Lower-Frequency Macroeconomic Fluctuations: Living Standards and Leisure

Benjamin A. Malin*
Stanford University

February 15, 2006

Abstract

Although it is well known that aggregate variables have slow-moving stochastic components, research on macroeconomic fluctuations has focused primarily on high-frequency movements of the data. I document some interesting lower-frequency facts in U.S. postwar data and investigate whether dynamic stochastic general equilibrium (DSGE) models can explain these facts. One fact of particular interest is that hours worked per capita is negatively correlated with both output per capita and total factor productivity (TFP) at lower frequencies, in stark contrast to the positive comovement of these three variables at high frequencies. I show that this lower-frequency fact is puzzling for many DSGE models and explore a variety of candidate solutions to this puzzle. I demonstrate that preferences which depend on a time-varying reference level of consumption ("living standards") can rationalize the observed patterns. Finally, I discuss the relative merits of the "living standards" interpretation of the model to alternative interpretations.

Keywords: Aggregate Fluctuations, Lower Frequency, Labor Hours
J.E.L. Classification: E32, E10

1 Introduction

Research on macroeconomic fluctuations has focused primarily on high-frequency movements of aggregate variables. As noted by King and Rebelo (1999), a fairly conventional definition of the business cycle is fluctuations in economic time series that have a periodicity of eight years or less. This definition has a strong intellectual tradition, following Burns and Mitchell (1946) and Prescott (1986), and an enormous amount of research has concentrated on these

*Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305 (email: bmalin@stanford.edu). I am extremely grateful to my advisors, Pete Klenow, Bob Hall, and Narayana Kocherlakota, for their guidance and support. I have also benefited from suggestions and comments from Soo Lee, Felix Reichling, Esteban Rossi-Hansberg, Todd Schoellman, Azeem Shaikh, John Taylor, Michele Tertilt, Mark Wright and participants of the Stanford Macro Lunch and Macroeconomics Seminar. I thank Ken Judd, Dirk Krueger, and Felix Kubler for sharing their knowledge of computational methods. Finally, the financial support of the Kohlhagen Fellowship Fund at Stanford University is gratefully acknowledged.

frequencies by using either the Hodrick-Prescott filter or a band pass filter to remove low-frequency fluctuations from the data.

Following Kydland and Prescott (1982), the dominant approach to modelling macroeconomic fluctuations has been the construction of dynamic stochastic general equilibrium (DSGE) models. One important feature of these models is that they unify business cycle and growth theory; the models are evaluated on their ability to explain high-frequency fluctuations but are constructed to be consistent with empirical regularities of long-run growth as well. In fact, these models have implications for the movements of economic variables at *all* frequencies, not just business cycle frequencies at one end of the spectrum and long-run growth at the other.

Although researchers have not completely ignored lower-frequency fluctuations, a large gap remains in our understanding of changes in the aggregate economy from, say, decade to decade. In this paper, I document some interesting facts about lower-frequency fluctuations in U.S. postwar data and investigate whether DSGE models can explain these facts. One fact of particular interest is that extended periods of rapid (stagnant) total factor productivity (TFP) growth are accompanied by fast (slow) growth in output per capita and prolonged decreases (increases) in hours worked per capita. This pattern manifests itself as a negative correlation between hours worked and TFP/output at lower frequencies.[1]

The high-frequency relationship between these three variables is very different: hours worked and output have a strong positive correlation, and TFP is positively correlated with both variables. Recessions are times of low output, low hours worked, and low productivity, and expansions are the reverse. Business cycle models have been constructed to capture this high-frequency comovement, but these models imply a positive correlation between labor input and TFP/output at lower frequencies as well. In other words, the lower-frequency behavior is a puzzle for many models of macroeconomic fluctuations.

In order to explore a variety of candidate solutions to this lower-frequency puzzle, I consider a general model formulation that allows for various specifications of primitives (i.e., preferences, technology, government, and the stochastic shock processes). One special case is a textbook real business cycle (RBC) model, e.g., Cooley and Prescott (1995), which is useful for illustrating why standard models perform poorly at lower frequencies. The challenge is then to identify the key ingredients that help explain the lower-frequency patterns and to provide an economically meaningful interpretation of these ingredients.

First, let me give some intuition for the difficulty faced by standard models. To capture stylized facts of long-run growth, these models (and mine) feature a steady state in which most aggregate variables (productivity, consumption, wages, etc.) grow at a common trend rate while hours worked per capita are stationary. Many models achieve this by assuming a preference specification which implies that the marginal rate of substitution between consumption and leisure (i.e., the intratemporal MRS) is linear in consumption (King, Plosser, and Rebelo 1988). The first-order condition for labor supply can then be expressed as

$$n = f\left(w - c\right),$$

[1] I provide a precise definition of high and lower frequency fluctuations in Section 2.

where n is log hours worked, w is the log wage, c is log consumption, and f is an increasing function. In steady state, wages and consumption grow at the same rate and hours worked are constant. Heuristically, the direct substitution effect of a wage increase is captured by w and the income (or wealth) effect by c. Thus, these preferences are said to have directly offsetting income and substitution effects.

Out of steady state, the fluctuations in hours worked depend on the nature of the stochastic driving force and the particulars of wage determination. For concreteness, consider a model with transitory shocks to TFP and a wage equal to the marginal product of labor. In response to a positive transitory shock, the increase in consumption is less than that of the wage because of consumption smoothing and hours worked increase. As a result, the model is able to produce the high-frequency positive comovement of productivity, output, and labor input. During extended periods of rapid productivity growth (i.e., a sequence of mostly positive shocks), the same forces are at work and contrary to the pattern in the data, hours worked increase. Put differently, the income effect is not strong enough to explain the negative lower-frequency comovement between hours worked and output/TFP.

The discussion of the standard model suggests three potential solutions: 1) separate the wage from the marginal product of labor, 2) consider more persistent driving shocks, and 3) relax the restriction that the intratemporal MRS is linear in consumption. Allowing for a wedge between the wage and the marginal productivity of labor severs the tight link between productivity and wage growth and could potentially explain the lower-frequency comovement. One wedge that immediately comes to mind is a distortionary labor income tax. Using and extending a measure of the average marginal income tax rate constructed by Barro and Sahasakul (1983), I find that taxes improve the performance of the model for some sub-periods (70s-80s) but worsen the model's predictions for other times (50-60s and 90s). Thus, taxes do not appear to be a sufficient explanation of the puzzle. The importance of other wedges considered in the business cycle literature, such as sticky wages, is much more difficult to justify at lower frequencies than at high frequencies.

As demonstrated by Campbell (1994), a positive shock to productivity that is *more* persistent than a random walk causes consumption to increase more than the wage as individuals dissave (i.e., borrow against higher expected future wage income). Although this force could help with the lower-frequency behavior of the labor supply, it is also a force for hours worked to counterfactually move in the opposite direction of output at high frequencies. Moreover, extended periods of gradual decline in hours worked would be accompanied by consumption increasing more rapidly than output, a pattern not found in the data. Thus, more persistent shocks will not, by themselves, be a solution to the puzzle. Still, a stochastic process that includes persistent and transitory shocks to productivity will play a significant role in the analysis that follows.

Finally, I consider preference specifications that are consistent with the stylized facts of long-run growth mentioned above but allow consumption to enter the intratemporal MRS non-linearly. The first-order condition for labor supply takes the form

$$n = f\left(w - \eta c - (1-\eta)\, x\right),$$

where $0 \leq \eta \leq 1$, and x is the log "living standard". In the long-run, the living standard grows at the same rate as other trending variables (including w and c), while in the short-run, its growth rate may vary from that of any particular economic variable. Note also that preferences are no longer restricted to imply directly offsetting income and substitution effects, but rather, η can be chosen based on estimates from microeconomic studies of labor supply.

The living standard is exogenous from the vantage point of any individual agent. It captures the average level of living in society as a whole, and individuals derive utility from their consumption and leisure compared to this norm. Its specific formulation is important both for its economic interpretation and the ability of the model to explain high and lower-frequency fluctuations. I consider specifications of the living standard that depend on recent levels of aggregate consumption and the "trend" productivity level. In this way, the living standard closely resembles a reference level of consumption as in Abel (1999), although with an elastic labor supply, it changes tastes for *both* consumption and leisure.

The formulation that best replicates the fluctuations in the data is one where the living standard grows smoothly, although its growth rate does change occasionally. In extended periods of fast productivity growth, the living standard grows (on average) as fast or faster than productivity, and in periods of slow growth, it grows (on average) as slowly or slower. Importantly, an increase in the living standard increases the MRS between consumption and leisure. This allows leisure to increase during prosperous times without requiring that consumption grow at a faster rate than output for an extended period. The living standard also changes the MRS between consumption at different dates, and thus, a general equilibrium model is necessary for understanding all its implications for aggregate fluctuations.

Models with state variables in preferences are common in macroeconomics. Examples include models with 'habit formation' [e.g., Abel (1990,1999); Campbell and Cochrane (1999)], home production [e.g., Benhabib, Rogerson and Wright (1991)], and leisure-enhancing production [Greenwood and Vandenbroucke (2005), Kopecky (2005)]. In fact, these models provide alternative interpretations for what I call the "living standard". I will discuss the relative merits of the various interpretations and, in the absence of more detailed empirical evidence, explain why I choose the living standard story. I should note, however, that this is certainly an area for future research.

Regardless of the explanation of the puzzle, economists have much to gain from understanding lower-frequency fluctuations. Because DSGE models have predictions for fluctuations at all frequencies, lower-frequency fluctuations could be used to discipline business cycle models and possibly distinguish between alternative theories. As DSGE models become more widely used as a tool for policy analysis, this discipline may be useful for highlighting limitations of the models and suggesting alternative explanations for episodes of interest. I view this paper as taking some useful first steps for documenting and explaining lower-frequency fluctuations of interest. The consideration of more persistent shock processes and the use of computational methods appropriate for solving models with these processes are novel aspects of this paper.

This paper is not the first to look at slow-moving fluctuations in aggregate variables, although the particular decomposition is, to my knowledge, new to the literature. To construct lower frequency fluctuations, I remove a linear trend (a fluctuation of infinite periodicity) and use a high-pass filter (Baxter and King 1999) to remove fluctuations more frequent than every 32 quarters. King, Plosser, and Rebelo (1988) look at linearly detrended data and mention the low correlation between hours and output. Hall (1997) uses a polynomial decomposition of time series to construct "medium frequency" fluctuations that are qualitatively similar to my lower frequency fluctuations. Neither of these papers focus on building a model to explain the fluctuations below business cycle frequencies. Comin and Gertler (2004) build a model to explain many features of "medium term" business cycles, defined as fluctuations of periodicity less than 200 quarters.

The rest of the paper is organized as follows. Section 2 documents some facts about lower-frequency fluctuations in postwar U.S. data, including the negative relationship between hours worked per capita and TFP/output per capita. Section 3 presents a general formulation of the DSGE model, which serves as a unifying framework for considering various explanations of the facts. It also provides a context for discussing the methods used to solve the model. Section 4 demonstrates that the lower-frequency relationship between hours worked and TFP/output is a puzzle from the view of a textbook RBC model. Moreover, variants of the textbook model found in the literature do not explain the puzzle. Section 5 demonstrates how a living-standard model resolves the puzzle and discusses alternative interpretations of the model. Section 6 offers some concluding remarks, and appendices describe the data construction and technical aspects in more detail.

2 High- and Lower-Frequency Fluctuations

In this section I first define high- and lower-frequency fluctuations and then develop some stylized facts from postwar U.S. data. Using a band pass filter (Baxter and King 1999) and a linear trend, I decompose the data into three parts: variations at frequencies between 2 and 32 quarters (high frequencies), those between 32 and an infinite number of quarters (lower frequencies), and a linear trend (the zero frequency). Figure 1 shows this frequency decomposition in the time domain. It plots log real GDP per capita, the trend associated with a high pass filter with cutoff frequency of 32 quarters [HP (32)], and a linear trend. The decomposition of the data into high and lower frequencies is given by

$$\begin{aligned}\text{Data} &= \underbrace{\text{Data} - \text{HP (32) Trend}} + \underbrace{\text{HP (32) Trend} - \text{Linear Trend}} + \text{Linear Trend} \\ &= \quad\text{High Frequencies} \quad + \quad \text{Lower Frequencies} \quad + \text{Linear Trend.}\end{aligned}$$

This definition of high and lower frequencies is useful for a couple of reasons. First at the high end of the frequency spectrum, the choice of 32 quarters as the cutoff between high- and lower-frequency cycles allows for comparison with many business cycle studies that filter the data with a Hodrick-Prescott filter, which, in practice, closely approximates a high pass (32) filter. Second at the low end, neoclassical growth theory predicts that variables will grow at a constant rate (i.e., with a linear trend) in the absence of shocks. Finally, with this three-way decomposition, none of the movements in the data are excluded from the analysis. DSGE models have implications for fluctuations at all frequencies, and the decomposition provides a compact, albeit stylized, way of comparing the fluctuations from model simulations with those in the data.

I focus on the volatility, persistence, and comovement of a set of aggregate variables including: output, consumption, investment, hours worked, labor and total factor productivity, and the net real return to capital. The data is quarterly from 1951:1 - 2001:1, the longest time period over which all variables are available.[2] Tables 1 and 2 display a set of statistics for the high- and lower-frequency fluctuations of each of these variables. All quantity variables are in real per capita terms, and all variables (except for returns to capital) are in logs. The moments are from 1954:1 - 1998:1 because the high pass trend is a moving average that depends on three years of past data and three years of future data.

Table 1: Moments of High-Frequency Fluctuations: U.S. 1954:1 - 1998:1

	Standard Deviation	SD relative to Y	Correlation with Y	Correlation with A	First-order Autocorrelation
Y	1.59	1	1	0.58	0.83
C	1.17	0.74	0.87	0.54	0.83
I	7.34	4.62	0.91	0.51	0.78
N	1.73	1.09	**0.86**	**0.10**	0.88
Y/N	0.89	0.56	0.11	0.83	0.71
r	0.35	0.22	0.69	0.68	0.75
A	0.95	0.60	0.58	1	0.75

The variables are GDP (Y), consumption (C), investment (I), hours worked (N), labor productivity (Y/N), the net realized return to capital (r), and TFP (A). All quantities are in real per capita terms, and all time series (except returns to capital) are in logs.

Table 2: Moments of Lower-Frequency Fluctuations: U.S. 1954:1 - 1998:1

	Standard Deviation	SD Relative to Y	Correlation with Y	Correlation with A	First-order Autocorrelation
Y	2.98	1	1	0.92	0.997
C	2.50	0.84	0.85	0.67	0.997
I	6.69	2.25	0.47	0.50	0.996
N	3.43	1.15	**-0.15**	**-0.46**	0.998
Y/N	4.87	1.64	0.72	0.89	0.999
r	0.68	0.23	0.43	0.26	0.997
A	4.18	1.40	0.92	1	0.999

The tables display some interesting statistics. First, the lower-frequency fluctuations are large; in fact, for all variables except investment, the lower frequencies are more volatile than the high frequencies. Second, the correlation between productivity (both labor and total factor) and output is higher at lower frequencies. This fact, along with the pattern in TFP over the postwar period, lead to the consideration of a model in which changes in productivity

[2] The data appendix describes the sample period, data measures and sources, and the construction of all the variables in detail. The restriction of the data to the period 1951:1-2001:1 is due to the availability of the capital stock data used in constructing the return to capital and TFP. The moments of variables that do not require this data do not change significantly when computed over the extended sample period 1947:1-2003:3.

play an important role in explaining lower-frequency fluctuations.

The most striking difference between the moments at different frequencies is the comovement between hours worked per capita and output per capita. These two variables are strongly positively correlated (0.86) at high frequencies but modestly negatively correlated (-0.15) at lower frequencies. The relationship between hours worked and TFP is similar: the two variables have a positive correlation (0.10) at high frequencies and a much stronger negative correlation (-0.46) at lower frequencies. A related feature of the data is that both labor productivity and TFP are less volatile than output at high frequencies but more volatile at lower frequencies. At high frequencies, the positive comovement of productivity and hours worked leads to amplified (relative to productivity) output fluctuations, while the negative comovement at lower frequencies is associated with dampened output. Explaining these differences between the high and lower frequencies is the main focus of this paper.

To better understand the nature of these differences, the top panel of figure 2 shows the deviations of log real GDP per capita and log hours worked per capita from their respective linear trends, while the bottom panel replaces log GDP with log TFP. Note that the trend in hours worked per capita over this sample period is basically zero (0.002%), so the raw data series for hours would look virtually identical to the detrended series in the figures.[3] Recall that removing a linear trend leaves *both* the high- and lower-frequency fluctuations, and thus, the linearly detrended series are a compact way to display all fluctuations at once. The positive relationship between TFP/GDP and hours worked from quarter to quarter (high frequencies) and the negative relationship from decade to decade (lower) are evident in figure 2.

Figure 3 shows the high and lower frequencies of GDP and hours worked. At high frequencies, hours worked and GDP move together, while at lower frequencies, hours worked are high when output is low and vice versa. Figure 4 replaces GDP with TFP. The plots reflect the weak positive correlation between TFP and hours at high frequencies and the negative correlation at lower frequencies. The high-frequency relationship between hours worked and output should come as no surprise: recessions (expansions) are times of low (high) market output and market work, and TFP is positively correlated with both variables although its correlation with hours worked is relatively weak. The lower-frequency relationship between these variables may be surprising, but it is influenced by two well-known patterns.[4] First, output and productivity grew faster than average during the 50s and 60s and then slowed over the next 25 years. Thus, relative to a linear trend, these variables were high in the 60s and 70s. Second, hours worked fell from 1951 to 1975 and then gradually rose, reaching their initial level in the late 90s.

These patterns suggest that the comovement of the growth rates of the lower-frequency fluctuations of output, hours worked, and TFP may be of some interest. These statistics include $Corr(\Delta Y, \Delta N) = 0.55$, $Corr(\Delta Y, \Delta A) = 0.72$, and $Corr(\Delta A, \Delta N) = -0.10$. The

[3] Moreover, the hours worked moments reported in table 2 do not change significantly if a linear trend of 0 is removed rather than the actual linear trend.

[4] The literature on the productivity and output growth slowdown is so large that I choose not to list references. The U-shaped pattern in hours worked can be found in King et. al. (1988) and Ingram et. al. (1997). Both studies use an hours worked series constructed using data from the Current Population Survey. McGrattan and Rogerson (2004) find a similar pattern in decennial census data.

negative correlation between the growth rates of TFP and hours worked will be useful for evaluating the performance of the theoretical models studied later in the paper. It is also interesting to note the growth rates of the lower-frequency fluctuations of these two variables have opposite signs in 55% of the quarters in the sample period. Lower-frequency TFP and hours worked move together at the major turning point of the productivity series (early 70s) and for brief periods of vigorous expansions (mid 60s) and severe recessions (early 80s), but generally, they move in opposite directions.

To summarize, the lower-frequency comovement of hours worked with output/TFP is very different than the comovement of these variables at business cycle frequencies. At lower frequencies, the fast productivity growth of the 50s and 60s was accompanied by rapid growth in output per capita and decreasing hours worked per capita. The slower productivity growth from the 70s through the 90s coincided with slower growth in output and, after a decline in the early 70s, a gradual increase in hours worked through the end of the sample period. As I will show in Section 4, many business cycle models have predictions for lower-frequency fluctuations that are at odds with the data. Because models of economic fluctuations have predictions for a number of aggregate variables, a satisfactory explanation of the lower-frequency comovements will also be required to be roughly consistent with other moments reported in tables 1 and 2.

2.1 Discussion

For analysis of lower-frequency fluctuations, one key issue of the data construction was how to define per capita. Throughout this paper, the population is defined as individuals of age 16-64. Two alternatives were also considered: individuals older than 16 and all ages. The rationale for taking the relevant population as 16-64 is as follows. First, some measure of the working age population seems more appropriate than using total population when one of the objectives is to explain labor input, especially if one's model abstracts from the decision to send children into the workforce. The 16-64 population was chosen rather than 16 and above because compositional changes in the age of the population have occurred over the sample period, while the primary age of retirement has remained 65 years of age.

Robustness checks were performed to check whether alternative definitions of the population change the facts reported above. They do not; the moments reported in Tables 1-2 are qualitatively similar, while the figures change only slightly. When total population is used rather than working-age population, hours worked per capita reaches its trough between 1967-1971 rather than 1975-1980. When retirement-aged individuals are included in the population, hours worked per capita is not as high in the 90s as it is in the figures above. The stylized fact that lower-frequency TFP/output per capita is high (low) when hours worked per capita are low (high), however, is robust to these changes.

Ramey and Francis (2005) consider more involved constructions of both the numerator and denominator of hours worked per capita. For the denominator, they argue that if one wants to limit the time endowment to the potential workforce (as opposed to including the total population), it should be done in a consistent manner that reflects the ability to engage in productive activities. Thus, they adjust the time endowment of individuals to reflect health as well as age. For the numerator, they adjust hours worked to include government work, time spent in formal schooling, and work at home. Their calculations imply significant changes to the time series for leisure, the biggest of which occur during the first half of the 20th century. For the time period under consideration in this paper (1951-2001), their time series for leisure and mine have qualitatively similar movements.[5]

3 A General Model

Formal general equilibrium models require an explicit characterization of technology, the stochastic impulses that shock the economy, preferences, the information available to economic agents, and the market structure. Because of the wide range of possible specifications for each of these components, multiple explanations of the stylized fact documented in the previous section may exist. In this section I describe a unifying framework for evaluating many potential explanations and discuss the methods used to solve the model.

3.1 Model Environment

I consider a one-sector stochastic neoclassical growth model. There is a single output good that is used for both consumption and investment, and the aggregate resource constraint for the economy is $C_t + I_t = Y_t$. The model abstracts from imports, exports, and government spending, although distortionary taxes on income that are rebated lump-sum to households will be considered. All economic agents are price takers and have perfect information about the state of the economy.

3.1.1 Technology

The aggregate production technology is given by

$$Y_t = A_t^{1-\alpha} F(K_t, N_t) = A_t^{1-\alpha} K_t^{\alpha} N_t^{1-\alpha} \tag{1}$$

where Y_t is output, K_t is the physical capital stock, N_t is labor input, and $A_t^{1-\alpha}$ is an index of TFP.[6]

The capital accumulation technology is

$$K_{t+1} = (1-\delta)K_t + I_t. \tag{2}$$

[5] My measure of leisure is simply the (constant) per period time endowment minus time spent in market work.

[6] A_t is raised to an exponent of $1-\alpha$ simply to allow for a cleaner normalization of the model's variables in the process of making them stationary for the solution procedure.

3.1.2 Stochastic Driving Force

Following Kydland and Prescott (1982), random shocks to productivity are taken as the driving force of the model's fluctuations. TFP evolves according to a Markov process

$$A_t = \gamma_t^A A_{t-1},$$

where the growth rate of TFP γ_t^A is exogenous and drawn from distribution Γ_t^A. Various specifications of Γ_t^A are used in the business cycle literature, and I will consider some alternatives below. Note that throughout the paper I use the notation γ^V to denote the growth rate of a variable V.

3.1.3 Preferences

There is a representative household whose preferences are given by

$$\max_{\{C_t,N_t\}} E_0 \sum_{t=0}^{\infty} \beta^t u(C_t, N_t; X_t),$$

where the momentary utility function is

$$u(C_t, N_t; X_t) = \frac{\left[\frac{\left(\frac{C_t}{X_t}\right)^{1-\eta}}{1-\eta} - b\frac{N_t^{1+\nu}}{1+\nu}\right]^{1-\frac{1}{\sigma}} - 1}{1-\frac{1}{\sigma}} \tag{3}$$

with $\nu, b > 0$, $\sigma \geq 0$, and $0 \leq \eta \leq 1$, and in the case $\eta = 1$ or $\sigma = 1$, the appropriate power function is replaced by the log function. X is an exogenous state variable from the perspective of the household.

The parameters of the utility function govern the household's choice of consumption and labor supply as follows: $\frac{1}{\nu}$ is the intertemporal (Frisch) elasticity of labor supply; η governs the income elasticity of labor supply; b is a scaling parameter used to set the steady state level of labor supply; and conditional on the other parameters, σ governs the (average) willingness of the household to substitute consumption through time.[7]

This momentary utility function nests several specifications used in the literature. For example, if $\eta = 1$ and $\sigma = \infty$,

$$u(C_t, N_t; X_t) = \ln(C_t) - b\frac{N_t^{1+\nu}}{1+\nu} - \ln(X_t).$$

With these parameter values, X enters the utility function in an additively separable way and has no effect on the household's decisions. A further special case is the reduced-form of

[7] For parametrizations in which preferences are nonseparable, the intertemporal elasticity of substitution (IES) will depend upon allocations and will vary over time. σ is chosen so that the model's steady-state IES is consistent with estimates of the IES from the literature.

the indivisible labor model of Rogerson (1988) as used by Hansen (1985), which is obtained by setting $\nu = 0$.[8] If instead, $\eta = 0$ and $\sigma = \infty$, utility takes the functional form used by Greenwood, Hercowitz, and Huffman (1988).

In Section 5, I consider alternative functional forms for and economic interpretations of the state variable X. For now, just note that, as long as X does not enter the utility function in a separable way, it affects both the marginal rate of substitution (MRS) between consumption at different dates and the MRS between contemporaneous consumption and leisure. It can be shown that an increase in the living standard increases both marginal rates of substitution: increasing the taste for current consumption relative to future consumption and for leisure relative to contemporaneous consumption.

Empirical observations motivate restrictions on the functional forms that X can take. In the long-run, many aggregate variables grow at roughly the same rate (Kaldor 1957) while leisure is relatively constant (Ramey and Francis 2005). For the model to have a balanced growth path (BGP) in which leisure is constant $\left(\gamma^N = 0\right)$ while trending variables grow at the same rate $\left(\text{i.e., } \gamma^C = \gamma^K = \gamma^I = \gamma^Y = \gamma^A \equiv \gamma\right)$, either $\eta = 1$ or $\gamma^X = \gamma$.[9]

$\eta = 1$ is the well-known requirement pointed out by King, Plosser and Rebelo (1988a) for obtaining a BGP when preferences are a stable function of only consumption and leisure. This requires that the income and substitution effects of a change in the wage are directly offsetting. Based on estimates of income and substitution effects from the labor supply literature (Blundell and MaCurdy 1999),[10] the relaxation of this requirement appears to have some merit. By including a trending variable in preferences that grows at the same rate as other trending variables in the long-run, the model can exhibit balanced growth without requiring $\eta = 1$.

3.1.4 Government

The government has a very limited role in this economy, levying distortionary taxes which it simply rebates lump-sum to the households. It is fully characterized by its budget constraint

$$T_t = \tau_{N,t} W_t N_t + \tau_{K,t} R_t K_t,$$

where T_t are lump-sum transfers to households, $\tau_{N,t}$ and $\tau_{K,t}$ are distortionary taxes on labor and capital income, and W_t and R_t are the wage and return to capital. From the view of private agents, the tax rates and transfers are exogenous processes, and for simplicity, agents expect that future taxes and transfers (relative to total output) will remain at their current rates.

[8] A utility function of the form $u(C_t, N_t) = \ln(C_t) + \ln\left(\overline{N} - N_t\right)$, where $\overline{N}$ is the per period time endowment, is closely approximated by the momentary utility function under consideration in this paper with parameters $\eta = 1$, $\sigma = \infty$ and $\nu = 0.25$.

[9] The requirement that either $\eta = 1$ or $\gamma^X = \gamma$ is derived from evaluating the first-order condition for labor supply, $-\frac{u_N(C_t, N_t; X_t)}{u_C(C_t, N_t; X_t)} = A_t^{1-\alpha} F_N\left(K_t, N_t\right)$, along a balanced growth path of the model.

[10] Tables 1 and 2 of Blundell and MaCurdy (1999) report estimates of the uncompensated wage elasticity and income elasticity from a number of studies. The income elasticity (income effect) is typically substantially smaller in magnitude than the uncompensated wage elasticity (substitution effect).

3.2 Equilibrium

It is straightforward, but tedious, to define a competitive equilibrium for this economy. Therefore, I simply state the necessary conditions that any competitive equilibrium allocation $\{C, K, N\}_{t=0}^{\infty}$ must satisfy.[11] These are derived by maximizing the expected present discounted value of utility subject to the per-period resource constraint and include

1. the intertemporal Euler equation

$$1 = \beta E_t \left\{ \frac{u_C\left(C_{t+1}, N_{t+1}; X_{t+1}\right)}{u_C\left(C_t, N_t; X_t\right)} \left[\left(1-\tau_{K,t+1}\right) A_{t+1}^{1-\alpha} F_K\left(K_{t+1}, N_{t+1}\right) + 1 - \delta\right] \right\}, \tag{4}$$

2. the intratemporal equilibrium condition

$$-\frac{u_N\left(C_t, N_t; X_t\right)}{u_C\left(C_t, N_t; X_t\right)} = \left(1-\tau_{N,t}\right) A_t^{1-\alpha} F_N\left(K_t, N_t\right) \tag{5}$$

3. and the resource constraint

$$C_t + K_{t+1} - (1-\delta)K_t = A_t^{1-\alpha} F\left(K_t, N_t\right). \tag{6}$$

3.3 Solution Method

I use recursive methods to solve for the equilibrium, and thus, it is necessary to transform the model so that all variables are stationary. This is done by scaling all trending variables by the variable X_t. A "hat" over a variable represents the stationary version of that variable (e.g. $\widehat{C}_t = \frac{C_t}{X_t}$).

The recursive formulation for the problem is

$$V\left(\widehat{K}_t, \left\{A_t^{state}\right\}, \left\{X_t^{state}\right\}\right) = \max_{\left\{\widehat{C}_t, N_t, \widehat{K}_t'\right\}} \left\{u\left(\widehat{C}_t, N_t\right) + \beta E_t h\left(\gamma_{t+1}^X\right) V\left(\widehat{K}_{t+1}, \left\{A_{t+1}^{state}\right\}, \left\{X_{t+1}^{state}\right\}\right)\right\}$$

$$\text{subject to } \widehat{C}_t + \widehat{K_t'} - (1-\delta)\widehat{K}_t = \widehat{A}_t^{1-\alpha} \widehat{K}_t^{\alpha} N_t^{1-\alpha}$$

where $\widehat{K_t'} = \gamma_{t+1}^X \widehat{K}_{t+1}$. $\left\{A_t^{state}\right\}$ and $\left\{X_t^{state}\right\}$ are sets of state variables that are needed for forming expectations and determining current levels of productivity A and the preference variable X. $h\left(\gamma_{t+1}^X\right) = 1$ given the preference specification in equation (3).

The model is solved numerically using policy function iteration, solving for nonlinear policy functions $\widehat{C}(S_t)$, $N(S_t)$, and $\widehat{K'}(S_t)$ over the multi-dimensional, continuous state space $S_t =$

[11] A full characterization of the equilibrium allocation would also include initial and transversality conditions. Note that equilibrium prices are given by the marginal product of the production technology with respect to the relevant inputs.

$\left\{\widehat{K}_t, \left\{A_t^{state}\right\}, \left\{X_t^{state}\right\}\right\}$. Specific examples of $\left\{A_t^{state}\right\}$ and $\left\{X_t^{state}\right\}$ will be given in following sections when I consider various specifications of the model's primitives. For now, note that the specifications considered in this paper require between 2 and 6 state variables. As the number of state variables increases, one faces a curse of dimensionality. To address this computational challenge, I use Smolyak's algorithm for approximating functions of high-dimensionality (Krueger and Kubler 2004). This algorithm has been used to solve models with up to 30 state variables, and it proves sufficient for the models considered below. For readers who are interested, a brief overview of the solution procedure and further references are provided in the technical appendix.

An alternative approach for solving models with multiple state variables is to take a linear approximation around the deterministic steady state of the economy. I choose to use a nonlinear solution method because it is more accurate in general and noticeably so for some specifications of the model. As an example, in models with very persistent shocks such as those considered in section 5, state variables γ^X and $\widehat{K}$ vary significantly from their steady state values for plausible realizations of productivity growth. Because the accuracy of a linearized approximation decreases as the model's state variables move further away from their steady state values, a nonlinear solution method is preferable.

4 A Puzzle

In this section I demonstrate that the lower-frequency comovement of hours worked with output/productivity is a puzzle from the view of many DSGE specifications found in the business cycle literature. I first use a textbook RBC model to illustrate the primary difficulty encountered by standard specifications. I then add distortionary labor income taxes to the model to demonstrate that taxes alone are not sufficient for resolving the puzzle. Finally, I show that models with more persistent productivity shocks can not resolve the puzzle without having other counterfactual predictions. This analysis points towards a DSGE specification, fleshed out in section 5, that can rationalize the lower-frequency patterns in the data.

4.1 Textbook RBC Model

The unifying framework described in Section 3 allows for various specifications of the distribution of productivity growth shocks, preferences, and taxes. A textbook RBC model (e.g., Cooley and Prescott (1995)) specifies these ingredients as follows. First, the stochastic process for TFP growth consists of transitory shocks around a constantly growing trend. Formally,

$$\gamma_t^A = \frac{A_t}{A_{t-1}} = \frac{G_t}{G_{t-1}} \frac{Z_t}{Z_{t-1}} = \gamma^G \frac{Z_t}{Z_{t-1}},$$

where γ^G is the constant growth rate of trend productivity G_t, while Z_t is a stationary cyclical component that reflects transitory shocks. Letting lowercase letters denote logged variables (i.e., $z_t = \ln(Z_t)$), log productivity growth γ^a follows a Markov process:

$$\begin{aligned} \gamma_t^a &= \mu_\gamma + z_t - z_{t-1} \\ z_t &= \rho_z z_{t-1} + \varepsilon_{z,t}, \text{ where } \varepsilon_z \backsim \text{i.i.d. } N(0, \sigma_z^2). \end{aligned} \tag{7}$$

where $\mu_\gamma = \ln\left(\gamma^G\right)$. Second, preferences are defined over consumption and time spent in market work only; that is, parameters will be chosen so that X enters preferences in an additively separable way. Finally, this specification abstracts from time-varying taxes altogether.

The state space for this model specification is $S_t = \left\{\widehat{K}_t, \widehat{A}_t\right\}$ with $\widehat{A}_t = z_t$. The current value of the transitory shock to productivity is all that is needed to form expectations of future productivity growth. Since X_t is additively separable in the utility function, it does not appear in either the intertemporal MRS for consumption or the intratemporal MRS between leisure and consumption. X_t is needed to normalize the other variables in the model, but simply letting X grow at a constant rate requires no additional state variables.

4.1.1 Choosing Parameter Values

Standard Parameters Many of the model parameters are selected to ensure agreement with observed long-run values for key postwar U.S. aggregates. These parameters will be set the same way in all versions of model considered in this paper. The mean growth rate of technological progress, μ_γ, is chosen to imply a 1.9 percent annual average growth rate of real per capita output, and the discount factor, β, is then set to imply an average real return to capital of 7.4 percent per year. Capital's share, α, is set to 0.333, consistent with broad evidence, while the depreciation rate, δ, implies average deprecation of 6 percent per year. The parameter governing the steady-state level of labor supply, b, is set to imply an average of 23 percent of available time spent in market work. The methods used to set these parameters are quite standard and are described in Cooley and Prescott (1995).

Preference Parameters A couple of the preference parameters are pinned down by long-run restrictions or functional form assumptions. As discussed in Section 3.1.3, $\eta = 1$ is required for the model to be consistent with a balanced growth path in which hours worked are invariant to the level of productivity. $\sigma = \infty$ makes utility log-separable in consumption and hours worked

$$u(C_t, N_t; X_t) = \ln(C_t) - b\frac{N_t^{1+\nu}}{1+\nu} - \ln(X_t) \tag{8}$$

and implies an intertemporal elasticity of substitution (IES) of 1.[12]

The other parameter that governs the household's preferences for consumption and leisure is taken from estimates from microeconomic studies of labor supply. The first-order condition

[12] Many business cycle models consider preferences that do not require IES = 1, e.g., Cobb-Douglass preferences $u(C,N) = \frac{\left(C^\varsigma(1-N)^{1-\varsigma}\right)^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}}$. These preferences, however, still have the feature that the intratemporal MRS is linear in consumption ($\eta = 1$), which will be shown to be the key feature that prohibits the model from explaining the lower-frequency comovement of labor input with output and productivity.

governing labor input (equation 5) takes the form

$$bN_t^{\nu}C_t = (1 - \tau_N)\, W_t.$$

Letting lowercase letters denote logged variables and dropping uninteresting constants,

$$n_t = \frac{w_t - c_t}{\nu}. \tag{9}$$

Using this equation, the intertemporal (Frisch) elasticity of labor supply $\frac{1}{\nu}$ is estimated to be in the 0.5 -1 range (Blundell and MaCurdy 1999), so I set $\nu = 1.5$. Many business cycle researchers have argued that the elasticity of labor supply should be much larger in a representative agent model than estimated by microeconomic studies because of the importance of the extensive margin (employment vs. unemployment). Note from equation (9) that a larger elasticity will affect the magnitude of the labor supply response to a shock, but it will not affect the sign. Thus, the model's implications for the lower-frequency comovement of labor input with output/TFP is robust to the choice of ν.

Productivity Parameters The parameters for the stochastic process for productivity shocks (equation 7) are chosen using standard procedures, e.g., section 4.1 of King and Rebelo (1999). The parameters are $\rho_z = 0.95$ and $\sigma_z = 0.0105$.[13]

The parameters for the textbook RBC specification are given in Table 3.

Table 3: Parameters for Textbook RBC Model

σ	β	η	ν	b	α	δ	μ_γ	ρ_z	σ_z
∞	0.987	1.0	1.5	30	0.333	0.015	1.0047	0.95	0.0105

4.1.2 Results

I conduct two experiments to assess the model's ability to capture the lower-frequency fluctuations observed in postwar U.S. data. First, I simulate the model many times by drawing stochastic shocks from a random number generator and report the average moments across these simulations, producing tables similar to Tables 1 and 2. I also calculate the variance (across simulations) of the moments, which is useful for making statements about how (un)likely it is that the model could produce the patterns seen in the data. Second, I calculate the Solow residual from actual data and then simulate the model using this residual as the forcing process. Implicitly, I assume that the model and data are the same in terms of their Solow residual, and ask whether they are similar along other dimensions (i.e., hours worked and output). Because this experiment only uses a subset of the stochastic shocks that hit the economy, the fit between the simulated and actual data is unlikely to be extremely close. Still,

[13] Recall that the transitory shock enters the production function as $Z^{1-\alpha}$ rather than Z. Thus, the standard deviation of the transitory shocks σ_z is consistent with values typically used in the business cycle literature.

this experiment will be useful for showing how endogenous variables respond to the driving force.

Tables 4 and 5 show average moments of the high and lower-frequency fluctuations across 200 model simulations. This model does not capture the different relationship that exists between labor input and output/TFP at different frequencies, as reflected by the high correlation between these variables at both the high and lower frequencies. The difference between the average correlation of labor input and output (TFP) at high frequencies and that at lower frequencies is only 0.23 (0.13). Moreover, any particular simulation of the model is unlikely to produce the lower-frequency comovement. The lower-frequency correlation of hours worked and output (TFP) in the data lies more than 9 (16) standard deviations away from the average moment across simulations.[14]

Table 4: Average Moments of High-Frequency Fluctuations: Textbook RBC

	Standard Deviation	SD Relative to Y	Correlation with Y	Correlation with A	First-order Autocorrelation
Y	1.07	1	1	0.998	0.69
C	0.36	0.33	0.92	0.89	0.77
I	4.13	3.84	0.99	0.99	0.68
N	0.31	0.28	**0.98**	**0.99**	0.68
Y/N	0.78	0.72	0.997	0.99	0.70
r	0.04	0.03	0.97	0.98	0.68
A	0.87	0.81	0.998	1	0.68

Table 5: Average Moments of Lower-Frequency Fluctuations: Textbook RBC

	Standard Deviation	SD Relative to Y	Correlation with Y	Correlation with A	First-order Autocorrelation
Y	1.95	1	1	0.98	0.996
C	1.37	0.70	0.88	0.77	0.997
I	5.51	2.83	0.88	0.95	0.995
N	0.39	0.21	**0.76**	**0.86**	0.995
Y/N	1.67	0.86	0.99	0.94	0.997
r	0.05	0.02	0.48	0.63	0.998
A	1.35	0.70	0.98	1	0.9996

The results of simulating the model with the actual U.S. Solow residual as the driving force are shown in Figure 5, which displays the lower frequencies of the Solow residual (TFP) series along with the hours worked and output series from the model and the data. The top panel shows that the model's lower-frequency pattern for hours worked varies significantly from the data. Simulated hours worked actually increase for the first 15 years of the sample period, and over the last 15 years, hours worked are basically flat. In contrast, the actual data shows

[14]The standard deviation (across simulations) of the correlation of lower-frequency labor input with output (TFP) was 0.11 (0.08).

a sizable decrease in hours worked in the 50s and 60s and a strong increase during the 90s. Moreover, the correlation of the growth rates of the lower-frequencies of model hours worked and productivity is 0.46 in comparison to -0.10 in the data. The bottom panel shows the related pattern that the model's GDP series is more volatile at lower frequencies than the TFP series, while actual GDP is less volatile than TFP. The lower-frequency positive correlation between TFP and hours worked in the simulation leads to the extra volatility of GDP.

The model's inability to capture the lower-frequency movements in labor input can be understood by examining the intratemporal equilibrium condition

$$n_t = \frac{w_t - c_t}{\nu} = \frac{y_t - c_t}{\nu + 1} = \frac{(1-\alpha)a_t + \alpha k_t - c_t}{\nu + \alpha}, \tag{10}$$

where the second equality comes from the fact that the wage equals the marginal product of labor. The textbook RBC model has only one force that works against hours worked and output moving in the same direction at all frequencies: the negative wealth effect associated with a wage increase. Given the preference specification in equation (8), the wealth effect is summarized by c_t. Because consumption is smoothed in response to transitory shocks, it does not increase (decrease) fast enough to offset the direct substitution effect of higher (lower) wages on hours worked during extended periods of rapid (stagnant) productivity growth.

Although the textbook RBC model abstracts from many ingredients that are relevant for explaining features of high-frequency fluctuations (e.g., nominal and real rigidities, government spending shocks, a financial accelerator, etc.), these ingredients are typically embedded in DSGE models in ways that do not significantly change the determination of labor supply (equation 10) at lower frequencies. For example, models with sticky wages do separate the wage from the marginal product of labor, but the separation lasts for at most a few years, not for decades. There are, however, some elaborations of the textbook model, to which I now turn my attention, that could be of first-order importance for explaining lower frequencies.

4.2 Distortionary Labor Taxes

In a study of U.S. and European economies, Prescott (2004) argues that different tax codes can explain much of the disparity between the performance of these countries' labor markets over the medium-run. His study focuses on two time periods: 1970-1974 and 1993-1996. For the U.S., he finds that tax rates were lower in the mid 90s than in the early 70s, and hours worked were higher. This raises the question of whether taxes can explain the lower-frequency comovement of hours worked with output/TFP over the entire postwar sample period.

Figure 6 shows an average marginal income tax rate series for the U.S. postwar period, constructed following Barro and Sahasakul (1983), and Figure 7 shows the results of feeding this series through the RBC model. Although I use a different measure of the tax rate than Prescott, our results are consistent for the early 70s and the mid 90s. The timing of tax changes, however, does not help explain the movement of hours worked during some other sub-periods of the sample. The slight decrease in tax rates between 1950 and the mid 60s causes an increase in hours worked in the model. Moreover, increasing taxes in the 90s push hours worked down. All told, a comparison of Figure 7 and Figure 2 suggests that taxes are not sufficient for explaining the puzzle.

4.3 Persistent Shocks

As mentioned in the introduction, an alternative way to generate the negative correlation between lower-frequency hours worked and output/TFP is to consider shocks to productivity that are extremely persistent. That is, rather than characterizing productivity shocks as transitory fluctuations around a stable trend, shocks could be to trend growth as in Pakko (2002) and Aguiar and Gopinath (2005). Formally,

$$\gamma_t^A = \gamma_t^G \frac{Z_t}{Z_{t-1}},$$

where γ_t^G is the stochastic growth rate of trend productivity G_t, and Z_t reflects transitory shocks. Note, however, that regardless of the functional form for the γ^G process, equation (10) still governs the labor supply. In particular, the labor supply is inversely related to consumption's share of output $\left(n_t = -\frac{c_t - y_t}{\nu+1}\right)$. Thus, if the model generates a negative correlation between hours worked and output at lower frequencies, it would also generate a positive correlation between the consumption share and output. This is not something that is seen in postwar U.S. data. Figure 8 shows lower-frequency GDP and the non-high-frequency consumption share (i.e., a linear trend is still in the consumption share). At lower frequencies, the consumption share typically moves in the opposite direction of output.

4.4 Discussion

The above analysis points the way to a potential solution of the lower-frequency puzzle. The key relationship discussed in this section has been the intratemporal equilibrium condition equating the marginal product of labor (MPL) to the MRS between consumption and leisure, where the MRS is a time-invariant, increasing function of market consumption and hours worked. Rewriting equation (10),

$$(1-\alpha)\frac{A_t^{1-\alpha}K_t^{\alpha}}{N_t^{\alpha}} \equiv MPL_t = MRS_t \equiv bN_t^{\nu}C_t. \qquad (11)$$

For this equation to hold during an extended period of fast TFP growth while hours worked decline requires that consumption (counterfactually) grows extremely quickly. Alternatively, there could be a wedge in the equilibrium condition; that is, $MPL_t = Wedge_t * MRS_t$. Distortionary labor taxes provide such a wedge although they did not move as needed to explain the lower-frequency fluctuations in postwar U.S. data.

Business cycle researchers have considered many stories for explaining the presence of this wedge at high frequencies. Examples include taxes, sticky wages, unions, home production (Benhabib, Rogerson, and Wright 1991), leisure-enhancing production [e.g., Greenwood and Vandenbroucke (2005), Kopecky (2005)] and habit formation [Lettau and Uhlig (1999); Boldrin, Christiano, Fisher (2001)]. In the next section I build on a habit formation story to develop a DSGE specification that can rationalize the lower-frequency patterns in the data.

5 A Solution

The solution to the lower-frequency puzzle presented herein considers the hypothesis that individuals derive utility from consumption and leisure relative to a "living standard". As in models with (external) habit formation, individuals' consumption and labor supply decisions are influenced by a state variable in preferences. Thus, the intratemporal MRS will no longer be a time-invariant function of only market consumption and hours worked; put differently, there will be a wedge between the "true" MRS and the "traditional" MRS given in equation (11). This "living standard", denoted by X_t in section 3, will also affect the MRS between consumption at different dates, and thus, a complete general equilibrium analysis is needed to understand all its implications for aggregate fluctuations.

After describing and parameterizing a baseline living-standard model, I repeat the same experiments performed in Section 4 to measure the model's ability to explain the lower-frequency puzzle. To demonstrate the model's key features, I then consider various formulations for the living standard and some alternative parameterizations. Finally, I discuss the model's interpretation of the postwar U.S. economy and the merits of alternative interpretations of the model, such as home or leisure-enhancing production stories, relative to the living standard interpretation.

5.1 "Living Standard" Model

The "living standard" model deviates from the textbook RBC model in two ways: its specifications of the stochastic process for productivity and of preferences. First, to capture sustained periods of rapid (stagnant) productivity growth, log TFP growth γ^a has the following stochastic process:

$$\begin{aligned} \gamma_t^a &= \gamma_t^g + z_t - z_{t-1} \\ z_t &= \rho_z z_{t-1} + \varepsilon_{z,t}, \text{ where } \varepsilon_z \backsim \text{ i.i.d. } N(0,\sigma_z^2) \\ \gamma_t^g &= \begin{Bmatrix} \varepsilon_{\gamma,t} \sim \text{i.i.d. } \Gamma\left(\mu_\gamma, \sigma_\gamma^2\right) \text{ with probability } P \\ \gamma_{t-1}^g \text{ with probability } 1-P \end{Bmatrix}, \end{aligned} \tag{12}$$

where γ_t^g is the stochastic (log) growth rate of trend productivity G_t, z_t reflects transitory shocks, and the distribution $\Gamma\left(\mu_\gamma, \sigma_\gamma^2\right)$ is a truncated normal distribution with mean μ_γ and variance σ_γ^2. The shocks to the trend growth rate are occasional – $P = \frac{1}{80}$ in a quarterly model implies an expected time between trend breaks of 20 years – but persistent, whereas the transitory shocks occur frequently.

For preferences, I consider specifications of the momentary utility function (equation (3)) that allow the living standard X_t to impact agents' decisions. This requires $\eta < 1$. Following Abel (1999), the living standard is a function of productivity and recent levels of consumption per capita. Specifically,

$$X_t = C_{t-1}^{\phi_1} G_t^{1-\phi_1} \tag{13}$$

where $0 \leq \phi_1 \leq 1$. The special case with $\phi_1 = 1$ corresponds to the "catching up with the Joneses" formulation in Abel (1990). The living standard is a function of G_t for a couple of reasons. First, a stochastic trend G_t is used rather than a constantly growing trend, say μ_γ^t, because it seems reasonable that the standard of living in a country may grow faster at some times and slower at others. Second, G_t is used rather than total productivity A_t to capture the idea that changes in the living standard depend on permanent changes in productivity. Finally, including trend productivity in addition to lagged consumption is a reduced-form way of capturing a forward-looking aspect of the living standard.

The state space for this model specification is $S_t = \left\{ \widehat{K}_t, \gamma_t^g, z_t, \gamma_t^X, \widehat{A}_t, \widehat{C}_{t-1} \right\}$. The current values of both the transitory shock to productivity z_t and the trend growth rate of productivity γ_t^g are needed to form expectations of future productivity growth. In general, the final three state variables are all necessary for forming expectations of future growth of the living standard γ_{t+1}^X, although for some parameterizations (e.g., $\phi_1 = 0$), a subset of these variables is sufficient.

5.2 Choosing Parameter Values

Many of the parameters are set as they were for the textbook RBC specification (see Section 4.1), while others, namely η and σ, which were previously pinned down by functional-form assumptions, can now be chosen based on estimates from the labor supply literature. To choose the productivity process parameters, I use techniques for estimating parameters of state-space models with regime switches. Finally, I investigate the implications of various functional forms for the living standard by considering alternative values for ϕ_1.

5.2.1 Preference Parameters

The first-order condition for labor input (equation 5) takes the form

$$bN_t^\nu C_t^\eta X_t^{(1-\eta)} = (1 - \tau_N) W_t,$$

which can be written, dropping constant terms, as

$$n_t = \frac{w_t - \eta c_t - (1-\eta) x_t}{\nu}$$

or

$$n_t = \frac{1}{\nu}\widehat{w}_t - \frac{\eta}{\nu}\widehat{c}_t. \tag{14}$$

Recall that lowercase letters denote logged variables and a "hat" denotes the variable has been normalized by X_t; that is, $\widehat{c} = \ln\left(\frac{C}{X}\right)$. From equation (14), one sees that η would govern the relative magnitude of the income elasticity of labor supply to the uncompensated wage elasticity in a static labor supply context. Tables 1 and 2 of Blundell and MaCurdy (1999) report estimates of the uncompensated wage elasticity and income elasticity from a number of studies. The income elasticity is typically substantially smaller in magnitude than the uncompensated wage elasticity, although its relative magnitude varies. η is set to 0.33 in the baseline model, and I will illustrate how alternative values affect the results. $\sigma = 1$ implies a steady-state elasticity of intertemporal substitution of 0.64, within the range of estimates from the literature.

In the baseline model, $\phi_1 = 0$ so that the living standard is only a function of trend productivity. I will also consider a "catching up with the Joneses" living standard, $\phi_1 = 1$, and living standards that depend on both lagged consumption and trend productivity, $0 < \phi_1 < 1$.

5.2.2 Productivity Parameters

The parameters for the productivity distribution are obtained by maximum likelihood estimation. A description of the estimator taken from Kim and Nelson (1999) is provided in the technical appendix. The estimates include $\mu_\gamma = 1.0047$ and $\sigma_\gamma = 0.0022$. These values imply a long-run growth rate of per-capita output of 1.9% per year, and because the distribution for γ^g is truncated at ± 3 standard deviations, a trend growth rate that is bounded between -0.8% and 4.6% per year. The parameters for the Z process, $\rho_z = 0.88$ and $\sigma_z = 0.0099$, imply less unconditional variance of the cyclical shock than in the textbook RBC model. Finally, the estimate for P implies that the trend growth of productivity changes, on average, once every 70 quarters.

The parameters for the baseline "living standard" model are given in Table 6.

Table 6: Parameters for Baseline "Living Standard" Model

σ	β	η	ν	b	ϕ_1	α	δ	μ_γ	σ_γ	P	ρ_z	σ_z
1.0	0.987	0.33	1.5	20	0	0.333	0.015	1.0047	0.0022	0.0144	0.88	0.0099

5.3 Results

I once again conduct the two experiments described in Section 4.1.2. Tables 7 and 8 report average moments of high and lower-frequency fluctuations from 200 model simulations. The baseline "living standard" specification performs noticeably better than the textbook RBC model. Although the correlation between hours worked and output/TFP at lower frequencies is not negative, the difference from high to lower frequencies is substantial. The correlation between TFP and hours worked drops by 0.69 when moving from high to lower frequencies (slightly larger than the change in the data), while $Corr(Y, N)$ drops by 0.55 (less than the change in the data but sizable). That these correlations are still positive at the lower frequencies reflects the fact that I consider shocks only to TFP. Other researchers have included a wide variety of alternative shocks (government spending, monetary shocks, etc.) within the DSGE framework. Adding such realistic shocks would reduce the correlations between TFP and other variables at all frequencies.[15]

[15] The choice of a simple model was made to most clearly illustrate the effects of the living standard. Given this modelling decision, there are clear discrepancies between the moments produced by the model and the data. Some of these can be fixed by considering alternative values for parameters. For example, a higher intertemporal elasticity of the labor supply $\left(\frac{1}{\nu}\right)$ would make labor input more volatile [Hansen (1985), Rogerson (1988)]. Other discrepancies, like the decrease in the correlation of TFP and output from high to lower frequencies, could be fixed by considering a richer model with more shocks.

Table 7: Average Moments of High-Frequency Fluctuations

	Standard Deviation	SD Relative to Y	Correlation with Y	Correlation with A	First-order Autocorrelation
Y	1.05	1	1	0.99	0.65
C	0.45	0.43	0.86	0.85	0.69
I	3.88	3.69	0.96	0.96	0.65
N	0.37	0.35	**0.98**	**0.97**	0.65
Y/N	0.69	0.66	0.99	0.99	0.66
r	0.04	0.03	0.97	0.98	0.65
A	0.81	0.77	0.99	1	0.65

Table 8: Average Moments of Lower-Frequency Fluctuations

	Standard Deviation	SD Relative to Y	Correlation with Y	Correlation with A	First-order Autocorrelation
Y	2.67	1	1	0.93	0.997
C	2.21	0.83	0.93	0.82	0.998
I	5.46	2.04	0.86	0.91	0.996
N	0.48	0.18	**0.43**	**0.27**	0.995
Y/N	2.51	0.94	0.98	0.95	0.998
r	0.06	0.02	0.44	0.67	0.996
A	2.03	0.76	0.93	1	0.997

The second experiment of simulating the model with the actual U.S. Solow residual as the driving force is more complicated than it was for the textbook RBC model. Specifically, it requires a decomposition of productivity growth into trend and transitory components. Alternative decompositions will change the model's predictions through two channels: the living standard and agents' expectations. Because only the trend (and not the transitory) part of productivity growth enters the living standard, any decomposition is an assumption about how the living standard evolved over the postwar period. As for expectations, the higher the trend component, the higher is the expected level of future productivity.

To illustrate how the living standard model works, I show the simulation results for the case of one change in trend productivity. I use Bai and Perron's (1998) methods for identifying trend breaks in univariate time series to provide discipline for the decomposition. At the 10% confidence level, the hypothesis that no trend breaks occurred in the productivity growth series can be rejected in favor of the alternative hypothesis that one break occurred. Moreover, these tests identify the date of the break as the second quarter of 1966. Given the break date, I construct the series for the trend γ^G and transitory z components in a way that imposes orthogonality between the two series (to be consistent with the assumptions on the stochastic processes) and minimizes the sum of the squared transitory shocks. Details are provided in the technical appendix.[16]

[16] I thank John Fernald and his research assistant, Andrew McCallum, for running the Bai-Perron statistical tests on my TFP series. It is difficult to make the case for the correctness of one particular decomposition of productivity growth over another. Although I report the results of my model simulations for the case of one

Figure 9 displays the lower frequencies of the TFP series along with the hours worked and output series from the model and the data. The hours worked series generated by the model is roughly similar to the actual series as the 50s and 90s are times of relatively high labor input. The one major discrepancy between the series is the increase in simulated hours worked during the late 60s and early 70s. The effect of this increase in hours worked can also be seen in the simulated output series in the bottom panel of Figure 9. This discrepancy could be reduced in a couple of ways. The first is by adding more shocks to the model. Figure 10 shows the impact of adding distortionary labor income taxes, which counteract some of the increase in hours worked during the late 60s.

The second way to improve the model's fit would be to consider the impacts of learning in the model. Figure 9 displays hours worked in a model where individuals perfectly recognize a downward shift in the trend rate of productivity in the second quarter of 1966. The impact of such a drop is for hours worked to increase due to the negative effect on permanent income of lower expected future productivity growth. As shown by Edge, Laubach, and Williams (2004), if individuals instead attributed part of this drop in productivity to low temporary shocks, they would not have as much of an incentive to increase hours worked. The technical appendix (8.4) discusses some details for adding learning (i.e., signal extraction) to the model, although results of this exercise are not yet available.

To understand the model's ability to capture the decrease in hours worked at times of fast TFP growth and the importance of the living standard for driving this result, it is useful to consider the response of hours to a permanent shock to the trend growth rate of productivity γ^g. Figure 11 displays the response of hours worked to a one-standard-deviation increase in γ^g for different values of the income elasticity η. For the baseline living standard model ($\eta = 0.33$), hours worked decline in the initial period and continue to decline as the economy transitions to a new steady-state. Thus, hours worked can have an extended decline while output and TFP grow at a fast rate. For a model without a living standard ($\eta = 1$), the initial response of hours is greater due to the stronger income effect, but in the long-run, hours worked actually increase.[17] Thus, without a living standard, a prolonged decrease in hours worked could only be generated by many positive shocks to γ^g hitting in succession. Not only would a string of successive shocks cause consumption to increase more rapidly than output (as discussed in section 4.3), it would also make hours worked negatively correlated with productivity at high frequencies.

In contrast, the living standard model can generate the positive high-frequency correlation

break, note that these tests can be used to check for an arbitrary number of breaks. In fact, the hypothesis of no breaks was also rejected in favor of the alternatives of 3 or 4 breaks as well. Moreover, even conditioning on the number of breaks, Fernald (2005) explains that there is considerable uncertainty about the exact break dates. Thus, this experiment should be seen as illustrating the effects of a living standard, rather than literally arguing for a living standard that grew at one rate before 1966:2 and at another rate afterwards.

[17] I considered a version of the model with the trend-break productivity process but no living standard ($\eta = 1$). The results of simulating the model with actual Solow residuals as the driving force are very similar to the "catching up with the Joneses" model shown in Figure 12. In comparison to the living standard model, hours worked do not decline as much in the 50s or increase as quickly in the 90s. Moreover, the break in trend productivity growth in the second quarter of 1966 causes a stronger increase in labor input. For the monte-carlo exercise, the correlation between hours worked and output (TFP) only drops by 0.27 (0.19) when moving from high to lower frequencies.

between hours worked and TFP/output at the same time as it provides a force for the negative correlation at lower frequencies. Consider the intratemporal equilibrium condition

$$n_t = \frac{w_t - \eta c_t - (1-\eta)\, x_t}{\nu} = \frac{(\eta-\alpha)g_t + (1-\alpha)z_t + \alpha k_t - \eta c_t}{\nu+\alpha}, \tag{15}$$

where the second equality follows from $w = MPL$ and $x = g$. The z shocks will drive much of the high-frequency correlation but will be less important at the lower frequencies because of their temporary nature. This equation also provides another way of framing the discussion of the previous paragraph. Abstracting from temporary shocks to productivity ($z = 0$), effective units of capital ($k - g$) are more responsive to a change in the growth rate of productivity γ^g than are effective units of consumption ($c - g$). (This must be true for hours worked to decline even when $\eta = \alpha = 0.33$.) For smaller values of η, the long-run decline in hours worked associated with an increase in the growth rate of productivity will be even greater.

5.4 Alternative "Living Standard" Formulations

The living standard in the baseline specification of the model is simply the trend part of productivity:

$$x_t = g_t = \sum_{s=1}^{t} \gamma_s^g,$$

but alternative formulations of the living standard can also be considered. This is useful for comparing the model to other 'habit formation' models in the business cycle literature and for demonstrating the key features of the living standard for explaining the lower-frequency puzzle. I first describe various alternatives and then present some summary statistics from simulations of these models.

By setting $\phi_1 = 1$, the living standard is fully backwards-looking, a function only of lagged aggregate consumption:

$$x_t = c_{t-1}.$$

This is the "catching up with the Joneses" specification considered by Abel (1990). More generally, the living standard could be a weighted average of trend growth and past levels of consumption, with ϕ_1 denoting the relative weight of lagged consumption:

$$x_t = \phi_1 c_{t-1} + (1-\phi_1) \sum_{s=1}^{t} \gamma_s^g.$$

To capture the idea that individuals may overreact, becoming very optimistic (pessimistic) during times of fast (slow) growth, I allow the living standard to increase more rapidly than the trend growth of TFP during prosperous times and to not grow as quickly during times of relative stagnation. This is done by setting the parameter $\lambda > 1$ for the living standard specification

$$x_t = \lambda \sum_{s=1}^{t} \left(\gamma_s^g\right) + t\left(1-\lambda\right)\ln\left(\mu_\gamma\right).$$

Note that $\lambda = 1$ delivers the baseline version of the living standard, while $\lambda = 0$ implies a constantly growing living standard.[18]

Table 9 reports the *difference* between the average correlation of hours worked with output/TFP at high and lower frequencies for these various formulations of the living standard. The formulations produce similar high-frequency correlations of around 0.90, but the correlations generated at lower-frequencies are quite different. Living standards which grow (on average) at least as quickly (slowly) as productivity during extended periods of fast (slow) productivity growth are best able to generate the drop in the correlation at lower frequencies.

Table 9: High-Frequency minus Lower-Frequency Labor Moments

	ϕ_1	λ	Correlation with Y	Correlation with A
Data			1.01	0.56
Baseline: Stochastic Trend	0	1	0.55	0.68
Consumption and Trend	0.5	1	0.37	0.41
"Catching up with Joneses"	1	1	0.26	0.18
Constant Trend	0	0	0.02	0.01
Accelerated Stochastic Trend	0	1.05	0.61	0.76

Figure 12 shows the lower-frequency fluctuations in the hours worked series that results from using the actual Solow residuals as the driving force for each of these specifications. The specifications with living standards that are mostly a function of the stochastic trend of productivity do best at replicating the movement in hours worked seen in the data. The key is that the "true" MRS, $bN_t^{\nu}C_t^{\eta}X_t^{(1-\eta)}$, grows quickly at times of fast productivity growth. Note that when the living standard simply grows at a constant rate, the model performs even more poorly than the textbook RBC model: the income effect is smaller in the living standard model ($\eta = 0.33$) and the wedge in the MRS (X) does not grow quickly.

[18] Another formulation for the living standard is one that is forward-looking, depending on expectations of future productivity. It could take the form

$$x_t = E_t\left\{\frac{a_{t+j}}{(1-\alpha)}\right\} - j\ln\left(\mu_\gamma\right)$$

where $E_t\left(a_{t+j}\right)$ is the expected productivity level j periods in the future and the other terms are simply for scaling purposes. The living standard could then be expressed as

$$x_t = g_t + f_1\left(j, P, \gamma_t^g - \ln\left(\mu_\gamma\right)\right) + f_2\left(j, \rho_z, z_t\right)$$

where f_1 is an increasing function of j, P, and the trend growth rate γ_t^g, and f_2 goes to 0 as j grows. For j sufficiently large, the transitory shock has little weight in the living standard. This forward-looking specification is similar to the baseline specification, $x_t = g_t$, although changes in the trend growth rate have an additional impact through f_1.

5.5 Alternative Interpretations of the Model

The model considered above takes changes in productivity as the exogenous driving force of the economy and specifies that persistent changes in productivity enter into a living standard. Individuals derive utility not just from the absolute levels of goods and leisure they consume, but from these levels relative to the standard. The high average productivity (wage) growth in the 50s and 60s is thus seen as driving the contemporaneous increase in leisure time, while the slow productivity growth in the last quarter of the 20th century is interpreted as a factor that led to increased market work through the 80s and 90s. As discussed before, the usual income effect of a wage change on labor supply is not strong enough to generate these patterns, and thus, a living standard story is proposed.

As mentioned in the introduction, other macroeconomic models share the mathematical structure of the living standard model, including models with (external) habits and various forms of non-market production. I will briefly discuss why I have chosen to interpret the state variable in preferences as a living standard. First, the distinction between a living standard and an external habit is slight but important. The main difference is that the living standard depends on growth in productivity (which is closely tied to income and wages) and not solely on lagged consumption. The idea of using trend productivity was to capture the idea that there may be a forward-looking aspect of agents' reference level. As shown in the previous subsection, allowing for productivity in the reference level, and not solely lagged consumption, is critical for the model's performance.

A leisure-enhancing production interpretation of the model is similar to the living standard interpretation, although the mechanism that drives the increased desire for leisure in prosperous times is different [Greenwood and Vandenbroucke (2005), Kopecky (2005)]. The increased desire for leisure would result from an increase in the consumption of goods that are complements of leisure time. Examples of such goods include books, radios, televisions, admissions to concerts, and travel. In periods of rapid growth, declines in the relative price of leisure goods, rather than a behavioral story for the increased desire of leisure, could explain the negative correlation between hours worked and output/TFP.

Figure 13 plots the relative price of a basket of leisure goods over the postwar period, as constructed by Kopecky (2005).[19] Periods of fast (slow) productivity growth do not appear to coincide with drops (increases) in this measure of the relative price of leisure goods, making this story less appealing than the living standard explanation. There is, however, an important caveat. The measure of the relative price in Figure 13 does not correct for changes in quality or the variety of leisure goods. If, for example, the quality improvements in leisure goods are greater than those for the average good in the CPI basket at times of fast productivity growth, a corrected relative price series could display the movements needed for this story to work. Pursuing these corrections is left for future research.

Models with home production provide yet another interpretation for the lower-frequency movement in hours worked. If the consumption of home-produced goods increase rapidly during times of fast market productivity growth, the marginal utility of consumption can

[19] I thank Karen Kopecky for making this data available.

be low enough to explain the contemporaneous decrease in market hours worked. In this interpretation, the 50s and 60s were times when individuals consumption of non-market goods grew rapidly, either because they spent more time in homework or because the productivity of homework grew rapidly. Figure 14 shows an estimate of the number of annual hours per capita spent in homework over the postwar period, as constructed by Ramey and Francis (2005).[20] Since homework apparently decreased over the first few decades of the sample period, the home production story would require that home productivity grew extremely fast in this period, much faster than market productivity. It is, of course, possible that the efficiency of home productivity does contemporaneously grow faster than market productivity during times of fast market productivity growth, although one may think that new technologies first impact the market before diffusing to the non-market sector.

6 Concluding Remarks

Long-run economic growth and business cycles have been two of the most active areas of macroeconomic research. This paper focuses on the often overlooked gap in between, lower-frequency fluctuations. One striking fact in U.S. postwar data is that hours worked per capita is negatively correlated with both output per capita and TFP at lower frequencies, while these variables are positively correlated at high frequencies. The hypothesis proposed in this paper is that extended periods of above average growth in productivity lead to rapid output growth and decreases in hours worked, while temporary increases in productivity lead to temporary increases in both output and hours worked. The mechanism that delivers this behavior in a DSGE model is a "living standard" that changes smoothly over time, growing quickly (sluggishly) in times of fast (slow) average productivity growth.

There are a number of directions in which future work on this topic is likely to proceed. As described in the previous section, more work can be done to distinguish between alternative interpretations of the time-varying state variable in preferences that I have, for now, labeled a living standard. Extending the model to include learning about trend breaks in productivity growth may improve the model's performance around breaks and could also be relevant for thinking about other interesting patterns in the data, such as asset-pricing puzzles. Finally, the tools used in this paper can be used for considering other lower-frequency patterns, including the tremendous medium-run variation seen in cross-country growth experience.

7 Data Appendix

7.1 Sample Period

Annual data for the private capital stock is from 1951 to 2001; population data is from 1947:1 to 2003:3; and data for all other variables is from 1947:1 to 2005:1. The band-pass filter used to separate high- and lower-frequency fluctuations is a moving average that depends on three

[20] I thank Valerie Ramey and Neville Francis for sharing their data with me.

years of past data and three years of future data, and thus, observations must be dropped at the beginning and end of the sample. Therefore, the moments reported in Tables 1 and 2 are from 1954:1 to 1998:1, whereas the figures display time series from 1951:1 to 2001:1.

7.2 Measures

- Per-capita output (Y)
  - Real GDP divided by population 16-64
- Per-capita consumption (C)
  - Real personal consumption expenditures divided by population 16-64
- Per-capita investment (I)
  - Real gross private domestic investment (fixed plus inventories) divided by population 16-64
- Per-capita hours worked (N)
  - Hours of all persons (business sector) divided by population 16-64
- Per-capita private capital stock (K)
  - Real net stock of private produced assets (fixed assets plus inventories) divided by population 16-64
- Total Factor Productivity (A)
  - Constructed from per-capita output, hours worked and private capital stock series using a time-invariant capital share of 0.333
- Labor Productivity (Y/N)
  - Per-capita output divided by per-capita hours worked
- (Net) Return to capital (r)
  - Capital share of income multiplied by the output to capital stock ratio minus the depreciation rate

All data series are converted from nominal to real units using the GDP deflator.

7.3 Sources

The specific sources for the data listed above are as follows:

National account, post-1947, quarterly

www.bea.gov, NIPA Table 1.1.5 (in bil. $)

Gross domestic product (GDP)

Personal consumption expenditures (PCE)

Gross private domestic investment (GDPI)

www.bea.gov, NIPA Table 1.1.6 (in bil. chained 2000 $)

Gross domestic product (GDP)

www.bea.gov, NIPA Table 5.9 (in bil. $) (annually, 1951 - 2001)

Private Produced Assets (Fixed Assets and Inventories)

www.bea.gov, NIPA Table 1.1.10 (in bil. $)

Private consumption of fixed capital (i.e., depreciation)

Net operating surplus of private enterprises

Gross domestic income (GDI)

www.bea.gov, NIPA Table 1.12 (in bil. $)

Proprietor's Income with IVA and CCAdj

Business current transfer payments (net)

Population, post-1947, quarterly

Economic Report of President, Table B-34

Hours worked, post-1947, quarterly

Global Insights Basic Economics Database, series LBMN (Total Hours Worked in Business Sector)

7.4 Construction of Capital Stock and Returns to Capital

The quarterly capital stock is constructed as follows: the private capital stock in the first quarter of each year is taken from NIPA Table 5.9 and converted to real terms using the GDP deflator. Gross private fixed investment and private consumption of fixed capital (both in real terns at quarterly rates) are used to construct the capital stock for quarters 2-4 and quarter 1 of the next year. The discrepancy between the constructed quarter 1 capital stock and the quarter 1 capital stock taken from Table 5.9 is then equally divided between the 4 quarters of the previous year.

The (net) return to capital is capital's share of income multiplied by the ratio of output to capital stock minus the depreciation rate $\left(\alpha\frac{Y}{K}-\delta\right)$. Capital's share of income is constructed as

$$\alpha = \frac{\text{Priv net operating surplus + Cons of fixed cap - Proprietor's inc. - Bus. trans. payments}}{\text{GDI - Proprietor's income - Business transfer payments}}.$$

8 Technical Appendix

Throughout this appendix, any notation not explicitly defined is consistent with the notation used in the body of the paper.

8.1 Solution Method

To solve the model, I use a policy-function iteration approach (for example, see sections 17.5-17.9 of Judd (1998)) that operates directly on the necessary and sufficient conditions for an equilibrium. In theory, the equilibrium can be described as functions, $\widehat{C}(s)$, $N(s)$, $\widehat{K'}(s)$, that satisfy the equilibrium conditions, equations (4)-(6), at all points in the state space, $s \in S$. Policy-function iteration is implicitly defined by

$$u_C\left(\widehat{C}^{j+1}(s), N^{j+1}(s)\right) = \beta E\left\{\frac{h\left(\gamma^{X^+}\right)}{\gamma^{X^+}} u_C\left(\widehat{C}^j(s^+), N^j\left(s^+\right)\right)\left[\left(\widehat{A}^+\right)^{1-\alpha} F_K\left(\widehat{K}^+, N^j\left(s^+\right)\right) + 1 - \delta\right]\right\},$$

$$-\frac{u_N\left(\widehat{C}^{j+1}(s), N^{j+1}(s)\right)}{u_C\left(\widehat{C}^{j+1}(s), N^{j+1}(s)\right)} = \widehat{A}^{1-\alpha} F_N\left(\widehat{K}, N^{j+1}(s)\right),$$

$$\widehat{C}^{j+1}(s) + \widehat{K'}^{j+1}(s) - (1-\delta)\widehat{K} = \widehat{A}^{1-\alpha} F\left(\widehat{K}, N^{j+1}(s)\right).$$

Given the functions $\widehat{C}^j$, N^j, $\widehat{K'}^j$ and some $s \equiv (\widehat{K}, z, \gamma^G, \widehat{A}, \gamma^X)$, I solve for the values of $\widehat{C}^{j+1}(s)$, $N^{j+1}(s)$, $\widehat{K'}^{j+1}(s)$ that solve the above system; since this can be done for each s, I have the functions $\widehat{C}^{j+1}$, N^{j+1}, $\widehat{K'}^{j+1}$. The procedure is repeated until convergence.

To implement this procedure on the computer, one needs to choose how to approximate the policy functions and how to approximate the expectation in the intertemporal euler equation. Because the state space is multi-dimensional, approximating the policy function could require solving the equilibrium conditions at a large number of points in the state space. Fortunately, Smolyak's algorithm provides an efficient way of approximating smooth functions of multiple variables. The algorithm consists of both a specification of the grid of points, $\mathcal{H} \in S$, at which to solve for the policy functions and the procedure for interpolating between these points. I use complete polynomials of degree 4 to approximate the equilibrium. For a more detailed description of Smolyak's method, please see Krueger and Kubler (2004), especially sections 3.3 - 3.5. For approximating the expectation in the intertemporal euler equation,I use the appropriate quadrature formulas from chapter 7 of Judd (1998).

To assess the quality of the solution, I compute relative Euler equation errors as in Judd (1992). For each of the three equilibrium conditions, I compute

$$err_t = \left| \frac{(\text{Right-Hand Side})_t}{(\text{Left-Hand Side})_t} - 1 \right|.$$

I simulate the economy for 10,000 periods and record the maximal and average error along the simulated path. For the intertemporal euler equation and resource constraint, the maximal error is typically on the order of 10^{-3} and the average error on the order of 10^{-4}. The intratemporal euler equation can actually be solved exactly for $N(s)$ given the solution for $\widehat{C}(s)$ and $\widehat{K'}(s)$. Thus, by construction, the error for this equation is 0.

8.2 Markov Chain for Persistent Shocks

In order to facilitate the choice of parameters (discussed in the next subsection) and to allow for a 'learning' extension of the model, the stochastic process for trend productivity growth γ^G in the "living standard" model is characterized by the following time-invariant Markov chain:

$$\overline{\gamma} = \begin{bmatrix} \overline{\gamma}^1 \\ \vdots \\ \overline{\gamma}^M \end{bmatrix}, \ \pi_\gamma = \begin{bmatrix} \pi_\gamma^1 \\ \vdots \\ \pi_\gamma^M \end{bmatrix}$$

$$p = \begin{bmatrix} 1-P & \frac{P\pi_\gamma^2}{1-\pi_\gamma^1} & \cdots & \frac{P\pi_\gamma^M}{1-\pi_\gamma^1} \\ \frac{P\pi_\gamma^1}{1-\pi_\gamma^2} & 1-P & & \\ \vdots & & \ddots & \\ \frac{P\pi_\gamma^1}{1-\pi_\gamma^M} & & & 1-P \end{bmatrix}$$

where $\overline{\gamma} \in R^M$ records the possible values for γ^G, π_γ records the unconditional probability of being in each state, and p is a transition matrix which records the probabilities of moving from one value of γ^G to another in one period. $\overline{\gamma}$ and π_γ are chosen so that the distribution from which trend productivity growth is drawn has approximately a normal distribution with mean μ_γ and variance σ_γ^2. Following Tauchen (1986),

$$\overline{\gamma}^m = \mu_\gamma + \left(\frac{2(m-1)}{M-1} - 1 \right) * Trunc * \sigma_\gamma, \; m = 1, ..., M.$$

$Trunc = 3$ and $M = 13$ for all specifications of the model considered in this paper. Thus, there are 13 possible values for γ^G that are equispaced over the interval $\left[\mu_\gamma - 3\sigma_\gamma, \mu_\gamma + 3\sigma_\gamma\right]$. Let $w = \overline{\gamma}^m - \overline{\gamma}^{m-1}$ and $\Phi(\cdot)$ be the standard normal cumulative distribution function. Then, if m is between 2 and $M-1$, set

$$\pi_\gamma^m = \Phi\left(\frac{\overline{\gamma}^m - \mu_\gamma + w/2}{\sigma_\gamma} \right) - \Phi\left(\frac{\overline{\gamma}^m - \mu_\gamma - w/2}{\sigma_\gamma} \right),$$

otherwise,

$$\pi_\gamma^1 = \Phi\left(\frac{\overline{\gamma}^1 - \mu_\gamma + w/2}{\sigma_\gamma} \right) \text{ and } \pi_\gamma^M = 1 - \Phi\left(\frac{\overline{\gamma}^M - \mu_\gamma - w/2}{\sigma_\gamma} \right).$$

8.3 Estimating Parameters of Persistent Shock Process

To estimate the parameters of the productivity process of the "living standard" model, I use the maximum likelihood estimation procedure described by Kim and Nelson (1999), which is a summary of the methods developed in Kim (1994). I first represent the productivity process in state-space form, then describe the Kim filter (an extended version of the Kalman filter that allows for regime switches), and finally detail the approximate maximum likelihood estimation of the parameters of the productivity process.

8.3.1 State-space Form

The productivity process can be represented as a state-space model with regime switching. There are M regimes: one for each possible value of γ^G.

Letting $\gamma_t^A = \ln\left(\frac{A_t}{A_{t-1}}\right)$, a state-space representation of the growth rate of TFP can be written compactly in the form:

$$\begin{aligned} \text{Measurement Equation:} \quad & \gamma_t^A = \mathbf{x}'\mathbf{a}_t + B_{S_t} \\ \text{Transition Equation:} \quad & \mathbf{a}_t = \mathbf{T}\mathbf{a}_{t-1} + \mathbf{R}\eta_t \\ \text{Shocks:} \quad & \eta_t \backsim N(0, Q) \end{aligned}$$

with

$$\begin{aligned}
\mathbf{a}_t &= \begin{bmatrix} z_t \\ z_{t-1} \end{bmatrix}, \quad \mathbf{x} = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \\
\mathbf{B}_{S_t} &= \ln\left(\overline{\gamma}_{S_t}\right) \text{ for } S_t = 1, ..., M \\
\mathbf{T} &= \begin{bmatrix} \rho_z & 0 \\ 1 & 0 \end{bmatrix}, \ \mathbf{R} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \text{ and } \eta = \varepsilon_z.
\end{aligned}$$

The variance-covariance matrix of the underlying shocks is

$$Q = \sigma_z^2.$$

Transition probabilities for the Markov-switching variable S_t are given by

$$p \equiv \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1M} \\ p_{21} & & & \\ \vdots & & \ddots & \\ p_{M1} & & & p_{MM} \end{bmatrix} = \begin{bmatrix} 1-P & \frac{P\pi_\gamma^2}{1-\pi_\gamma^1} & \cdots & \frac{P\pi_\gamma^M}{1-\pi_\gamma^1} \\ \frac{P\pi_\gamma^1}{1-\pi_\gamma^2} & 1-P & & \\ \vdots & & \ddots & \\ \frac{P\pi_\gamma^1}{1-\pi_\gamma^M} & & & 1-P \end{bmatrix}$$

where $p_{ij} = \Pr[S_t = j | S_{t-1} = i]$.

8.3.2 Kim's (1994) filter for state-space models with Markov-switching (Kim and Nelson, section 5.2)

The filter for the state-space model with Markov switching is a combination of extended versions of the Kalman filter and the Hamilton filter, along with appropriate approximations. Here I describe the filter for the state-space representation discussed above.

First, I will define some notation. Let ψ_{t-1} denote the vector of observations available as of time $t-1$, $\psi_{t-1} = \{\gamma_1^A, ..., \gamma_{t-1}^A\}$. The goal is to form a forecast of the unobserved state vector $\mathbf{a}_t$ based not just on ψ_{t-1} but also conditional on the random variable S_t taking on the value j and on S_{t-1} taking on the value i:

$$\mathbf{a}_{t|t-1}^{(i,j)} = E[\mathbf{a}_t | \psi_{t-1}, S_t = j, S_{t-1} = i].$$

The proposed algorithm calculates a battery of M^2 such forecasts for each date t, corresponding to every possible value for i and j. Associated with these forecasts are M^2 different mean squared error matrices:

$$\mathbf{P}_{t|t-1}^{(i,j)} = E[(\mathbf{a}_t - \mathbf{a}_{t|t-1})(\mathbf{a}_t - \mathbf{a}_{t|t-1})' | \psi_{t-1}, S_t = j, S_{t-1} = i].$$

The key part of the algorithm is to reduce the $(M \times M)$ posteriors ($\mathbf{a}_{t|t}^{(i,j)}$ and $\mathbf{P}_{t|t}^{(i,j)}$) into M posteriors ($\mathbf{a}_{t|t}^{j}$ and $\mathbf{P}_{t|t}^{j}$) to complete the Kalman filter described below. Kim's (1994) filter contains the following steps:

1. Run the Kalman filter given in equations (16)-(21) for $i, j = 1, 2, ..., M$.

$$\mathbf{a}_{t|t-1}^{(i,j)} = \mathbf{T}\mathbf{a}_{t-1|t-1}^{i} \tag{16}$$

$$\mathbf{P}_{t|t-1}^{(i,j)} = \mathbf{T}\mathbf{P}_{t-1|t-1}^{i}\mathbf{T}' + \mathbf{R}Q\mathbf{R}' \tag{17}$$

$$\upsilon_{t|t-1}^{(i,j)} = \gamma_t^A - \mathbf{x}'\mathbf{a}_{t|t-1}^{(i,j)} - B_j \tag{18}$$

$$f_{t|t-1}^{(i,j)} = \mathbf{x}'\mathbf{P}_{t|t-1}^{(i,j)}\mathbf{x} \tag{19}$$

$$\mathbf{a}_{t|t}^{(i,j)} = \mathbf{a}_{t|t-1}^{(i,j)} + \mathbf{P}_{t|t-1}^{(i,j)}\mathbf{x}[f_{t|t-1}^{(i,j)}]^{-1}\upsilon_{t|t-1}^{(i,j)} \tag{20}$$

$$P_{t|t}^{(i,j)} = (\mathbf{I} - \mathbf{P}_{t|t-1}^{(i,j)}\mathbf{x}[f_{t|t-1}^{(i,j)}]^{-1}\mathbf{x}')\mathbf{P}_{t|t-1}^{(i,j)} \tag{21}$$

2. Calculate $\Pr[S_t, S_{t-1}|\psi_t]$ and $\Pr[S_t|\psi_t]$, for $i, j = 1, 2, ..., M$.

$$\Pr[S_t = j, S_{t-1} = i|\psi_{t-1}] = \Pr[S_t = j|S_{t-1} = i] * \Pr[S_{t-1} = i|\psi_{t-1}] \tag{22}$$

$$f(\gamma_t^A|S_{t-1} = i, S_t = j, \psi_{t-1}) = (2\pi)^{-\frac{N}{2}}|f_{t|t-1}^{(i,j)}|^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}\upsilon_{t|t-1}^{(i,j)\prime}\left(f_{t|t-1}^{(i,j)}\right)^{-1}\upsilon_{t|t-1}^{(i,j)}\right\} \tag{23}$$

$$f(\gamma_t^A, S_{t-1} = i, S_t = j|\psi_{t-1}) = f(\gamma_t^A|S_{t-1} = i, S_t = j, \psi_{t-1}) * \Pr[S_t = j, S_{t-1} = i|\psi_{t-1}] \tag{24}$$

$$f(\gamma_t^A|\psi_{t-1}) = \sum_{j=1}^{M}\sum_{i=1}^{M} f(\gamma_t^A, S_{t-1} = i, S_t = j|\psi_{t-1}) \tag{25}$$

$$\Pr[S_t = j, S_{t-1} = i|\psi_t] = \frac{f(\gamma_t^A, S_{t-1} = i, S_t = j|\psi_{t-1})}{f(y_t|\psi_{t-1})} \tag{26}$$

$$\Pr[S_t = j|\psi_t] = \sum_{i=1}^{M} \Pr[S_t = j, S_{t-1} = i|\psi_t] \tag{27}$$

3. Using these probability terms, collapse $M \times M$ posteriors in (20) and (21) into $M \times 1$ using the following equations:

$$\mathbf{a}_{t|t}^{j} = \frac{\sum_{i=1}^{M} \Pr[S_t = j, S_{t-1} = i|\psi_t]\mathbf{a}_{t|t}^{(i,j)}}{\Pr[S_t = j|\psi_t]}$$

$$\mathbf{P}_{t|t}^{j} = \frac{\sum_{i=1}^{M} \Pr[S_t = j, S_{t-1} = i|\psi_t]\left\{\mathbf{P}_{t|t}^{(i,j)} + \left(\mathbf{a}_{t|t}^{j} - \mathbf{a}_{t|t}^{(i,j)}\right)\left(\mathbf{a}_{t|t}^{j} - \mathbf{a}_{t|t}^{(i,j)}\right)'\right\}}{\Pr[S_t = j|\psi_t]}$$

4. To start the filter, the following initial values are needed:

$$\mathbf{a}_{0|0}^{j} = \mathbf{0}$$

$$vec\left(\mathbf{P}_{0|0}^{j}\right) = (\mathbf{I} - \mathbf{T} \otimes \mathbf{T})^{-1} vec\,(\mathbf{RQ})$$

$$\Pr(S_0 = j) = \pi_\gamma^j$$

Note that $vec(\mathbf{RQ})$ is not comformable to the other matrices. For this step, in a slight abuse of notation, let $\mathbf{R}$ and $\mathbf{Q}$ be 2×2 matrices, where the last column of $\mathbf{R}$ and last row and last column of $\mathbf{Q}$ are zeros.

8.3.3 Approximate Maximum Likelihood Estimation

The filter computes the density of γ_t^A conditional on past information, $f(\gamma_t^A|\psi_{t-1})$, $t = 1, 2, ..., T$, from equation (25). The approximate log likelihood function is given by

$$LL = \ln\left[f\left(\gamma_1^A, ..., \gamma_T^A\right)\right] = \sum_{t=1}^{T} \ln\left[f\left(\gamma_t^A|\psi_{t-1}\right)\right]. \tag{28}$$

To estimate the parameters of the model, I use a nonlinear optimization procedure to maximize (28) with respect to the underlying unknown parameters, $\left\{\rho_z, \sigma_z, \mu_\gamma, \sigma_\gamma, P\right\}$.

8.4 Decomposing Productivity Growth into Trend and Cycle

8.4.1 Bai-Perron

For a model with both persistent and transitory shocks to productivity growth, feeding the actual U.S. Solow residuals through the model requires a decomposition of productivity growth into the two components. Bai and Perron's (1998) methods for identifying trend breaks in univariate time series are used to identify break dates. Let B be the number of breaks and $t_1, ..., t_B$ denote the dates of the breaks.

Given the logged Solow residuals (a_t) and the break dates $\{t_1, ..., t_B\}$, I construct the series $\{\gamma_t^g\}$ and $\{z_t\}$ by running the following OLS regression with serially correlated errors:

$$a_t = \varsigma + \left[\Gamma_1 D_1 + ... + \Gamma_{B+1} D_{B+1}\right] t + \varepsilon_t,$$

where $D_1 = 1$ for $t \geq 1$, $D_b = 1$ for $t \geq t_{b-1}$, $b = 2, ..., B+1$, and ε_t is an AR(1) process. Then, for $t_{b-1} < t \leq t_b$, $\gamma_t^g = \sum_{i=1}^{b} \Gamma_i$ and $z_t = \varepsilon_t$. This construction imposes the orthogonality of γ^g and z and minimizes the sum of the squared transitory shocks.

8.4.2 Kim Filter

An alternative way to discipline the decomposition of productivity growth into trend and transitory components is to apply the Kim filter as described in section 8.3. If the economic agents in the model know the underlying processes for γ^g and z but only see the realization of productivity growth γ^a, the Kim filter could be used to solve the signal extraction problem.

References

Abel, A. B. (1990): "Asset Prices under Habit Formation and Catching up with the Joneses," *American Economic Review*, 80(2), 38–42.

——— (1999): "Risk Premia and Term Premia in General Equilibrium," *Journal of Monetary Economics*, 43, 3–33.

Bai, J., and P. Perron (1998): "Estimating and testing linear models with multiple structural changes," *Econometrica*, 66, 47–78.

Barro, R., and C. Sahasakul (1983): "Measuring the Average Marginal Tax Rate from the Individual Income Tax," *Journal of Business*, 56(4), 419–452.

Baxter, M., and R. G. King (1999): "Measuring Business Cycles: Approximate Band-Pass Filters for Macroeconomic Time Series," *Review of Economics and Statistics*, 81(4), 575–93.

Benhabib, J., R. Rogerson, and R. Wright (1991): "Homework in macroeconomics: household production and aggregate fluctuations," *Journal of Political Economy*, 99, 1166–1187.

Blundell, R., and T. MaCurdy (1999): "Labor Supply: A Review of Alternative Approaches," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 3, pp. 1559–1695. Elsevier Science, Amsterdam.

Boldrin, M., L. J. Christiano, and J. Fisher (2001): "Habit Persistence, Asset Returns and the Business Cycle," *American Economic Review*, 91, 149–166.

Burns, A. F., and W. C. Mitchell (1946): *Measuring Business Cycles*. National Bureau of Economics Research, New York.

Campbell, J. Y. (1994): "Inspecting the Mechanism: An Analytical Approach to the Stochastic Growth Model," *Journal of Monetary Economics*, 33(3), 463–506.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–251.

Comin, D., and M. Gertler (2004): "Medium-Term Business Cycles," Working Paper, New York University.

Cooley, T. F., and E. C. Prescott (1995): "Economic Growth and Business Cycles," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, pp. 1–38. Princeton University Press, Princeton, NJ.

Fernald, J. (2005): "Trend Breaks, Long Run Restrictions, and the Contractionary Effects of Technology Improvements," Working paper, Federal Reserve Bank of San Francisco.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988): "Investment, Capacity Utilization and the Real Business Cycle," *American Economic Review*, 78(3), 402–417.

Greenwood, J., and G. Vandenbroucke (2005): "Hours Worked: Long-Run Trends," Working paper, University of Rochester.

Hall, R. E. (1997): "Macroeconomic Fluctuations and the Allocation of Time," *Journal of Labor Economics*, 15(1), S223–S250.

Hansen, G. D. (1985): "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 16, 309–327.

Ingram, B. F., N. R. Kocherlakota, and N. Savin (1997): "Using Theory for Measurement: Measuring the Cyclical Behavior of Home Production," *Journal of Monetary Economics*, 40, 435–456.

Judd, K. L. (1992): "Projection Methods for Solving Aggregate Growth Models," *Journal of Economic Theory*, 58, 410–52.

——— (1998): *Numerical Methods in Economics*. MIT Press, Cambridge, MA.

Kaldor, N. (1957): "A model of economic growth," *Economic Journal*, 67, 591–624.

Kim, C.-J. (1994): "Dynamic Linear Models with Markov-Switching," *Journal of Econometrics*, 60, 1–22.

Kim, C.-J., and C. R. Nelson (1999): *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*. MIT Press, Cambridge, MA.

King, R. G., C. I. Plosser, and S. T. Rebelo (1988): "Production, growth, and business cycles: I. The basic neoclassical model," *Journal of Monetary Economics*, 21, 195–232.

King, R. G., and S. T. Rebelo (1999): "Resuscitating Real Business Cycles," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, pp. 927–1007. Elsevier, Amsterdam.

Kopecky, K. (2005): "The Trend in Retirement," Working Paper, University of Rochester.

Krueger, D., and F. Kubler (2004): "Computing Equilibrium in OLG Models with Stochastic Production," *Journal of Economic Dynamics and Control*, 28(7), 1411–1436.

Kydland, F. E., and E. C. Prescott (1982): "Time to build and aggregate fluctuations," *Econometrica*, 50, 1345–1370.

Lettau, M., and H. Uhlig (2000): "Can Habit Formation Be Reconciled with Business Cycle Facts," *Review of Economic Dynamics*, 3, 79–99.

McGrattan, E. R., and R. Rogerson (2004): "Changes in Hours Worked, 1950-2000," *Federal Reserve Bank of Minneapolis Quarterly Review*, 28(1), 14–33.

Prescott, E. C. (1986): "Theory ahead of business-cycle measurement," *Federal Reserve Bank of Minneapolis Quarterly Review*, 10(4), 9–22.

——— (2004): "Why Do Americans Work So Much More Than Europeans," *Federal Reserve Bank of Minneapolis Quarterly Review*, 28(1), 2–13.

Ramey, V. A., and N. Francis (2005): "A Century of Work and Leisure," Working Paper, UCSD.

Rogerson, R. (1988): "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, 21, 3–16.

Tauchen, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20, 177–181.

Figure 1: Frequency Decomposition for GDP

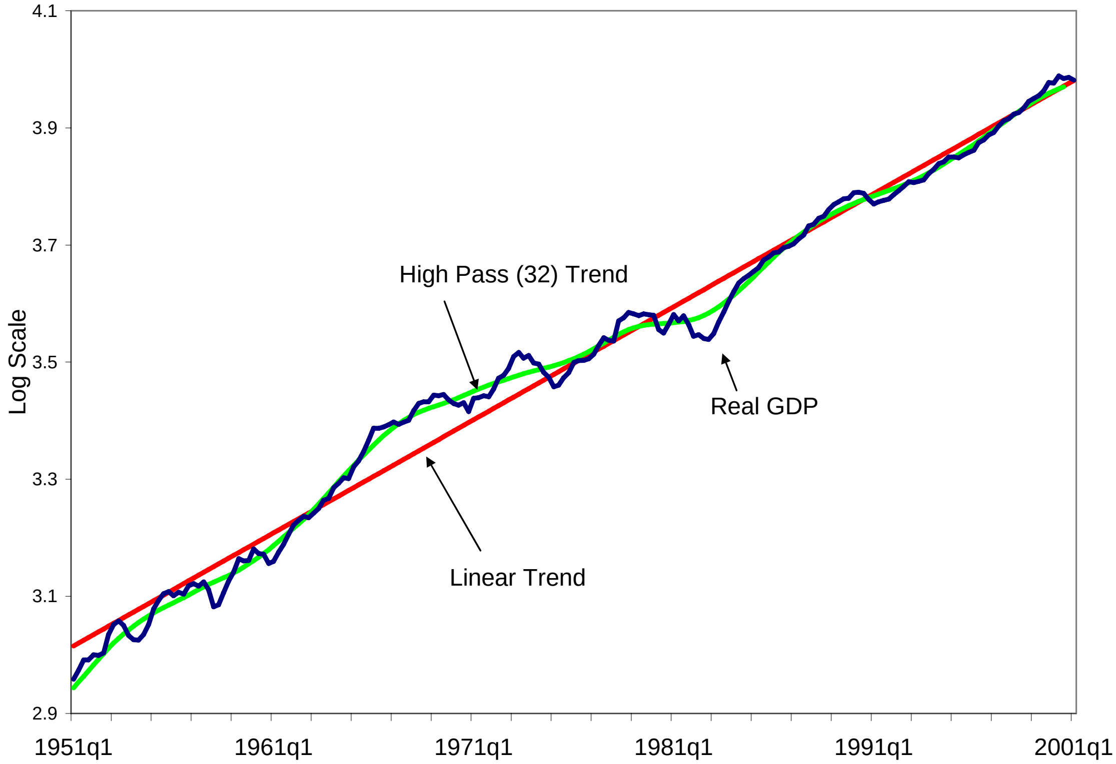

Figure 2: Linearly Detrended Variables

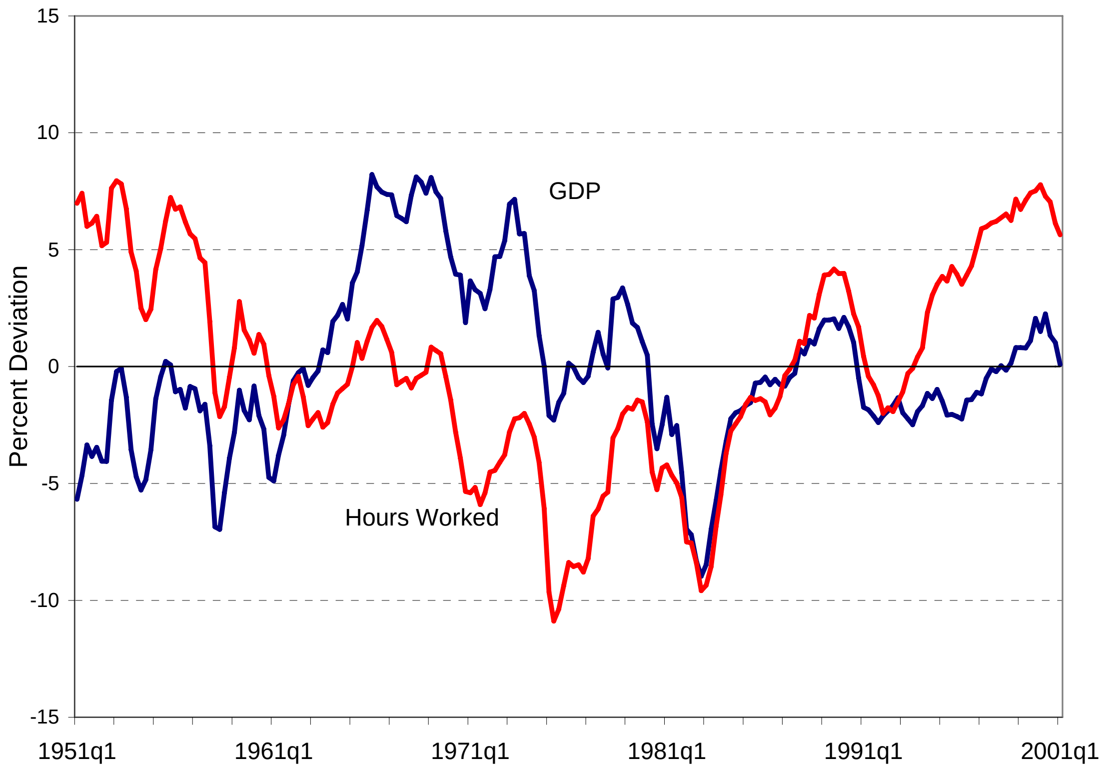

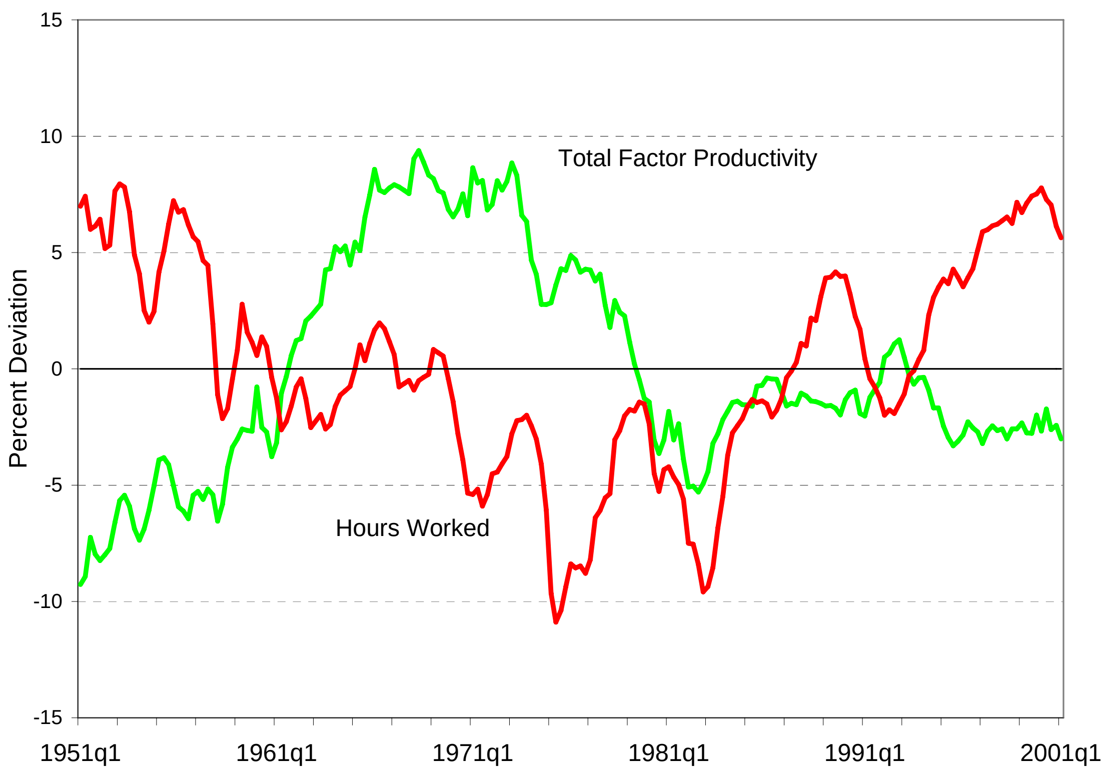

Figure 3: Output and Hours Worked

High Frequencies

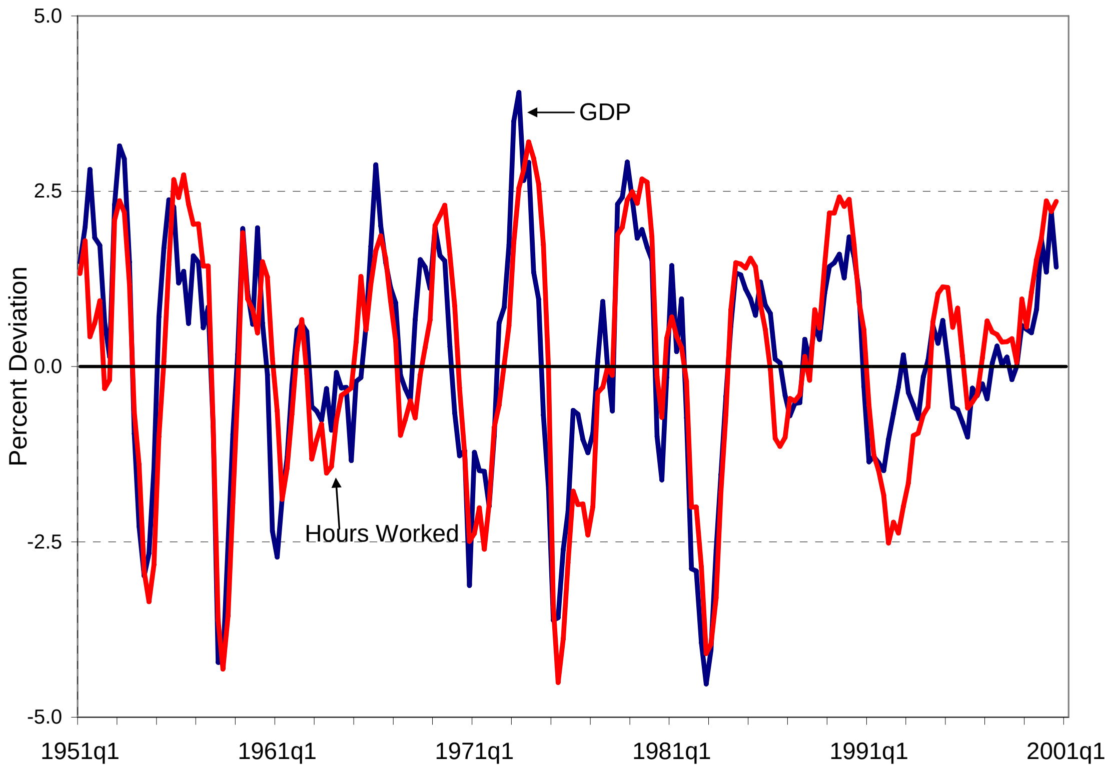

Lower Frequencies

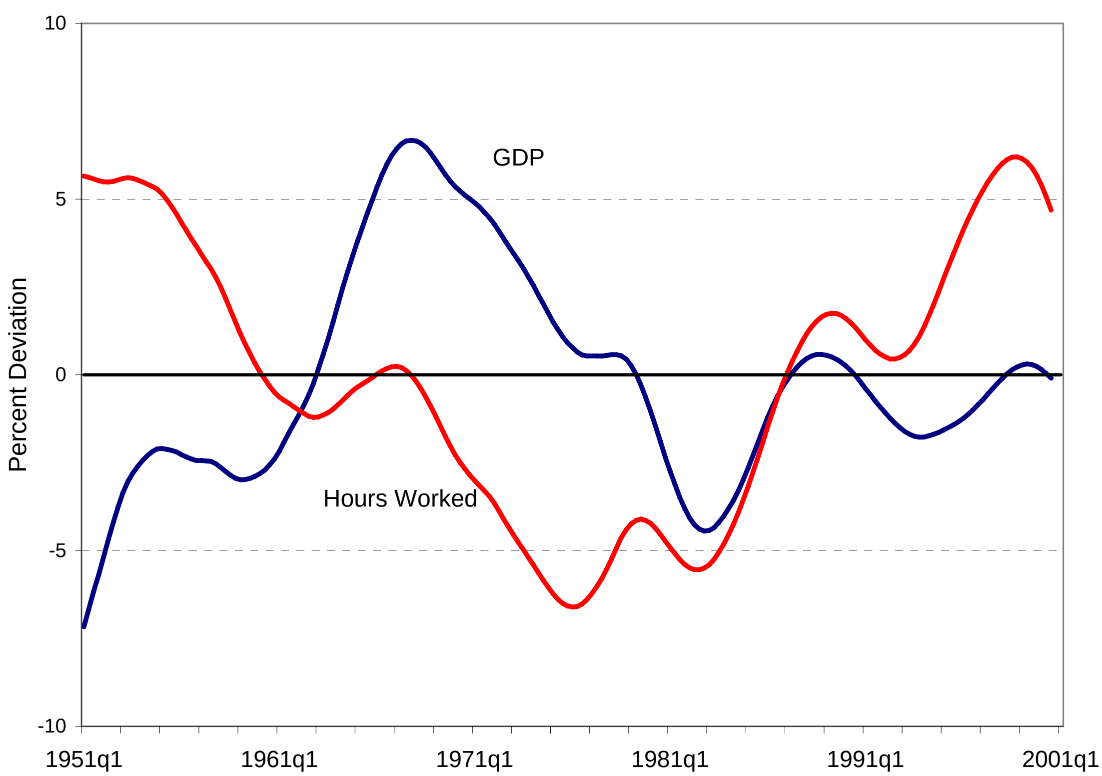

Figure 4: TFP and Hours Worked

High Frequencies

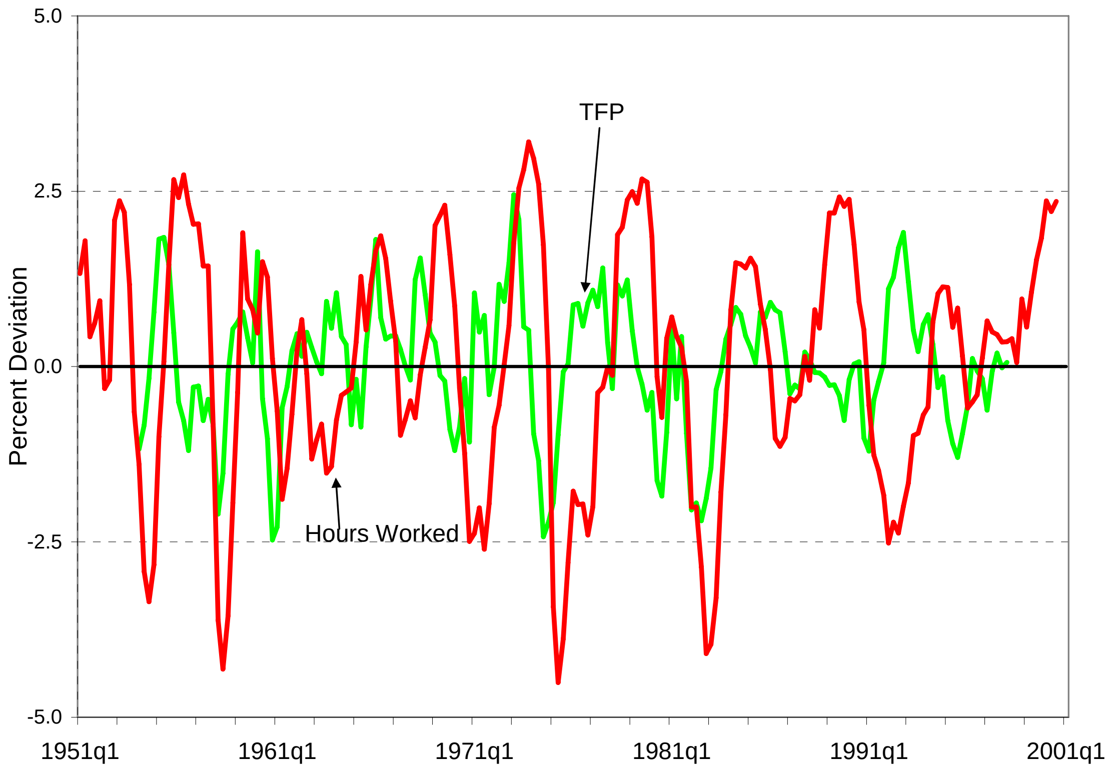

Lower Frequencies

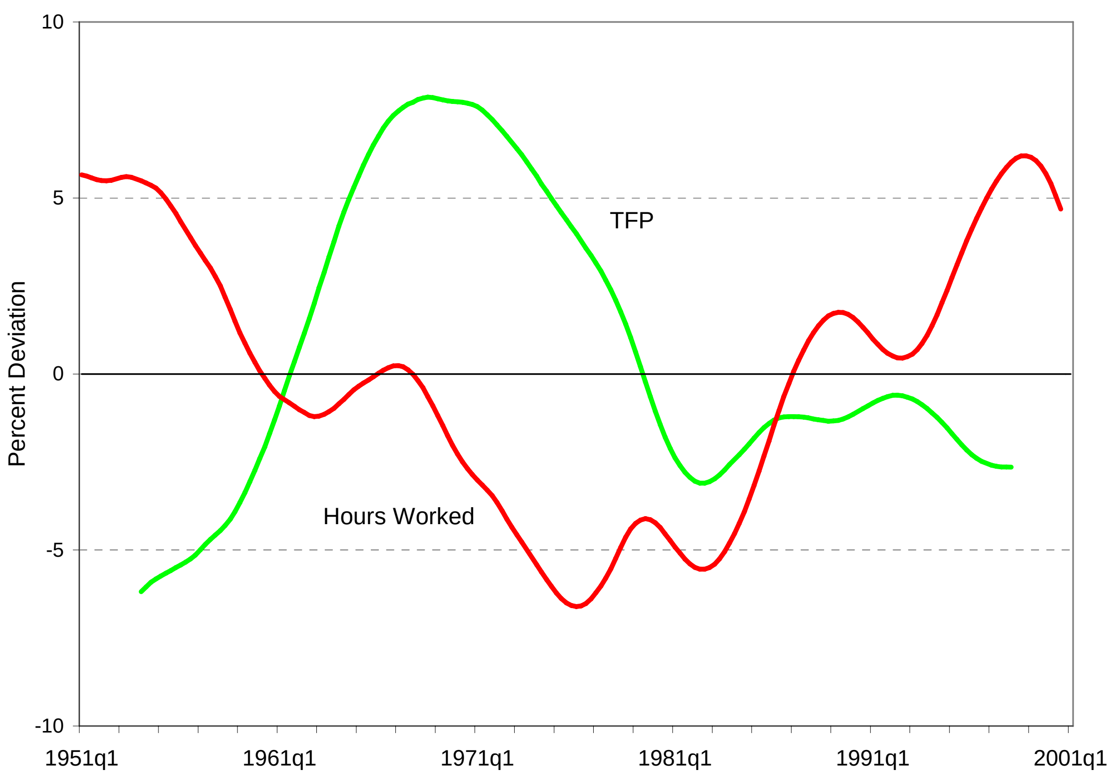

Figure 5: Textbook RBC Model vs. Data
TFP and Hours Worked

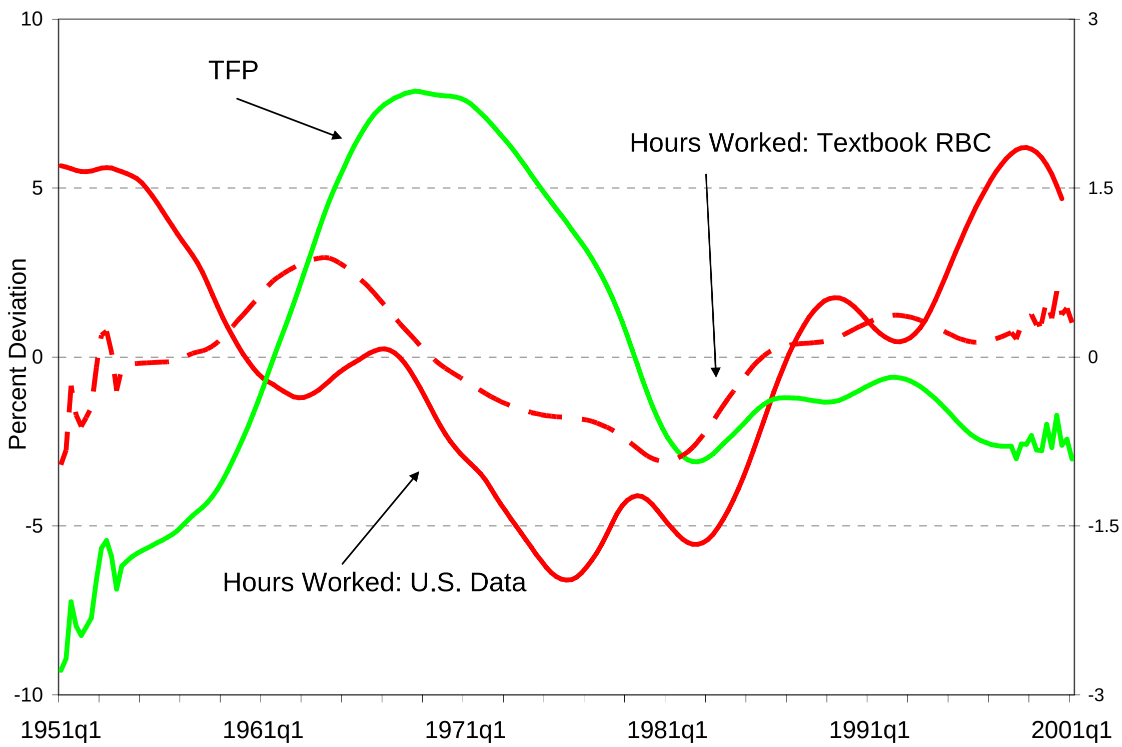

TFP and Output

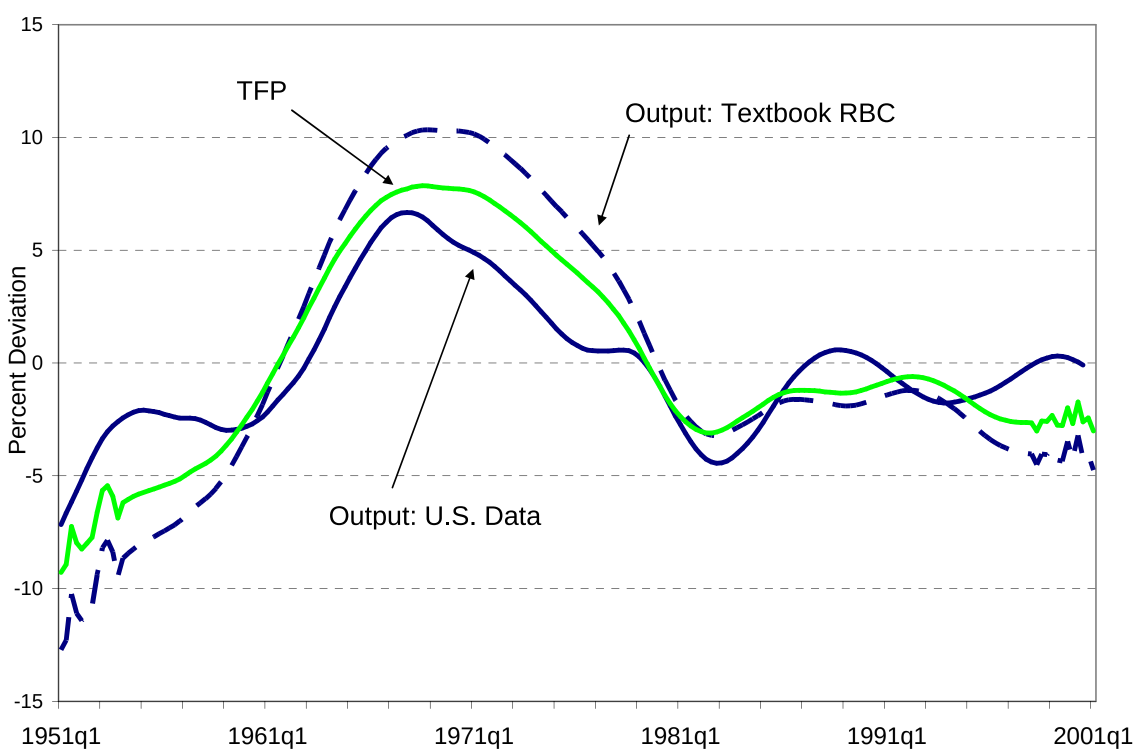

Figure 6: Average Marginal Income Tax Rate

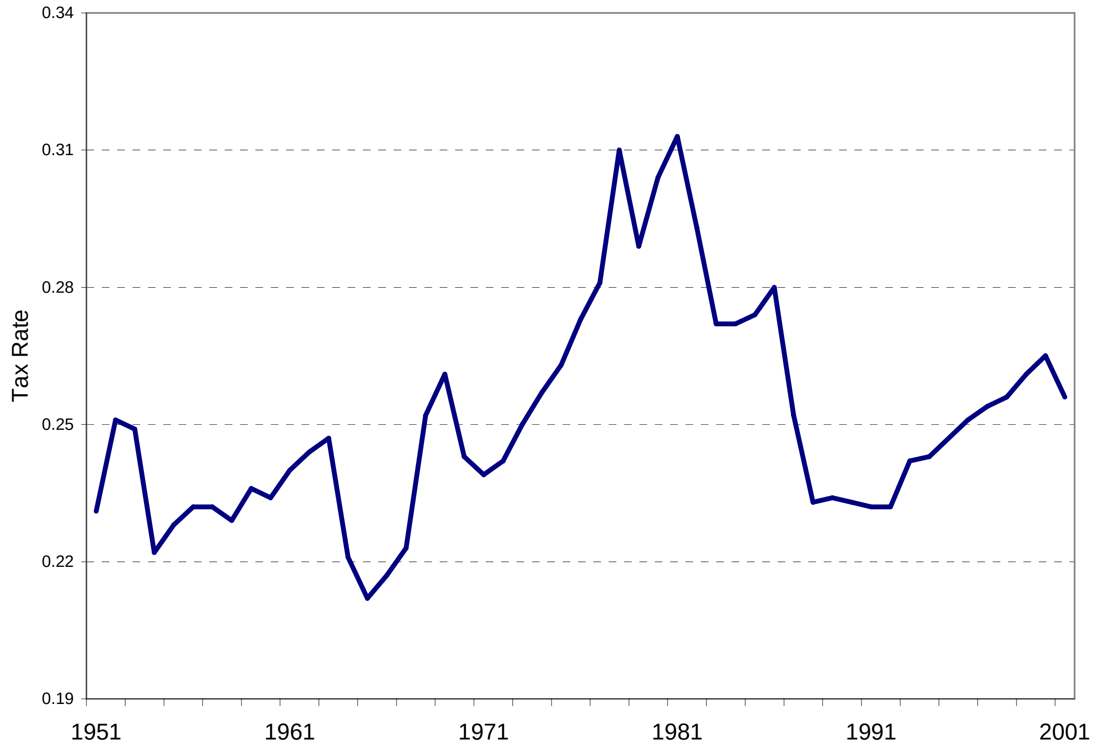

Figure 7: Textbook RBC Model with Distortionary Taxes

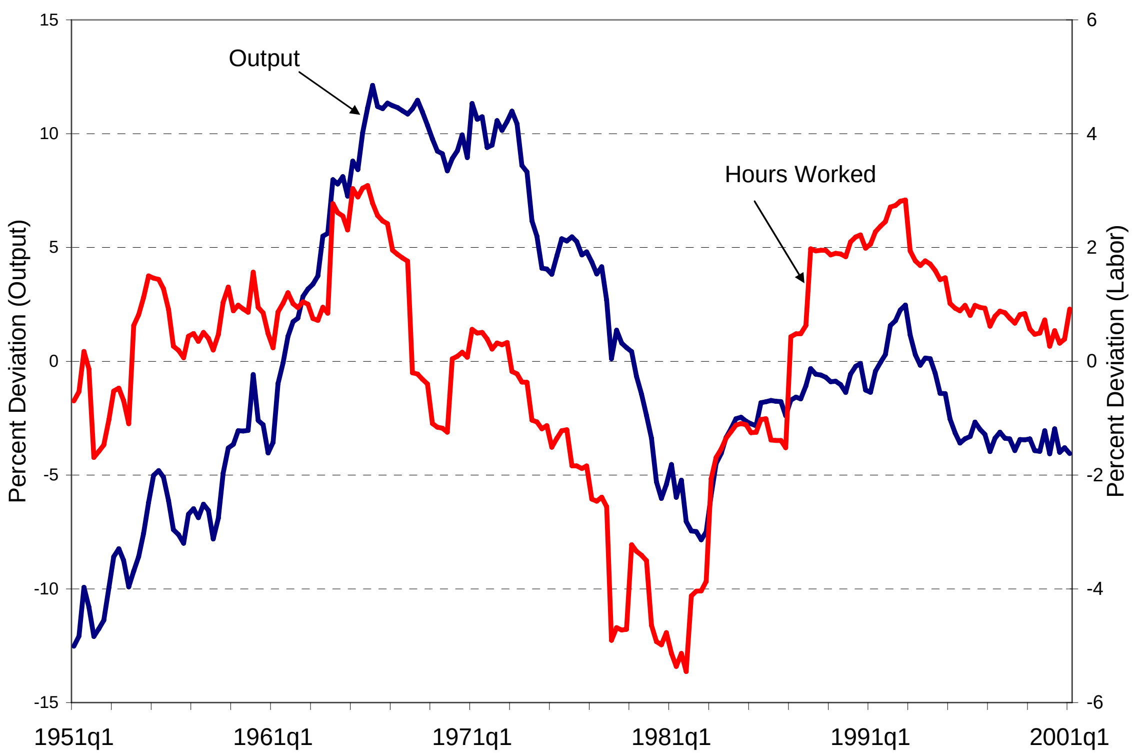

Figure 8: Lower-Frequency GDP and Consumption Share

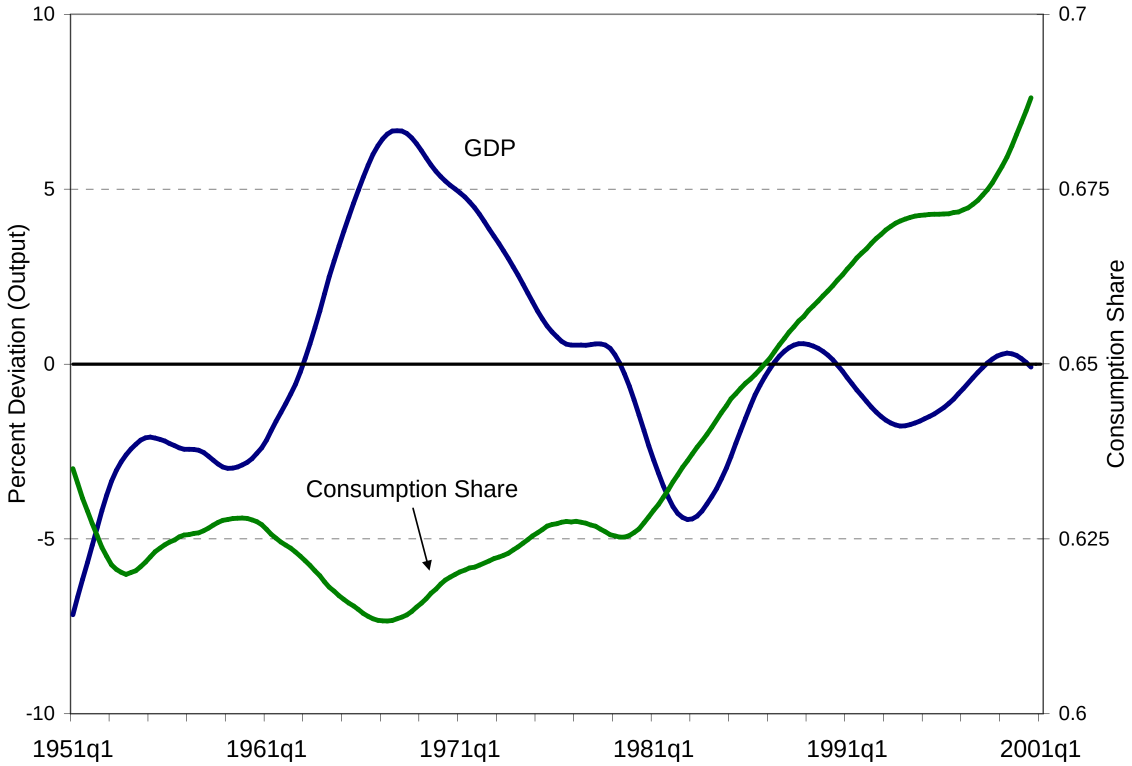

Figure 9: Living Standard Model vs. Data

TFP and Hours Worked

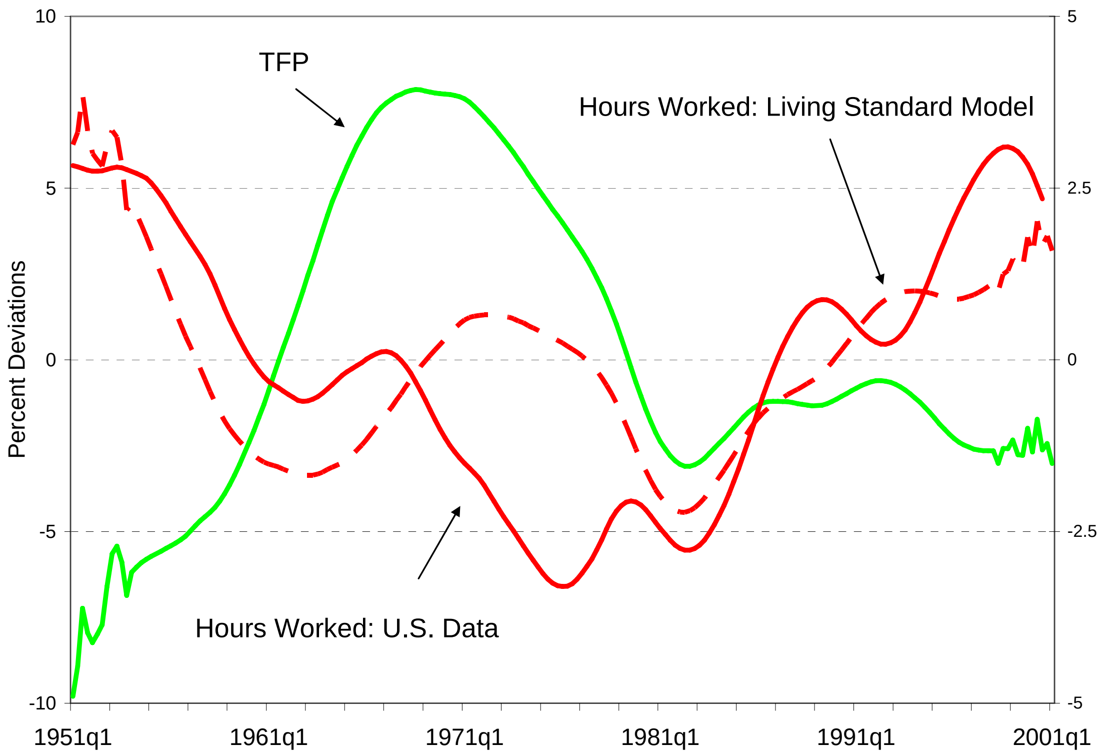

TFP and Output

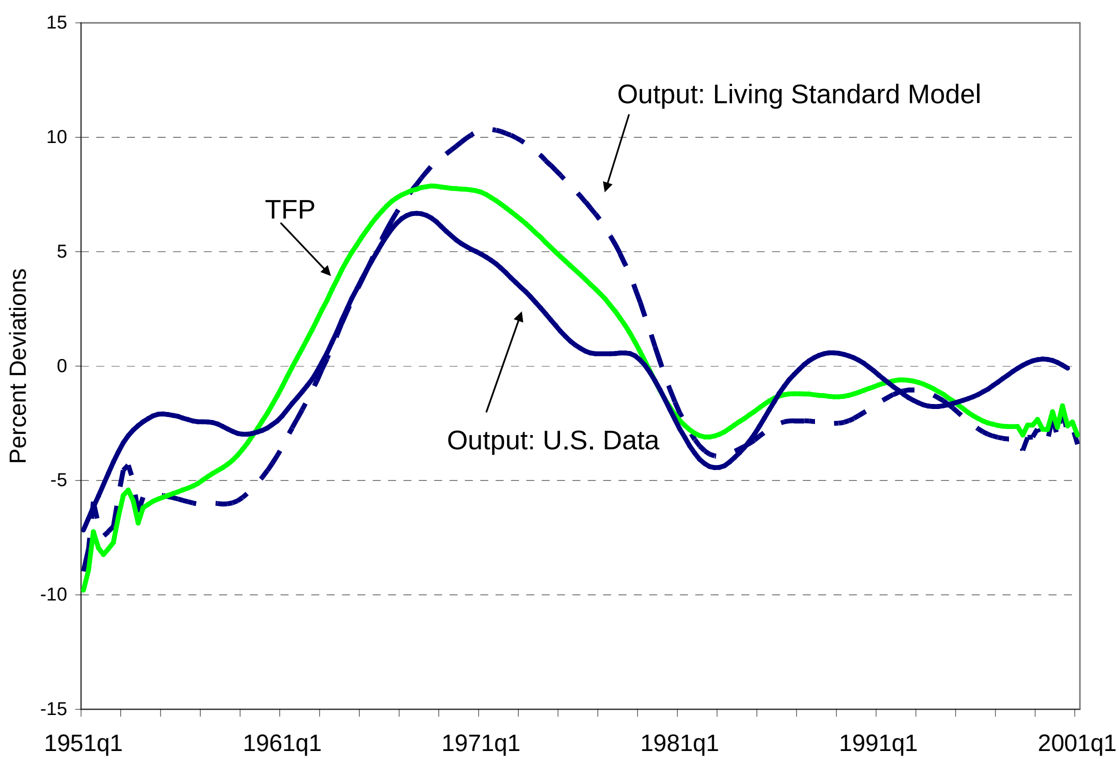

Figure 10: Living Standard Model (w/ taxes) vs. Data
TFP and Hours Worked

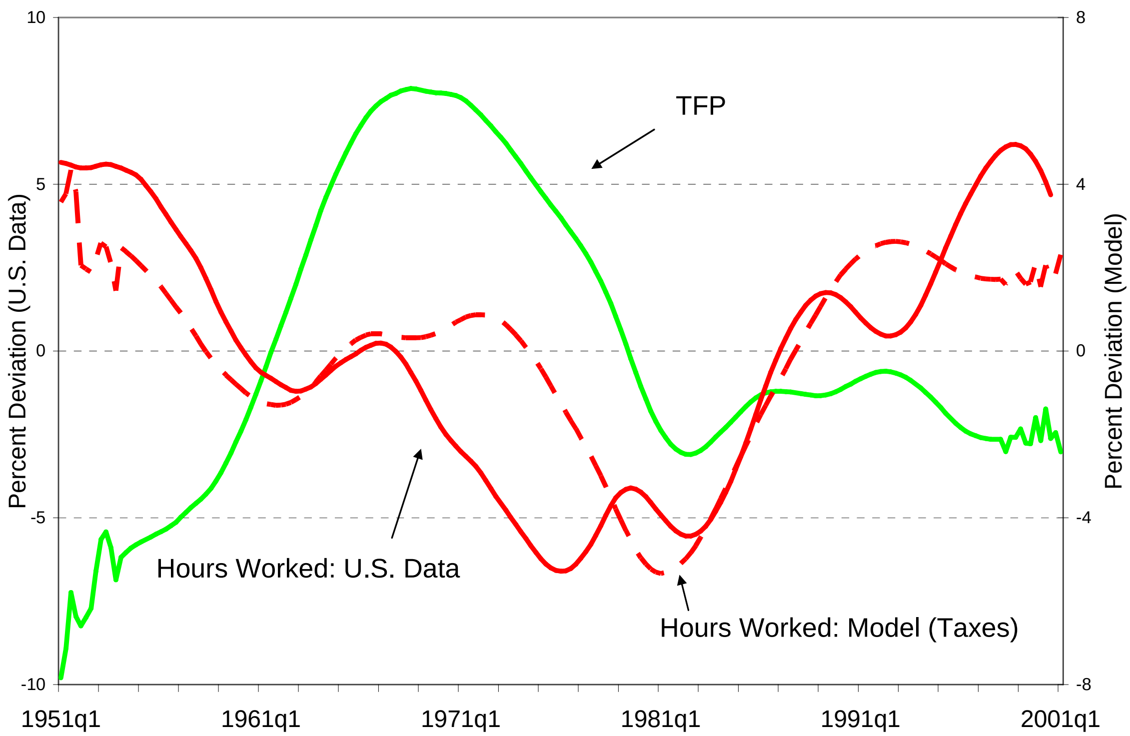

Figure 11: Response of Hours Worked to an increase in trend productivity growth

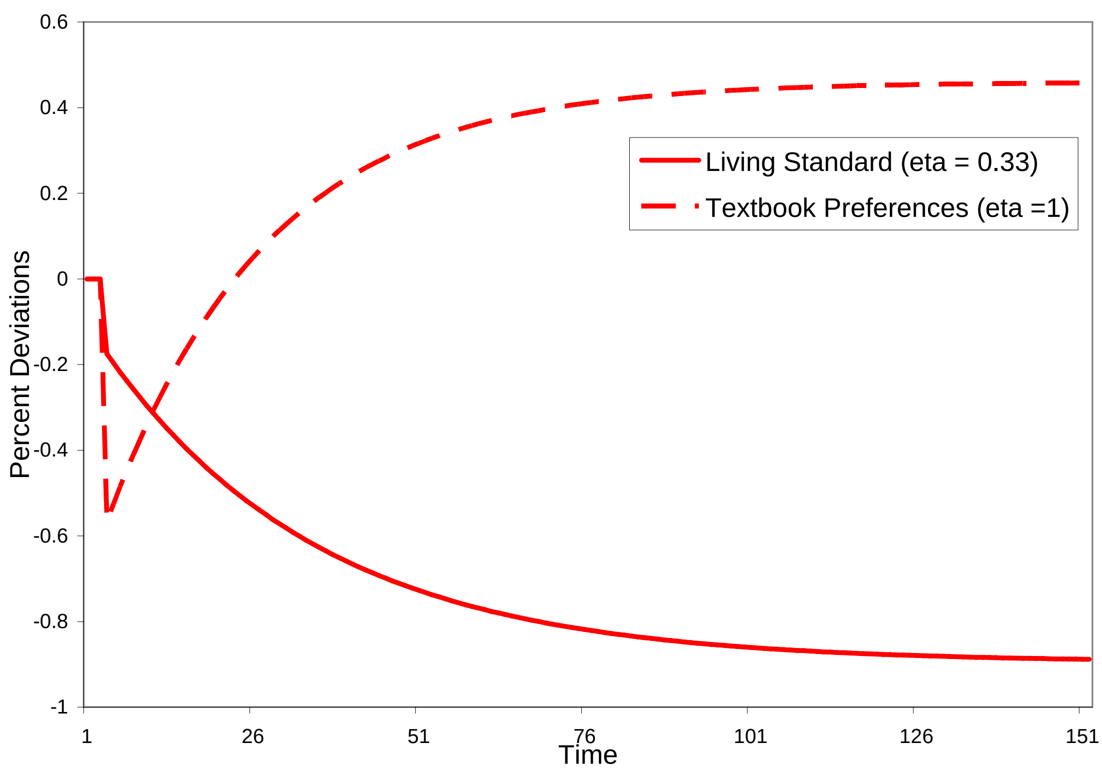

Figure 12: Lower-Frequency Hours Worked for Alternative Formulations of Living Standard

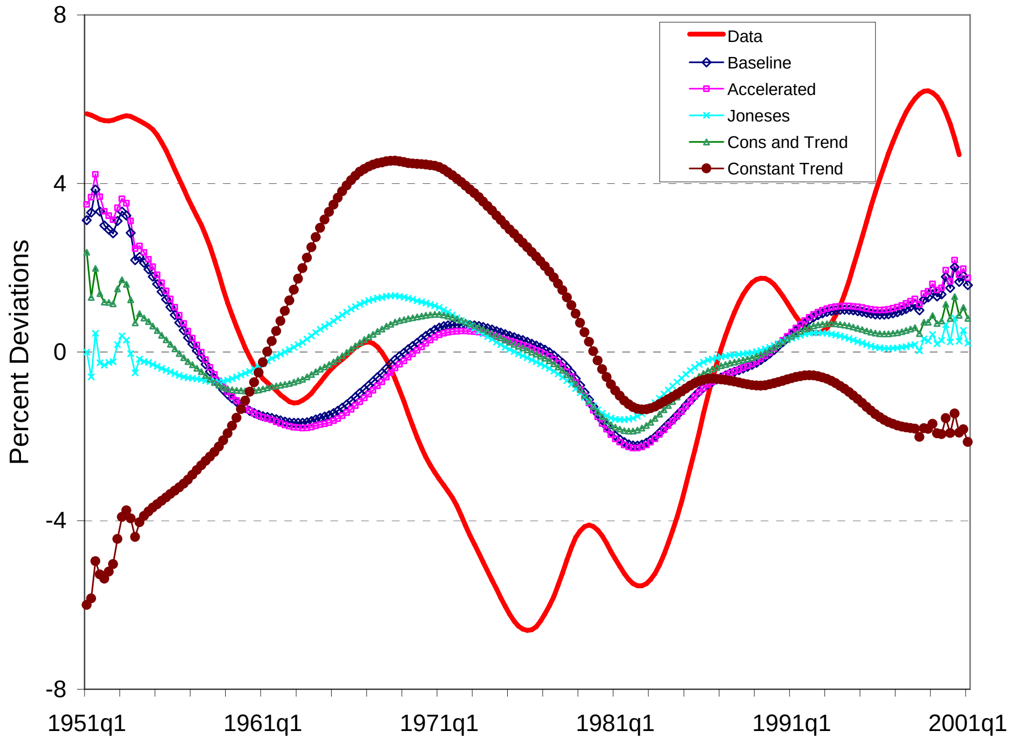

Figure 13: Relative Price of Leisure Goods

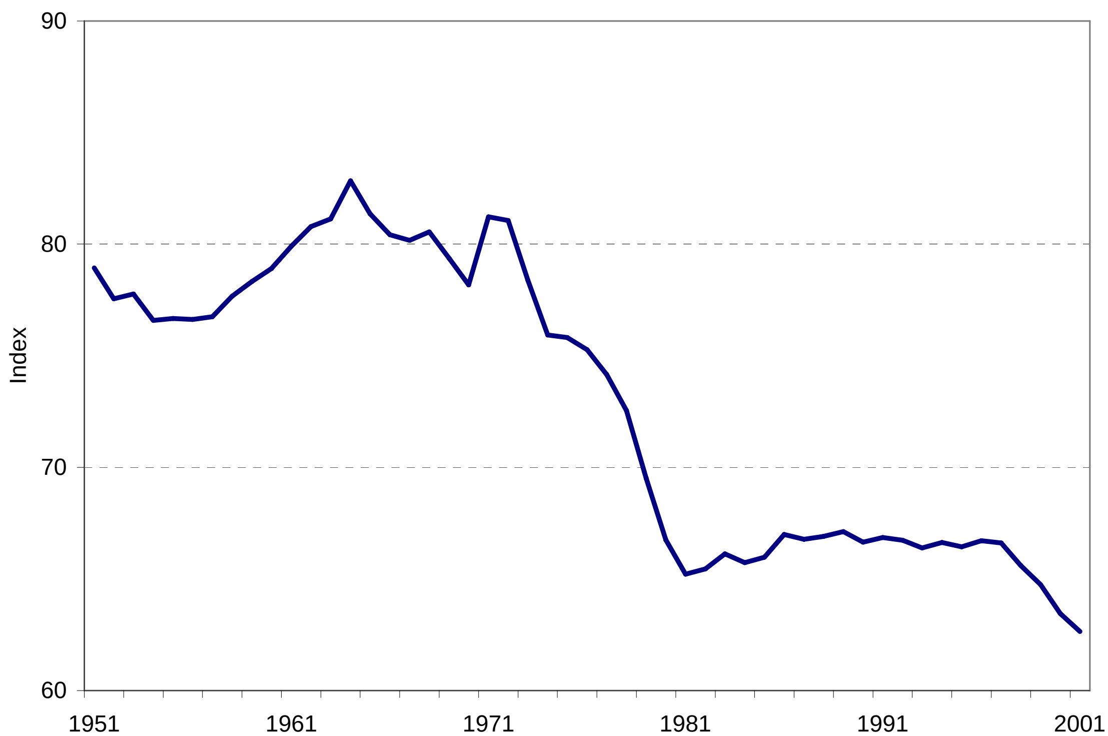

Figure 14: Hours Spent in Home Production