression states are rare – after the initial 10-year burn-in period from 1947 to 1957, there are only really two observations that place even modest probability on the depression state – the recession in 1981 and the financial crisis at the end of 2008. This implies that depression states are nearly 'Peso' events in the post WW2 sample.

The agent's beliefs are quite volatile early in the sample in all of the models. This is not surprising. Since initial parameter beliefs are highly uncertain, the agent has a difficult time discerning the current state as parameter uncertainty exacerbates state uncertainty. As the agent learns, parameter uncertainty decreases and state identification is easier. It is important to note that even with full knowledge of the parameters, the agent will never be able to perfectly identify the state.[21] In sum, even when we allow uncertainty about parameters, the state uncertainty principally reflects uncertainty about the stationary fluctuations of the economy, associated with business cycles. This is a result from filtering the data given the prior beliefs and not immediate ex ante.

Next, consider beliefs over parameters. Due to the large number of parameters and in the interests of parsimony, we focus on a few of the more economically interesting and important parameters. For the 2-state model, the top panels of Figure 2 display posterior means

[21] The posterior variance of the state, $var\left[s_t|\mathcal{M}_k, y^t\right]$, does decline over time due to decreasing parameter uncertainty. This will be discussed further when we use GDP growth as an additional observation to help identify the state.

Figure 1 - Evolution of Mean State Beliefs

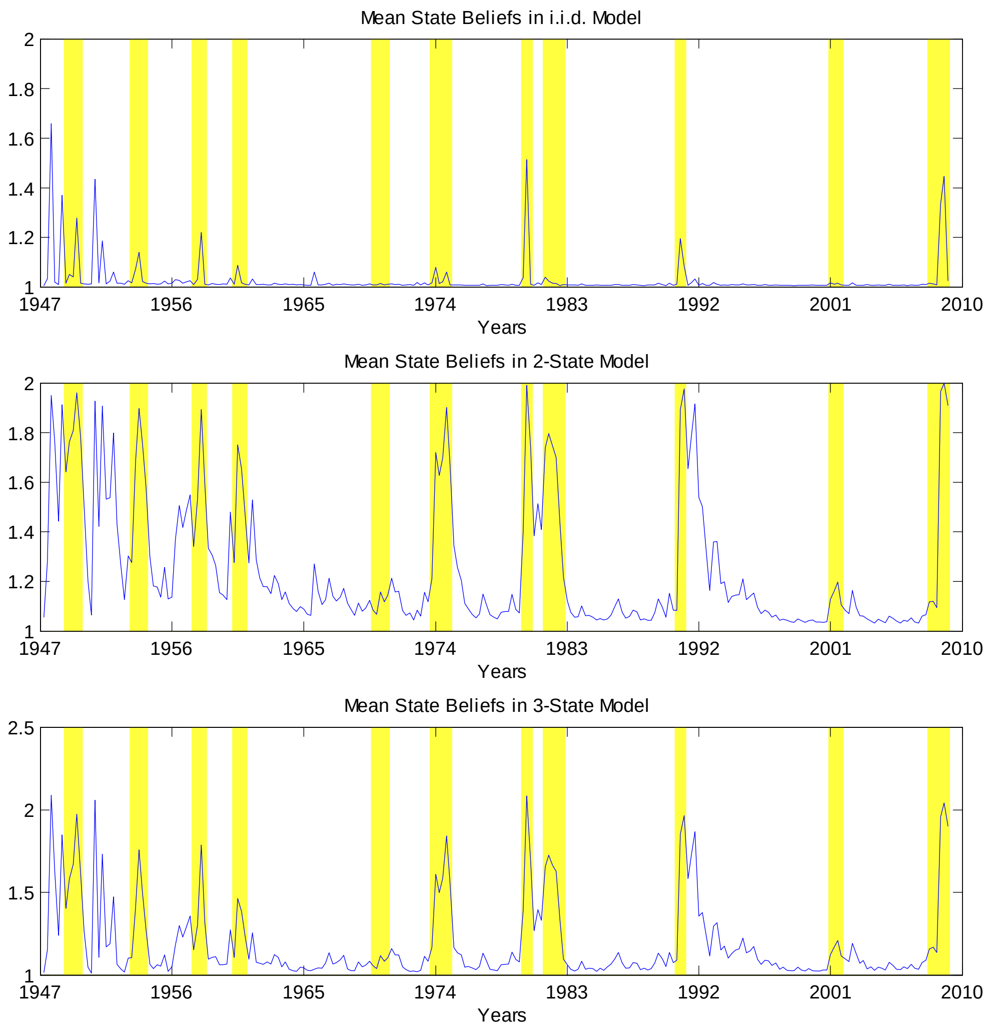

Figure 1: The plots show the means of agents' beliefs about the state of the economy at each point in time, $E_t[s_t]$. Note that s_t =1,2 in the i.i.d. and the 2-state model, while s_t =1,2,3 in the 3-state model. "1" is an expansion state, "2" is a recession state, and "3" is a disaster state. The models have either 2 or 3 states as indicated on each plot, and the time t state beliefs are formed using the history of consumption only up until and including time t. The "i.i.d. Model" is a model with i.i.d. consumption growth but that allows for jumps ("2" is a jump state). The sample is from 1947:Q2 until 2009:Q1.

of the beliefs over σ_1^2 and σ_2^2. Notice that these conditional variance parameters overall decrease slowly throughout the sample. This is a combination of the Great Moderation (realized consumption volatility did decrease over the post-war sample) and the initial beliefs, which based on the historical experience expected higher consumption growth volatility. The decline in consumption volatility conditional on being in the good state is quite large over the sample (from about 0.7% per quarter to about 0.4% per quarter).

The lower panels in Figure 2 display the mean beliefs over the transition probabilities, π_{11} and π_{22}. After the 10-year burn-in period, the former is essentially increasing over the sample, while the latter is decreasing. That is, 50 years of, on average, long expansions and high consumption growth leads to revisions in beliefs that are manifested in higher probabilities of staying in the good state and lower probabilities of staying in recession state. The probability of staying in an expansions, conditional on being in an expansion, goes down up from 0.88 to 0.93. Such positive shocks to the agents' perception of the data generating process lead all else equal to higher ex post equity returns than compared to ex ante expectations.

As previously stated, the mean beliefs follow a unit root process, from the law of iterated expectations. Thus, the non-stationary parameter drifts apparent in Figure 2 are a source of permanent shifts in beliefs that of this reason will strongly affect asset prices.

Figure 3 shows the marginal model probabilities, $p\left(\mathcal{M}_k|y^t\right)$.[22] The prior probability of each model as of 1947Q1 is set to 1/3. Note first that the posterior probability of the i.i.d. model decreases relatively quickly towards zero upon having observed both the high growth period of the 1960s followed by two strong recessions in the 1970s, indicating persistence in consumption growth. Thus, i.i.d. consumption growth is rejected by a Bayesian agent that updates by observing past realized consumption growth. Although not reported for brevity, this conclusion is robust even if the prior probability of the i.i.d. model is set to 0.95 - in this case it takes somewhat longer (but still only slightly more than half the sample) for the probability of the i.i.d. model to drop very close to zero. The 3-state model also sees a reduction in its likelihood and ends at about a 1% probability level at the end of the sample. As mentioned in the introduction, a single large negative consumption shock would quickly change these probabilities. In sum, we observe large, non-stationary changes in the model

[22] Note that marginal model probabilities (i.e., where parameter uncertainty is integrated out) penalizes extra parameters as more sources of parameter uncertainty tends to flatten the likelihood function. Thus, it is not the case, as we see an example of here, that a 3-state model always dominates a 2-state model in Bayesian model selection.

Figure 2 - Evolution of Mean Parameter Beliefs

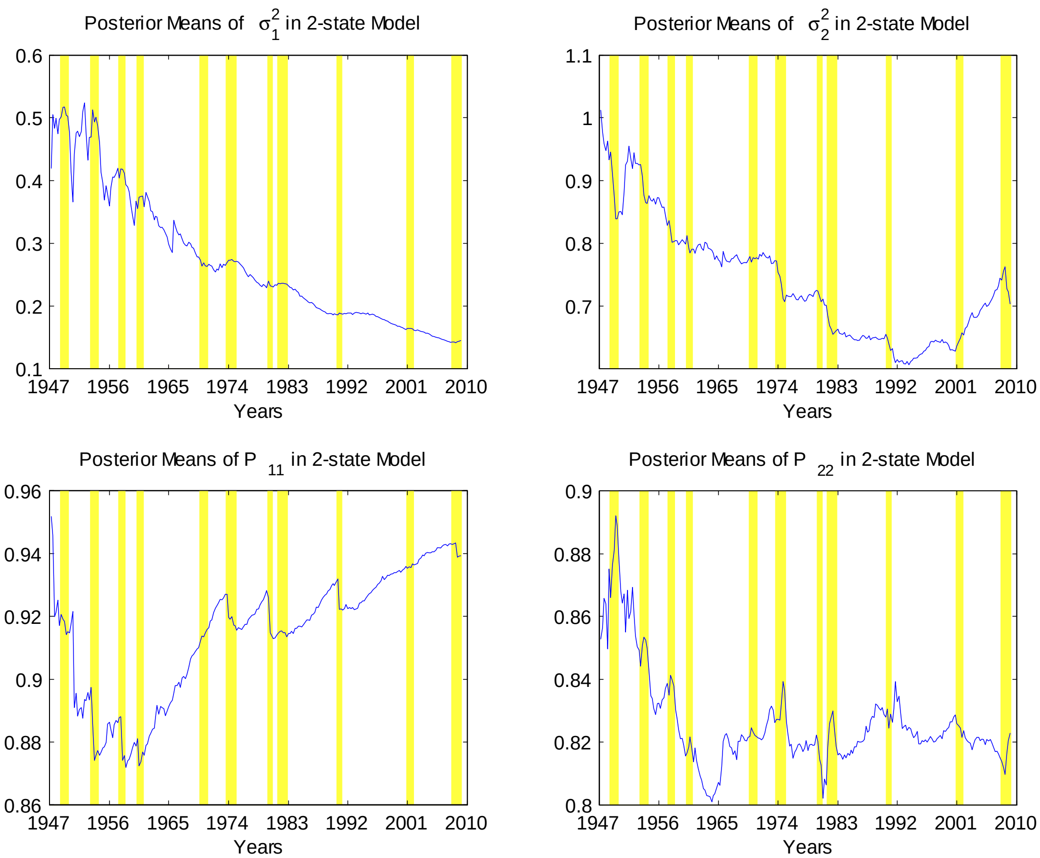

Figure 2: The two top plots in this figure show the mean beliefs about the volatility parameters within each state for the 2-state model. The two lower plots show for the same model the mean beliefs of the probabilities of remaining in the current state. The sample is from 1947:Q2 until 2009:Q1.

uncertainty over the sample.

The fact that the agent can learn that consumption growth is not i.i.d. is important. Many asset pricing models specify i.i.d. consumption growth with the implicit assumption that it is not possible or difficult to detect non-i.i.d. dynamics in consumption. Our results show that agents, using only consumption growth data, can detect non-i.i.d. dynamics, and can do so in real time, which is an even stronger result. The agent does not need to wait until the end of the sample. This result holds for various prior specifications and is robust to time-aggregation.[23]

3.2 Speed of learning

Bayesian learning is efficient and therefore typically fast. Previous literature has focused on learning about one dimension at a time (e.g., the current state, or one parameter), whereas we consider a general, high-dimensional learning problem. Learning about multiple dimensions of uncertainty leads to confounding effects that slows down the learning process in two particularly interesting ways.

First, learning is faster for parameters that govern the dynamics in states that are more common. Table 1 shows the end-of-sample posterior mean and standard deviation for each parameter. Two patterns emerge. In terms of mean beliefs, the agent updates beliefs in the overall direction of the world being a less risky place—expansions last longer, recessions are less severe, and volatilities are lower in the end-of-sample beliefs relative to the beginning-of-sample priors. In terms of posteriors, the end-of-sample posterior standard deviation of beliefs, as a fraction of prior standard deviation of beliefs, are smaller for parameters associated with the good state, than for parameters associated with the recession state, than for the parameters associated with the depression state. In fact, since the depression state did not really materialize in the post-WW2 sample, there is almost no learning about the parameters of the depression state, whereas the reduction in the standard deviation of beliefs about, say, the mean in the good state is 82%. While this is a quite intuitive finding, it has important asset pricing implications. For instance, updates in beliefs will be larger in recession and depression states than in the good state, which will be reflected in a higher volatility of returns in these states, consistent with the data.

[23] In the Appendix, we show that taking out an autocorrelation of 0.25 from the consumption growth data, which is what time-aggregation of i.i.d. data predicts (see Working (1960)), does not qualitatively change these results - if anything it makes the rejection of the i.i.d. model occur sooner. The same is true if we purge the data of its full sample first order autocorrelation.

Figure 3 - Marginal Model Probabilities

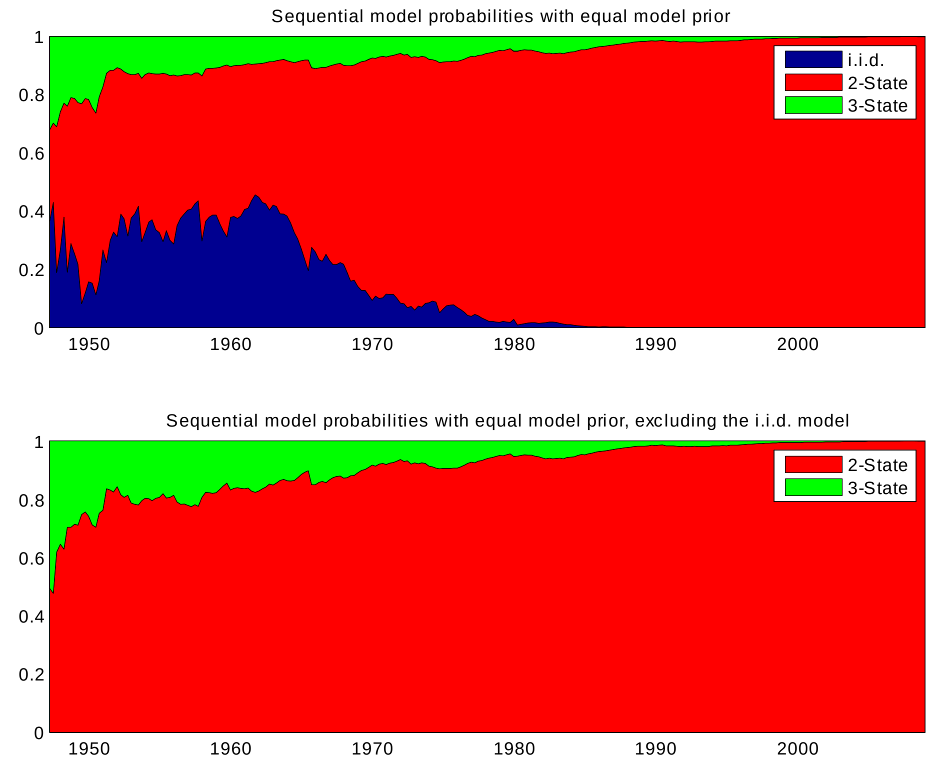

Figure 3: The top panel shows the evolution of the probability of each model being the true model, where the models at the beginning of the sample are set to have an equal probability, and where state and parameter uncertainty have been integrated out. The lower plot shows the same when the agent considers only the general 2-state model and the 3-state model as possible models of consumption dynamics. The sample period is 1947:Q2 - 2009:Q1.

[Table 2 about here]

Second, there is a strong confounding effect that arises from the joint learning about parameters, θ, and the state of the economy, s_t. Table 2 shows the results of the following Monte-Carlo experiment based on the 2-state model with true parameters set to the maximum likelihood estimates obtained from post-WW2 consumption data. In particular, we simulate 1,000 economies of length 209 quarters, as in the 1957 to 2009 sample, where the agent is assumed to have unbiased beliefs with dispersion equal to that given in Table 1. Next, we let in one case the agent learn the parameters and the state of the economy by filtering from consumption growth only. In the other case, we assume the agent also observes the current state and, therefore, only needs to learn about the parameters within each state.

The average difference in the end-of-sample posterior variance of parameter beliefs between these two cases is striking. For all the parameters, the posterior variance decreases significantly more in the case where the states are observed relative to when they are unobserved. This effect is strongest for the transition probabilities. If the states are observed, the filtering problem can in fact be solved analytically. In terms of the transition probabilities, the parameter learning is simply a matter of counting the frequency of transitions from one state to itself or the other state. With unobserved states, however, the number of state variables, or sufficient statistics, that characterize beliefs is proportional to 2^T, where T is the sample length. Thus, it is for pricing purposes an infinite-dimensional problem.

The reason the problem is so high-dimensional is that parameter learning depends on the identification of the state, while the identification of the state depends on the parameter beliefs. It is this feedback mechanism in the learning problem that strongly confounds inference. This finding provides a rationale for using the particle filter, which is a numerical approximation algorithm designed to solving such high-dimensional, sequential inference problems.

3.3 Beliefs about consumption growth moments

The results of the previous section show that beliefs about parameters, states, and models vary through the sample, but it is not clear from this how much variation in the conditional moments of consumption growth is present.[24] To provide asset-pricing relevant measures, we

[24] As an example, consider the conditional volatility of consumption growth. A decrease in the probability of the bad state, which has higher consumption growth volatility, could be offset by an increase in the

Figure 4 - Quarterly Mean and Standard Deviation of Consumption Growth

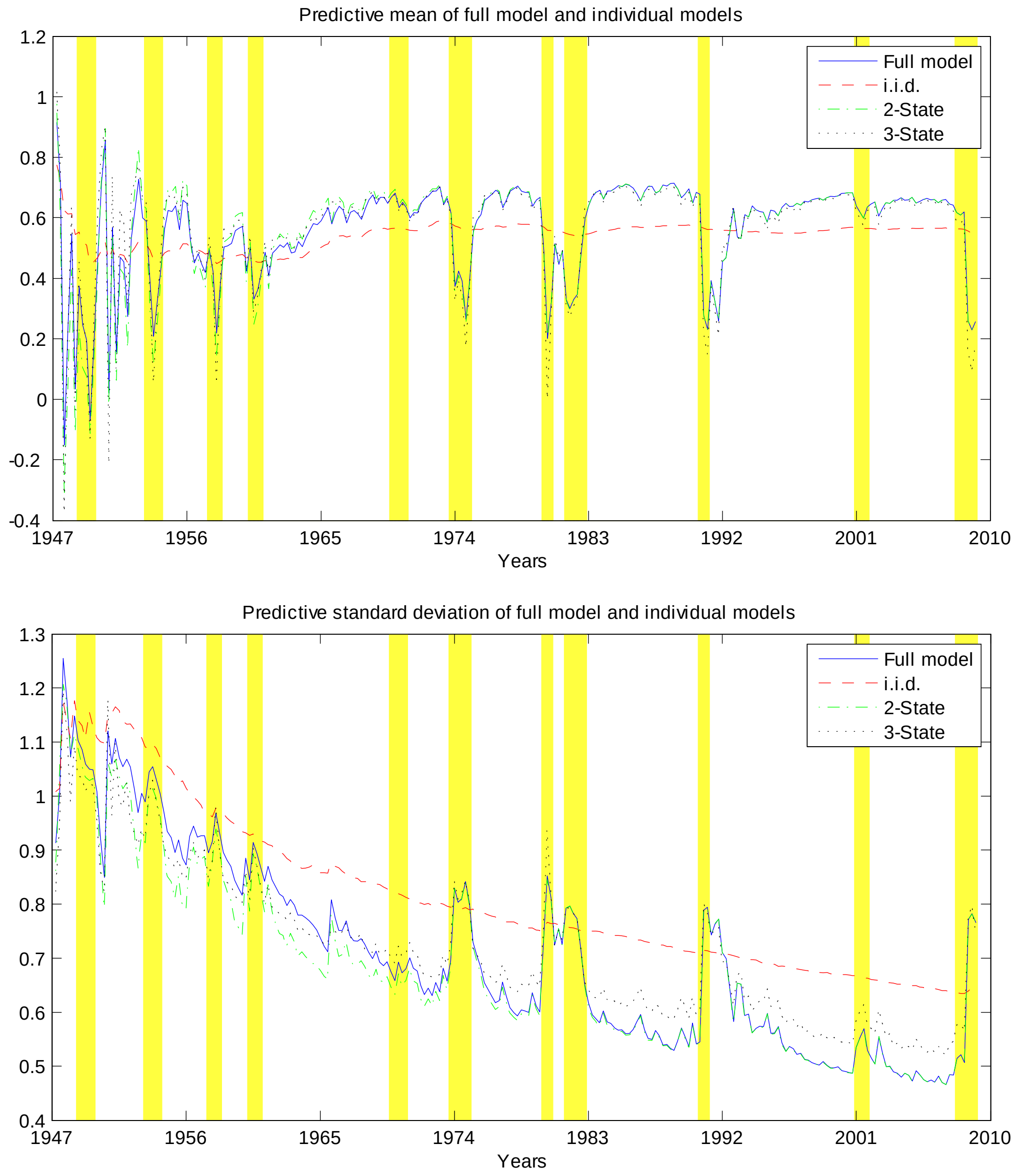

Figure 4: The top panel shows the quarterly conditional expected consumption growth, where state and parameter uncertainty have been integrated out, from each of the three benchmark models: the "i.i.d.", and the general 2- and 3-state switching regime models. The solid line shows the conditional expected consumption growth rate for the 'full' model, where also model uncertainty has been integrated out. The lower plot shows the expected quarterly conditional consumption growth for the same cases. The sample period is 1947:Q2 - 2009:Q1.

report the agent's beliefs regarding the first four moments of conditional consumption growth and model probabilities. All of these quantities are marginal, integrating out parameter, state, and/or model uncertainty. For example, the predictive mean for a given model, $\mathcal{M}_k$, is

$$E\left[\Delta c_{t+1}|\mathcal{M}_k, y^t\right] = \int \Delta c_{t+1} p\left(\Delta c_{t+1}|\theta, s_t, \mathcal{M}_k, y^t\right) p\left(\theta, s_t|\mathcal{M}_k, y^t\right) d\theta ds_t.$$

In describing these moments, we in our discussion refer to the first ten years as a 'burn-in' period, in order to allow the beliefs some time to adjust to the data.

The top panel in Figure 4 displays the conditional expected quarterly consumption growth for each model, as well as the full model where model uncertainty is also integrated out. The 2- and 3-state models generate relatively modest differences in this moment – both pick up business cycle fluctuations in expected consumption growth, with the 3-state model identifying the recessions in the early 80's and the financial crisis in '08 as severe. Persistent recessions are missing from the i.i.d. model, as expected. All three models exhibit a low frequency increase in expected consumption growth over the first half of the sample, due to parameter learning. When averaging across the three models by weighting by the model probabilities give the beliefs from the full learning problem. Overall, recessions are associated with a mean quarterly consumption growth of about 0.3%, while the mean consumption growth in expansions is about 0.6%. Since business cycles are relatively persistent, these fluctuations in conditional consumption growth are a source of long-run consumption risk, akin to that of Bansal and Yaron (2004). However, the lower frequency fluctuations we observe in expected consumption growth, which is due to parameter learning, constitute "truly" long-run risk, as shocks to parameter beliefs are permanent.

Averaging beliefs across models using the model probabilities, we get the beliefs from the 'full' learning model. These beliefs do not deviate much from those arising in the 2-state mode, so model uncertainty does not play a major role in determining expected consumption growth dynamics.

Turning to the conditional volatility of quarterly consumption growth, the bottom panel of Figure 4 shows that there is a downward trend in consumption growth volatility through the sample, with marked increases during recessions for the non-i.i.d. models. The secular decline in the conditional standard deviation of consumption growth is largely driven by downward revisions in estimates of the volatility parameters as realized consumption growth

---

consumption volatility in the good state, σ_1, keeping the total conditional volatility of consumption growth constant.

was less volatile in the second half of this century. In the agent's beliefs, the conditional volatility of consumption growth in expansions decreases from a little more than 1% per quarter to about 0.5%. This is the Great Moderation - the fact that consumption volatility has decreased also over the post-war sample. In the models considered here, the agent learning in real-time perceives this decrease to happen gradually, in contrast to studies that find *ex post* evidence of structural breaks or regime shifts at certain dates. For this moment, however, model learning increases the downward drift. The 3-state model has overall higher conditional consumption volatility than the 2-state model due to the presence of the 'depression' state'. At the same time, the probability of the 3-state model is decreasing over the sample. This shows how model learning, just like parameter learning, contributes to non-stationary changes in beliefs.

Every recession in the sample is associated with higher consumption growth volatility, although the size of the increase varies. The largest increase, on a percentage basis, occurs with the financial crisis of 2008. The increase is largest in the 3-state model, as the mean state belief at this time approaches the third state, which has a very high volatility. There is little updating about the volatility of the disaster state through the sample, since there have been no prolonged visits to this state. Thus, this reflects the fear that prevailed in the fall of 2008 that the economy was potentially headed into a depression not seen since the 1930s. This econometric result squares nicely with anecdotes from the crisis.

The top panel of Figure 5 shows the time-series of conditional consumption growth skewness, again with the model averaged estimates reported as the 'full' learning model. The time-variation in the conditional skewness is dominated by business cycle variation for the two and 3-state models, and there is a slight downward trend, as the probability of a disaster and recession decrease. When the economy is in a recession, consumption growth is naturally less negatively skewed for two reasons: (1) there is a high probability that the economy jumps to a higher (i.e. better) state and (2) expected consumption volatility is high, which tends to decrease skewness. Note that in terms of skewness, the 3-state model, with its severe recession (disaster) state, has more negative skewness than the 2-state model. Since the probability of the 3-state model is decreasing over the sample, the model learning creates less drift in the conditional skewness than any of the individual models do.

The bottom panel of Figure 5 shows the time-series of conditional consumption growth excess kurtosis. Conditional kurtosis is lower in bad states as these states are the least persistent and volatility is highest. Large, rare, outcomes are more likely when the econ-

Figure 5 - Quarterly Consumption Growth Skewness

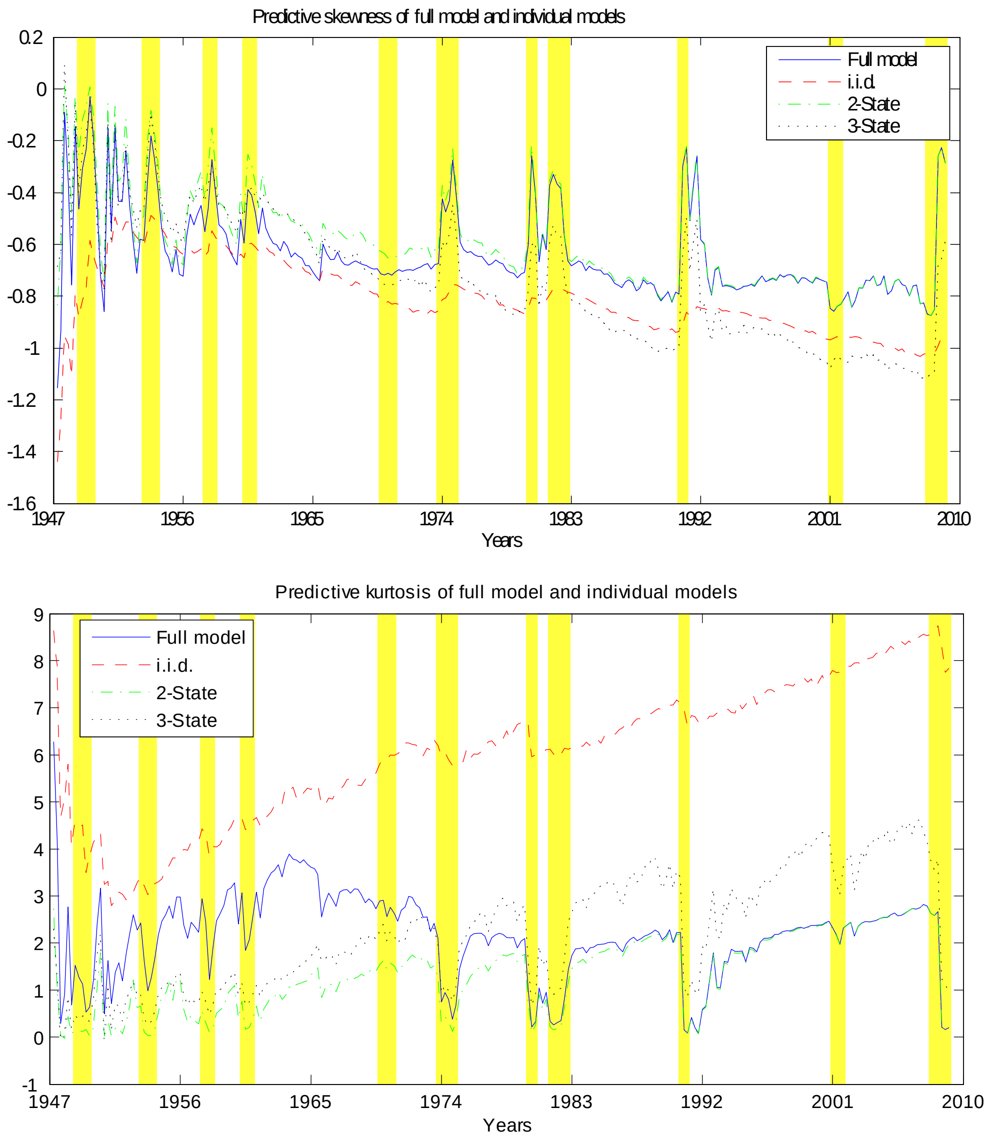

Figure 5: The top panel shows the quarterly conditional skewness of consumption growth, where state and parameter uncertainty have been integrated out, from each of the three benchmark models: the "i.i.d.", and the general 2- and 3-state switching regime models. The solid line shows the conditional expected consumption growth rate for the 'full' model, where also model uncertainty has been integrated out. The lower plot shows the conditional kurtosis of consumption growth for the same cases. The sample period is 1947:Q2 - 2009:Q1.

omy is in the good state. This has potentially interesting option pricing implications (see, e.g., Backus, Chernov, and Martin (2009)), as the skewness and kurtosis will be related to volatility smiles. For each model the kurtosis is overall increasing over the sample as the persistence of the good state increases creating a more non-normal predictive consumption growth distribution.

Parameter uncertainty gives an extra 'kick' to conditional skewness and kurtosis measures relative to the case of fixed parameters, where the skewness and kurtosis both move little over time (the fixed parameter case is not reported here for brevity). Both for skewness and kurtosis, there is clear evidence of parameter learning over the business cycle: the skewness becomes more negative and the kurtosis higher the longer an expansion last, reflecting updating of the transition probabilities, which reflect business cycle dynamics. Similar to skewness, there are now relatively large differences between the 2- and 3-state models, but for kurtosis the i.i.d. model stands out with its high kurtosis. The 3-state model has significantly higher conditional kurtosis than the 2-state model, due to the presence of the disaster-state. In terms of the conditional kurtosis after model uncertainty is integrated out (the 'full' learning model), model learning here contributes to a less overall drift in kurtosis over the sample relative to the individual models. Thus, among the models considered here, model uncertainty and its dynamic behavior have the strongest implications for assets such as out-of-the-money options that are more sensitive to the tail behavior of consumption growth, as given here by the negative skewness and high excess kurtosis of consumption growth.

With Bansal and Yaron (2004), long-run consumption risks have come into prominence in the asset pricing literature. Since shocks to parameter and model beliefs are permanent, such shocks constitute a source of 'truly' long-run risks. To illustrate this, we in Figure 6 plot shocks to long-run expected consumption growth in the 2-state model with *unknown* parameters, as well as shocks to long-run expected consumption growth in the 2-state model with *known* parameters. Long-run expected consumption growth is defined as $E_t\left[\sum_{j=1}^{\infty}\rho^j\Delta c_{t+j}\right]$, where ρ is set to 0.99 and where the expectation is taken with respect to the agent's beliefs about the state, parameters, and models at time t. From the figure it is clear that the long-run shocks to consumption growth are much more volatile – in fact, 3.4 times more volatile – in the case of unknown parameters relative to the case where there is only state learning. Such permanent shocks to beliefs will have a large impact on aggregate valuation ratios and thus help generated excess return volatility. Further, to the extent the preference

Figure 6 - Long-run consumption shocks

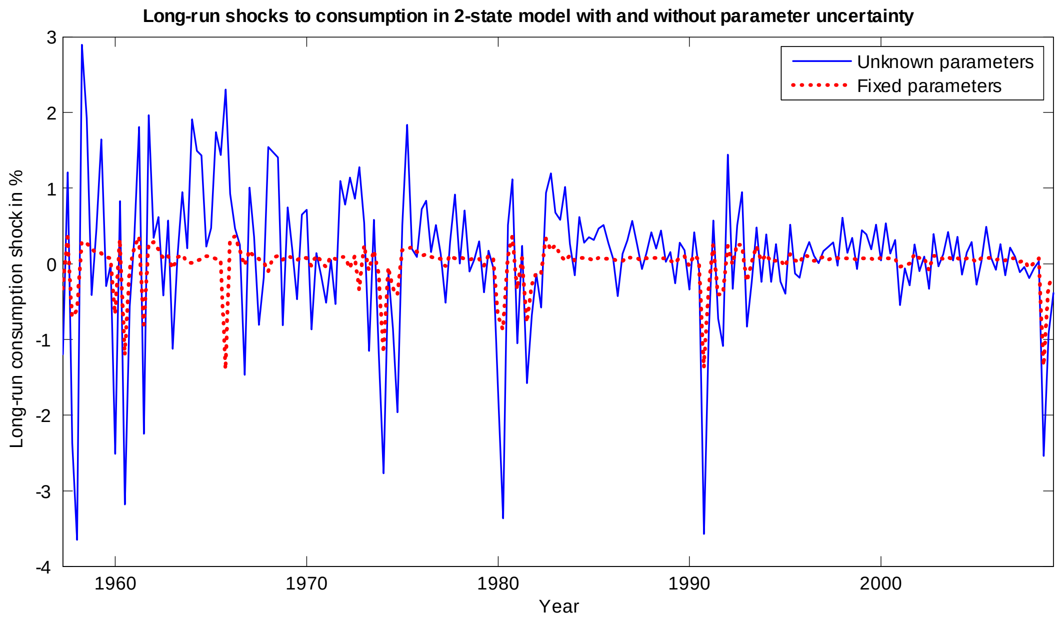

Figure 6: The figure shows shocks to long-run expected consumption growth for the 2-state model with fixed parameters as well as the case of unknown parameters. The state and parameter uncertainty are integrated out when forming expectations. Long-run expected consumption growth is calculated as the expected consumption growth from time $t+1$ to infinity, where period $t+j$'s quarterly consumption growth is discounted by 0.99^j. The sample period is 1947:Q2 - 2009:Q1.

specification prices such long-run risk, parameter uncertainty can be a significant additional source of macro risk (see Collin-Dufresne, Johannes, and Lochstoer (2013)). Finally, it is worth noting that the biggest negative shocks to long-run consumption are associated with recessions, indicating that recessions are a significant source of such risks.

4 Does learning matter for asset prices?

4.1 A new test for the importance of learning

The previous results indicate that the agent's beliefs – about parameters, moments, and models – vary substantially at both very low frequencies and over the business cycle. If learning is an important determinant of asset prices, changes in beliefs should be a significant determinant of asset returns. This is a fundamental test of the importance of learning about the consumption dynamics. For example, if agents learn that expected consumption growth is higher than previously thought, this revision in beliefs will be reflected in the aggregate wealth-consumption ratio (if the elasticity of intertemporal substitution is different from one). In particular, if the substitution effect dominates, the wealth-consumption ratio will increase when agents revise their beliefs about the expected consumption growth rate upwards (see, e.g., Bansal and Yaron (2004)). As another example, if agents learn that aggregate risk (consumption growth volatility) is lower than previously thought, this will generally lead to a change in asset prices as both the risk premium and the risk-free rate are affected. In the Bansal and Yaron (2004) model, an increase in the aggregate volatility leads to a decrease in the stock market's price-dividend ratio.

To test this, we regress excess quarterly stock market returns (obtained from Kenneth French's web site) on changes in beliefs about expected consumption growth and expected consumption growth variance. This is a particularly stringent test of learning, which to our knowledge has not been done in the previous literature. We use the beginning of period timing for the consumption data here and elsewhere in the paper.[25] The regressors are the shocks, $E_t\left(\Delta c_{t+1}\right) - E_{t-1}\left(\Delta c_{t+1}\right)$ and $\sigma_t\left(\Delta c_{t+1}\right) - \sigma_{t-1}\left(\Delta c_{t+1}\right)$. Notice that the only thing that is changing is the conditioning information set as we go from time $t-1$ to time t; the regressors are revisions in beliefs. We calculate these conditional moments integrating out state, model and parameter uncertainty. The first 10 years of the sample are used as a burn-in period to alleviate any prior misspecification.

We control for contemporaneous consumption growth and lagged consumption growth (the direct cash flow effect). By controlling for realized consumption growth, we ensure that the results are driven by model-based revisions in beliefs, and not just the fact that realized consumption growth (a direct cash flow effect) was, for example, unexpectedly high. To separate out the effects of parameter from state learning, we use revisions in expected consumption growth beliefs computed from the 3-state model with fixed parameters (set to their full-sample values) as an additional control.[26]

[25] Due to time-averaging (see Working, 1960), Campbell (1999) notes that one can use either beginning of period or end of period consumption in a given quarter as the consumption for that quarter. The beginning of period timing yields stronger results than using the end of period convention (although the signs are the same in the regressions). In principle, the results should be the same, so this is consistent with some information being impounded in stocks before the consumption data is revealed to the Bureau of Economic Analysis.

[26] Using the fixed parameter 2-state model as the control instead does not change the results.

[Table 3 about here]

Specifications 1 and 2 in Table 3 show that increases in expected conditional consumption growth are positively and strongly significantly associated with excess contemporaneous stock returns for both priors. This result holds controlling for contemporaneous and lagged consumption growth (the direct cash flow effect), and so we can conclude that revisions in beliefs are significantly related to shocks to the price-dividend ratio. This is a very strong result, pointing to the importance of a learning-based explanation for realized stock returns. These results could be driven by parameter or state learning.

Specification 3 shows that the updates in expected consumption growth derived from the model with fixed parameters (that is, a case with state learning only) are also significantly related to realized stock returns. The R^2, however, is lower than for the case of the full learning model, and when we include the revisions in beliefs about expected consumption growth from both the full learning model and the fixed parameters benchmark model in the regression (specification 4), the updates in expected consumption growth that arise in a model with fixed parameters are insignificant (and have the wrong sign), while the belief revisions from the full learning model remain significant. That is, updates in expectations when learning about about parameters, states, and models are more closely related to realized stock market returns than the corresponding updates in expectations based on a single model with known parameters but hidden states estimated on the full sample. To our knowledge, this is the first direct comparison of learning about models and parameters versus the traditional implementation of the rational expectation explanations in terms of explaining the time-series of realized stock returns using the actual sequence of realized macro shocks.

This result is driven by the nonlinear process of jointly learning about parameters and states. However, returns may be a nonlinear function of the updates in beliefs. To control for this possible effect, we in specification 5 also put in as a control the change in the price-dividend ratio as it appears in returns for the fixed parameter model (with preference parameters that are standard and will be discussed in the next section). Again, the updates in beliefs from the 'full' learning case with state, parameter, and model learning dominates that of the model with only state learning.

For the variance (regression specifications 6 and 7 in Table 3) we get the opposite result, as one would expect (at least with a high elasticity of intertemporal substitution, as we will use later in the paper): unexpected increases in conditional consumption growth variance are associated with negative contemporaneous stock returns. This result is not significant at the 5% level when including contemporaneous and lagged consumption growth in the regressions (specification 7). This does not mean there is no effect; we just cannot distinguish it from the direct cash flow effect when learning from consumption data alone.

To summarize, we find strong evidence that the updates in beliefs elicited from our model/prior combinations are associated with actual updates in agent beliefs at the time, as proxied by stock market returns. Again, it is important to recall that no asset price data was used to generate these belief revisions.

4.2 Learning from additional macro variables

Agents have access to more than just aggregate consumption growth data when forming beliefs. Here we provide one approach for incorporating this additional information and apply this methodology to learning from quarterly GDP growth, in addition to consumption. Suppose x_t represents the common growth factor in the economy and evolves via:

$$x_t = \mu_{s_t} + \sigma_{s_t}\varepsilon_t, \tag{3}$$

where $\varepsilon_t \overset{i.i.d.}{\sim} \mathcal{N}(0,1)$, and s_t is the state of the economy, which follows the same Markov chains specified earlier. Consumption growth Δc and J additional variables $Y_t = [y_t^1, y_t^2, ..., y_t^J]'$ are assumed to follow:

$$\Delta c_t = x_t + \sigma_c \varepsilon_t^c, \tag{4}$$

where

$$y_t^j = \alpha_j + \beta_j x_t + \sigma_j \varepsilon_t^j, \quad for \quad j = 1, 2, ..., J \tag{5}$$

and $\varepsilon_t^c \overset{i.i.d.}{\sim} \mathcal{N}(0,1)$, and $\varepsilon_t^j \overset{i.i.d.}{\sim} \mathcal{N}(0,1)$ for any j. Note that the coefficients in equation (5) are not state dependent, which implies that the additional variables will primarily aid in state identification. The specification allows for the additional observation variables to be stronger or weaker signals of the underlying state of the economy than consumption growth. For the case of GDP growth, this captures the idea that investment is more cyclical than consumption, which makes GDP growth a better business cycle indicator. The linearity of the relationship is an assumption that is needed for conjugate priors.

The similar conjugate priors for the parameters are applied. For each state $s_t = i$,

$p(\mu_i|\sigma_i^2)p(\sigma_i^2) \sim \mathcal{NIG}(a_i, A_i, b_i, B_i)$, where $\mathcal{NIG}$ is the normal/inverse gamma distribution. σ_c is assumed to follow an inverse gamma distribution $\mathcal{IG}(b_c, B_c)$, and for each $j = 1, 2, ..., J$, $p([\alpha_j, \beta_j]'|\sigma_j^2)p(\sigma_j^2) \sim \mathcal{NIG}(a_j, A_j, b_j, B_j)$, where $p([\alpha_j, \beta_j]'|\sigma_j^2)$ is a bivariate normal distribution $\mathcal{N}(a_j, A_j\sigma_j^2)$, a_j is a 2×1 vector and A_j is a 2×2 matrix. Particle filtering is straightforward to implement in this specification by modifying the algorithm described in the Appendix.

To analyze the implications of additional information, we consider learning using real, per capita U.S. GDP growth as an additional source of information. This exercise generates a battery of results: time series of parameter beliefs, conditional moments, and model probabilities. We report only a few interesting statistics in the interests of parsimony. Basically, adding GDP growth mainly aids in the state identification. In particular, adding GDP growth results in a greater difference in expected consumption growth across the states due. Figure 7 shows that the difference in the expected consumption growth rate in recessions versus expansions is about 0.6% per quarter, versus about 0.3% in the case of consumption information only (see Figure 4). The dynamic behavior of the conditional standard deviation of consumption growth is not significantly changed (not reported for brevity).

Figure 8 shows that the model specification results are similar, as the data again favors the 2-state model, leaving the 3-state model with a very low probability at the end of the sample. Overall, however, the 3-state model has a higher probability than earlier. \

Table 4 shows the regressions of contemporaneous stock returns and updates in agent beliefs about conditional expected consumption growth and consumption growth variance, as calculated from this extended model. The results are similar, but in fact overall *stronger* than the results using only consumption growth. Updates in agent expectations about these moments from the full learning model are significantly related to stock returns, also after controlling for contemporaneous and lagged consumption growth and updates in expected consumption growth derived from a model with fixed parameters. Again, this evidence indicates that learning about parameters and models is an important feature of the data.

[Table 4 about here]

Figure 7 - Conditional expected consumption growth (GDP)

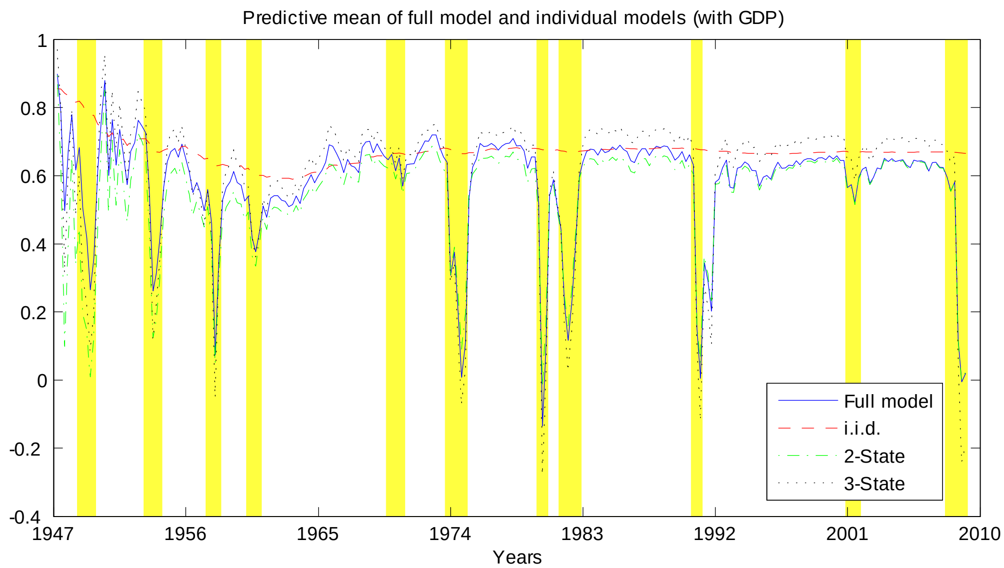

Figure 7: The top panel shows the quarterly conditional expected consumption growth, where state and parameter uncertainty have been integrated out, from each of the three benchmark models: the "i.i.d.", and the general 2- and 3-state switching regime models. The solid line shows the conditional expected consumption growth rate for the 'full' model, where also model uncertainty has been integrated out. In this case, GDP growth is used in addition to consumption growth in the agent's learning problem, as explained in the text. The sample period is 1947:Q2 - 2009:Q1.

Figure 8 - Model Probabilities (learning also from GDP data)

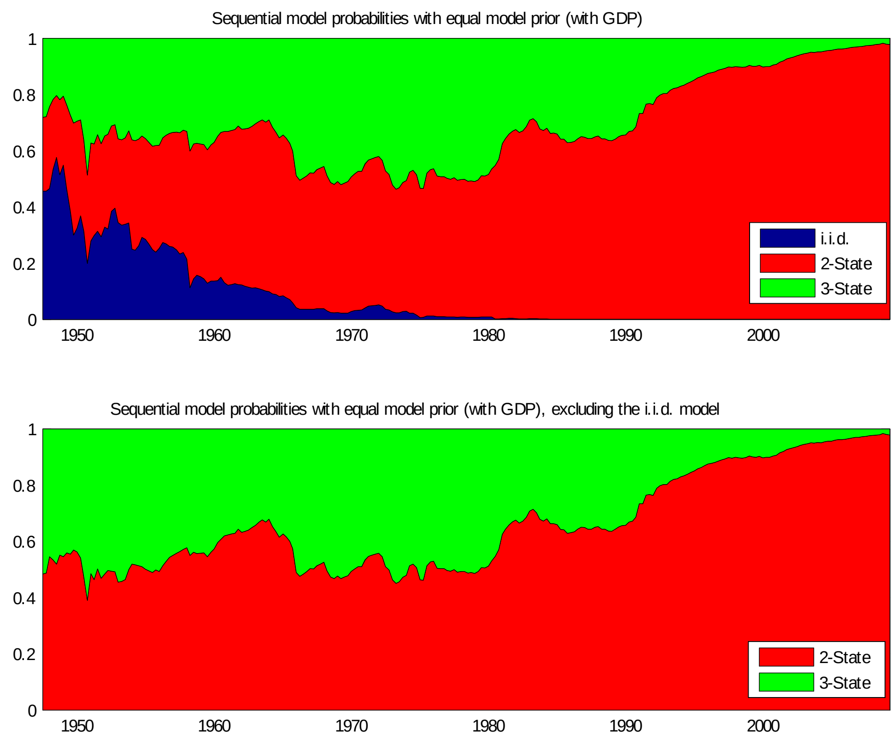

Figure 8: The top panel shows the evolution of the probability of each model being the true model, where the models at the beginning of the sample are set to have an equal probability, and where state and parameter uncertainty have been integrated out. The lower plot shows the same when the agent considers only the general 2-state model and the 3-state model as possible models of consumption dynamics. In this case, GDP growth is used in addition to consumption growth in the agent's learning problem, as explained in the text. The sample period is 1947:Q2 - 2009:Q1.

4.3 Asset pricing implications

We now embed the beliefs of our learning agent in a general equilibrium asset pricing model. There are considerable computational and technical issues that need to be dealt with when considering such an exercise. First, the state space is prohibitively large. The 3-state model, as an example, have 12 parameters governing the exogenous consumption process, and the beliefs over each parameter are governed by 2 hyper-parameters. In addition, there are the beliefs over the state of the economy and the corresponding parameter and state beliefs for the i.i.d. and the general 2-state models. Second, as pointed out by Geweke (2001) and Weitzmann (2007), some parameter distributions must be truncated in order for utility to be finite. This introduces additional nuisance parameters.

Given the computational impediments, we follow Sargent and Cogley (2008) and Piazzesi and Schneider (2010) and apply the principle of "anticipated utility" (originally suggested by Kreps (1998)) to our main pricing exercise. Under this assumption, the agents maximize utility at each point in time assuming that the parameters and model probabilities are equal to the agents' current mean beliefs and will remain constant forever. Of course, at time $t+1$ the mean parameter beliefs will in general be different due to learning. While parameter and model uncertainty are not *priced* risk factors in this framework, they are nonetheless important for the time-series of asset prices as updates in mean parameter and model beliefs lead to changes in prices. We do integrate out state uncertainty in the pricing exercise, so state uncertainty is a priced risk factor (as in, e.g., Lettau, Ludvigson, and Wachter (2008)). The anticipated utility approach reduces the number of state variables to three (the belief about the state in the general 2-state model, and the 2-dimensional belief about the state in the 3-state model).

The purpose of the pricing exercise is to examine what features of the post-WW2 U.S. aggregate consumption and asset price data a realistic, general learning problem can help explain. Since we do not integrate out the parameter and model uncertainty in the pricing exercise, we focus on two aspects of the model that we show later are robust to the introduction of *priced* parameter uncertainty.

1. *Ex-ante versus ex post*

   With parameter and model learning *ex ante* expectations can in general be quite far away from average *ex post* outcomes in samples like those we typically have available. This is due to the non-stationary aspect of parameter and model learning and different from the typical rational expectations model implementation. In the following, we argue that substantial components of the observed equity premium, excess return volatility, the degree of in-sample excess return predictability, and the time-series of the aggregate price-dividend ratio can be explained by the (nonstationary) time-path of mean parameter beliefs.

2. *The joint time-series path of beliefs and asset prices*

   The time-series of beliefs about consumption dynamics will have direct implications, under some mild additional assumptions, for the time-series of the aggregate price-dividend ratio. Since the shocks to mean parameter beliefs are *permanent shocks* to investor information sets, even a small shock to beliefs will be reflected in the price-dividend ratio with a large multiplier. We show in the following that the covariance between the aggregate price-dividend ratio in the data and that implied by the models with learning are about an order of magnitude larger than the covariance between the price-dividend ratio in the data and the price-dividend ratio generated by the models with fixed parameters.

4.3.1 The model

The model is solved at the quarterly frequency, and the representative agent is assumed to have Epstein and Zin (1989) preferences, which are defined recursively as:

$$U_t = \left\{(1-\beta)C_t^{1-1/\psi} + \beta\left(E_t\left[U_{t+1}^{1-\gamma}\right]\right)^{\frac{1-1/\psi}{1-\gamma}}\right\}^{\frac{1}{1-1/\psi}}, \tag{6}$$

where C_t is the consumption, $\psi \neq 1$ is the intertemporal elasticity of substitution (IES) in consumption, and $\gamma \neq 1$ is the coefficient of relative risk aversion. These preferences imply the stochastic discount factor:

$$M_{t+1} = \beta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\left(\beta\frac{PC_{t+1}+1}{PC_t}\right)^{\frac{1/\psi-\gamma}{1-1/\psi}}, \tag{7}$$

where PC_t is the wealth-consumption ratio – that is, the price-dividend ratio for the claim to the stream of future aggregate consumption. The first component of the pricing kernel is that which obtains under standard power utility, while the second component is present if the agent has a preference for the timing of the resolution of uncertainty (i.e., if $\gamma \neq 1/\psi$). As mentioned earlier, we consider an anticipated utility approach to the pricing problem in terms of parameter and model uncertainty, while state uncertainty is priced.[27] This corresponds to a world where investors understand and account for business cycle fluctuations, but where they simply use their best guess for the parameters governing these dynamics.

Our goal in this section is to, for reasonable preference parameters, understand how learning affects pricing relative to the benchmark case of fixed parameters. Given that the consumption dynamics are not *ex post* calibrated but estimated in real-time, we also do not calibrate preference parameters to match any particular moment(s). Instead, we simply use the preference parameters of Bansal and Yaron (2004). Thus, $\gamma = 10$, $\psi = 1.5$, and $\beta = 0.998\hat{}3$.

Following both Bansal and Yaron (2004) and Lettau, Ludvigson, and Wachter (2008), we price a levered claim to the consumption stream with a leverage factor λ of 4.5. The annual consumption volatility over the post-war sample is only 1.34%, and so the systematic annual dividend volatility is therefore about 6%. Quarterly log dividend growth is defined as:

$$\Delta d_t = \mu + \lambda \left(\Delta c_t - \mu \right) + \varepsilon_{d,t}, \tag{8}$$

where $\varepsilon_{d,t} \overset{i.i.d.}{\sim} \mathcal{N}\left(-\frac{1}{2}\sigma_d^2, \sigma_d^2\right)$ is the idiosyncratic component of dividend growth. σ_d is chosen to match the observed annual 11.5% volatility of dividend growth reported in Bansal and Yaron (2004). With these choices of λ and σ_d we also in fact closely match the sample correlation they report between annual consumption and dividend growth (0.55).[28]

Unconditional Moments

[Table 5 about here]

Table 5 reports realized asset pricing moments in the data, and also those generated by our learning models over the same sample period. The priors of each model are given in

[27] The model is solved numerically through value function iteration *at each time t* in the sample, conditional on the mean parameter beliefs at time t, which gives the time t asset prices. The state variables when solving this model are the beliefs about the hidden states of the economy for each model under consideration. For a detailed description of the model solution algorithm, please refer to the Online Appendix.

[28] The dividend dynamics imply that consumption and dividends are not cointegrated, which is a common assumption (e.g., Campbell and Cochrane (1999), and Bansal and Yaron (2004)). One could impose cointegration between consumption and dividends, but at the cost of an additional state variable. Further, it is possible to also learn about λ and σ_d^2. However, quarterly dividends are highly seasonal, which would severely complicate such an analysis. Further, data on stock repurchases is mainly annual. We leave a rigorous treatment of these issues to future research.

Table 1 and the sample period is 1957Q2 to 2009Q1 (the 10 years from 1947 through 1956 were used as a burn-in period for the priors). We consider cases with and without parameter learning.

The models with parameter uncertainty match the sample equity premium reasonably well: 4.7% in the data versus 3.7% for the full model with both parameter and model uncertainty and learning only from consumption data, and 5.2% for the full model with learning from both consumption and GDP data. The models where GDP is used as an additional signal has, as reported earlier, a more severe recession state, which is why the average return is higher in this model. The benchmark fixed parameters 2- and 3-state models have sample equity premiums 1.1% and 1.7%, respectively. Thus, allowing for parameter and model uncertainty more than triples the sample risk premiums, despite the fact that parameter and model uncertainty are not priced risk factors in the anticipated utility pricing framework.

The reason for this difference can be gleaned from the reported average *ex ante* equity risk premium ($E_T \left[E \left(R^{excess}_{m,t+1} | I_t \right) \right]$, where I_t denotes the information set (beliefs) of agents at time t and $E_T [\cdot]$ denotes the sample average). Since parameter and model uncertainty are not priced risks with anticipated utility, one may expect the average ex ante risk premiums to be similar to the corresponding model with known parameters. For the 2-state model with parameter uncertainty and learning from consumption only, however, the average ex ante sample risk premium is 1.8%, whereas the average ex ante sample risk premium for the 2-state model with known parameters is 0.9%. A similar difference is present between the 3-state models. The beginning-of-sample beliefs are, as discussed earlier, that expansions are less persistent and that recessions are more severe than the end-of-sample posterior beliefs. This leads to a higher ex ante risk premium early in the sample, At the end of sample, on the other hand, the conditional risk premium of the models with parameter uncertainty is lower and similar to that of the fixed parameters models.

These unexpected, overall positive surprises in belief updates not only decreases the ex ante risk premium, but also increases the price-dividend ratio over the sample. Thus, ex post average returns are much higher than the ex ante expected returns, which explains the remaining difference between the sample risk premium in the parameter and model learning case versus the fixed parameter cases. In particular, the sequence of shocks realized over the post-war sample generate a times series of beliefs that have a systematic time series pattern: the initial low mean and high volatility of consumption growth causes an upward revision in the mean growth rates and a negative revision in the volatility parameters, as described in

Section 3. Fama and French (2002) reach a similar conclusion in terms of the *ex post* versus the *ex ante* risk premium when looking at the time-series of the aggregate price-earnings and price-dividend ratios. Sargent and Cogley (2008) assume negatively biased beliefs in their model to highlight the same mechanism. The results we present here are consistent with their conclusions, but our models are estimated from fundamentals alone.

The equity return volatility is, in all the cases permitting parameter and model uncertainty about 15%, close to the 17.1% annual return volatility in the data. In contrast, the equity return volatility in the models with fixed parameters is about 12%, which is almost all cash flow volatility as the annual dividend growth volatility is 11.5%. Thus, the sample variation in discount and, especially, dividend growth rates arising from updates in agents' beliefs cause excess return volatility (Shiller, 1980). This is reflected in the sample volatility of the log price-dividend ratio, which is 0.38 in the data. In the cases with parameter and model uncertainty the volatility of the log price-dividend ratio is about 0.25, depending somewhat of the exact model specification.[29] While this is only about three quarters of its volatility in the data, it is 4 to 5 times the volatility of the log price-dividend ratio in the benchmark fixed parameters models (here the volatility of the log price-dividend ratio is 0.06 for the 2-state model and 0.07 for the 3-state model).

The sample correlation between the log price-dividend ratios from the model versus the data, is 0.67 for the full learning model using both GDP and consumption to estimate beliefs and 0.42 for the full learning model using consumption only to estimate beliefs. The models with fixed parameters have lower correlations, 0.25 for the 2-state model and 0.26 for the 3-state model. Thus, the covariance between the price-dividend ratio in the data and the full learning model using both consumption and GDP growth is 0.0354, more than an order of magnitude higher than the highest covariance between the price-dividend ratio in the data and the models with fixed parameters (0.0013). Thus, with parameter and model learning the model tracks the aggregate stock market price level (normalized by dividends) much more closely than either of the models we consider with fixed parameters, providing further evidence for the empirical relevance of parameter and model learning over the post-WW2 sample.

[29] The price-dividend ratio in each model is calculated as the corresponding in the data by summing the last four quarters of payouts to get annual payout. The price-dividend ratio from the data includes share repurchases in its definition of total dividends.

Anticipated utility versus priced parameter uncertainty. To assess the robustness of the asset pricing implications of learning that we document to fully rationally priced parameter uncertainty, we contrast the anticipated utility pricing results for the 2-state model with unknown parameters and anticipated utility pricing and learning from consumption only with the corresponding metrics for a 2-state model where all the parameters are uncertain but where the agent prices the estimation risk rationally. In order to have analytical expressions for the evolution of beliefs, which are necessary when solving such a model, we have to assume that the state in the Markov chain is observed. Otherwise, the priors are the same across the two models.

[Table 6 about here]

The details of the numerical solution for this exercise are given in the Online Appendix. We here just mention that the problem is extremely computationally intensive with 9(!) state variables. We solve the 2-state model with priced parameter uncertainty both with $\gamma = 10$, the same value as used for the anticipated utility models, and with $\gamma = 5$. As explained in Collin-Dufresne, Johannes, and Lochstoer (2013), updates in beliefs are a source of long-run consumption risks that carry a high price of risk when agents have a preference for early resolution of uncertainty, as is the case here.

Table 6 shows that the sample average excess returns for the equity claim in this model is 7.5% relative to 3.7% for the anticipated utility 2-state model with learning from consumption only. With $\gamma = 5$ the risk premium is as in the data (4.6%) . The average ex ante sample risk premiums are 5.6% and 2.5%, respectively, versus 1.8% for the anticipated utility case. Thus, also in this case a large fraction of the realized average equity returns are due to unanticipated positive shocks. However, the fraction is a little smaller than for the anticipated utility cases. This occurs since discount rates are higher with priced parameter uncertainty, which in turn means that the shocks to the growth rate affects the price-dividend ratio less.

Return volatility and the volatility of the price-dividend ratio are somewhat lower than in the anticipated utility case, due to the same higher ex ante discount rate effect. It is worth nothing, however, that the correlation between the price-dividend ratio in the data and that of the models are by far highest for the cases with priced parameter uncertainty, than that from the anticipated utility case or that with fixed parameters.

The time-series of asset prices. As a formal test of the learning model's match of the aggregate stock price level (the log D/P ratio) we regress the price-dividend ration in the data on the price-dividend ratio implied by the different models in levels and in changes.

Panel A of Table 7 shows the level regressions. The best 'fit' is achieved by the anticipated utility model with learning about both parameters and models where the agent learns from both consumption and GDP. This model has an R^2 of 46% and one cannot reject that the intercept is zero and that the slope is one. The price-dividend ratio from the 'fixed parameter' model is insignificant in a regression that also includes the learning model's pd ratio. The table also shows that even though the 2-state model quickly becomes the one with the highest model probability, the full model adds explanatory power relative to this model mainly as the 3-state model's depression state is important for asset prices even though it is quite unlikely. The increase in fit from the full learning models stems from a better match of the business cycle fluctuations in the dividend yield, as well as low-frequency fluctuations. In particular, with parameter learning the dividend yield displays a downward trend over the sample, similar to that found in the data as documented by, for instance, Fama and French (2002).

Panel B of Table 7 shows the same regressions in changes. Again, the models with anticipated utility fare the best and has a significant relation to the changes of the price-dividend ratio in the data. In sum, including parameter and model uncertainty leads to not only better fit of the unconditional asset pricing moments, but a significantly better fit of the realized aggregate stock price level in the post-WW2 era. Learning with both states and parameters uncertain matches the data best.

[Table 7 about here]

Permanent shocks and the volatility of long-run yields. With parameter and model uncertainty, the updates in mean beliefs constitute permanent shocks to expectations about consumption growth rates, consumption growth volatility, and higher order moments. This is a distinguishing feature of models with learning about constant quantities relative to learning about or observing a stationary underlying process (such as our state of the Markov chain, long-run risk in Bansal and Yaron (2004), or the surplus consumption ratio in Campbell and Cochrane (1999)). The latter models have transitory variables only in marginal utility growth. Shocks to a transitory state variable eventually die out, and so (very) long-run expectations are constant. Shocks to, for instance, the mean belief about the unconditional growth rate of consumption are, on the other hand, permanent, leading to permanent shocks to marginal utility growth. This has implications for all asset prices, but can be most clearly seen when considering the volatility of long-run default-free real yields, which can be readily calculated from our model. Table 8 shows the volatility of annualized yields for default-free real, zero-coupon bonds at different maturities. The data column gives the volatility of yields on U.S. TIPS, calculated from monthly data for the longest available sample, 2003 to 2011, from the Federal Reserve Board, along with the standard error of the volatility estimates. In the remaining columns, the corresponding model-implied yield volatilities, calculated from each of the models considered in this paper over the post-WW2 sample, are given.

First, the yield volatilities for the models with parameter and model uncertainty are substantially higher than the yield volatilities from the models with fixed parameters. The 2-year yields are twice as volatile, while the 10-year yields are an order of magnitude more volatile. This is a direct consequence of the permanent shocks to expectations resulting from parameter learning, whereas the models with fixed parameters have constant long-run consumption growth mean and volatility. Notably, the long maturity yields in the data have about the same yield volatility as in the models with parameter uncertainty, and so this is another dimension along which learning about parameters and models can help explain historical asset pricing behavior.

[Table 8 about here]

Conditional Moments Figure 9 shows the time-series of the conditional risk premium and Sharpe ratio for the 2-state model with anticipated utility pricing as well as the 2-state model with priced parameter uncertainty. The latter has a strongly counter-cyclical risk premium and Sharpe ratio, due to the higher parameter uncertainty in recessions. The anticipated utility model also has a counter-cyclial risk premium and Sharpe ratio, but the magnitude is much smaller as all the ex ante priced risk in this case is state uncertainty.

Lastly, we consider excess market return forecasting regression using the dividend yield as the predictive variable. These regressions have a long history in asset pricing and remain a feature of the data that asset pricing models typically aim to explain (e.g., Campbell

Figure 9 - Anticipated utility vs. priced parameter uncertainty: Conditional risk premium and Sharpe ratio in the 2-state model

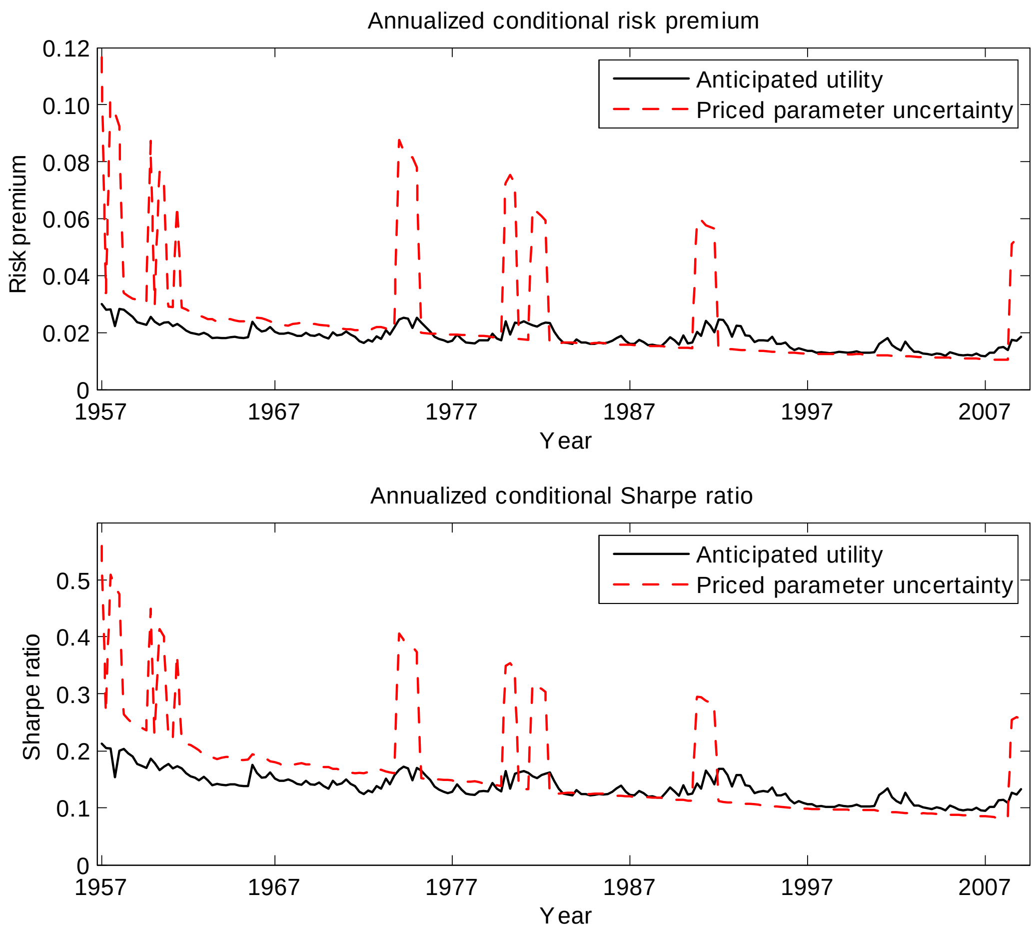

Figure 9: The top plot shows the annualized conditional risk premium on the dividend claim from the 2-state model with anticipated utility pricing (the black, solid line), as well as the case where the parameter uncertainty is fully rationally priced (the red, dashed line). In the former case, the risk aversion parameter $\gamma = 10$, whereas in the latter case $\gamma = 5$, such that the average risk premium over the sample is similar. The sample period is 1957:Q2 - 2009:Q1. The bottom plot shows the annualized conditional Sharpe ratio of returns to the dividend claim for the same two cases over the same sample period.

and Cochrane (1999), Bansal and Yaron (2004)). However, the strength of the empirical evidence is under debate (see, e.g., Stambaugh (1999), Ang and Bekaert (2007), Boudoukh, Richardson and Whitelaw (2008), and Goyal and Welch (2008) for critical analyses). Here we run standard forecasting regressions overlapping at the quarterly frequency using the sample of market returns and dividend yields as implied by each of the models. Note that, as before, we are not looking at population moments or average small-sample moments, but the single sample generated by feeding the models the actual sample of realized consumption growth.

Table 9 shows the forecasting regressions over different return forecasting horizons from the data. We use both the market dividend yield and the approximation to the consumption-wealth ratio, *cay*, of Lettau and Ludvigson (2001) to show the amount of predictability implied by these regressions in the data. We then run the same regressions using model implied returns and dividend yields. The benchmark models with fixed parameters (bottom right in the table) show no evidence of return predictability at the 5% significance level and the R^2's are very small. These models do, in fact, feature time-variation in the equity risk premium, but the standard deviation of the risk premiums are only about 0.5% per year and so the signal-to-noise ratio in these regressions is too small to result in significant predictability in a sample of the length we consider here. The models with parameter uncertainty, however, display significant in-sample return predictability and the regression coefficients and the R^2's are large and increasing in the forecasting horizon similar to those in the data. The *ex ante* predictability in these models is in fact similar to that in the fixed parameters cases, but since the parameters are updated at each point in time, there is significant *ex post* predictability. For instance, an increase in the mean parameters of consumption growth leads to high returns and lower dividend yield. Thus, a high dividend yield *in sample* forecasts high excess returns *in sample*. This is the same effect of learning as that pointed out in Timmermann (1993) and Lewellen and Shanken (2002). The models here show that the significant regression coefficients in the classical forecasting regressions show up in the sample only in the model where there is parameter learning which generates a significant difference between ex ante expected returns and ex post realizations. Thus, the model predicts that the amount of predictability is much smaller out-of-sample, consistent with the empirical evidence in Goyal and Welch (2008) and Ang and Bekaert (2007).

[Table 9 about here]

5 Conclusion

This paper studies the statistical problem and asset pricing implications of learning about parameters, states, and models in a standard class of models for consumption dynamics. Our approach is empirical, focuses on the specific implications generated by learning about U.S. consumption dynamics during the post World War II period, and contributes to a growing empirical literature documenting the importance of learning for asset prices (e.g., Malmendier and Nagel (2011), and Pastor and Veronesi (2003)).

We find broad support for the importance of learning about parameters and models. Agents' beliefs about consumption growth dynamics are strongly time-varying, nonstationary, and help explain the realized time-series of equity returns and price-dividend ratio. In particular, the new and significant relationship we document between contemporaneous realized returns and revisions in beliefs is strong support for the importance of learning. Incorporating learning and our estimated time-series of beliefs in a general equilibrium model uniformly improves the model fit with respect to the standard asset pricing moments.

Taken together, this evidence questions the typical implementations of rational expectations consumption-based exchange economy models, in which agents know with certainty the data generating process for aggregate consumption growth. Further, the nonstationary dynamics induced by learning about fixed quantities such as parameters and models translates to nonstationary dynamics in marginal utility growth and asset valuation ratios. This, in turn, implies that standard econometric approaches to model tests and parameter estimation should be used with caution (see also Cogley and Sargent (2008)).

The procedure implemented in this paper can in a straightforward way be implemented for other countries or markets, or extended to multi-country or multi-asset settings. For instance, learning about the joint dynamics of dividends and consumption is an interesting exercise abstracted away from in this paper. In terms of other countries, it is clear that the post World War II experience of Japan would lead to a very different path of beliefs. Learning about the joint dynamics of, say, the U.S. and Japan's economies would have interesting implications, not only for their respective equity markets, but also for the real exchange rate dynamics. It will in future research be interesting to consider priced parameter uncertainty with Epstein-Zin preferences. Parameter and model uncertainty will be major sources of anxiety for agents with preferences for early resolution of uncertainty as these risks are nonstationary and thus truly "long-run." As in Bansal and Yaron (2004), these sources of uncertainty will likely command high risk prices.

References

[1] Ai, H. (2010), "Information about Long-Run Risk: Asset Pricing Implications," forthcoming, *Journal of Finance.*

[2] Ang, A. and G. Bekaert (2007), "Return predictability: Is it there?", *Review of Financial Studies*, 20, 3, 651-707.

[3] Backus, D., M. Chernov and I. Martin (2009), "Disasters Implied by Equity Index Options", *NYU Working Paper*

[4] Bakshi, G. and G. Skoulakis (2010), "Do Subjective Expectations Explain Asset Pricing Puzzles?", *Journal of Financial Economics,* December 2010, 117 - 140.

[5] 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

[6] Bansal, R. and I. Shaliestovich (2010), "Confidence Risk and Asset Prices," *American Economic Review P&P*, 100, 537 – 541.

[7] Barberis, N. (2000), "Investing for the Long Run When Returns Are Predictable", *Journal of Finance* 55(1), 225 - 264

[8] Barillas, F., Hansen, L. and T. Sargent, "Doubts or Variability?", *Journal of Economic Theory*, forthcoming.

[9] Barro, R. (2006), "Rare Disasters and Asset Markets in the Twentieth Century", *Quarterly Journal of Economics* 121, 823 - 866

[10] Barro, R., Nakamura, E., Steinsson, J., and J. Ursua (2009), "Crises and Recoveries in an Empirical Model of Consumption Disasters," Columbia Business School working paper

[11] Barro, R. and J. Ursua (2008), "Consumption Disasters since 1870", *Brookings Papers on Economic Activity*

[12] Bhamra, H., L. Kuehn and I. Strebulaev (2009), "The Levered Equity Risk Premium and Credit Spreads: A Unified Framework", *Review of Financial Studies* (forthcoming)

[13] Boguth, O. and L. Kuehn (2009), "Consumption Volatility Risk," Working Paper Carnegie Mellon University

[14] Boudoukh, J., M. Richardson, and R. Whitelaw (2008), "The Myth of Long Horizon Predictability," *Review of Financial Studies* 21, 1577 - 1605.

[15] Brandt, M. , Q. Zeng, and L. Zhang (2004), "Equilibrium stock return dynamics under alternative rules of learning about hidden states," *Journal of Economic Dynamics and Control*, 28, 1925 – 1954.

[16] Brennan, M. and Y. Xia (2002), "Dynamic Asset Allocation Under Inflation," *Journal of Finance* 57, 1201 – 1238.

[17] Bray, M. M. and N. E. Savin (1986), "Rational Expectations Equilibria, Learning, and Model Specification," *Econometrica* 54, 1129 - 1160.

[18] Cagetti, M., L. Hansen, T. Sargent and N. Williams (2002), "Robustness and Pricing with Uncertain Growth", *Review of Financial Studies* 15, 363 - 404

[19] Carvalho, Johannes, Lopes, and Polson (2010a), "Particle Learning and Smoothing," *Statistical Science*, 25, 88-106.

[20] Carvalho, Johannes, Lopes, and Polson (2010b – forthcoming), "Particle learning: Simulation-based Bayesian inference," *Bayesian Statistics* 9.

[21] Cecchetti, S., P. Lam and N. Mark (1990), "Mean Reversion in Equilibrium Asset Prices", *American Economic Review* 80, 398 - 418

[22] Cecchetti, S., P. Lam and N. Mark (1993), "The Equity Premium and the Risk Free Rate: Matching the Moments", *Journal of Monetary Economics* 31, 21 - 46

[23] Chen, H. (2009), "Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure", *Journal of Finance* (forthcoming)

[24] Chen, H., S. Joslin, and N. Tran (2010), "Rare Disasters and Risk Sharing with Heterogeneous Beliefs," MIT Working paper

[25] Cogley, T. and T. Sargent (2008), "The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?", *Journal of Monetary Economics* 55, 454 - 476

[26] Cogley, T. and T. Sargent (2009), "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making", *International Economic Review* 49, 185 - 221

[27] Collin-Dufresne, P., M. Johannes, and L. Lochstoer (2013), "Parameter learning in general equilibrium: the asset pricing implications," Columbia Business School working paper.

[28] David, A. and P. Veronesi (2009), "What ties return volatilities to price valuations and fundamentals?" University of Calgary Working Paper.

[29] Detemple, J. (1986), "Asset pricing in a production economy with incomplete information." *Journal of Finance*, 41, 383–390.

[30] Dothan, M. U. and D. Feldman (1986), "Equilibrium interest rates and multiperiod bonds in a partially observable economy." *Journal of Finance*, 41, 369 – 382.

[31] Drechsler, I. and A. Yaron (2008), "What's Vol Got to Do with It?", Working Paper, Wharton School of Business, University of Pennsylvania

[32] Epstein, L. and S. Zin (1989), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework", *Econometrica* 57, 937 - 969

[33] Fama, E. and K. French (1988), "Dividend yields and expected stock returns," *Journal of Financial Economics* 22, 3 - 25.

[34] Fama, E. and K. French (2002), "The Equity Premium," *Journal of Finance* 57, 637 – 659.

[35] Froot, K. and S. Posner (2002), "The Pricing of Event Risks with Parameter Uncertainty", *GENEVA Papers on Risk and Insurance - Theory* 27, 153 - 165

[36] Hall, R. (1978), "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, December 1978, 86(6), pp. 971-987.

[37] Gabaix, X. (2009), "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance", *NYU Working Paper*

[38] Gennotte, G. (1986), "Optimal Portfolio Choice under Incomplete Information," *Journal of Finance*, 61, 733-749.

[39] Geweke, J. (2001), "A note on some limitations of CRRA utility," *Economics Letters* 71, 341 - 345.

[40] Goyal A. and I. Welch, "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction," July 2008, *Review of Financial Studies* 21(4) 1455-1508.

[41] Hansen, L. (2007), "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk," Richard T. Ely Lecture, *The American Economic Review* 97, No. 2., 1 - 30

[42] Hansen, L. and T. Sargent (2009), "Robustness, Estimation and Detection," Working paper.

[43] Hansen, L. and T. Sargent (2010), "Fragile Beliefs and the Price of Uncertainty," *Quantitative Economics*, Vol. 1, Issue 1, pp. 129-162

[44] Johannes, M. and N. Polson (2006), "Particle Filtering," *Springer Verlag Handbook of Financial Time Series*, edited by Torben G. Andersen Richard A. Davis, Jens-Peter Kreiss, and Thomas Mikosch, September 2006.

[45] Kandel, S. and R. Stambaugh (1990), "Expectations and Volatility of Consumption and Asset Returns", *Review of Financial Studies* 2, 207 - 232

[46] Kandel, S. and R. Stambaugh (1996), "On the Predictability of Stock Returns: An Asset-Allocation Perspective", *Journal of Finance* 51, 66 - 74

[47] Kreps, D. (1998), "Anticipated Utility and Dynamic Choice", *Frontiers of Research in Economic Theory*, (Cambridge: Cambridge University Press) 242 - 274

[48] Lettau, M. and S. Ludvigson (2001), "Consumption, Aggregate Wealth and Stock Returns," *Journal of Finance* 56, 815 - 849

[49] Lettau, M., S. Ludvigson and J. Wachter (2008), "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?", *Review of Financial Studies* 21(4), 1653 - 1687

[50] Lewellen, J. and J. Shanken (2002), "Learning, Asset-Pricing Tests, and Market Efficiency," *Journal of Finance* (57(3), 1113 - 1145.

[51] Lucas, R. (1978), "Asset Prices in an Exchange Economy," *Econometrica* 46, 1429-1446

[52] Lucas, R. and T. Sargent (1979), "After Keynesian Macroeconomics," The Federal Reserve Bank of Minneapolis, *Quarterly Review* 321.

[53] Malmendier, U. and S. Nagel (2011), "Depression babies: Do Macroeconomic Experience Affect Risk-Taking?", *Quarterly Journal of Economics* 126, 373 – 416.

[54] Mehra, R. and E. Prescott (1985), "The Equity Premium: A Puzzle", *Journal of Monetary Economics* 15, 145 - 161

[55] Moore, B. and H. Schaller (1996), "Learning, Regime Switches, and Equilibrium Asset Pricing Dynamics", *Journal of Economic Dynamics and Control* 20, 979 - 1006

[56] Pastor, L. (2000), "Portfolio Selection and Asset Pricing Models," *Journal of Finance* 55, 179 - 223

[57] Pastor, L. and P. Veronesi (2003), "Stock valuation and learning about profitability," *Journal of Finance*, 58, 1749 – 1789.

[58] Pastor, L. and P. Veronesi (2006), "Was there a NASDAQ bubble in the late 1990s?," *Journal of Financial Economics*, 81, 61 – 100.

[59] Pastor, L. and P. Veronesi (2009), "Learning in Financial Markets," *Annual Review of Financial Economics.*

[60] Piazzesi, M. and M. Schneider (2010), "Trend and Cycle in Bond Risk Premia," Working Paper Stanford University

[61] Rietz, T. (1988), "The equity risk premium: A solution?", *Journal of Monetary Economics*, Volume 22, Issue 1, July 1988, Pages 117-131

[62] Romer, C. (1989), "The Prewar Business Cycle Reconsidered: New Estimates of Gross National Product, 1869-1908", *Journal of Political Economy* 97, 1 - 37

[63] Shaliastovich, I. (2008), "Learning, Confidence and Option Prices," Working Paper, Duke University

[64] Stambaugh, R. (1999), "Predictive Regressions." *Journal of Financial Economics,* 1999, 54, pp. 375–421.

[65] Timmermann, A. (1993), "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices." *Quarterly Journal of Economics*, 1993, 108, 1135-1145.

[66] Veronesi, P. (1999), "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model," *Review of Financial Studies*, 12, 5, Winter 1999

[67] Veronesi, P. (2000), "How does information quality affect stock returns?" *Journal of Finance*, 55, .

[68] Weitzman, M. (2007), "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, 97(4), 1102-1130.

[69] Whitelaw, R. (2000), "Stock Market Risk and Return: An Equilibrium Approach", *Review of Financial Studies* 13, 521 - 547.

[70] Working, H. (1960), "Note on the Correlation of First Differences of Averages in a Random Chain," *Econometrica* 28(4), 916 - 918.

[71] Xia, Y. (2001), "Learning about Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation", *Journal of Finance* 56(1), 205 - 246

6 Appendix

6.1 Existing literature and alternative approaches for parameter, state, and model uncertainty.

Our paper is related to a large literature studying the asset pricing implications of parameter or state learning. Most of this literature focuses on learning about a single unknown parameter or state variable (assuming the other parameters and/or states are known) that determines dividend dynamics and power utility. For example, Timmerman (1993) considers the effect of uncertainty on the average level of dividend growth, assuming other parameters are known, and shows in simple discounted cash-flow setting that parameter learning generates excess volatility and patterns consistent with the predictability evidence (see also Timmerman 1996). Lewellen and Shanken (2002) study the impact of learning about mean cash-flow parameters with exponential utility with a particular focus on return predictability.

Veronesi (2000) considers the case of learning about mean-dividend growth rates in a model with underlying dividend dynamics with power utility and focuses on the role of signal precision or information quality. Pastor and Veronesi (2003, 2006) study uncertainty about a fixed dividend-growth rate or profitability levels with an exogenously specified pricing kernel, in part motivated in order to derive cross-sectional implications. Weitzman (2007) and Bakshi and Skoulakis (2009) consider uncertainty over volatility.

Cogley and Sargent (2008) consider a 2-state Markov-switching model, parameter uncertainty over one of the transition probabilities, tilt beliefs to generate robustness via pessimistic beliefs, and use power utility. After calibrating the priors to the 1930s experience, they simulate data from a true model calibrated to the post War experience to show how priced parameter uncertainty and concerns for robustness impact asset prices, in terms of the finite sample distribution over various moments.

A number of papers consider state uncertainty, where the state evolves discretely via a Markov switching model or smoothing via a Gaussian process. Moore and Shaller (1996) consider consumption/dividend based Markov switching models with state learning and power utility. Brennen and Xia (2001) consider the problem of learning about dividend growth which is not a fixed parameter but a mean-reverting stochastic process, with power utility. Veronesi (2004) studies the implications of learning about a peso state in a Markov switching model with power utility. David and Veronesi (2010) consider a Markov switching model with learning about states.

In the case of Epstein-Zin utility, Brandt, Zeng, and Zhang (2004) consider alternative rules for learning about an unknown Markov state, assuming all parameters and the model is known. Lettau, Ludvigson, and Wachter (2008) consider information structures where the economic agents observe the parameters but learn about states in Markov switching consumption based asset pricing model. Chen and Pakos (2008) consider learning about the mean of consumption growth which is a Markov switching process. Ai (2010) studies learning in a production-based long-run risks model with Kalman learning about a persistent latent state variable. Bansal and Shaliastovich (2008) and Shaliastovich (2010) consider learning about the persistent component in a Bansal and Yaron (2004) style model with sub-optimal Kalman learning.

Additionally, some papers consider combinations of parameter or model uncertainty and robustness, see, e.g., Hansen and Sargent (2000,2009) and Hansen (2008).

6.2 Econometrics

This section briefly reviews the mechanics of sequential Bayesian learning and introduces the econometric methods needed to solve the high-dimensional learning problem. For ease of exposition, we abstract here from the problem of model uncertainty and drop the dependence on the model specification. Model uncertainty can be dealt with easily in a fashion analogous to the problem considered here.

The agent begins with initial beliefs over the parameters and states, $p(\theta, s_t) = p(s_t|\theta)\, p(\theta)$, and then updates via Bayes' rule. If at time t the agent holds beliefs $p(\theta, s_t|y^t)$, then updating occurs in a two step process by first computing the predictive distribution, $p(\theta, s_{t+1}|y^t)$, and then updating via the likelihood function, $p(y_{t+1}|s_{t+1}, \theta)$:

$$p\left(\theta, s_{t+1}|y^{t+1}\right) \propto p\left(y_{t+1}|\theta, s_{t+1}\right) p\left(\theta, s_{t+1}|y^t\right).$$

The predictive distribution is

$$p\left(\theta, s_{t+1}|y^t\right) = \int p\left(s_{t+1}|s_t, \theta\right) p\left(\theta, s_t|y^t\right) ds_t,$$

which shows the recursive nature of Bayesian updating, as $p(\theta, s_{t+1}|y^{t+1})$ is functionally dependent on $p(\theta, s_t|y^t)$.

The main difficulty is characterizing $p(\theta, s_t|y^t)$ for each t, which is needed for sequential learning. Unfortunately, even though s_t is discretely valued, there is no analytical form for $p(\theta, s_t|y^t)$, as it is high-dimensional and the dependence on the data is complicated and nonlinear. We use Monte Carlo methods called particle filters to generate approximate samples from $p(\theta, s_t|y^t)$. Johannes and Polson (2008) developed the general approach we use, and it was extended and applied to Markov switching models by Carvalho, Johannes, Lopes, and Polson (2010a, 2010b) and Carvalho, Lopes and Polson (2009). Details of the algorithms are given in those papers.

The first step of the approach, data augmentation, introduces a conditional sufficient statistics, T_t, for the parameters. Sufficient statistics imply that the full posterior distribution of the parameters conditional on entire history of latent states and data takes a known functional form conditional on a vector of sufficient statistics: $p(\theta|s^t, y^t) = p(\theta|T_t)$, where $p(\theta|T_t)$ is a known distribution. The conditional sufficient statistics are given by $T_{t+1} = \mathcal{T}(T_t, s_{t+1}, y_{t+1})$, where the function $\mathcal{T}(\cdot)$ is analytically known, which implies the sufficient statistics can be recursively updated. For Markov switching models, the sufficient statistics contain random variables such as the number of times and duration of each state visit, the mean and variance of y_t in those visits, etc. This step requires conjugate priors.

The key is that it is easier to sample from $p(\theta, T_t, s_t|y^t)$ than $p(\theta, s_t|y^t)$, where

$$p\left(\theta, T_t, s_t|y^t\right) = p\left(\theta|T_t\right) p\left(T_t, s_t|y^t\right). \tag{9}$$

By the definition of sufficient statistics and the use of conjugate priors, $p(\theta|T_t)$ is a known distribution (e.g., normal). This transforms the problem of sequential learning of parameters and states into one of sequential learning of states and sufficient statistics, and then standard updating by drawing from $p(\theta|T_t)$. The dimensionality of the target distribution, $p(\theta, T_t, s_t|y^t)$, is fixed as the sample size increases.

An $N-$ particle approximation, $p^N(\theta, T_t, s_t|y^t)$, approximates $p(\theta, T_t, s_t|y^t)$ via 'particles' $\left\{(\theta, T_t, s_t)^{(i)}\right\}_{i=1}^{N}$ so that:

$$p^N(\theta, T_t, s_t|y^t) = \frac{1}{N}\sum_{i=1}^{N}\delta_{(\theta,T_t,s_t)^{(i)}},$$

where δ is a Dirac mass. A particle filtering algorithm merely consists of a recursive algorithm for generating new particles, $(\theta, T_{t+1}, s_{t+1})^{(i)}$, given existing particles and a new observation, y_{t+1}. The approach developed in Johannes and Polson (2008) and Carvalho, Johannes, Lopes, and Polson (2009a, 2009b) generates a direct or exact sample from $p^N(\theta, T_t, s_t|y^t)$, without resorting to importance sampling or other approximate methods. The algorithm is straightforward to code and runs extremely quickly so that it is possible to run for large values N, which is required to keep the Monte Carlo error low. These draws can be used to estimate parameters and states variables.

In addition to sequential parameter estimation, particle filters can also be used for Bayesian model comparison. Bayesian model comparison and hypothesis testing utilizes the Bayes factor, essentially a likelihood ratio between competing specifications. Formally, given a number of competing model specifications, generically labeled as model $\mathcal{M}_k$ and $\mathcal{M}_j$, the Bayesian approach computes the probability of model k as:

$$p\left(\mathcal{M}_k|y^t\right) = \frac{p(y^t|\mathcal{M}_k)p\left(\mathcal{M}_k\right)}{\sum_{j=1}^{N} p(y^t|\mathcal{M}_j)p\left(\mathcal{M}_j\right)},$$

where $p(\mathcal{M}_k)$ is the prior probability of model k,

$$p(y^{t+1}|\mathcal{M}_k) = p(y_{t+1}|y^t, \mathcal{M}_k)p\left(y^{t-1}|\mathcal{M}_k\right),$$

and

$$p(y_{t+1}|y^t, \mathcal{M}_i) = \int p\left(y_{t+1}|\theta, s_t, \mathcal{M}_i\right) p\left(\theta, s_t|y^t, \mathcal{M}_i\right) d\left(\theta, s_t\right)$$

is the marginal likelihood of observation y_{t+1}, given data up to time t in model k. Marginal likelihoods are not known analytically and are difficult to compute even using MCMC methods. Since our algorithm provides approximate samples from $p(s_t, \theta|y^t)$, it is straightforward to estimate marginal likelihoods via

$$p^N(y_{t+1}|y^t, \mathcal{M}_k) = \frac{1}{N}\sum_{i=1}^{N} p\left(y_{t+1}|\left(\theta, s_t\right)^{(i)}, \mathcal{M}_k\right).$$

For all of our empirical results, we ran particle filtering algorithms with $N = 10K$ particles. We performed extensive simulations to insure that this number of particles insured a low and negligible Monte Carlo error.

Table 1 - Priors (1957Q1) and end-of-sample posteriors (2009Q1)

Table 1: The table shows the 1957Q1 priors for the parameters of the three different models of log, real per capita, quarterly consumption gorwth considered in the paper, as well as the end-of-sample posteriors (as of 2009Q1). The parameters within a state (mean and variance) have Normal/Inverse Gamma distributed priors, while the transition probabilities have Beta distributed priors. Note that $\hat{\pi}_{ij} \equiv \frac{\pi_{ij}}{1-\pi_{ij}}$.

Panel A: Priors and end-of-sample posteriors for i.i.d. model

Parameter	μ_1	μ_2	σ_1^2	σ_2^2	π_{11}
Prior mean	0.60%	-1.56%	0.59 $(\%)^2$	2.78 $(\%)^2$	5.51%
Prior st.dev.	0.11%	0.73%	0.14 $(\%)^2$	1.31 $(\%)^2$	3.80%
Posterior mean	0.61%	-0.96%	0.25 $(\%)^2$	2.39 $(\%)^2$	3.79%
$\frac{\text{Posterior st.dev}}{\text{Prior st.dev}}$	31%	73%	21%	73%	49%

Panel B: Priors and end-of-sample posteriors for 2−state model

Parameter	μ_1	μ_2	σ_1^2	σ_2^2	π_{11}	π_{22}
Prior mean	0.81%	-0.03%	0.40 $(\%)^2$	0.84 $(\%)^2$	0.88	0.83
Prior st.dev.	0.18%	0.28%	0.16 $(\%)^2$	0.23 $(\%)^2$	0.07	0.09
Posterior mean	0.70%	0.13%	0.15 $(\%)^2$	0.63 $(\%)^2$	0.93	0.83
$\frac{\text{Posterior st.dev}}{\text{Prior st.dev}}$	18%	41%	11%	51%	34%	70%

Panel C: Priors and end-of-sample posteriors for 3−state model

Parameter	μ_1	μ_2	μ_3	σ_1^2	σ_2^2	σ_3^2	π_{11}	$\hat{\pi}_{12}$	$\hat{\pi}_{21}$	π_{22}	$\hat{\pi}_{31}$	π_{33}
Prior mean	0.74%	-0.23%	-1.84%	0.52 $(\%)^2$	0.55 $(\%)^2$	0.56 $(\%)^2$	0.92	0.86	0.86	0.75	0.35	0.65
Prior st.dev.	0.17%	0.29%	0.47%	0.15 $(\%)^2$	0.21 $(\%)^2$	0.28 $(\%)^2$	0.05	0.13	0.01	0.08	0.24	0.10
Posterior mean	0.72%	0.01%	-1.77%	0.18 $(\%)^2$	0.45 $(\%)^2$	0.56 $(\%)^2$	0.94	0.93	0.92	0.77	0.34	0.65
$\frac{\text{Posterior st.dev}}{\text{Prior st.dev}}$	22%	42%	98%	14%	46%	95%	40%	54%	60%	85%	97%	99%

Table 2 - The speed of learning:
A Monte-Carlo experiment of observed vs. unobserved states

Table 2: The table shows the results of the following Monte-Carlo experiment. Assume a 2-state Markov switching regime model, like that considered in the body of the paper, with true parameters as estimated by MCMC over the post-WW2 sample. These true parameters are reproduced in the row 'True values.' Next, we simulate 500 economies of 209 quarters from this model, assuming the first state is drawn randomly according to the unconditional probability of each state. Finally, using the particle filter, we run sequentially through each sample with unbiased prior means and with prior variances as used for the 2-state model and given in Table 1, Panel B. For each sample, we run the particle filter assuming either that the current state is observed or, like in the actual main empirical exercise, that the state is also unobserved. Thus, the latter problem embodies the joint problem of learning about both states and parameters, whereas the former only has parameter learning. Finally, we report the average end-of-sample posterior variances for each parameter for the case of known states as well as unknown states to investigate whether the joint learning about both states and parameters confounds inference and slows down parameter learning.

Parameter	μ_1	σ_1^2	μ_2	σ_2^2	π_{11}	π_{22}
True values	0.68%	$0.13\,(\%)^2$	0.21%	$0.49\,(\%)^2$	0.95	0.83
Posterior mean						
Unknown states	0.69%	$0.12\,(\%)^2$	0.22%	$0.45\,(\%)^2$	0.93	0.80
Known states	0.68%	$0.13\,(\%)^2$	0.21%	$0.47\,(\%)^2$	0.95	0.82
Posterior variance						
Unknown states	$0.0012\,(\%)^2$	$0.0004\,(\%)^2$	$0.0238\,(\%)^2$	$0.0146\,(\%)^2$	$0.0012\,(\%)^2$	$0.0050\,(\%)^2$
Known states	$0.0008\,(\%)^2$	$0.0002\,(\%)^2$	$0.0114\,(\%)^2$	$0.0109\,(\%)^2$	$0.0003\,(\%)^2$	$0.0024\,(\%)^2$
Reduction of posterior variance when states are known relative to unknown:	33%	46%	52%	25%	77%	52%

Table 3 - Updates in Beliefs versus Realized Stock Returns

Table 3: The table shows the results from regressions of innovations in agents' expectations of future consumption growth ($E_{t+1}[\Delta c_{t+2}] - E_t[\Delta c_{t+2}]$) and conditional consumption growth variance ($\sigma^2_{t+1}[\Delta c_{t+2}] - \sigma^2_t[\Delta c_{t+2}]$) versus contemporaneous excess stock market returns. Expectations integrate out parameter, state and model uncertainty, unless otherwise noted. The controls are lagged and contemporaneous realized log consumption growth, as well as the innovation in expected consumption growth derived from the 3-state model with fixed parameters (i.e., no model or parameter uncertainty), as well as the i.i.d. model with uncertain parameters. Heteroskedasticity and autocorrelation adjusted (Newey-West; 3 lags) standard errors are reported in paranthesis. * denotes significance at the 10% level, ** denotes significance at the 5% level, and *** denotes significance at the 1% level. The sample is from 1947:Q2 until 2009:Q1. In the below regressions, we have removed the first 40 observations (10 years), as a burn-in period to alleviate misspecification of the priors.

Dependent variable: $r_{m,t+1} - r_{f,t+1}$ (log excess market returns)

	1	2	3	4	5	6	7
$E_{t+1}[\Delta c_{t+2}] - E_t[\Delta c_{t+2}]$	40.61***	32.42**		55.49***	37.27***		
	(8.75)	(12.28)		(17.00)	(10.51)		
$\sigma^2_{t+1}[\Delta c_{t+2}] - \sigma^2_t[\Delta c_{t+2}]$						−40.73***	−18.30
						(11.13)	(11.55)
Controls:							
Δc_{t+1}		0.92					3.24**
		(1.71)					(1.42)
Δc_t		2.60*					2.17
		(1.43)					(1.44)
$[E_{t+1}[\Delta c_{t+2}] - E_t[\Delta c_{t+2}]]^{\text{3-state model}}_{\theta\text{ known}}$			23.75***	−12.63			
			(7.59)	(10.66)			
$\left[\ln\left(\frac{P_{t+1}/D_{t+1}+1}{P_t/D_t}\right)\right]^{\text{3-state model}}_{\theta\text{ known}}$					8.44		
					(11.03)		
R^2_{adj}	10.0%	11.7%	5.9%	10.0%	9.8%	6.3%	9.9%

Table 4 - Updates in Beliefs versus Realized Stock Returns (GDP)

Table 4: The table shows the results from regressions of innovations in agents' expectations of future consumption growth ($E_{t+1}[\Delta c_{t+2}] - E_t[\Delta c_{t+2}]$) and conditional consumption growth variance ($\sigma^2_{t+1}[\Delta c_{t+2}] - \sigma^2_t[\Delta c_{t+2}]$) versus contemporaneous excess stock market returns. Expectations integrate out parameter, state and model uncertainty, unless otherwise noted. The controls are lagged and contemporaneous realized log consumption growth, as well as the innovation in expected consumption growth derived from the 3-state model with fixed parameters (i.e., no model or parameter uncertainty). Both consumption and GDP data is used to estimate the models, as described in the main text. Heteroskedasticity and autocorrelation adjusted (Newey-West; 3 lags) standard errors are reported in paranthesis. * denotes significance at the 10% level, ** denotes significance at the 5% level, and *** denotes significance at the 1% level. The sample is from 1947:Q2 until 2009:Q1. In the below regressions, we have removed the first 40 observations (10 years), as a burn-in period to alleviate misspecification of the priors.

Dependent variable: $r_{m,t+1} - r_{f,t+1}$ (log excess market returns)

	1	2	3	4	5	6
$E_{t+1}[\Delta c_{t+2}] - E_t[\Delta c_{t+2}]$	31.16***	28.42***	34.15***	29.43***		
	(6.15)	(7.97)	(8.34)	(6.85)		
$\sigma^2_{t+1}[\Delta c_{t+2}] - \sigma^2_t[\Delta c_{t+2}]$					−48.39***	−38.66***
					(8.72)	(10.84)
Controls:						
Δc_{t+1}		0.08				1.23
		(1.51)				(1.52)
Δc_t		2.09				2.58*
		(1.32)				(1.38)
$[E_{t+1}[\Delta c_{t+2}] - E_t[\Delta c_{t+2}]]^{\text{3-state model}}_{\theta \text{ known}}$			−4.75			
			(6.79)			
$\left[\ln\left(\frac{P_{t+1}/D_{t+1}+1}{P_t/D_t}\right)\right]^{\text{3-state model}}_{\theta \text{ known}}$				7.90		
				(9.94)		
R^2_{adj}	14.2%	14.7%	13.9%	14.0%	10.9%	13.0%

Table 5 - Asset Price Moments

Table 5: The table reports the asset pricing implications of the models with an anticipated utility version of the Epstein-Zin preferences under different priors, as well as the fixed parameters cases. For all the models, $\gamma = 10$, $\beta = 0.994$, $\psi = 1.5$, $\lambda = 4.5$. The volatility of the idiosyncratic component of dividend growth ($\epsilon_{d,t}$) is calibrated to match the historical standard deviation of dividend growth, as reported in Bansal and Yaron (2004). The statistics are annualized. The expectation operator with a T subscript, E_T, denotes the sample average, while the volatility operator, σ_T denotes the sample standard deviation. The sample period is from 1957:Q2 until 2009:Q1, with 1957Q1 priors as given in Table 1 for the case of learning from observing consumption growth only, and as given in an Online Appendix for the case of learning from both consumption and GDP growth.

	Data	Learning from consumption				Learning from cons. & GDP				Known parameters	
Moments	1957:Q2-2009:Q1	Full model	i.i.d. model	2-state model	3-state model	Full model	i.i.d. model	2-state model	3-state model	2-state model	3-state model
The real risk-free rate:											
$E_T(r_t^f)$	1.6%	3.7%	3.7%	3.7%	3.5%	3.5%	3.8%	3.5%	3.4%	3.8%	3.7%
$\sigma_T(r_t^f)$	1.6%	0.7%	0.3%	0.8%	1.0%	1.1%	0.2%	0.9%	1.4%	0.6%	0.9%
The dividend claim: $\Delta d_t = \mu + \lambda\left(\Delta c_t - \mu\right) + \varepsilon_{d,t}$											
ex post:											
$E_T(r_t - r_t^f)$	4.7%	3.7%	2.9%	3.7%	6.0%	5.2%	2.6%	5.2%	6.7%	1.1%	1.7%
$\sigma_T(r_t - r_t^f)$	17.1%	15.1%	14.5%	15.2%	14.1%	14.9%	14.9%	14.2%	15.3%	11.9%	12.2%
Sharpe ratio	0.27	0.25	0.20	0.24	0.43	0.35	0.17	0.37	0.44	0.09	0.14
$\sigma_T(pd_t)$	0.38	0.22	0.24	0.23	0.21	0.23	0.31	0.25	0.15	0.06	0.07
$Corr_T(pd_t^{Model}, pd_t^{Data})$	n/a	0.42	0.28	0.39	0.50	0.67	0.35	0.43	0.62	0.25	0.26
ex ante:											
$E_T[E_t(r_{t+1} - r_{t+1}^f)]$	n/a	1.8%	1.4%	1.8%	3.2%	3.2%	1.4%	2.4%	5.2%	0.9%	1.5%

Table 6 - Priced parameter uncertainty vs. Anticipated utility:

The asset price moments

Table 6: The table reports the asset pricing implications of the 2-state Markov switching regime model for different assumptions on investors information set, as well as the pricing methodology. In particular, 'P.P.U. E-Z' stands for priced parameter uncertainty Epstein-Zin utility, and corresponds to the fully rational case where investors account for the parameter uncertainty ex ante in the pricing of asset claims. In this case, the state is assumed to be observed. 'Anticipated E-Z' corresponds to the anticipated utility Epstein-Zin agent as given in Table 5, where the agent at each point in time prices assets using their current best guess of the parameters as the true parameter value. In this case the state is unobserved, and the agent takes into account state uncertainty when pricing the claim. Finally, 'Known parameters' correspond to the case where parameters are known but where the state is unobserved but there is no parameter uncertainty, also reproduced from Table 5. In all cases, the agent learns from observing consumption growth only. For all the models, $\gamma = 10$, $\beta = 0.994$, $\psi = 1.5$, $\lambda = 4.5$. The volatility of the idiosyncratic component of dividend growth ($\epsilon_{d,t}$) is calibrated to match the historical standard deviation of dividend growth, as reported in Bansal and Yaron (2004). The statistics are annualized. The expectation operator with a T subscript, E_T, denotes the sample average, while the volatility operator, σ_T denotes the sample standard deviation. The sample period is from 1957:Q2 until 2009:Q1, with 1957Q1 priors as given in Table 1 for the case of learning from observing consumption growth only, and as given in an Online Appendix for the case of learning from both consumption and GDP growth.

Moments	Data 1957:Q2-2009:Q1	P.P.U. E-Z 2-state model $\gamma = 5$	P.P.U. E-Z 2-state model $\gamma = 10$	Anticipated E-Z 2-state model $\gamma = 10$	Known parameters 2-state model $\gamma = 10$
The real risk-free rate:					
$E_T(r_t^f)$	1.6%	3.6%	3.0%	3.7%	3.8%
$\sigma_T(r_t^f)$	1.6%	0.9%	1.3%	0.8%	0.6%
The dividend claim: $\Delta d_t = \mu + \lambda(\Delta c_t - \mu) + \varepsilon_{d,t}$					
ex post:					
$E_T(r_t - r_t^f)$	4.7%	4.6%	7.5%	3.7%	1.1%
$\sigma_T(r_t - r_t^f)$	17.1%	13.8%	13.8%	15.2%	11.9%
Sharpe ratio	0.27	0.34	0.54	0.24	0.09
$\sigma_T(pd_t)$	0.38	0.11	0.11	0.23	0.06
$Corr_T(pd_t^{Model}, pd_t^{Data})$	n/a	0.62	0.69	0.39	0.25
ex ante:					
$E_T[E_t(r_{t+1} - r_{t+1}^f)]$	n/a	2.5%	5.6%	1.8%	0.9%

Table 7 - Dividend Yield Regressions

Table 7: The table reports the results of regressions where the U.S. log aggregate stock market dividend price ratio is the independent variable. Panel A shows regressions of this on the contemporaneous log dividend price ratios from the Anticipated utility 'full' model with parameter and model uncertainty (dp^{AU_full}), the Anticipated utility 2-state model (dp^{AU_2state}), the fixed parameters 3-state model ($dp^{FixedPar\_3state}$), as well as the 2-state model with observed states but fully rationally priced parameter uncertainty ($dp^{PricedPU\_2state}$). Panel B shows the corresponding regressions using changes in the log dividend-price ratios. The standard errors are corrected for heteroskedasticity and given in parantheses under the coefficient estimates. The 'dagger' symbol seen in the final column means that the particular regressor has been orthogonalized with respect to the other regressors. * denotes significance at the 10% level, ** denotes significance at the 5% level, and *** denotes significance at the 1% level. The full sample period is from 1957:Q2 until 2009:Q1, with 1957Q1 priors as given in Table 1 for the case of learning from observing consumption growth only, and as given in an Online Appendix for the case of learning from both consumption and GDP growth.

Panel A:	Learning from consumption				Learning from consumption and GDP			
Indep. var.: dp^{data}	1	2	3	4	5	6	7	8
constant	−0.87	1.86	0.57	3.44***	−0.09	1.86	−0.71	−1.43***
	(0.64)	(1.98)	(1.67)	(0.91)	(0.46)	(1.98)	(1.40)	(0.34)
dp^{AU_full}	0.77***		0.71***		1.13***		1.15***	†2.71***
	(0.19)		(0.17)		(0.15)		(0.15)	(0.33)
dp^{AU_2state}								0.68***
								(0.11)
$dp^{FixedPar\_3state}$		1.47***	0.45			1.48***	−0.19	
		(0.54)	(0.42)			(0.54)	(0.39)	
$dp^{PricedParUnc\_2state}$				2.38***				
				(0.31)				
R^2	18.7%	6.2%	18.8%	46.3%	45.5%	6.2%	45.3%	63.4%

Panel B:	Learning from consumption				Learning from consumption and GDP			
Indep. var.: Δdp^{data}	1	2	3	4	5	6	7	8
constant	−0.00	−0.00	0.00	−0.00	−0.00	−0.00	0.00	0.00
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
dp^{AU_full}	0.34***		0.72**		0.38***		0.77***	†0.83*
	(0.09)		(0.30)		(0.09)		(0.28)	(0.47)
dp^{AU_2state}								0.37***
								(0.09)
$dp^{FixedPar\_3state}$		0.34***	−0.51			0.34***	−0.56*	
		(0.10)	(0.36)			(0.10)	(0.32)	
$dp^{PricedParUnc\_2state}$				0.62***				
				(0.22)				
R^2	6.5%	3.9%	7.2%	3.4%	8.4%	3.9%	10.1%	8.7%

Table 8 - Real risk-free yield volatilities

Table 8: The table reports the sample standard deviation of annualized real risk-free yields at different maturities as computed from each of the models with anticipated utility pricing considered in the paper over the post-WW2 sample (1957 – 2009). The data column reports the standard deviation of annualized yields from the available data on TIPS from the Federal Reserve, which is monthly from January 2003 to February 2011.

TIPS (2003 – 2011)	*Data* (*s.e.*)	*Learning from consumption*	*Learning from consumption, GDP*	*Fixed Parameters* 2-*state*	3-*state*
5-yr yield	0.75% (0.18%)	0.35%	0.54%	0.16%	0.21%
10-yr yield	0.45% (0.11%)	0.33%	0.45%	0.08%	0.10%
20-yr yield	0.30% (0.06%)	0.31%	0.43%	0.05%	0.06%
30-yr yield	*n/a*	0.31%	0.42%	0.03%	0.03%

Table 9 - Return Forecasting Regressions

Table 9: This table presents quarterly excess market return forecasting regressions over various forecasting horizons (q quarters; 1 to 16). The top left panel shows the results when using market data and a measure of the log aggregate dividend yield; the *cay*-variable of Lettau and Ludvigsson (2001) and the CRSP aggregate log dividend yield ($\ln \frac{D_t}{P_t}$ where dividends are measured as the sum of the last four quarters' dividends. The rest of the table shows the results using the returns and dividend yield generated within the models. "Cons. only" denotes the model results in the case where only consumption growth is used to update beliefs, while "Cons. and GDP" denotes the model results in the case where both consumption and GDP growth are used to update beliefs. The results for 'Anticipated E-Z utility' refers to the anticipated utility case of the full model with state, parameter, and model uncertainty. The 'P.P.U. E-Z utility' results refer to the 2-state model with fully rational pricing of the parameter uncertainty where the states are observed. Finally, the fixed parameters case correspond to the case of state uncertainty where the parameters are known. Newey-West autocorrelation and heteroskedasticity adjusted standard errors are given in parentheses (the number of lags is equal to the number of overlapping observations). * denotes significance at the 10% level, ** denotes significance at the 5% level, and *** denotes significance at the 1% level. The sample period is from 1957:Q2 until 2009:Q1, with 1957Q1 priors as given in Table 1 for the case of learning from observing consumption growth only, and as given in an Online Appendix for the case of learning from both consumption and GDP growth.

$$r_{t,t+q} - r_{f,t,t+q} = \alpha_q + \beta_{q,dp} \ln(D_t/P_t) + \varepsilon_{t,t+q}$$

	Data				Anticipated E-Z utility			
	$\ln(D_t/P_t) := cay_t$		$\ln(D_t/P_t) := \ln \frac{\Sigma_{j=0}^{3} D_{t-j}^{Mkt.}}{P_t^{Mkt.}}$		Cons. only		Cons. and GDP	
q	β_{dp} (s.e.)	R^2_{adj}	β_{dp} (s.e.)	R^2_{adj}	β_{dp} (s.e.)	R^2_{adj}	β_{dp} (s.e.)	R^2_{adj}
1	1.19^{***}	4.67%	0.03^{*}	1.6%	0.03	1.0%	0.03	0.8%
	(0.31)		(0.02)		(0.03)		(0.02)	
4	4.29^{***}	15.65%	0.11^{**}	6.6%	0.18^{**}	6.4%	0.14^{**}	5.1%
	(1.18)		(0.05)		(0.07)		(0.06)	
8	7.60^{***}	28.1%	0.17^{*}	8.5%	0.37^{***}	14.7%	0.29^{***}	13.4%
	(1.72)		(0.10)		(0.10)		(0.09)	
16	12.31^{***}	41.6%	0.22^{**}	9.5%	0.59^{***}	20.7%	0.42^{***}	16.6%
	(1.82)		(0.11)		(0.16)		(0.14)	

	P.P.U. E-Z utility				Fixed parameters			
	$\gamma = 5$		$\gamma = 10$		2-state model		3-state model	
q	β_{dp} (s.e.)	R^2_{adj}	β_{dp} (s.e.)	R^2_{adj}	β_{dp} (s.e.)	R^2_{adj}	β_{dp} (s.e.)	R^2_{adj}
1	0.10^{**}	2.1%	0.12^{***}	3.4%	-0.01	0.0%	0.005	0.0%
	(0.05)		(0.04)		(0.06)		(0.06)	
4	0.32^{*}	6.0%	0.42^{***}	11.2%	0.19	1.1%	0.21	1.3%
	(0.17)		(0.15)		(0.17)		(0.16)	
8	0.54^{*}	9.1%	0.68^{**}	15.9%	0.40	2.5%	0.44^{*}	3.3%
	(0.32)		(0.28)		(0.24)		(0.23)	
16	1.03	12.6%	1.16^{**}	18.0%	0.29	0.8%	0.32	1.1%
	(0.73)		(0.53)		(0.31)		(0.31)	