

Credit Constraints, Learning and Aggregate Consumption Volatility*

Daniel L. Tortorice[†]

First Version: November 2007

This Version: November 2009

Abstract

This paper documents three empirical facts. First, consumption volatility relative to income volatility rose from 1947-1960 and then fell dramatically by 75 percent from the 1960s to the 1990s. Second, the correlation between consumption growth and personal income growth fell by about 75 percent over the same time period. Finally, absolute deviations of consumption changes from their mean exhibit two breaks in U.S. data, and the mean size of the absolute deviations has again fallen by about 75 percent. First, I find that a standard benchmark permanent income hypothesis model is unable to explain these facts. Then, I examine the ability of two hypotheses: a fall in credit constraints and changing beliefs about the permanence of income shocks to explain these facts. I find evidence for both explanations and find that these facts can be almost completely explained by a model with learning about the nature of income shocks and a reduction in credit constraints. Importantly, I find that estimated changes in beliefs about the permanence of income shocks have substantial explanatory power for consumption changes.

JEL Codes: E21, E32, D83

Keywords: Consumption, Business Fluctuations, Learning

*A previous version of this paper was titled "Income Process Uncertainty, Learning and Aggregate Consumption Volatility". I thank Robert Barro, John Campbell, Gita Gopinath, George Hall, David Laibson, N. Gregory Mankiw, Catherine Mann, Blake LeBaron, James Stock, Aleh Tsyvinski, and members of the Harvard macroeconomics lunch seminar and Brandeis faculty workshop for helpful comments, suggestions, and advice. I am grateful to the Jacob K. Javits Fellowship Program at the U.S. Department of Education and Harvard University for financial support. All errors are mine.

[†]Department of Economics and International Business School, Brandeis University, tortoric@brandeis.edu

1 Introduction

This paper establishes three facts about aggregate U.S. consumption. First, after rising from the 1950s to the 1960s the standard deviation of consumption growth *relative* to income growth fell 50% from the 1960s to the 1970s and then another 50% from the 1980s to the 1990s. Second, the response of consumption to income, or more exactly the correlation between consumption growth and income growth, has fallen 75% over this same time period. Finally, there appears to be two structural breaks in the size of average consumption changes about their mean. One break occurs in 1955 and another break occurs in 1992. The total fall in the average absolute deviation of consumption from its mean is again about 75%.

These facts complement the Great Moderation facts by noting that consumption volatility has fallen more than we would expect based on the observed fall in income volatility. This paper is not the first paper to suggest a changing relationship between consumption and income. Stock & Watson (2002) estimate structural breaks in consumption. Dynan *et al.* (2006) and Dynan *et al.* (2009) also note that the marginal propensity to consume appears to have fallen substantially in U.S. data. However, this paper is the first to document all these facts and the first paper to study the ability of different hypothesis (a reduction in credit constraints and changing beliefs about the nature of income shocks) to explain them.

I model credit constraints using the Campbell & Mankiw (1990) "rule of thumb" framework. In this model a fraction of consumers are credit constrained and their consumption change equals their income change. I estimate a fraction of credit constrained consumers that varies over time and use the credit constraint model to simulate a consumption series. I examine how well this series replicates the consumption facts in the data.

From a theoretical standpoint it is not clear if reduced credit constraints would reduce the volatility of consumption. If individuals want to smooth income shocks then reducing financial constraints would make consumption smoother. However, if individuals want to borrow to respond more than one to one with income shocks, as they should if shocks are permanent, then reduced financial constraints could increase consumption volatility. I

therefore model the nature of income shocks explicitly allowing for uncertainty as to whether or not shocks are permanent or temporary. This model has two advantages: it allows one to study the credit constraint model without making assumptions about the permanence of income shocks and changing beliefs about the nature of income shocks may themselves potentially explain the observed patterns in consumption volatility.

As pointed out by Deaton (1992), if income is stationary about a stable time trend, the case where shocks are persistent but transitory, consumption changes exhibit less volatility than income changes. However, if income is non-stationary and shocks are permanent, consumption exhibits greater volatility than income. It is also known (Stock (1991) and Cochrane (1988)) that distinguishing between these two models in samples the length of the U.S. macroeconomic time series is very difficult. Since these two income processes are so difficult to distinguish, one might expect consumers to have uncertainty as to which is the true model. Their beliefs may change over time resulting in time varying consumption volatility to income volatility. (The second to last paragraph of section 4.2.1 provides more motivation to why consumers share this econometric uncertainty.)

To examine the plausibility of this hypothesis, I study a model based on the learning model from Cogley & Sargent (2005). In my model, the agent believes the income process is non-stationary with some probability and stationary otherwise. At each point in time the individual, using Bayes's Rule, first updates her beliefs about the parameters of the two models, and then updates her beliefs about the probability that each model is true. Based on these beliefs she then chooses her optimal level of consumption. Next, I examine a special case of the model, where the individual knows the parameters of the stationary and non-stationary model, but does not know which is the true model. From this model I can simulate a consumption series and examine the ability of the model to capture the previously outlined facts in U.S. consumption data.

First I find that two benchmark permanent income hypothesis models fail to predict the fall in consumption volatility. Then, I find that only the learning model is able to explain

the rise and fall in consumption volatility early in the sample. The credit constraints model is consistent with an overall decline in the response of consumption to income, but cannot explain the early rise in consumption volatility and its quick fall after 1960. Importantly, I find only a model with learning about the income process can replicate the two breaks estimated in the U.S. data. In contrast, simulated consumption from the credit constraints model shows evidence of only one break. Finally, I find that the learning model's implied probability weights on the different income processes have substantial explanatory power for consumption changes even controlling for changes in income, permanent income, and time trends. In sum, both the learning model and the credit constraints model capture the decline in consumption volatility relative to income volatility. However, learning is more consistent with variation in consumption volatility early in the sample and is necessary to explain the magnitude of the drop in consumption volatility. Finally, I find a model with learning and credit constraints fits the data best. One of the virtues of this paper is the simplicity of the learning model. Estimation of the model does not use any information on the observed patterns in consumption volatility. It is remarkable then that the model is able to match the observed patterns in consumption data so well.

This paper relates to three important strands of the macroeconomics literature. First it relates to the literature on learning in macroeconomic models (Evans & Honkapohja (2001)). Like Cogley & Sargent (2005) it is an example of how relaxing the rational expectations assumption in favor of learning can result in a significantly different choice for the optimal policy. However, there is a distinction between my work and theirs. In their model, the optimal choice differs because while the agent is fairly sure one model is correct, they harbor some very small belief that another model could be true. The optimal policy under the model that they believe is true would cause a disastrous outcome under the model they believe might be true. Therefore this model, though mostly ruled out by the data, still influences optimal policy. In my learning model, the two possible income processes are difficult to distinguish in small samples. Hence, learning about which process is true takes a long time. However,

because consumption volatility is sensitive to small deviations from stationarity, the two processes imply very different optimal policies. Learning then introduces a very different choice of consumption relative to the no-learning model.

The paper also relates to some of the most recent research in the consumption literature. Guvenen (2005) finds that introducing learning about income into a life-cycle consumption model changes many of the model's predictions. Specifically, optimal consumption choices are very different when one learns about the trend in income than when it is known. While my paper also demonstrates how introducing learning can affect model predictions, it focuses on the relative likelihood of distinct models of the income process. Guvenen focuses on the differences in model predictions with learning about one parameter of the income process versus when there is no learning. The paper also relates to Aguiar & Gopinath (2005). They explain the differences in consumption volatility relative to income volatility in different countries. In their model, different countries have different ratios of consumption volatility to income volatility because they have different income processes. In my paper, the variance of consumption relative to income varies over time depending on the relative likelihood of two different income processes.

Finally, this paper relates to the extensive literature on the Great Moderation. While an extensive literature review would be out of place here, I note three papers most closely related, Cecchetti *et al.* (2005), Dynan *et al.* (2006) and Cecchetti *et al.* (2006). Cecchetti *et al.* (2005) show that the reduction in the standard deviation of consumption was accompanied by a rise in debt level in the U.S., supporting the financial innovation explanation. Dynan *et al.* (2006) show that estimated marginal propensities to consume seem to have fallen over time. Cecchetti *et al.* (2006) relates change in consumption volatility across countries to changes in estimated fractions of rule-of-thumb consumers. This paper differs for three reasons. First, the question is different. I ask if financial innovation can account for the fact that consumption volatility has declined more than personal income volatility in U.S. data. Secondly, I estimate a model of consumption with credit constraints and test the simulated

consumption series directly to see if it is consistent with the decline in consumption volatility. Finally, I solve an innovative learning model to see if it can explain features of the data that the credit constraints model does not.

Other explanations have been proposed for the Great Moderation. McConnell & Perez-Quiros (2000) argue that better inventory management has led to a decline in the volatility of output. It is unclear that this model makes predictions for the decline in consumption volatility relative to income volatility, so I do not study this explanation. Clarida *et al.* (2000) argue that better monetary policy has led to the moderation in economic activity. Stable inflation, by stabilizing ex-post real interest rates, could lead to smoother consumption. My models abstract from variation in the real interest rate, but in section seven I show that the largest fall in consumption volatility relative to income volatility was accompanied by a substantial *rise* in the volatility of ex-post real interest rates, demonstrating the need for an explanation not based solely on improved monetary policy.

The rest of the paper proceeds as follows. Section two describes the data I use. Section three describes the empirical facts I attempt to explain. Four describes the credit constraint and learning models. Five tests these models and section six provides more intuition for the learning model's predictions. Seven examines robustness to different parameter choices and to allowing the real interest rate to vary. Section eight concludes.

2 Data

The data I analyze come from the National Income and Product Accounts (NIPAs) available from the Bureau of Economic Analysis (BEA). The consumption data are estimates of aggregate consumption of nondurable goods and services and the measure of income is the personal income series. From the consumption series, I remove housing services. They do not represent quarterly consumer decisions but imputations based on observed rental rates.

Data begin in 1947 and end in 2005¹. Series are transformed into a real, per-capita series by deflating with the nondurable consumption deflator, the service consumption deflator, and the GDP deflator respectively, and dividing by the total US population.

I pay particular attention to the distinction between accrued wages and distributed wages. In general, firms distribute most of the wages accrued or earned within a quarter. However, at times firms distribute wages early or late. Occasionally this difference is large. For example, in 1992:Q4 firms distributed 63 billion dollars more in wages than employees earned, and in 1993:Q1 employees earned 72.1 billion dollars more in wages than firms distributed. Similarly, in 1993:Q4 firms distributed 50 billion dollars more in wages than employees earned, and in 1994:Q4, employees earned 56.4 billion dollars more in wages than firms distributed. These discrepancies appear due to an increase in income tax rates in 1993 and 1994.² These distributions are large enough to affect income variance estimates. They appear as large shocks to income though they do not represent uncertainty but the effect of known tax rates. As a result, I assign to income all wages earned in a quarter regardless of how much is distributed. However, this modification does not effect the conclusions of the paper.

3 Empirical Facts

This section establishes three facts about the aggregate U.S. consumption series. First, the relative volatility of consumption to income rose early in the sample and then fell substantially over the last 45 years. Second, the correlation of consumption growth and income growth has fallen over time. Finally, there are substantial breaks in the mean absolute

¹As usual in the literature, consumption analysis is carried out using only data from 1955 onwards avoiding the Korean War. As Campbell & Mankiw (1990) point out the National Service Life Insurance (NSLI) benefits paid to WWII veterans in 1950:1 distort tests of the permanent income hypothesis.

In theory 1950:1 could effect the estimates of the income process. To investigate this possibility, I re-estimated the income process setting the 1950:1 observation to the mean of the previous quarter and the following quarter, using only the 1955 onwards data, and with data removing personal current transfer receipts, the component including the insurance payment. Results were very similar with these robustness checks. Since the 1950:1 observation does not influence the results, I use all the income data to estimate the processes. Because beliefs about the income process are central to this paper, I tie them to the data as much as possible.

²I thank Kurt Kunze, an economist at the Bureau of Economic Analysis, for this insight.

deviation of consumption growth from its mean, falling about 75% over the last 60 years.

3.1 Consumption Volatility Relative to Income Volatility

In Figure 1, I plot the rolling standard deviation of the normalized consumption change divided by the rolling standard deviation for income growth. I use a window over the next 10 years and nondurables and services consumption.³ In my empirical analysis I focus on the normalized consumption change $\frac{\Delta c_t}{y_{t-1}}$ because Campbell & Deaton (1989) show that the permanent income hypothesis makes sharp predictions for this quantity, especially its variance relative to income. Normalizing by consumption is problematic for the model because consumption is a random walk and does not trend as it does in the data. However, for the empirical analysis it does not matter if one normalizes by income or consumption (therefore studying consumption growth). I found that the same facts hold for consumption growth.

As one can clearly see, consumption variance relative to income variance rises over time reaching a peak around 1960 and then falls, quite dramatically by 50% from 1960 to 1970. It remains constant for about 20 years before beginning to fall again in 1990 by an additional 50%. I obtained similar results (omitted) with nondurables alone, finding a larger, but more gradual decline beginning in 1960 and ending in 1980.

3.2 Time Varying Response of Consumption to Income

The data demonstrate that consumption changes appear to have moderated over time relative to income changes. To get a parametric representation of this fact I estimate a time

³Standard errors are calculated using the GMM delta method formulas. See, for example, Cochrane (2001), pp. 207. In short, let $\mu = [E(\frac{\Delta c_t}{y_{t-1}})^2 \ E(\frac{\Delta c_t}{y_{t-1}}) \ E(\frac{\Delta y_t}{y_{t-1}})^2 \ E(\frac{\Delta y_t}{y_{t-1}})]'$ and $\phi(\mu) = \frac{\sqrt{E(\frac{\Delta c_t}{y_{t-1}})^2 - E(\frac{\Delta c_t}{y_{t-1}})^2}}{\sqrt{E(\frac{\Delta y_t}{y_{t-1}})^2 - E(\frac{\Delta y_t}{y_{t-1}})^2}}$. Then $var(\phi) = \frac{1}{T} \left[\frac{d\phi}{d\mu} \right]' \sum_{j=-\infty}^{j=\infty} cov(x_t, x_{t-j}) \left[\frac{d\phi}{d\mu} \right]$ where $x_t = \left[(\frac{\Delta c_t}{y_{t-1}})^2 \ (\frac{\Delta c_t}{y_{t-1}}) \ (\frac{\Delta y_t}{y_{t-1}})^2 \ (\frac{\Delta y_t}{y_{t-1}}) \right]$. I estimate the covariance matrix using a Newey-West estimator with 5 lags.

varying coefficient in the following regression

$$\frac{\Delta c_t}{y_{t-1}} = \alpha + \beta_t \Delta \ln y_t + \varepsilon_t \quad (1)$$

where c_t is the NIPA value of real, per-capita, nondurable and service consumption and y_t is real, per-capita, personal income. I let β_t take the form $\beta_t = \beta_0 + \beta_1 t$. I found that higher order terms (quadratic and cubic) were not significant. The results, in column 1 of table 1, indicate that there is a significant time varying component to β as β_1 is statistically significant and negative. At the beginning of the sample period $\beta_t = 0.4$; by the end of the sample period β_t is around 0.1. This result again shows a considerable moderation of consumption *relative* to personal income. While income has moderated over this time, consistent with the Great Moderation facts, consumption has moderated even more.

This is of course a simplistic way of measuring a marginal propensity to consume. Dynan *et al.* (2009) estimate a marginal propensity to consume controlling for potential output, interest rates, the unemployment rate and the index of consumer sentiment. They also find a substantial fall in the response of consumption to income.

3.3 Breaks in the Variance of Consumption

Since consumption has moderated more than income volatility it is natural to examine how much consumption has moderated. To investigate, I follow the methodology of Stock & Watson (2002) and estimate a break in the mean of the absolute value of the residual from the regression of the normalized consumption change on a constant.⁴

$$\frac{\Delta c_t}{y_{t-1}} = \alpha + \eta_t \quad (2)$$

$$|\eta_t| = \beta + \beta_1 \tau + \varepsilon_t \quad (3)$$

⁴I do not use an AR(1) specification since the permanent income hypothesis implies consumption changes should be uncorrelated. Also, I use all the consumption data, not just the post 1955 data as is the norm in the literature, because the estimated break is very close to 1955, and Bai and Perron recommend looking for a break only after the first 15% of the sample.

To estimate τ (the break dates) I use the methodology of Bai & Perron (1998) and the algorithms and GAUSS code available in Bai & Perron (2003). The methodology, in addition to providing a feasible way to estimate the structural break model, describes a sequential method for estimating the number of breaks (allowing for the possibility of more than one break). These methods consistently estimate the number of breaks and the proportion of the sample that occurs before the break date occurs.

To estimate the number of breaks, Bai & Perron (1998) recommend the sequential application of a test of l breaks versus the alternative of $l + 1$ breaks, called the $\sup F(l + 1|l)$ test. The procedure first tests one break versus the alternative of zero breaks, and then two versus one and so on until the statistic fails to reject the null hypothesis. In addition, Bai and Perron provide a test statistic, called the UD max statistic, that tests the hypothesis of zero breaks against the alternative of $k > 0$ breaks where k is unknown. The estimates of the break dates are those that minimize the sum of squared residuals in equation (3).

Table 2a presents the results from these tests⁵. The consumption series has two breaks: one at 1955:4 and one at 1992:1. The means for each segment are .007, .004, and .002. This implies a moderation in consumption volatility of 70%.

4 Explanations

This section outlines the two models I use to explain the fall in consumption volatility relative to the fall in income volatility. The first model is a model of credit constraints with a time varying fraction of credit constrained consumers. The second is a model of learning where the true income process is unknown and therefore agents' beliefs about the true process are changing over time.

⁵Estimated standard errors and confidence intervals allow for heterogeneity and autocorrelation in the residuals using Andrews (1991) automatic bandwidth with AR(1) approximation and a quadratic kernel. The residuals are AR(1) pre-whitened and the variance-covariance matrix is allowed to vary across segments. See Bai & Perron (2003) for details.

4.1 Credit Constraints

This section of the paper outlines a model with time varying credit constraints as an attempt to explain the previously described empirical facts. Consumption may be very volatile early in the sample because individuals are credit constrained but become less volatile as more assets are accumulated or better consumption insurance becomes available. This explanation would seem to have difficulty explaining why consumption volatility rises at the beginning of the sample until the 10 years beginning in 1960. However, to investigate more carefully, I use the approach of Campbell & Mankiw (1990).

In their setup, a fraction λ of consumers are credit-constrained and increase consumption one for one with increases in income. Therefore, I estimate this regression:

$$\frac{\Delta c_t}{y_{t-1}} = \alpha + \lambda_t \Delta \ln y_t + \varepsilon_t \quad (4)$$

where Δc_t is consumption change, $\Delta \ln y_t$ is the log income change, and ε_t is the innovation in permanent income. Because $\Delta \ln y_t$ is correlated with ε_t one must use an instrumental variables estimation procedure since ordinary least squares will be biased. I choose as instruments the lagged income growth rate and the lagged change in the S&P 500 index. I found that these instruments lead to the largest first stage F-statistics and inclusion of additional instruments suggested by Campbell & Mankiw (1990)– for example lagged consumption changes, savings rates, or interest rates– does not increase the power of the instruments. I also found that using twice lagged instruments resulted in instruments that were too weak.⁶ I estimate several different polynomials in λ , up to a cubic. For the cases where λ varies by time I instrument with the time interactions of the instruments as well.

The results of the estimation are in table 3. I find limited evidence of time variation in λ . In all cases from a linear trend to a cubic trend the time interactions have p-values greater

⁶As pointed out by Christiano *et al.* (1991), time aggregation bias can result in spurious rejection of the permanent income hypothesis when one uses once lagged instruments. To the extent this bias is a problem though, it will tend to overestimate the importance of credit constraints.

than 0.1. However, they do exhibit a downward trend. Next I simulate the credit constraints consumption series setting:

$$\Delta c_t^{cc} = \lambda_t \Delta y_t + (1 - \lambda_t) \Delta c_t^{pi} \quad (5)$$

where Δy_t is the change in current income and Δc_t^{pi} is the change in consumption of a permanent income consumer. In section five I report results for consumption changes with a linear, quadratic and cubic trend in λ . To estimate Δc_t^{pi} , I use the income processes described in table 4 and the learning model described in the next section.

Admittedly the "rule of thumb" framework is a simplistic view of credit constraints. However, it has one main advantage. It gives a very precise prediction of how credit constraints evolved over time. An alternative approach would write down a model where credit constraints vary over time and estimate the model using consumption data. While this approach has its merits, one might worry that such a model is so flexible it could not fail to match the data. That model would not give a true description of the evolution of credit constraints over time but just the pattern of credit constraints one needs to replicate consumption patterns.

4.2 The Learning Model

4.2.1 Motivation

To motivate the learning model I analyze, I begin with an example from Deaton (1992). He studies a basic version of the rational expectations, permanent income hypothesis model and shows how sensitive its predictions for the volatility of consumption relative to income are to the process assumed for income. In this model the representative agent maximizes:

$$\sum_{j=0}^{\infty} \beta^j u(c_t) \quad (6)$$

subject to:

$$A_{t+1} = (1 + r_{t+1})(A_t + y_t - c_t) \quad (7)$$

where A_t is the asset position, y_t is labor income received, c_t is consumption, and r_{t+1} is the risk free real interest rate. Under the assumption of quadratic utility and a constant real interest rate equal to the rate of time preference, consumption is a random walk equal to the annuity value of total wealth.

$$c_t = E_t(c_{t+1}) \quad (8)$$

$$c_t = \frac{r}{1+r}A_t + \frac{r}{1+r} \sum_{j=0}^{\infty} (1+r)^{-j} E_t y_{t+j} \quad (9)$$

Now if we assume that the income process is of the form $\alpha(L)y_t = \beta(L)\varepsilon_t$, where $\alpha(L)$ and $\beta(L)$ are lag polynomials and ε_t is mean zero and serially uncorrelated, then as shown by Deaton (1992) and Hansen & Sargent (1981)

$$\Delta c_t = \frac{r}{1+r} \frac{\beta(\frac{1}{1+r})}{\alpha(\frac{1}{1+r})} \varepsilon_t \quad (10)$$

This formula holds if applied to a detrended income series with a deterministic time trend or if there is a unit root in the $\alpha(L)$ polynomial, so that the income process is non-stationary and income shocks have permanent effects.

Deaton (1992) then estimates a trend stationary and a difference stationary model on the real per capita income series and obtains:

$$y_t = \alpha + 1.42y_{t-1} - .45y_{t-2} + \gamma t + \varepsilon_t \quad (11)$$

$$\Delta y_t = \alpha + .44\Delta y_{t-1} + \varepsilon_t \quad (12)$$

The first equation describes the dynamics of income around a stable, deterministic time trend. In this case, shocks to income are persistent but eventually transitory. The second equation imposes a unit root and sets the time trend to zero. Here, income shocks are permanent. Using the formula (10) and the estimated income processes (11) and (12) and setting the quarterly real interest rate r equal to 0.01, we find that under the first income

process

$$\Delta c_t = .28\varepsilon_t$$

and under the second

$$\Delta c_t = 1.77\varepsilon_t$$

These models clearly predict an enormous difference in the way consumption responds to an income shock. They are also very difficult to tell apart in time series of the length common in macroeconomics. If we add y_{t-1} to equation (12) the two income processes become

$$y_t = \alpha + 1.42y_{t-1} - .45y_{t-2} + \gamma t + \varepsilon_t \quad (13)$$

$$y_t = \alpha + 1.44y_{t-1} - .44y_{t-2} + \varepsilon_t \quad (14)$$

As one can see, if γ is small, they have almost identical parameters on the lag values of income. However, the slightly larger parameters in equation (14) are sufficient for income shocks to change from being transitory to permanent and sufficient for consumption to change from being very smooth to very volatile.

Given that the parameter values are so similar in the two equations it would seem that these two models would be very difficult to distinguish in small samples. This assertion is confirmed by Stock (1991), among others, who finds, using 60 years of macroeconomic data, that the confidence interval for the largest autoregressive root in the GNP series (recall a root above one implies the process is non-stationary) is (.6, 1.04) at the 90% level and (.634, 1.029) at the 80% level.

Of course uncertainty on the part of the econometrician does not necessarily imply that the agent shares this uncertainty. However, given that the uncertainty for the econometrician is so large, it seems warranted to explore the assumption that the agent shares some of this uncertainty. Consequently, she puts a non-zero probability on each of these models being true at any given point in time. Moreover, these weights will change over time as more

evidence is amassed for the stationary or non-stationary model. Since these models imply very different reactions to income shocks the model will make important predictions for consumption variance that differ markedly from a model without learning.

While this model may seem based on a difficult econometric problem, too far detached from the decision making of individuals, I believe it captures two important features of consumption decisions. First, imagine a situation like the 1990s where the economy begins to grow faster. Two views of the world may emerge: the first, the new era view, espouses that the gains are permanent. A second, more pessimistic view, would be that the gains are temporary and the economy will eventually return to trend. If uncertainty about these two views is important, when the recession of 2000 comes along and gives evidence against the new era view, one could imagine an additional effect on consumption from changes in beliefs, over and above the direct effect of the income shock. Second, note that the U.S. has done comparatively well over the last 50 years returning often to a stable trend. One could imagine that, after seeing this repeated pattern, agents will become more confident in the U.S. economy, and react less to recessions, seeing them as temporary deviation from trend.

To demonstrate that these types of changing beliefs are important for consumption decisions, I now develop a model that will allow the agent to have changing beliefs about the nature of the income process. Then I demonstrate that the model captures several features of the consumption data outlined in the previous empirical section.

4.2.2 The Full Model

The learning model in this paper comes from Cogley & Sargent (2005). While their paper is concerned with how uncertainty about the trade-off between inflation and unemployment affects the monetary authority's optimal policy, their model is equally applicable to inferences about the income process in a model of consumption choice.

In the learning model, there are two models of the income process indexed $i = s, ns$

which can be written in regression form $y_t = x_t' \theta_t + \varepsilon_t$. For the stationary model we have

$$y_t = \alpha + \theta_1^s y_{t-1} + \dots + \theta_p^s y_{t-p} + \gamma t + \varepsilon_t^s \quad (15)$$

for the non-stationary model we have

$$\Delta y_t = \alpha + \theta_1^{ns} \Delta y_{t-1} + \dots + \theta_{p-1}^{ns} \Delta y_{t-(p-1)} + \varepsilon_t^{ns} \quad (16)$$

Letting Z^t represent the joint history of Y_t and X_t up to time t , the agents prior beliefs on the parameters for each model are given by:

$$p(\theta | \sigma^2, Z^{t-1}) = N(\theta_{t-1}, \sigma^2 P_{t-1}^{-1}) \quad (17)$$

$$p(\sigma^2 | Z^{t-1}) = IG(s_{t-1}, v_{t-1}) \quad (18)$$

θ_{t-1} is the parameter estimate based on $t-1$ data, $\sigma^2 P_{t-1}^{-1}$ is the estimate of the variance-covariance matrix of $\hat{\theta}_{t-1}$, s_{t-1} is the residual sum of squares and v_{t-1} is the degrees of freedom for estimating variance of the residuals ($t-k$). IG is the inverse gamma distribution.

Maximum likelihood implies the parameters are updated recursively according to

$$P_t = P_{t-1} + x_t x_t' \quad (19)$$

$$\theta_t = P_{t-1}^{-1} (P_{t-1} \theta_{t-1} + x_t y_t) \quad (20)$$

$$s_t = s_{t-1} + y_t' y_t + \theta_{t-1}' P_{t-1} \theta_{t-1} - \theta_t' P_t \theta_t \quad (21)$$

$$v_t = v_{t-1} + 1 \quad (22)$$

Next the agent updates the probability weights on each model. For each model the marginalized likelihood is

$$m_{it} = \int \int \prod_{s=1}^t p(y_s | x_s, \theta_i, \sigma_i^2) p(\theta_i, \sigma_i^2) d\theta_i d\sigma_i^2 \quad (23)$$

and the probability weight is $w_{it} = m_{it}p_{i,0}$, where $p_{i,0}$ is the prior probability on model i .

To calculate the marginalized likelihood Cogley & Sargent (2005) note that Bayes's rule implies, for any θ_i, σ_i^2 :

$$p(\theta_i, \sigma_i^2 | Z_t) = \frac{\prod_{s=1}^t p(y_s | x_s, \theta_i, \sigma_i^2) p(\theta_i, \sigma_i^2)}{m_{it}} \quad (24)$$

$$m_{it} = \frac{\prod_{s=1}^t p(y_s | x_s, \theta_i, \sigma_i^2) p(\theta_i, \sigma_i^2)}{p(\theta_i, \sigma_i^2 | Z_t)} \quad (25)$$

Therefore,

$$\frac{w_{i,t+1}}{w_{i,t}} = \frac{m_{i,t+1}p_{i,0}}{m_{i,t}p_{i,0}} \quad (26)$$

$$= p(y_{t+1} | x_{t+1}, \theta_i, \sigma_i^2) \frac{p(\theta_i, \sigma_i^2 | Z_t)}{p(\theta_i, \sigma_i^2 | Z_{t+1})} \quad (27)$$

As Cogley & Sargent (2005) show, this expression can be evaluated analytically since:

$$p(y_{t+1} | x_{t+1}, \theta_i, \sigma_i^2) = N(y_{t+1} - x'_{t+1}\theta_i, 0, \sigma_i) \quad (28)$$

$$p(\theta_i, \sigma_i^2 | Z_{t+1}) = N(\theta_i, \hat{\theta}_{t+1}, \sigma_i P_{t+1}^{-1}) IG(\sigma_i, s_{t+1}, v_{t+1}) \quad (29)$$

where N and IG indicate the normal and inverse gamma probability density functions respectively. There are analytical expressions for both distributions. Finally, to get the actual probabilities for each model one normalizes the weights to sum to one:

$$p_{NS,t} = \frac{w_{NS,t}}{w_{S,t} + w_{NS,t}} \quad (30)$$

$$p_{S,t} = 1 - p_{NS,t} \quad (31)$$

4.2.3 Model For Empirical Analysis

While the dynamic path of $p_{NS,t}$ implied by the model can be calculated from income data once $p_{NS,0}$ (the prior belief that the non-stationary model is true) and the lag length p are specified, the model is difficult to use for empirical analysis for two reasons. Firstly, the constant reevaluation of parameters introduces a substantial source of variation in consumption. New information changes the parameter estimates and the forecast of permanent income introducing an important source of variation. This source can be very large especially early in the sample when the number of observations are small, obscuring the effect of changing inferences about the type of income process.

The second difficulty with implementing this model directly is the presence of small sample bias in estimates of the stationary model. As noted by Andrews (1981), Hurwicz (1950), Kendall (1954), and Marriott & Pope (1954), ordinary least squares and maximum likelihood estimates are biased downwards when the largest autoregressive root is near one, i.e. the model is very close to being non-stationary. Since the stationary model is close to being non-stationary, even after detrending, estimates of its parameters are biased in small samples. This effect makes the estimates of $p_{NS,t}$ unreliable.

To overcome these challenges I sidestep learning about the parameters of each model and focus on learning about which type of model is the correct model. Under this assumption, the agent places probability one on one set of parameters $\{\theta_i, \sigma_i\}$, and probability zero on all others, equation (27) simplifies to

$$\frac{w_{i,t+1}}{w_{i,t}} = p(y_{t+1}|x_{t+1}, \theta_i, \sigma_i^2) \quad (32)$$

and the dynamics of the probability are given by

$$\frac{w_{NS,t+1}}{w_{S,t+1}} = \frac{p(y_{t+1}|x_{t+1}, \theta_{NS}, \sigma_{NS}^2)w_{NS,t}}{p(y_{t+1}|x_{t+1}, \theta_S, \sigma_S^2)w_{S,t}} \quad (33)$$

$$p_{S,t} = \frac{1}{1 + \frac{w_{NS,t}}{w_{S,t}}} \quad (34)$$

The parameter values the agent uses to evaluate the likelihood of each model (given in table four) are the OLS estimates on the full income series. One can think of these parameters as the most likely parameter values for the non-stationary and stationary model (conditional on the lag length).

Given these probabilities, consumption is defined by a modification of (9).

$$c_t^l = \frac{r}{1+r} A_t + \frac{r}{1+r} \sum_{j=0}^{\infty} (1+r)^{-j} [p_{s,t}(E_t y_{t+k}|S) + p_{ns,t}(E_t y_{t+k}|NS)] \quad (35)$$

with the budget constraint given by (7).

The consumption model makes simplifications. I assume constant interest rates and quadratic utility to get a closed form solution. Many current consumption models allow for constant relative risk aversion utility, idiosyncratic shocks to labor income, variation in the interest rate, and a zero lower bound on assets. These modifications improve the model's realism. However, to highlight the role for changing beliefs I use the simplest model possible. While these additional complications to the consumption model may further improve the fit of the model, there is no reason to think that they would reduce the ability of learning to explain the data. And while variation in the real interest rate could be responsible for the changing volatility of consumption, I show in section 7 that the real interest rate became more volatile over the time that relative consumption volatility was falling.

While it may seem unorthodox to have learning in the type of model, but not the parameters of the model, the assumption can be justified. Start with the strict rational expectations model of consumption. Here the agent knows the true model of income and the parameters of that model. As argued above, there is considerable uncertainty about which model is the correct model of the U.S. income time series even after 60 years of data (Stock (1991)). This fact motivates me to study a model with uncertainty about which model is the true model. In addition, I could step further from the rational expectations benchmark and introduce uncertainty about the parameters of each model. However there is a less compelling case to

do so. In a long time series it is still difficult to tell these two processes apart, however it is not hard to get precise estimates of the parameters of each model. Of course, by ignoring learning in the parameters I will miss potentially interesting variation in consumption that come from changes in the parameters of the model, e.g. the trend in the stationary model. While not denying that this is an interesting avenue for future research, it is not the main focus of this paper. This paper’s main focus is the interesting variation in consumption that stems from uncertainty about the stationarity of the income process and whether or not adding this uncertainty can explain the observed patterns in consumption volatility.

4.2.4 Overall Volatility and GMM Estimates

It should be clear that the learning model is tightly parameterized. The parameters for the income processes come only from estimating the income processes on income data. No information about consumption is used. The only free parameters then are the lag length and the prior belief at the beginning of the sample in 1947. I choose the lag-length based on the Schwarz-Bayesian information criterion (results are in table 4b.). I choose the prior in 1947 by matching the variance of normalized consumption changes in the data with consumption changes from the model. I show in section 7 that the results of this paper are robust to different prior and lag-length choices. I examine nondurable consumption and services⁷.

Formally, to match the variance in the model to that in the data, let the errors of the model take the form

$$u_t(p_{ns,0}) = \left(\frac{\Delta c_t^d}{y_{t-1}} - \overline{\frac{\Delta c_t^d}{y_{t-1}}} \right)^2 - \left(\frac{\Delta c_t^{sim}(p_{ns,0})}{y_{t-1}} - \overline{\frac{\Delta c_t^{sim}(p_{ns,0})}{y_{t-1}}} \right)^2 \quad (36)$$

where the bar above the variable indicates the mean, the *sim* superscript denotes consump-

⁷Since not all consumption is included in these categories, I normalize the consumption from the model measure. This transformation involves scaling consumption in the model by the ratio of the mean level of consumption in the data with the mean level of income. One could normalize by the share in total consumption as well. Doing so would result in a different estimate of the prior, but would not effect any other aspects of the analysis.

tion from the learning model, and the d superscript denotes the consumption data. Define

$$g_T(p_{ns,0}) = E_T[u_t(p_{ns,0})] \quad (37)$$

$$S = \sum_{j=-\infty}^{\infty} E[u_t(p_{ns,0})u_{t-j}(p_{ns,0})'] \quad (38)$$

Then

$$\widehat{p_{ns,0}} = \arg \min_{p_{ns,0}} g_T(p_{ns,0})' \widehat{S}^{-1} g_T(p_{ns,0}) \quad (39)$$

$$var(\widehat{p_{ns,0}}) = \frac{1}{T} \left(\frac{\partial g_T(p_{ns,0})'}{\partial p_{ns,0}} \widehat{S}^{-1} \frac{\partial g_T(p_{ns,0})}{\partial p_{ns,0}} \right)^{-1} \quad (40)$$

Results of this GMM estimation are in table 5b. I estimate the prior for the learning model and the credit constraints model where the permanent income consumer is modeled as a consumer who learns. The prior, in the learning model is estimated to be 0.22 with a Newey-West standard error, with 5 lags, of 0.045. The resulting confidence interval is [0.133, 0.309]. The prior, in the credit constraints model with learning is estimated to be 0.20 with a Newey-West standard error, with 5 lags, of 0.055. The resulting confidence interval is [0.093, 0.309]. Interestingly a prior of zero is the standard permanent income hypothesis model with a stationary income process, while a prior of one is the standard permanent income hypothesis with a non-stationary income process. These alternatives are clearly outside the confidence bands.

Table 5a presents the standard deviations of normalized non-durable and services consumption changes. The standard deviation in the data is 1.272 which the learning and credit constraints model with learning matches. The credit constraints model without learning predicts a standard deviation of 1.035. The standard permanent income hypothesis models fair poorly. (The predictions of these models are calculated using equations (7) and (9).) The model with a stationary income process predicts a standard deviation of 0.44 and the model with a non-stationary income process predicts a standard deviation of 4.45.

5 Testing the Models Empirically

This section reports the results from empirical analysis of the credit constraints model, the learning model, the credit constraints model with learning, and two versions of the standard permanent income hypothesis model, one with a stationary income process and one with a non-stationary income process. I take these versions of the standard permanent income hypothesis models as benchmark models to improve upon and show that the learning and credit constraints modifications are necessary to general time varying volatility in consumption relative to income. To generate the predictions of the learning model, I calculate the consumption series from equations (35) and (7). To calculate the predictions of the no-learning models I use equations (9) and (7).⁸ To calculate the prediction of the credit constraints model with learning I use equation (5) and permanent income forecasted from the learning model. For the credit constraints model without learning, I focus on the model with a linear trend in credit constraints and income described by a stationary income process. I report results for other credit constraint trends and for the non-stationary income process. However, since the higher order terms do not improve the fit of the model much and the version with a non-stationary income process fits so poorly, I do not dwell on these models. Estimates of the non-stationary and stationary models for income are in table 4a. These estimates are based on a lag length p equal to five.⁹ For the learning model I calibrate the prior belief as described in the previous section. As I will show in section 7, the results are robust to a range of choices for the prior and the lag length.

⁸Since the income processes is estimated using log income, forecasting permanent income is complicated by the fact that $E[y_t] \neq \exp(E(\ln y_t))$. However, using the estimated parameters and shocks, and simulating the possible paths of income, I found that the error in calculating permanent assuming the above inequality is an equality income was 1% of permanent income for the non-stationary process and .01% for the stationary process. Moreover, it does not vary over time. I therefore use $\exp(E(\ln y_t))$ to calculate $E[y_t]$.

⁹To choose the lag length I applied the Schwarz-Bayesian Information Criterion to detrended income and the log change in income. Results, in table 4b, show the information criterion is minimized at $p = 5$. Results were similar using the Akaike Information criterion or the Hannan and Quinn information criterion.

5.1 Dynamics of the Relative Variances of Consumption and Income Growth

Figure 2 plots the agent’s probability weight on the non-stationary income model implied by the learning model. Recall, the initial prior is 0.22 in 1947. As is evident from the graph the probability rises reaching a peak in 1962 and then falls, non-monotonically, putting more and more weight on the stationary model as time passes. This graph suggests that consumption variance relative to income variance would rise early in the sample and fall later on. We saw that this was in fact true in U.S. data (figure 1), so now we examine how well the model actually captures this pattern.

Figure 3 plots the predictions for the relative variance of consumption to income from the models without learning. The no-learning, non-stationary income model predicts the ratio should be roughly constant. It also overestimates the level of the ratio. The no-learning, stationary model again predicts this ratio should be flat and substantially underestimates its level. In contrast, the learning model, shown in the last panel of Figure 3, does much better than these benchmark models. It predicts a rise and fall in consumption volatility. There is a clear rise and fall in consumption variance relative to income variance, peaking at the same time as the peak in the data. However, the model predicts too much variance at the beginning of the sample and too little at the end of the sample. It is important to underscore that the learning model replicates the rise and fall in the variance of consumption without using information on consumption data. It is quite remarkable that the model replicates this distinctive consumption pattern using probabilities calculated only from income data.

Figure 4 examines the ability of the credit constraints model with and without learning to match this fact. The credit-constraint model with a declining linear trend and a stationary income process can explain the magnitude of the overall decline but can not explain the rise and fall of volatility centered around 1960. (Figure 5 reports the results for cubic and quadratic trends and for a non-stationary income process. The cubic and quadratic trends do little to improve the fit; with a non-stationary income process the model predicts an

increase in consumption volatility.) It also cannot explain the steepness of the drop after 1960; it predicts a more gradual decline.

These limitations of the model are solved by adding in the learning model to the credit-constraints model. As one can see in the bottom panel of figure 4, this model tracks the early rise and fall of consumption volatility almost exactly. Its sole discrepancy is predicting that volatility should continue to fall gradually after 1970 when in fact the fall was more abrupt, and occurred later around 1990.

The last panel also highlights an important contribution of the learning model. If loosening of credit constraints are an important contributor to the reduction in consumption volatility it is important to have a good model of consumption choice absent credit constraints. The credit constraints model is sensitive to the choice of income process. If the income process is stationary, the rule of thumb model gives the correct qualitative result. However, if the income process is non-stationary the reverse is true – consumption becomes more volatile as credit constraints loosen. One can view the learning model as completing the credit constraints model. It allows one not to impose an income process in face of uncertainty as to which process is correct and brings the model closer to the data.

5.2 Time Varying Response of Consumption to Income

In section three I found that in the, following regression

$$\frac{\Delta c_t}{y_{t-1}} = \alpha + \beta_t \Delta \ln y_t + \varepsilon_t \quad (41)$$

β_t fell over time ($\beta_t = \beta_0 + \beta_1 t$ with $\beta_0 = 0.4$ and $\beta_1 = -0.0013$). Table 1 reports the ability of the different models to match this fact. We first see that the models without time varying credit constraints or learning (labeled NS and S) cannot match this fact (β_1 is not significantly different from zero). Both the learning model and the credit constraints model, improve over the benchmark models; these models predict a fall in β over time. For the

learning model $\beta_0 = 0.7$ and $\beta_1 = -0.003$ and for the credit constraint model $\beta_0 = 0.5$ and $\beta_1 = -0.0012$. The learning model predicts a larger fall than is in the data. However, adding the time varying credit constraints to the learning model brings it much closer to the data. In this case $\beta_0 = 0.68$ and $\beta_1 = -0.0019$. The basic time varying credit constraints model comes closest to matching the data. This is perhaps not surprising as credit constraints are modeled linearly as is the fall in consumption correlation in this regression.

5.3 Breaks in the Variance of Consumption

In section three I measured breaks using the following set of regressions.

$$\frac{\Delta c_t}{y_{t-1}} = \alpha + \eta_t \quad (42)$$

$$|\eta_t| = \beta + \beta_1 \tau + \varepsilon_t \quad (43)$$

Table 2 presents the results from this estimation. The nondurables and services series was estimated to have two breaks one at 1955:4 and the other at 1992:1. The means for each segment are 0.007, 0.004, and 0.002. The simulated consumption series from the learning model is estimated to have two breaks the first occurring in the third quarter of 1965.¹⁰ While several years from the break estimated in the data, it is within the estimated confidence intervals for the break dates in the data, 1952:4-1966:1. The second break occurs at 1982 which is outside the confidence interval for the break date of 1992:1 in the data. The respective means for each segment from the consumption model with learning are 0.005, 0.002, and 0.001. The model matches both the number of breaks and the relative magnitudes of the change. However, the break dates do not line up exactly. None of the consumption measures from the benchmark standard no-learning permanent income models show evidence of any breaks. Consumption data from the model with credit constraints show only one break at 1958. It needs an additional factor to create variation in its consumption volatility. That

¹⁰One can think of these breaks as "measured breaks" since the simulated data does not have discrete breaks in consumption variance.

additional factor is provided by adding learning dynamics to the credit constraints model. This model shows evidence of two breaks similar to the plain learning model.

5.4 Relation Between Consumption Changes and Probability Changes

The last few subsections argued that a learning model with changing probability weights on the non-stationary model is consistent with the observed change in the variance of consumption over time. This subsection extends that argument. In the data large changes in consumption are associated with large changes in the estimated probability that the non-stationary model is true. I find that this fact is replicated by the learning model.

To motivate this section's analysis, imagine we wanted to investigate the hypothesis that changes in beliefs influenced consumption choice. Then to test this theory we might want to identify the points in time that resulted in the largest changes in beliefs and see if consumption changed more at those times than would be predicted based on changes in permanent income. Since my model identifies how much beliefs should change at each point in time, I can use my estimated beliefs to examine if changes in beliefs help explain consumption changes controlling for other factors that influence consumption. To this end, I estimate the following regression on the consumption data¹¹:

$$\left| \frac{\Delta c_t}{y_{t-1}} - \frac{\overline{\Delta c_t}}{\overline{y_{t-1}}} \right| = \alpha + \beta \left| \Delta p_{s,t} - \overline{\Delta p_s} \right| + \gamma X_t + \varepsilon_t \quad (44)$$

where c_t is the NIPA measure of nondurable and service consumption, $p_{s,t}$ is the estimated probability the stationary model is true (note this is a function only of the income data and not of the consumption data), and X_t is a vector of control variables including a time variable, the time variable squared, $|\Delta \ln PI_t - \overline{\Delta \ln PI}|$ where PI stands for permanent income measured from either the stationary or non-stationary model and $|\Delta \ln y_t - \overline{\Delta \ln y}|$ where y_t is personal income.

¹¹Results are similar for using squared deviations instead of the absolute value of deviations. However, I choose the absolute value to mitigate the potential for outliers to influence the results.

The results in table 6 indicate that the absolute value of the probability change is a significant predictor of the absolute value of the consumption deviation from its mean. The coefficient $\beta = 0.06$ and is significant at the 1% level. In addition, the R^2 of the regression increases by 20% when the probability term is included. I obtain this significance even controlling for changes in income, changes in permanent income, and a quadratic time trend. In other words, large changes in the estimated probability are associated with large changes in consumption controlling for changes in income, permanent income, and time trends.

Table 6 reports the ability of the different models to replicate this fact. The benchmark models without learning or time-varying credit constraints are unable to replicate this fact. The model with learning naturally replicates this fact with a $\beta = 0.2$. However, the coefficient is too high. On the other hand, the model with credit constraints, while predicting a significant coefficient, has $\beta = 0.02$ which is too low. Putting the credit constraints model together with the learning model gets a coefficient that is close to correct, $\beta = 0.08$.

6 Mechanism of the Learning Model

To better explain the internal mechanisms of the learning model, I discuss two implications of the model in more detail. First, I examine the largest movements in the probability the stationary model is true and show that they roughly accord with intuition. When above trend growth continues to be persistent we move toward the non-stationary model and when above trend growth returns to trend we move towards the stationary model. Second, I decompose the consumption variance in the learning model into variance due to changes in permanent income and those due to changes in the estimated probability weights on each income model. I show that the large run up in consumption volatility early in the sample is due to the changing probability weights on each income processes.

6.1 Largest Movements in the Income Process Probability

Recall, from (33), the dynamics of the income process probability are given by

$$\frac{w_{NS,t+1}}{w_{S,t+1}} = \frac{p(y_{t+1}|x_{t+1}, \theta_{NS}, \sigma_{NS}^2)w_{NS,t}}{p(y_{t+1}|x_{t+1}, \theta_S, \sigma_S^2)w_{S,t}} \quad (45)$$

$$= \frac{p(\varepsilon_{t+1}^{NS})w_{NS,t}}{p(\varepsilon_{t+1}^S)w_{S,t}} \quad (46)$$

$$p_{S,t} = \frac{1}{1 + \frac{w_{NS,t}}{w_{S,t}}} \quad (47)$$

Table 7 lists the four dates when the $\frac{p(\varepsilon_{t+1}^{NS})}{p(\varepsilon_{t+1}^S)}$ are highest, and the four dates when they are the lowest, in order to build intuition for the mechanisms of the model. As one can see, the dates that provide the best evidence for the stationary model are the dates where robust income growth is followed by a strong contraction in income growth or vice versa. The largest movement towards the stationary model occurs during the first quarter of 1950. Very slow, slightly negative growth was followed by very robust growth of 26.8% at an annual rate.¹² Similarly income grew in 1980 at an annual rate of -6.5%, compared to a full sample mean of 2.2%. In the previous three quarters it had averaged 3.5%. This reversal in the growth rate of income provides evidence for the stationary model against the non-stationary model.

A similar story emerges from examining the periods of time that lend the strongest support to the non-stationary model. In the first quarter of 1949 income growth fell at an annual rate of -11.7% after two quarters of low growth. This observation supports the non-stationary model versus the stationary model which expects more of a return to trend. Similarly, in the fourth quarter of 1972 strong growth is followed by very strong growth of over 12%. This observation again favors the non-stationary model over the stationary model. Finally, above average growth is followed by very high growth of 6.56% in the first quarter

¹²As noted in the data section (2.1) this large growth is partially the result of a one time government transfer of life insurance benefits to WWII veterans, though there is evidence of real economic growth as well. As noted in the data section care has been taken to make sure this observation does not bias the results.

of 1989, again favoring the non-stationary model.

The only exception to this general pattern occurs in the third quarter of 1959 when robust growth followed by slow growth favors the non-stationary model. This occurs because, if one goes a bit father back in the data, there was a strong contraction in income growth. As a result, the stationary model still expected more robust growth in 1959 than occurred, while the non-stationary model expected less. Finally, it is interesting to note that these dates are slightly overrepresented during recessions. Of the top five movements towards the non-stationary model and top five movements towards the stationary model 30% occur during recessions (40% occur during or at the border of a recession), while only 20% of the sample dates are during recessions. This observation suggests that recessions, and big movements in income in general, tend to have large revelatory power about the income process.

6.2 Decomposition of Consumption Variance

Figure 6 decomposes the variance in consumption relative to income from the learning model into the variance due to permanent income and variance due to changes in probability. Equation (35) can be rewritten to describes the change in consumption:

$$\Delta c_t^l = \frac{r}{1+r} [p_{s,t}\Delta PIS_t + p_{ns,t}\Delta PINSt + (p_{s,t} - p_{s,t-1})(PIS_{t-1} - PINSt_{-1})] \quad (48)$$

where $\Delta PIS_t = \sum_{j=0}^{\infty} (1+r)^{-j} [E_t - E_{t-1}] (y_{t+j}|S)$, $\Delta PINSt = \sum_{j=0}^{\infty} (1+r)^{-j} [E_t - E_{t-1}] (y_{t+j}|NS)$, $PIS_{t-1} = \sum_{j=0}^{\infty} (1+r)^{-j} E_{t-1} (y_{t+j}|S)$ and $PINSt_{-1} = \sum_{j=0}^{\infty} (1+r)^{-j} E_{t-1} (y_{t+j}|NS)$.

Therefore we can decompose the change in consumption into changes due to permanent income shocks and changes due to changes in beliefs:

$$\Delta c_t = \Delta c_t^1 + \Delta c_t^2 \quad (49)$$

$$\Delta c_t^1 = \frac{r}{1+r} [p_{s,t}\Delta PIS_t + p_{ns,t}\Delta PINSt] \quad (50)$$

$$\Delta c_t^2 = \frac{r}{1+r} [(p_{s,t} - p_{s,t-1})(PIS_{t-1} - PINSt_{-1})] \quad (51)$$

A standard variance decomposition gives:

$$Var\left(\frac{\Delta c_t}{y_{t-1}}\right) = Var\left(\frac{\Delta c_t^1}{y_{t-1}}\right) + Var\left(\frac{\Delta c_t^2}{y_{t-1}}\right) + 2Cov\left(\frac{\Delta c_t^1}{y_{t-1}}, \frac{\Delta c_t^2}{y_{t-1}}\right) \quad (52)$$

Figure 6 contains the results from this variance decomposition, normalized by the variance of the income growth rate. Note that I report the variance of $\frac{\Delta c_t}{y_{t-1}}$ predicted by the model, divided by the variance of income. This is not directly comparable to Figure 3 which plots the ratio of standard deviations, however the standard deviation, being non-linear, does not decompose easily. All variances are calculated using the data from the starting date up to 10 years ahead of the date. As one can see, there is a large rise and fall in Δc_t^2 and the covariance term early in the sample. This drives much of the rise in consumption volatility early in the sample. However, by 1963, as learning has reduced the volatility of the changes in probability, most of the variance of consumption is given by the shocks to permanent income (Δc_t^1) and the continued fall in the variance of consumption is due to the increased weight put on the stationary model.

7 Robustness

7.1 Lag length and prior choice

One of the nice features about the learning model is its tight parameterization. There are only two main parameters that need to be set: the first is the lag length of the income process and the second is the initial weight put on the non-stationary model. Figure 7 explores the sensitivity of the results to different choices for the prior probability. I plot the consumption standard deviation ratio from the bottom panel of figure 3 with an initial prior of 0.22 and for a lag choice of 3, 4, 5, 6, and 7. Then I plot the same figure for priors on the non-stationary model of 0.1, 0.2, 0.3, 0.4, and 0.5 with a lag length of 5.

As seen in column one of Figure 7, the initial rise in consumption variance relative to

income variance and the subsequent fall are present for all lag lengths. Moreover, the fit gets better as the lag length increases; the fit is good for any $lag \geq 5$ and the fit deteriorates only slightly for lag lengths of 4 and 3 (This deterioration is due to the large change in the probability associated with 1950:1). Similarly the initial rise and fall in consumption variance relative to income variance is visible for all priors. Increasing the prior only serves to raise the mean level of consumption variance but does not effect the dynamics.

7.2 Real Interest Rate Volatility

All the previous analysis assumed that the real interest rate is constant. However, if the real interest rate varies then the consumer's log linearized Euler equation generalizes to (see, for example, Hall (1988)):

$$\Delta c_t = \alpha + \frac{1}{\gamma} r_t + \varepsilon_t \quad (53)$$

Therefore, interest rate volatility can be an independent source of consumption volatility. To see if interest rate volatility can explain the observed consumption variance, I replot the rolling standard deviation of consumption volatility of figure 2 along with the standard deviation of the rate on the three month t-bill minus the ex post rate of inflation measured with the GDP deflator. The plot, in figure 8, shows that interest rate volatility has the opposite pattern as consumption volatility. It was lowest in the 1960s, when consumption variance relative to income variance was highest, and was highest during the 1970s when consumption volatility was falling relative to income volatility. Therefore, variation in the real interest rate can not explain the observed pattern of consumption volatility. Consumption volatility fell despite increased real interest rate volatility.

This figure casts doubt on the ability of monetary policy to explain the relative fall in consumption volatility. The 1970s was a time of poor monetary policy, with high inflation, but the relative volatility of consumption actually fell. Something acted to smooth consumption. It was not better monetary policy and low inflation in the 1970s. Monetary policy may

be important for explaining absolute changes in consumption or income volatilities. It also may have a role in the relative decline of consumption volatility in the 1990s. However, this figure shows that monetary policy is not the whole story.

8 Conclusion

This paper studies three empirical observations. First, after an increase early in the sample, the standard deviation of consumption growth relative to income growth fell by 75%. Similarly the correlation between consumption growth and personal income growth has also fallen by about 75%. Finally, the consumption series has three estimated breaks in the absolute deviation of consumption from its mean. The fall from the highest mean deviation to lowest mean deviation is again about 75%.

I examine two explanations of these facts. The first is that a fall in the fraction of credit constrained consumers has lead to a reduction in consumption volatility relative to income volatility. While this model captured the fall in consumption volatility it was unable to capture the rise in consumption volatility early in the sample and the steepness of the decline after 1960. It also left open the question of how an unconstrained consumer would respond to an income shock, since if income shocks are permanent, relaxing credit constraints could make consumption more volatile. To address these shortcoming, I studied a model with an agent who learns about whether or not income shocks are permanent or transitory. This model was consistent with the overall shape of the relative volatility of consumption to income over time and adding this model to the credit constraints model resulted in a model that explained the empirical facts highlighted in this paper.

There are two topics for future research suggested by this paper. First, it would be interesting to embed this learning mechanism in a general equilibrium model to see the predictions for GDP volatility. Given the decline in GDP volatility documented by Stock & Watson (2002) among others, it would be interesting to discover if this learning mechanism

can account for some of this reduction. Second, the theoretical insights of the model are applicable to any situation where there is uncertainty about the future expectation of a variable and that uncertainty is large enough to justify different outcomes of the choice variable. Therefore, the type of learning mechanism proposed here can be applicable more broadly outside the aggregate consumption dynamics literature. For example, in the work by Guvenen (2005), he argues that it is difficult to distinguish at the individual level between a trend model in income with mildly persistent shocks and a model without trend and very persistent shocks. While he chooses to study a model with the first income process and learning about the parameters of the process, it would also be possible to study a model with uncertainty about which is the true process and use the learning mechanism of this paper to study beliefs, consumption choice and asset prices under that uncertainty. Finally, the credit constraints model in this paper, while highly tractable, is also simplistic. A more structural approach that uses a detailed model to estimate credit constraints by trying to match moments on financial data (e.g. household debt) may yield additional insight into the impact of credit constraints on consumption volatility.

References

- AGUIAR, MARK, & GOPINATH, GITA. 2005. Emerging market business cycles: The cycle is the trend. *Mimeo*, April.
- ANDREWS, DONALD W K. 1991. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, **59**(3), 817–58.
- ANDREWS, DONALD W.K. 1981. Exactly median-unbiased estimation of first order autoregressive/unit root models. *Econometrica*, **61**(1), 139–65.
- BAI, JUSHAN, & PERRON, PIERRE. 1998. Estimating and testing linear models with multiple structural changes. *Econometrica*, **66**(1), 47–78.
- BAI, JUSHAN, & PERRON, PIERRE. 2003. Computation and analysis of multiple structural change models. *Journal of applied econometrics*, **18**(1), 1–22.
- CAMPBELL, JOHN Y, & DEATON, ANGUS. 1989. Why is consumption so smooth? *Review of economic studies*, **56**(3), 357–73.
- CAMPBELL, JOHN Y, & MANKIW, N GREGORY. 1990. Permanent income, current income, and consumption. *Journal of business & economic statistics*, **8**(3), 265–79.
- CECCHETTI, STEPHEN G, FLORES-LAGUNES, ALFONSO, & KRAUSE, STEFAN. 2005. Assessing the sources of changes in the volatility of real growth. *In*: KENT, CHRISTOPHER, & NORMAN, DAVID (eds), *The changing nature of the business cycle*. RBA Annual Conference Volume. Reserve Bank of Australia.
- CECCHETTI, STEPHEN G, FLORES-LAGUNES, ALFONSO, & KRAUSE, STEFAN. 2006. Assessing the sources of changes in the volatility of real growth. *mimeo*, December.
- CHRISTIANO, LAWRENCE J, EICHENBAUM, MARTIN, & MARSHALL, DAVID. 1991. The permanent income hypothesis revisited. *Econometrica*, **59**(2), 397–423.
- CLARIDA, RICHARD, GALÍ, JORDI, & GERTLER, MARK. 2000. Monetary policy rules and macroeconomic stability: Evidence and some theory. *The quarterly journal of economics*, **115**(1), 147–180.
- COCHRANE, JOHN. 2001. *Asset pricing*. Princeton University Press.
- COCHRANE, JOHN H. 1988. How big is the random walk in gnp? *Journal of political economy*, **96**(5), 893–920.
- COGLEY, TIMOTHY, & SARGENT, THOMAS J. 2005. The conquest of us inflation: Learning and robustness to model uncertainty. *Review of economic dynamics*, **8**(2), 528–563.
- DEATON, ANGUS. 1992. *Understanding consumption*. Oxford University Press.

- DYNAN, KAREN E., ELMENDORF, DOUGLAS W., & SICHEL, DANIEL E. 2006. Can financial innovation help to explain the reduced volatility of economic activity? *Journal of monetary economics*, **53**(1), 123–150.
- DYNAN, KAREN E., EDELBERG, WENDY, & PALUMBO, MICHAEL. 2009. The effects of population ageing on the relationship among aggregate consumption, saving and income. *American economic review: Papers and proceedings*, May.
- EVANS, GEORGE, & HONKAPOHJA, SEPPÖ. 2001. *Learning and expectations in macroeconomics*. Princeton University Press.
- GUVENEN, FATİH. 2005. Learning your earning: Are labor income shocks really very persistent? *Mimeo*, July.
- HALL, ROBERT E. 1988. Intertemporal substitution in consumption. *Journal of political economy*, **96**(2), 339–57.
- HANSEN, LARS PETER, & SARGENT, THOMAS J. 1981. A note on wiener-kolmogorov prediction formulas for rational expectations models. *Economics letters*, **8**, 253–60.
- HURWICZ, L. 1950. Least-squares bias in time series. *in statistical inference in dynamic economics models*. ed. by T.C. Koopmans, Wiley:New York.
- KENDALL, M.G. 1954. Note on bias in the estimation of autocorrelation. *Biometrika*, **41**, 403–04.
- MARRIOTT, F.H.C., & POPE, J.A. 1954. Bias in the estimation of autocorrelations. *Biometrika*, **41**, 390–402.
- MCCONNELL, MARGARET M., & PEREZ-QUIROS, GABRIEL. 2000. Output fluctuations in the united states: What has changed since the early 1980's? *American economic review*, **90**(5), 1464–1476.
- STOCK, JAMES H. 1991. Confidence intervals for the largest autoregressive root in u.s. macroeconomic time series. *Journal of monetary economics*, **28**(3), 435–459.
- STOCK, JAMES H., & WATSON, MARK W. 2002 (Aug.). *Has the business cycle changed and why?* NBER Working Papers 9127. National Bureau of Economic Research, Inc.

Table 1: Normalized Consumption Change Regressed on the Income Growth Rate

	Benchmark Models				<u>Learning</u>	<u>Rule of Thumb</u>	<u>ROT and Learning</u>
	<u>Data</u>	<u>NS</u>	<u>S</u>				
$\Delta \ln(\text{income})$	0.400*** [0.080]	1.500*** [0.092]	0.145*** [0.012]		0.703*** [0.114]	0.514*** [0.007]	0.684*** [0.039]
$\Delta \ln(\text{income}) * \text{time}$	-0.0013*** [0.0005]	0.0002132 [0.0006508]	0.00004 [0.00008]		-0.003*** [0.001]	-0.00120*** [0.00006]	-0.00194*** [0.00025]
time	0.0001 [0.00000]	0.0000014 [0.0000058]	0.0000011 [0.0000009]		0.00002*** [0.00001]	-0.0000020*** [0.0000006]	0.0000018 [0.0000020]
Constant	0.0013743** [0.0006203]	-0.0085577*** [0.0008798]	-0.00100433*** [0.00012233]		-0.0040425*** [0.0009356]	-0.00019378*** [0.00006972]	-0.00108729*** [0.00032466]
	203	203	203		203	203	203

Newey-West Standard Errors with Five Lags in Brackets (* significant at 10%; ** significant at 5%; *** significant at 1%)

This table reports the result of regressing the normalized consumption change, and the predictions of the different models on the income growth rate, allowing the coefficient to vary by time.

Table 2a: Breaks in Consumption Volatility

	<u>Data</u>	Benchmark Models		<u>Learning</u>	<u>Rule of Thumb</u>	<u>Learning and ROT</u>
		<u>Non-Stationary</u>	<u>Stationary</u>			
supF(1 0)	9.44*	2.39	2.38	13.3*	10.1*	9.06*
supF(2 1)	24.12*	2.31	3.8	8.65*	6.51	10.03*
supF(3 2)	0.78	7.35	5.03	1.43	2.93	2.61
UD Max	19.42*	3.64	3.71	13.3*	10.1*	9.85*
Break Date 1	1955:4	--	--	1965:3	1958:1	1959:3
90% CI	1952:4-1966:1	--	--	1965:3-1980:1	1957:3-1966:2	1958:4-1974:3
Break Date 2	1992:1	--	--	1982:1	--	1984:3
90% CI	1991:2-2000:4	--	--	1977:2-1993:4	--	1978:1-1995:4
$\delta 1$	0.007	--	--	0.005	0.005	0.005
	(0.0007)	--	--	(0.0005)	(0.0004)	(0.0005)
$\delta 2$	0.004	--	--	0.002	0.002	0.003
	(0.0003)	--	--	(0.0005)	(0.0002)	(0.0003)
$\delta 3$	0.002	--	--	0.001	--	0.002
	(0.0005)	--	--	(0.0004)	--	(0.0004)

This Table reports the results from testing for a break in the residual of the normalized consumption change regressed on a constant. The means of the residual are given by δ , for before and after the breaks. supF and UDMax are tests for l+1 breaks vs. 1 and any breaks respectively and the critical values for the tests are given in Table 2b. Star denotes $\geq 10\%$ significance.

Table 2b: Critical Values for the Tests

	<u>10%</u>	<u>5%</u>	<u>1%</u>
supF(1 0)	7.04	8.58	12.29
supF(2 1)	8.51	10.13	13.89
supF(3 2)	9.41	11.14	12.66
UD Max	7.46	8.88	12.37

Critical values come from Bai and Perron (2003).

Table 3: Credit Constraint Regressions

	$\frac{\Delta c_t}{y_{t-1}}$	$\frac{\Delta c_t}{y_{t-1}}$	$\frac{\Delta c_t}{y_{t-1}}$	$\frac{\Delta c_t}{y_{t-1}}$
Income Growth	0.445***	0.804***	1.604**	1.302
	[0.115]	[0.289]	[0.768]	[1.56]
(time/100)*Income Growth		-0.003	-0.019	-0.005
		[0.0017]	[0.0127]	[0.043]
time/100		0.001	0.008	0.018
		[0.0009]	[0.0059]	[0.022]
(time/100)²*Income Growth			0.607	-0.831
			[0.450]	[3.618]
(time/100)²			-0.002	-0.010
			[0.002]	[0.019]
(time/100)³*IncomeGrowth				0.399
				[0.915]
(time/100)³				0.002
				[0.0047]
Constant	0.003***	0.001	-0.002	-0.006
	[0.00073]	[0.00175]	[0.00344]	[0.0075]
First Stage F-Statistic				
Income Growth	14.88	9.19	6.22	4.75
time*Income Growth		10.71	6.95	5.39
time ² *Income Growth			9.26	7.42
time ³ *IncomeGrowth				8.18
P-value (Time Interaction = 0)		0.11	0.27	0.16

Robust standard errors in brackets. (* significant at 10%; ** significant at 5%; *** significant at 1%)

This table gives the results from the regression of the consumption change divided by income on income growth instrumenting with lag values of S&P 500 and income growth rates and their time interactions.

Table 4a: Estimates of the Stationary and Non-Stationary Models

	<u>Stationary Model</u>	<u>Non-Stationary Model</u>
Lag 1	1.1859 [0.0654]	0.2096 [0.0645]
Lag 2	-0.0771 [0.0989]	0.1269 [0.0649]
Lag 3	-0.011 [0.0969]	0.1107 [0.0638]
Lag 4	-0.3348 [0.0938]	-0.2304 [0.0617]
Lag 5	0.205 [0.0630]	
time	0.0002 [0.0001]	
Constant	0.2946 [0.1554]	0.0043 [0.0008]
Observations	230	230
RMSE	0.00796	0.00799

Estimates of a stationary and non-stationary model using log real per-capita personal income.

Table 4b: Schwarz-Bayesian Information Criterion

<u>Lags</u>	<u>Stationary Model</u>	<u>Non-Stationary Model</u>
0	-3.939	-6.763
1	-6.673	-6.783
2	-6.709	-6.778
3	-6.71	-6.764
4	-6.697	-6.8003*
5	-6.722*	-6.7798
6	-6.699	-6.757
7	-6.676	-6.734
8	-6.669	-6.711

Estimation of lag length using the Schwartz-Bayesian Information Criterion. Star denotes the estimated lag length.

Table 5a: Full Sample Standard Deviations of 400*($\Delta c_t/y_{t-1}$)

	<u>Standard Deviation</u>	<u>Ratio of SD to Income SD</u>
Non-Durables and Services	1.272	0.461
Consumption Learning	1.272	0.461
NS Consumption	4.448	1.612
S Consumption	0.443	0.161
Rule of Thumb	1.035	0.375
Rule of Thumb and Learning	1.275	0.462

Table 5a reports the variance of the normalized consumption change for the non-durables and the non-durables series beginning in 1955 and the corresponding predictions from the learning, with an initial prior of .5, and non-learning models.

Table 5b: GMM Results

	<u>Learning Model</u>	<u>Rule of Thumb and Learning Model</u>
Estimated Prior	0.221	0.2012
Standard Error	0.045	0.055
Confidence Interval	[0.133 0.309]	[0.093 0.309]

Newey-West standard errors with 5 lags in parentheses

Table 5b estimates the prior weight on the non-stationary model by matching the variance of non-durable consumption using GMM.

Table 6: Regression of $|\Delta c_t/y_{t-1}-E(\Delta c_t/y_{t-1})|$ on Change in Probability

	Benchmark Models				<u>Rule of Thumb</u>	<u>Learning and ROT</u>
	<u>Data</u>	<u>Non-Stationary</u>	<u>Stationary</u>	<u>Learning</u>		
$ \Delta(ps_t)-E(\Delta(ps)) $	0.057*** [0.021]	-0.001 [0.001]	0.002 [0.003]	0.198*** [0.026]	0.019** [0.008]	0.079*** [0.011]
$ \Delta \ln(PIS_t)-E(\Delta \ln(PIS)) $	1.22 [3.627]		1.327*** [0.035]	-6.674*** [2.069]	-4.640*** [1.472]	-6.682*** [1.732]
$ \Delta \ln(PINS_t)-E(\Delta \ln(PINS)) $	-0.11 [0.350]	1.332*** [0.003]		0.685*** [0.193]	0.678*** [0.139]	0.863*** [0.161]
$ \Delta \ln(y_t)-E(\Delta \ln(y)) $	0.124* [0.067]					
t/100	0.0018* [0.0009]					
(t/100)²	-0.00081** [0.00034]					
Constant	0.000577 [0.000595]	-0.000001 [0.000017]	0.000030* [0.000017]	0.000341** [0.000137]	0.000351*** [0.000092]	0.000424*** [0.000106]
R²	0.26					
R² (without row one)	0.22					

Newey-West Standard Errors with Five Lags in Brackets (* significant at 10%, ** significant at 5%; *** significant at 1%)

Column 1 reports the results of regressing the absolute value of the normalized consumption change in the data on the estimated change in probability that the stationary model is true controlling for the changes in permanent income from the stationary model (PIS), changes in permanent income from the non-stationary model (PINS) and income (y). Columns 2-6 redoes the analysis of column 1 using simulated data from the non-stationary, no-learning model from the stationary, no-learning model, from them learning model, from the rule of thumb model, and from the rule of thumb with learning model.

Table 7: Largest Movements in the Non-Stationary Probability

Movements towards the Stationary Model

<u>Year</u>	<u>Quarter</u>	<u>PNS</u>	<u>$P(\epsilon^{NS})/P(\epsilon^S)$</u>	<u>Income Growth Rate</u>	<u>NS Forecast</u>	<u>S Forecast</u>
1949	2	0.288	0.916	-0.47	-2.33	-1.79
1949	3	0.287	0.991	-0.18	-0.50	-0.11
1949	4	0.299	1.062	-0.86	0.80	1.15
1950	1	0.172	0.485	26.79	4.14	4.56
1979	3	0.085	1.102	3.39	0.92	0.55
1979	4	0.094	1.117	4.22	2.03	1.56
1980	1	0.101	1.071	3.04	1.89	1.37
1980	2	0.061	0.581	-6.50	3.81	3.22
1973	2	0.153	1.030	3.90	3.35	2.90
1973	3	0.147	0.953	1.49	2.72	2.27
1973	4	0.178	1.260	4.64	-0.14	-0.60
1974	1	0.128	0.681	-2.65	2.94	2.18
2004	1	0.027	1.021	2.18	2.55	2.98
2004	2	0.029	1.078	0.59	2.00	2.49
2004	3	0.028	0.987	2.46	2.02	2.53
2004	4	0.021	0.746	8.26	2.07	2.59


Movements towards the Non-Stationary Model

<u>Year</u>	<u>Quarter</u>	<u>PNS</u>	<u>$P(\epsilon^{NS})/P(\epsilon^S)$</u>	<u>Income Growth Rate</u>	<u>NS Forecast</u>	<u>S Forecast</u>
1948	2	0.222	1.007	6.56	5.94	5.78
1948	3	0.221	0.993	1.71	1.42	1.53
1948	4	0.221	0.999	-2.11	4.72	4.70
1949	1	0.307	1.561	-11.67	0.81	1.12
1957	4	0.178	1.119	-2.51	0.41	0.78
1958	1	0.218	1.290	-4.68	1.94	2.30
1958	2	0.226	1.049	-0.26	0.31	0.91
1958	3	0.208	0.900	2.61	0.63	1.26
1958	4	0.208	1.002	2.47	2.28	2.86
1959	1	0.203	0.966	4.40	3.60	4.25
1959	2	0.163	0.763	7.32	3.29	4.05
1959	3	0.208	1.353	-0.41	3.47	4.19
1972	1	0.119	1.018	3.53	2.40	2.23
1972	2	0.122	1.034	3.81	2.02	1.83
1972	3	0.131	1.076	5.86	3.18	2.91
1972	4	0.168	1.347	12.25	3.06	2.77
1988	2	0.060	0.966	1.94	2.80	2.27
1988	3	0.061	1.016	2.46	2.29	1.79
1988	4	0.068	1.123	3.57	1.48	0.97
1989	1	0.089	1.325	6.56	2.34	1.72

Note: Average annual per-capita income growth rate is 2.2%

This Table lists the largest movements in the non-stationary probability along with income growth and its predictions.

Figure 1: Ratio of the Standard Deviation of Normalized Consumption Growth to Income Growth



Note: Dashed lines represent the 95% confidence interval.

Figure 2: Probability Weight Over Time



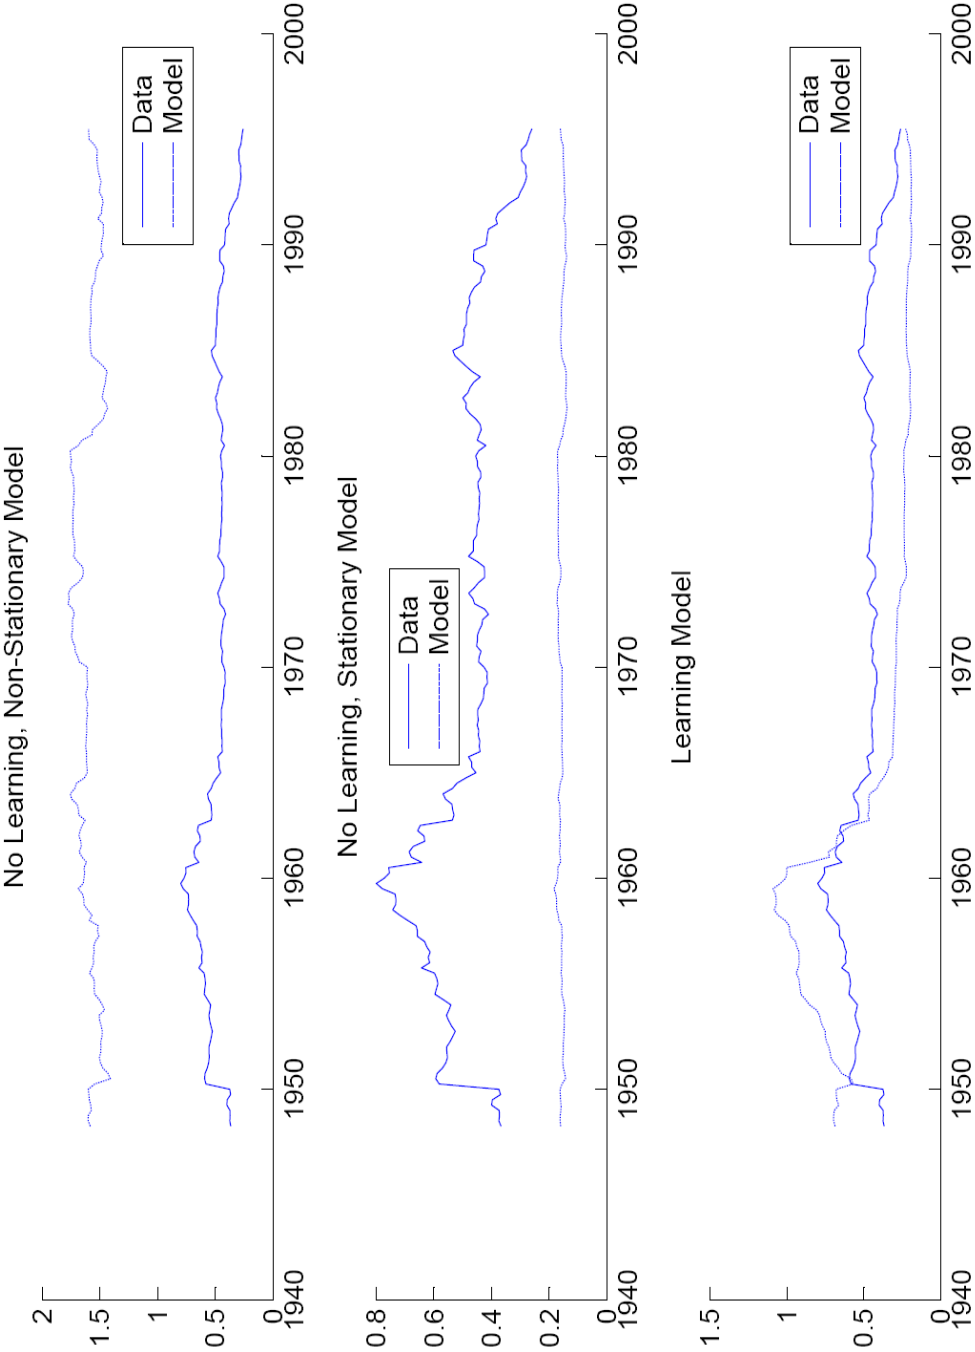



Figure 3: Ratio of Standard Deviation of Normalized Consumption Growth to Income Growth (Models)



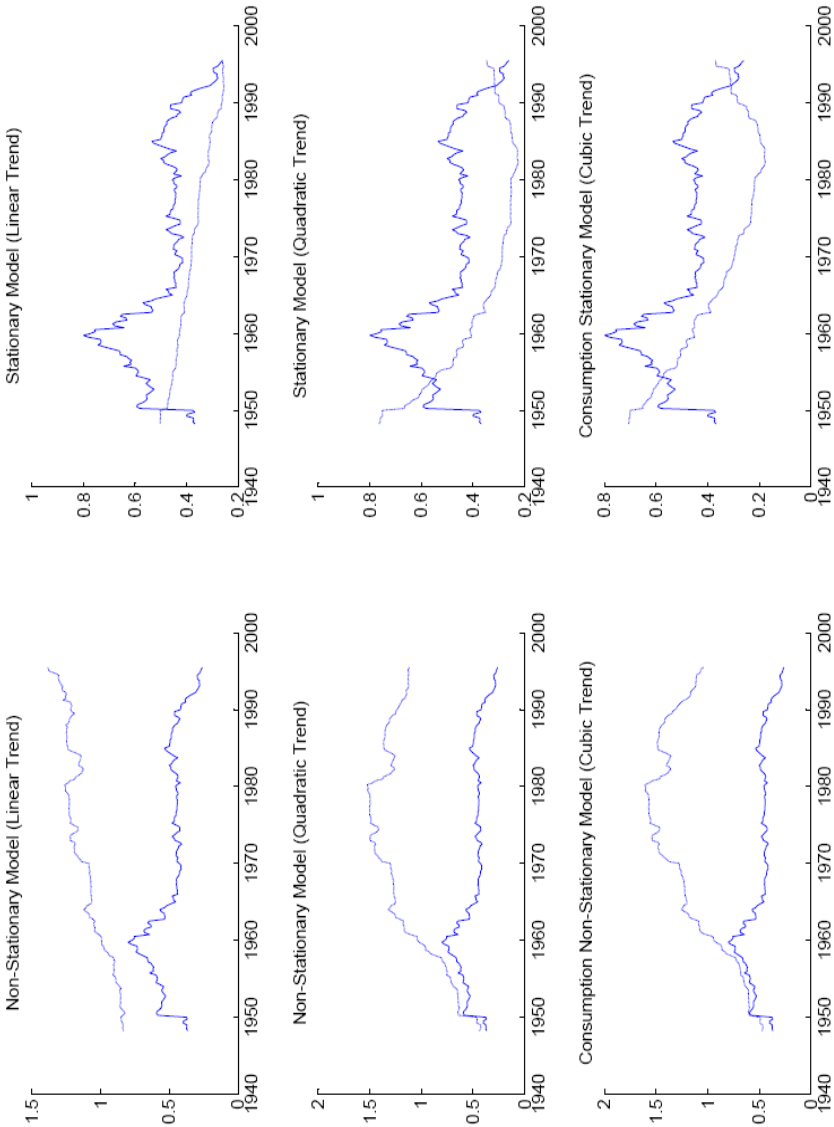
Note: Data (solid line) and model (dashed line).

Figure 4: Ratio of Standard Deviation of Normalized Consumption Growth to Income Growth (Models)



Note: Data (solid line) and model (dashed line).

Figure 5: Ratio of Standard Deviation of Normalized Consumption Growth to Income Growth
Credit Constraint Models



Note: Data (solid line) and model (dashed line).

Figure 6: Variance Decomposition Over Time




Figure 7: Robustness to Different Choices of Lag Length and Prior

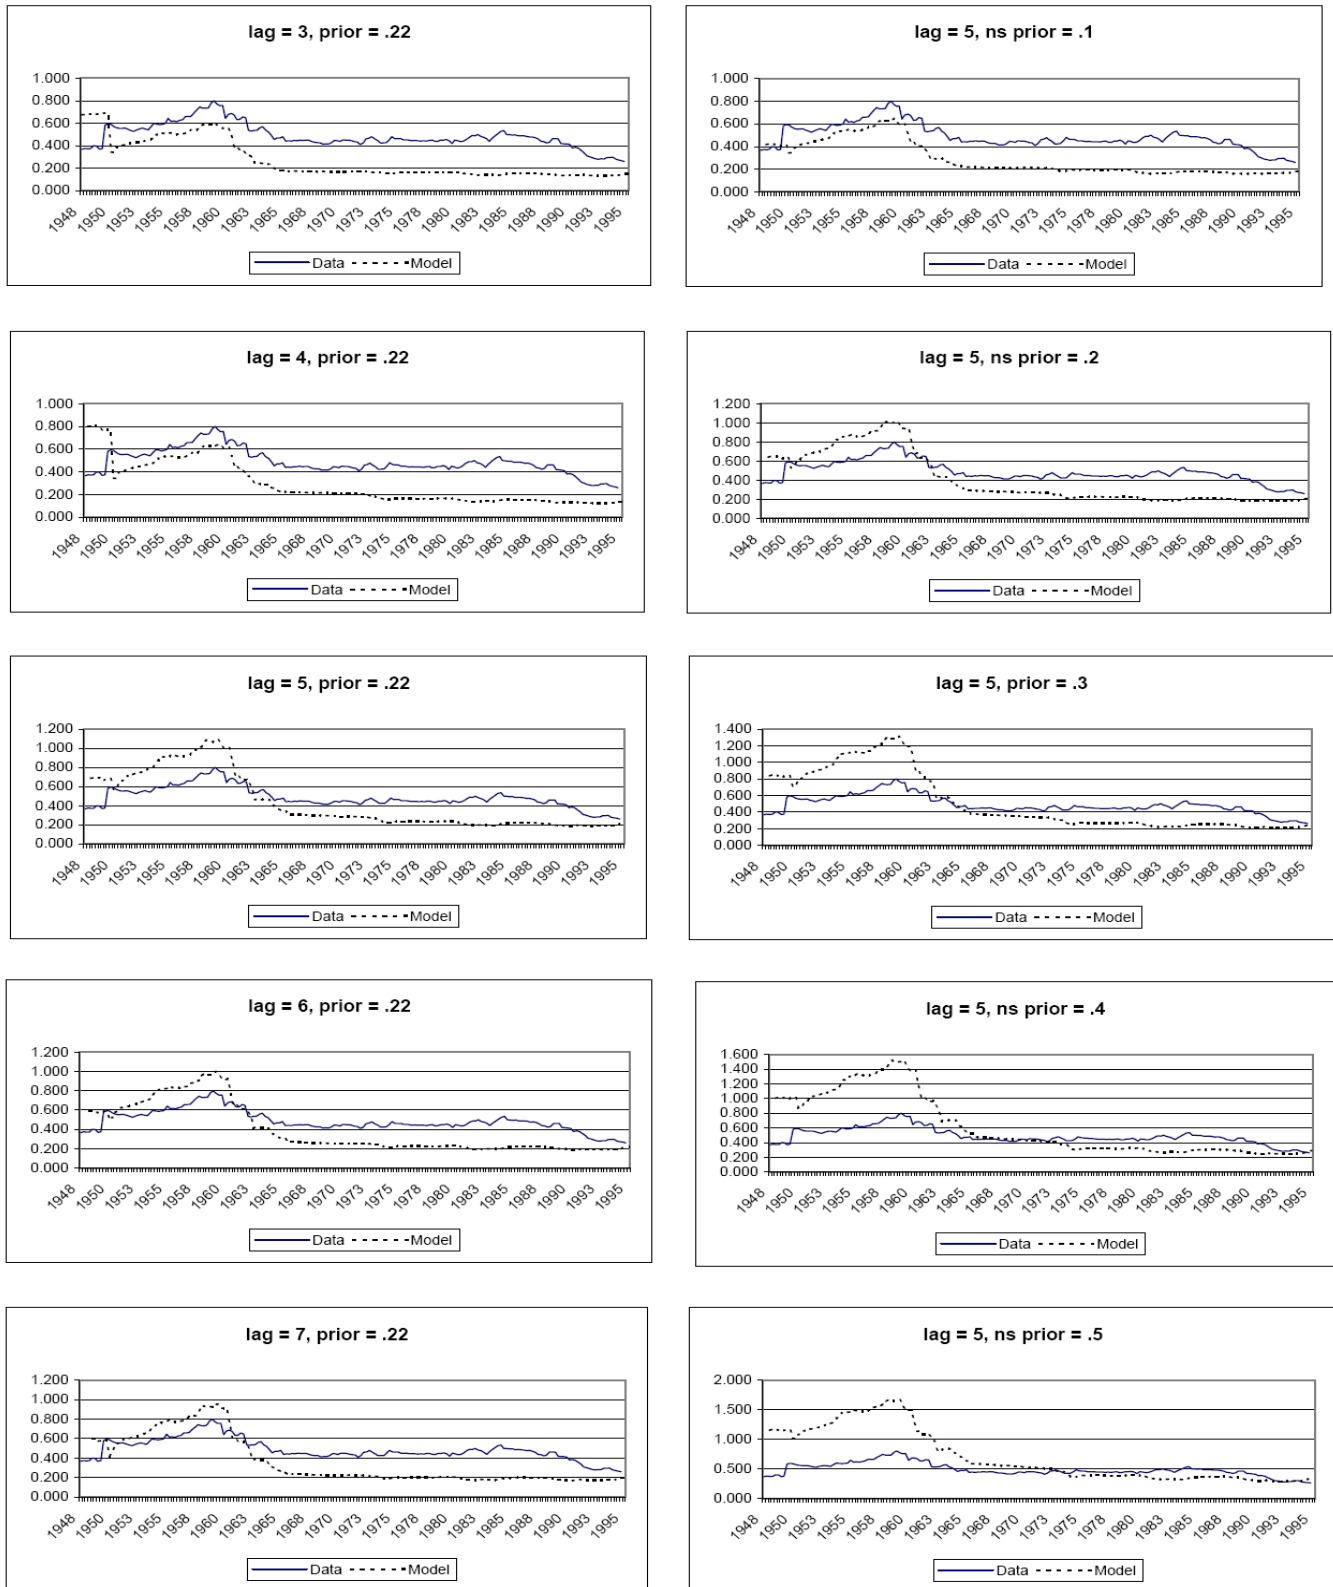



Figure 8: Volatility of the Real Interest Rate versus Volatility of Consumption



Note: Solid line is the consumption data and the dashed line is the real interest rate data.