

NBER WORKING PAPER SERIES

MEASURES OF PER CAPITA HOURS
AND THEIR IMPLICATIONS FOR THE
TECHNOLOGY-HOURS DEBATE

Neville Francis
Valerie A. Ramey

Working Paper 11694
<http://www.nber.org/papers/w11694>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2005

The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Neville Francis and Valerie A. Ramey. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Measures of Per Capita Hours and their Implications for the Technology-Hours Debate
Neville Francis and Valerie A. Ramey
NBER Working Paper No. 11694
October 2005
JEL No. E2, E3

ABSTRACT

Structural vector autoregressions give conflicting results on the effects of technology shocks on hours. The results depend crucially on the assumed data generating process for hours per capita. We show that the standard measure of hours per capita has significant low frequency movements that are the source of the conflicting results. HP filtered hours per capita produce results consistent with the those obtained when hours are assumed to have a unit root. We provide an alternative measure of hours per capita that adjusts for low frequency movements in government employment, schooling, and the aging of the population. When the new measure is used to determine the effect of technology shocks on hours using long-run restrictions, both the levels and the difference specifications give the same answer: hours decline in the short-run in response to a positive technology shock.

Neville Francis
Department of Economics
University of North Carolina, Chapel Hill
Gardener Hall, #3305
Chapel Hill, NC 24599
nrfanci@email.unc.edu

Valerie A. Ramey
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
and NBER
vramey@ucsd.edu

I. Introduction

The role of technology shocks in business cycle fluctuations has recently received considerable attention. A myriad of papers has emerged on this topic addressing the controversial conclusion reached by Galí (1999) that *technology shocks cannot be the main driving force behind cyclical movements in macroeconomic data*. This conclusion challenges the core of the long-standing Real Business Cycle (hereafter RBC) theory, thus, it comes as no surprise that so many recent papers have been written either in defense of or to challenge Galí's findings (see Galí and Rabanal (2004) for a review of the literature).

Standard RBC theory teaches that all factor inputs should rise when there is a positive technological innovation. However, recent empirical tests of the theory find that labor input falls in response to a positive shock to technology, a finding which has sparked a debate for the last five and a half years with little resolution. The crux of the debate has to do with the data generating process assumed for per capita labor input in empirical models. If one were to rely on econometrics, which fails to reject the presence of a unit root in per capita labor, one would be led to enter labor input in first differences when estimating a (structural) vector autoregression (VAR). Entered in differences the results of a typical VAR predict a fall in labor input in response to a positive shock to technology, opposite of that predicted by the standard RBC model. However, *common sense* tells us that per capita labor being a bounded series cannot have a unit root. For this reason, several papers have assumed per capita labor is stationary and, thus, should enter the VARs in levels. When entered in levels the standard result emerges that labor input rises when there is a positive innovation to technology.

In this paper, we show that there are significant (low frequency) demographic movements, over the postwar period, that are features of the commonly used measure of

population available to work, the civilian population 16+ and the denominator in per capita labor. These low frequency movements in the standard measure distort unit root tests (which have low power to begin with), make the time series for per capita labor inconsistent over time and with RBC theory, and are the source of conflicting results in the levels versus first difference debate. We begin by showing that extraction of these low frequency movements using an HP filter produces results similar to those obtained using first-differenced hours. We then devise a new measure of hours per capita that adjusts for demographic and institutional changes involving government employment, school enrollment and the aging of the population. The new series is virtually free of low frequency movements and comes closer to its model-generated counterpart than the standard measure used in the literature.

The absence of significant low frequency movements in our new measure has important implications for the technology hours debate. In particular, our new measures provide consistent implications for the role of technology shocks in business cycle fluctuations. Positive technology shocks, identified with long-run restrictions, lead to a short-run decrease in hours worked regardless of the stationary assumption made for per capita labor.

II. The Problem of Low Frequency Movements in the Standard Series of Hours Per Capita

Growth theory and RBC models are generally written in terms of a representative agent's consumption, work, and leisure. To match the representative agent in empirical applications, macroeconomic variables are measured in per capita terms. Theory rarely specifies how "per capita" should be measured, yet virtually all RBC empirical applications measure "per capita" as the BLS series on the civilian noninstitutional population ages 16 and over (e.g. King, Plosser, Stock and Watson (1991), Burnside and Eichenbaum (1996)). The omission of the military, children younger than 16, and the institutionalized population is based on the desire to measure

“per capita” as the available workforce rather than the entire population. To measure hours worked per capita, researchers typically use the BLS index of hours worked in private business or the more narrow index of hours worked in private nonfarm business. There are no quarterly series on hours worked in government.

According to RBC models with standard preference specifications, the hours per capita variable should be stationary in the absence of permanent shifts in government spending, labor income taxes, and preference shifts. Yet the most widely used measure of private hours per capita shows significant low frequency movements. Figure 1 shows the behavior of private hours divided by the civilian noninstitutional population ages 16 and over during the post-WWII period. Hours show a U-shape, with a downward trend until the mid-1970s, which partially reverses by 2004. The low frequency movements are so pronounced that the series does not return to its mean for decades at a time.

While these low frequency features are not an issue for analyses that HP-filter the data before analyzing it, they are very problematic for structural VARs in which assumptions about stationarity are key parts of the identification. In particular, these low frequency movements in hours per capita have important implications for empirical structural VAR models that identify technology shocks using long-run restrictions. Based on the results of standard unit root tests, Francis and Ramey (forthcoming) assume that hours per capita have a unit root, and thus enter hours in first differences in the model. They find that a positive technology shocks leads to a decline in hours worked. In contrast, Christiano, Eichenbaum and Vigfusson (2003) argue that hours per capita cannot logically have a unit root, and offer alternative empirical tests against a unit root. They enter hours in levels and find that a positive technology shock leads to a rise in hours worked.¹

¹ This literature has generated a further controversy about whether these VARs can capture the results from the model. In particular, Chari, Kehoe, and McGrattan (2005) show that they can generate data from a model in which technology shocks have a positive effect on hours yet the VAR shows them to have negative effects. Christiano,

To illustrate, we re-estimate the structural VAR used by Galí, Francis and Ramey, and Christiano, Eichenbaum and Vigfusson using the new measures of hours. In the baseline bivariate case, we estimate the following system:

$$\begin{bmatrix} \Delta x_t \\ n_t \end{bmatrix} = \begin{bmatrix} C^{11}(L) & C^{12}(L) \\ C^{21}(L) & C^{22}(L) \end{bmatrix} \begin{bmatrix} \varepsilon_t^z \\ \varepsilon_t^m \end{bmatrix}$$

x_t denotes the log of labor productivity, n_t denotes the log of hours per capita, ε^z denotes the technology shock, and ε^m denotes the non-technology shock. $C(L)$ is a polynomial in the lag operator. We maintain the usual assumption that ε^z and ε^m are orthogonal. Our assumption identifying the technology shock implies that $C^{12}(1) = 0$, which restricts the unit root in productivity to originate solely in the technology shock.

This system applies to the case in which hours are assumed to be stationary. We also estimate a system in which hours are assumed to have a unit root:

$$\begin{bmatrix} \Delta x_t \\ \Delta n_t \end{bmatrix} = \begin{bmatrix} C^{11}(L) & C^{12}(L) \\ C^{21}(L) & C^{22}(L) \end{bmatrix} \begin{bmatrix} \varepsilon_t^z \\ \varepsilon_t^m \end{bmatrix}$$

We impose the same restriction, that $C^{12}(1) = 0$, to identify the technology shock. In the baseline case, we use four lags and limit our attention to a bivariate system. The data are quarterly and extend from 1948 through 2004.

Eichenbaum and Vigfusson (2005), however, show that the Chari, Kehoe, McGrattan example is an anomaly, being at odds with the data. Erceg, Guerreri and Gust (2004) and Francis, Owyang, and Roush (FOR, 2005) also show that VARs applied to artificial data from RBC models are consistent with the simulated results of the underlying model. In addition to VARs, FOR (2005) also use a variant of long run restriction and identifies technology as that which maximizes the forecast error variance of productivity at a long but finite horizon. Their results come closer to the model generated results than the results using long run restrictions.

Recall the previous summary of the literature. Using standard measures of hours per capita, the specification with stationary hours implies that hours increase significantly in response to a technology shock. In contrast, the specification with a unit root in hours implies that hours fall significantly in response to a technology shock. This pattern can be seen in Figure 2 where we use the standard measure of hours per capita with the civilian population 16 and over as the population measure. The first column shows the results from the system with hours per capita in levels and the second column shows the results from the system estimated with hours per capita in first differences. The model is bivariate in the logs of labor productivity and hours, but we also show the implied effects for the log of output, since it is equal to the sum of the other two variables. The graphs display the same conflicting results from the literature.

Is the over-differencing of hours per capita leading to erroneous results or are the low-frequency movements in the level of hours per capita leading to misleading results in the levels specification? To investigate the plausibility of these explanations, we remove the very low frequency movements in hours per capita using a very conservative HP filter with a λ parameter set equal to 160,000 rather than the usual 1600 for quarterly data. Figure 3 shows the estimated trend. It displays a pronounced U-shape, with the highest part in the early part of the sample. We then use the detrended hours series in the bivariate SVAR model, both in levels and first-differenced.

Figure 4 shows the estimated impulse response functions using HP filtered hours per capita, with the results from the levels specification on the left and the results from the first-differenced specification on the right. Interestingly, both specifications imply that a positive technology shock leads to a decline in hours in the short-run, consistent with Galí's (1999) finding and Francis and Ramey's finding. Although one would suppose that the difference

specification is plagued by over-differencing when HP filtered hours are used, the results are quite similar to those when the filtered hours levels are used.²

These results support Fernald's (2004) contention that the coincidental U-shape in both productivity growth and the standard measure of hours per capita is driving CEV's finding of a positive response of hours. When Fernald removes the U-shape in productivity growth, but leaves the U-shape in the standard measure of hours per capita, he finds a negative effect of technology shocks on hours. Conversely, when we eliminate the U-shape in hours per capita by removing the low frequency component, but do not allow for structural breaks in labor productivity, we also obtain the same negative response.

One might worry, though, that HP filtering the data could distort the dynamics. Or, perhaps the HP filter is simply a crude way to correct the standard hours measure for demographic and institutional changes it does not capture. In the next section, we show that this indeed the case. We use data on institutional and demographic trends to construct a new measure of hours per capita that does not display these low frequency movements.

III. A Better Measure of Hours Per Capita

In this section, we highlight three important demographic and institutional changes that affect the available workforce during the post WWII period. In particular, we show the importance of changes in government employment, schooling, and the aging of the population. This section builds on earlier work using annual data for the entire 20th Century (Francis and Ramey (2004)). Here, we improve on those measures and discuss how to apply the insights to quarterly post-WWII data. Our aim is to develop a new measure of hours per capita that adjusts for demographic and institutional trends that are not captured by standard RBC models.

² The results are similar when an HP filter with standard parameter values is used.

A. Government Employment

As discussed above, the standard measure of hours includes only private hours because government hours are not available at the quarterly frequency. The omission of government hours from the standard hours series, however, distorts the estimate of actual hours worked and induces significant trends in the series. Figure 5A shows the behavior of annual government employment as a fraction of total employment. Government employment (both civilian and military) was around 10 percent of total employment in 1948. It rose to a peak of 17 percent in the late 1960s and early 1970s, and has fallen somewhat down to around 14 percent currently. Changes in military employment are a small part of the changes. Of the seven percentage point rise in government employment 1948 to 1971, six percentage points were accounted for by increases in civilian government employment. A third of the increase in government employment through the late 1960s stemmed from the increased employment of teachers to educate the baby boom.

Note that the behavior of government hours follows an inverted U-shape with similar timing to the U-shape in the private hours series. Clearly, one source of the low frequency movements in the standard measure of hours per capita is its focus on private hours. It is not possible, however, to include total hours at a quarterly frequency because of data limitations. Thus, we must instead adjust the population measure to capture the population available for work in the private sector.

In adjusting the denominator of the private hours per capita measure, it is not correct to simply subtract government employment from the available workforce. To see this, consider an economy of 100 two-person households in which one person engages in market work and the other engages in home production. The employment-population ratio is $50/100 = 0.5$. Suppose that suddenly 10 jobs are reclassified from private sector jobs to government jobs (TSA jobs, for instance). If we simply subtracted those government jobs from population in the denominator,

we would measure the private employment-population ratio as $40/(100 - 10) = 0.44$, which is misleading. Instead, we should adjust the denominator by a θ that equates the trend in total hours worked divided by available labor force to private hours divided by the adjusted available labor force, i.e., the θ that equates the following two ratios:

$$\frac{H_P N_P + H_G N_G}{\bar{N}} = \frac{H_P N_P}{\bar{N} - \theta N_G},$$

where H_P is average hours per full-time equivalent employee in the private sector, H_G is average hours per full-time equivalent employee in government, N_P is the number full-time equivalent employed in the private sector, N_G is the number full-time equivalent employed in government, and \bar{N} is the workforce available for private or government work. This value of θ is:

$$\theta = \frac{\bar{N}}{\frac{H_P}{H_G} N_P + N_G}$$

θ is therefore the inverse of the employment-work force ratio and is greater than unity.³ In our example above, if $H_G = H_P$, we would use a $\theta = 1/0.5 = 2$. This adjustment, however, not only removes the low frequency movements, but also potentially affects the cyclicity of our new measure of private hours per capita, since the cyclicity of private employment N_P makes θ very cyclical. In order to purge the denominator of the cyclical components, we use the estimated low frequency component of private employment using a standard HP filter rather than private employment itself when we construct θ . We also use the HP filtered trend of the ratio of private-

³ The ratio of hours in the denominator puts private sector and government sector workers on the same full-time equivalence basis.

to-government hours in the denominator of θ . These adjustments apply only to θ , not to private hours used in the numerator of the hours per capita measure.^{4,5}

B. Schooling Trends and the Baby Boom

A second important trend affecting the standard measure of hours per capita is the combination of the baby boom and the rise in the years spent in school. Figure 5B shows the number enrolled in school in grades 11, 12, and higher education as a fraction of the civilian noninstitutional population ages 16 and over. This series uses full-time equivalent enrollment, as discussed in the data appendix. The enrollment percentage rose from four percent in 1947 to almost 10 percent in the mid-1970s, and then stabilized around 8 percent during the last 20 years. The hump was caused by the baby boomers moving through high school and college age, whereas the overall upward trend comes from individuals attending more years of school.

Note that this series also has an inverted U-shape. Thus, these trends in schooling induce low frequency movements in the standard measure of hours per capita because they influence the population available for work. To adjust for these low frequency trends, we subtract full-time equivalent enrollment in school (grades 11 and up) from the population in order to measure the available labor force.

C. Older Population

One of the most important recent demographic trends is the aging of the population and the decrease in the labor force participation rates of older Americans. As Figure 5C shows, the fraction of the population 65 and older was below 8 percent in 1947 and has risen steadily to well

⁴ We do not apply a filter to government employment because we want temporary changes in government employment, such as the Korean War, to show up as a decline in θ , which measures the fraction of the population available for work in the private sector.

over 12 percent. At the same time, the labor force participation of this group has decreased. The labor force participation rate of those 65 and older relative to those ages 25-64 was 40 percent in 1947. The rate fell steadily to 13 percent in the mid-1980's and has risen slightly to 17 percent currently.

The standard measure counts anyone 65 and older who is not institutionalized as part of the available work force. In contrast, Prescott (2004) omits the entire population 65 and older in his measure of the labor force. We adopt an intermediate strategy by weighting the ages 65 and older population by its labor force participation relative to those ages 25 to 64. We use the 25 to 64 age group in the denominator to avoid any effects from schooling. This measure captures the institutional effects of mandatory retirement laws in place during part of the period and the subsidies for retirement offered by Social Security, Medicare, and private pension systems.

D. A New Measure of Hours Per Capita

Based on these considerations, we now construct a new measure of hours per capita. For the denominator, we use a new measure of an economy's *population available to carry out productive activity* in the private sector. The measure is defined as follows:

$$\text{Available workforce} = (\text{Noninstitutional Population 16-64}) + ([\text{Relative LFPR 65+}] * [\text{Noninstitutional Population 65+}]) - (\# \text{ enrolled in Grades 11 \& 12}) - (\text{Full-Time Equivalent College Enrollment}) - (\theta * \text{Full-Time Equivalent Government Employment}).$$

⁵ The results we present in the next section are very similar whether we use the actual value of θ or the one with the filtered private employment series. The results are also robust if we use the unfiltered ratio of private-to-government hours when calculating θ .

θ is the reciprocal of the employment population ratio (derived and discussed above) and *Relative LFPR65+* is the labor force participation of individuals 65 and over divided by similar participation rates for individuals 25 to 64.

Our alternative measure of the population available for work starts with individuals 16 years of age. The primary reason for this lower bound is to match the age range of the standard measure.⁶ Additionally, this is the common (minimum) age at which individuals would join the labor force after dropping out of high school. For example, see the Annual Progress Report on Students Who Drop Out 2000 – 2001 by Kristopher Kaase (www.wcpss.net/evaluation-research/reports/2002/0222_dropout.pdf) a report on the Wake County public school system and U.S. Bureau of the Census, Current Population Reports, P20-500, table No. 297.

We use this measure of workforce in the denominator to create a new measure of hours per capita with hours worked in private business as the numerator. Figure 6 presents the new measure of hours per capita. There is less evidence of low frequency movements in this measure. There is no longer a U-shape as in the standard measure. There is, however, an upward trend over time.

How much of the low frequency movement in the standard hours per capita series is expunged when we use our new population available measure? Figure 7 plots the HP trend (with $\lambda=160,000$) of the following ratios that decompose the standard hours per capita series:

$$\frac{\text{Private Hours}}{\text{Civilian Population 16+}} = \frac{\text{Private Hours}}{NBAR} * \frac{NBAR}{\text{Civilian Population 16+}}$$

NBAR is our new measure of the population available to work. A glance at Figure 7 reveals that dividing by the civilian population 16+ is the source of significant low frequency movements that has plagued recent empirical RBC studies. Evidence of this can be seen in the similarity of

⁶ In Ramey and Francis (2005) we use younger ages because child labor was quite common in the early part of the 20th Century.

the graphs and the high correlation of 0.97 between the HP-Trends of the ratios (private hours/civilian population 16+) and (NBAR/civilian population 16+). However, comparing the HP-Trends of (private hours/civilian population 16+) to (private hours/NBAR) further punctuates our point. The two trends are dissimilar having a correlation coefficient of -0.33 and the latter ratio displays much less low frequency. Therefore, it is obvious that schooling, aging and government employment are very important slow moving demographics in the civilian population measure.⁷

IV. The Effects of Technology Shocks Using the New Measure of Hours Per Capita

A. Baseline Impulse Responses

We now investigate how the use of our new measure of hours per capita changes the previous results on the effect of technology on hours. The first thing to note is that our improved measure show more evidence of stationarity. Table 1 shows standard ADF tests for both the standard measure and our alternative measure. The evidence against a unit root is very weak in the case of the standard measure of hours. The p-value is 0.12 against the stationary alternative and 0.27 against the deterministic linear trend alternative. In contrast, the p-values are very low for the new measure, with strong evidence against a unit root in favor of stationarity.

We then re-estimate the structural VAR used by Galí, Francis and Ramey, and Christiano, Eichenbaum and Vigfusson using the new measure of hours. Figure 8 plots the impulse responses from the system using the new measure of hours. The levels specification is on the left and the first-difference specification is on the right. In both specifications, per capita labor hours respond negatively in the short-run to the technology shock in both the levels and first difference specifications. Hours become positive, though not significantly so, after a year or more.

⁷ When $\lambda=1600$ the correlations of the HP trends are 0.97 and -0.13 respectively.

Thus, in contrast to the case with the standard measure, the negative effect of technology shocks on the new hours measures is robust across specifications, whether hours are assumed stationary or not. These results also shed light on the debate concerning the results with the standard measure. CEV claim that over-differencing of hours per capita leads to different estimated effects of technology shocks on hours. This is not true with our new measure, or with the HP filtered standard measure as shown earlier. Even though the standard ADF tests reject a unit root, assuming a unit root in hours does not change the qualitative nature of the impulse response functions.

B. Robustness Checks

How robust are the results? CEV initially argued that omitted variables were the source of the Galí finding. To check the robustness, we estimate the larger system used by CEV. This system adds four variables: the federal funds rate, the rate of inflation (measured using the GDP deflator), the log of the ratio of nominal consumption to nominal GDP (where consumption is measured as the expenditures on nondurables and services plus government expenditures), and the log of the ratio of nominal investment expenditures to nominal GDP (where investment is measured as expenditures on consumer durables and gross private investment). The $C(L)$ matrix of this system is now a block 6×6 matrix in the lag operator. If labor productivity is the first variable in the system, we identify the technology shock by imposing the restriction that $C^{lj}(1) = 0$ for $j = 2, 3, 4, 5, 6$. Because the federal funds rate is only available beginning in 1954, the model is estimated over a shorter sample.

Figure 9 shows the results when hours are specified in levels. The labor hours response continues to be significantly negative in the short run. Output and investment also dip slightly in the short-run, but they are not significantly below zero.⁸

⁸ Additionally, in the larger system using first differenced labor hours, the hours response is significantly negative.

We also checked robustness in two more ways. First, we estimated the bivariate system in which we allowed hours per capita to have a linear trend. We thought it was important to consider this possibility because of the slight upward trend in the new measure. The results (not shown) look similar to those for the first difference specification. Hours respond negatively for the first three quarters (significant only for the first two quarters) before turning positive, but not significant, for the remainder of the response period. The initial negative impact effect is estimated to be -0.23 , and is statistically significant. (The impact effects were -0.36 for the levels specification and -0.29 for the difference specification).

Cooley and Dwyer (1998) point out that the results from structural VARs may be sensitive to auxiliary assumptions with respect to lag length. Our baseline models all include four lags. To determine whether our results were due to too few lags, we re-estimated the bivariate system in levels and included 50 lags. The impulse responses were qualitatively similar to those from the system with only 4 lags.

Thus, when our new measure of hours per capita is used, hours always respond negatively to technology shocks. This is true for levels, first-differences, and trend specifications. It is true for specifications with more variables in the system. It is also true when we add 50 lags to the specification.

V. Conclusion

In this paper we propose that the conflicting results in the debate concerning the effects of technology shocks on hours stems from the low frequency movements in the standard measure of hours per capita. We first show that removal of the low frequency movements using a very conservative HP filter produces results suggesting a positive technology shock lowers hours in the short-run. We then argue that the HP filter is capturing demographic and institutional

changes that affect the available workforce. We produce of new measure of hours per capita that adjusts for these changes. Our measure, which adjusts for trends in government employment, aging, and schooling, removes most of the low frequency trend. We find that removing these slow moving components leads to consistent results on the effects of technology on per capita hours worked, regardless of the stationary assumption assumed for per capita labor hours. That is, in contrast to results using the standard measure, our new measures produce uniformly negative effects of technology on hours, whether labor hours enter the VARs in levels, first differences or linearly detrended.

Data Appendix

Hours in Private Business and Productivity:

For the series on labor productivity and labor input, we use the BLS series “Index of output per hour, business” and “Index of hours in business”.

Government and Private Employment

Bureau of Economic Analysis, Table 6.5, 6.8, and 6.9, and BLS data on unpaid family workers (<http://www.bls.gov/webapps/legacy/cpsatab5.htm>). The BLS hours in business index includes hours of employees, the self-employed and unpaid family workers in the private sector, as well as hours worked in government enterprises (such as the post office). We use the full-time equivalent employees in government (excluding government enterprises) as government employment and full-time equivalent persons engaged in the private sector plus government enterprises plus unpaid family workers as private employment. (The BEA “persons engaged” series includes the self-employed but not unpaid family workers.) We obtained annual measures of average hours of full-time equivalent workers in government versus the private sector by dividing the BEA series of total hours of employees (full time and part time) by the full-time equivalent number of employees.

Schooling Data:

Enrollment data were obtained by combining information from the *Digest of Education Statistics, 2002* and Claudia Goldin “A Brief History of Education in the U.S.” August 1999, NBER working paper H0119. Only data on public enrollment was available for grades 11 and 12. We assumed that the fraction of total enrollment in grades 9-12 that was accounted for by grades 11 and 12 was the same in public and private schools in order to impute total enrollment in grades 11 and 12.

Fraction of college students who are full-time: The fraction of college students who were enrolled fulltime was available only from 1963 – 1998 from the *Digest of Education Statistics*. We used the 1963 fraction for the years before 1963 and the 1998 fraction for the years after 1998. We assumed that part-time college students were equal to 30% of a full-time college student.

Population:

Data Sources: Data, including age breakdown, is from the U.S. Census, *Mini Historical Statistics*, Table HS-3 and *Economic Report of the President, 2003*, Table B-34.

Labor Force Participation Rates

The labor force participation rates by age group can be obtained for the civilian population from the BLS.

References

- Burnside, Craig and Martin Eichenbaum, "Factor Hoarding and the Propagation of Business Cycle Shocks," *American Economic Review*, 86 (1996): 1154-1174.
- Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan, "A Critique of Structural VARs Using Real Business Cycle Theory," May 2005 manuscript.
- Christiano, Lawrence, Martin Eichenbaum, and Robert Vigfusson, "What Happens After a Technology Shock?" 2003 NBER working paper 9819.
- , "Assessing Structural VARs," 2005 manuscript.
- Cooley, Thomas F. and Mark Dwyer, "Business Cycle Analysis without Much Theory: A Look at Structural VARs," *Journal of Econometrics* 83 (March-April 1998): 57-88.
- Erceg, Christopher J., Luca Guerrieri and Christopher Gust, "Can Long Run Restrictions Identify Technology Shocks?" International Finance Discussion Papers 792. Washington: Board of Governors of the Federal Reserve System, 2004.
- Fernald, John, "Trend Breaks, Long-Run Restrictions, and the Contractionary Effects of Technology Shocks," March 2004 manuscript.
- Francis, Neville and Valerie A. Ramey, "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited," forthcoming *Journal of Monetary Economics*.
- , "The Source of Historical Economic Fluctuations: An Analysis using Long-Run Restrictions," NBER working paper 10631, July 2004.
- Francis, Neville, Michael Owyang and Jennifer Roush, "A Flexible Finite-Horizon Identification of Technology Shocks," International Finance Discussion Papers 832. Washington: Board of Governors of the Federal Reserve System, 2005.
- Galí, Jordi, "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations," *American Economic Review*, 89 (March 1999): 249-271.
- Galí, Jordi and Pau Rabanal, "Technology Shocks and Aggregate Fluctuations: How Well does the RBC Model Fit Postwar US Data? 2004 NBER Macroeconomics Annual.
- Goldin, Claudia, "A Brief History of Education in the U.S." August 1999, NBER working paper H0119.
- King, Robert G., Charles I. Plosser, James H. Stock, and Mark W. Watson, "Stochastic Trends and Economic Fluctuations," *American Economic Review* 81 (1991): 819-840.
- Prescott, Edward C., "Why Do Americans Work So Much More Than Europeans?" *Federal Reserve Bank of Minneapolis Quarterly Review*, 28 (July 2004): 2-13.

Ramey, Valerie A. and Neville Francis, “A Century of Work and Leisure” Mimeo 2005.


Shapiro, Matthew D. and Mark Watson, “Sources of Business Cycle Fluctuations,” *NBER Macroeconomics Annual*, 1988, pp. 111 – 148.

Table 1: Augmented Dickey-Fuller Tests: Quarterly Data 1948-2004
P-values

Variable	Against H_a: stationarity	Against H_a: linear deterministic trend
Standard hours per capita measure	0.118 (3 lags)	0.270 (3 lags)
New hours per capita measure	0.003 (3 lags)	0.011 (3 lags)


Note: Lags were chosen optimally.

Figure 1: Private Hours Per Capita in Post-WWII Quarterly Data: Based on Civilian Non-institutional Population Aged 16 and Over



**Figure 2: Impulse Responses to a Technology Shock: Quarterly 1948-2004
(Civilian Population 16+, Bivariate System with 95% standard error bands)**

Hours in Levels



Hours in First Differences

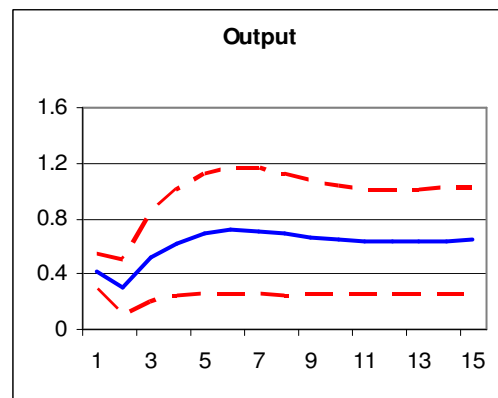



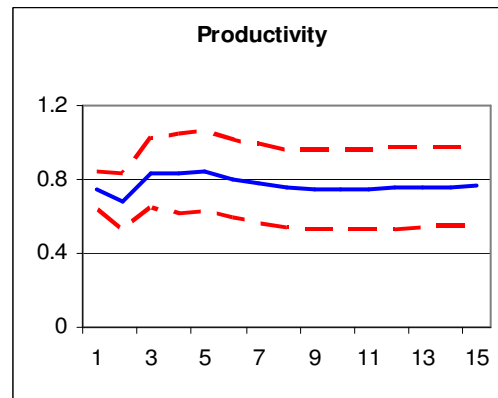
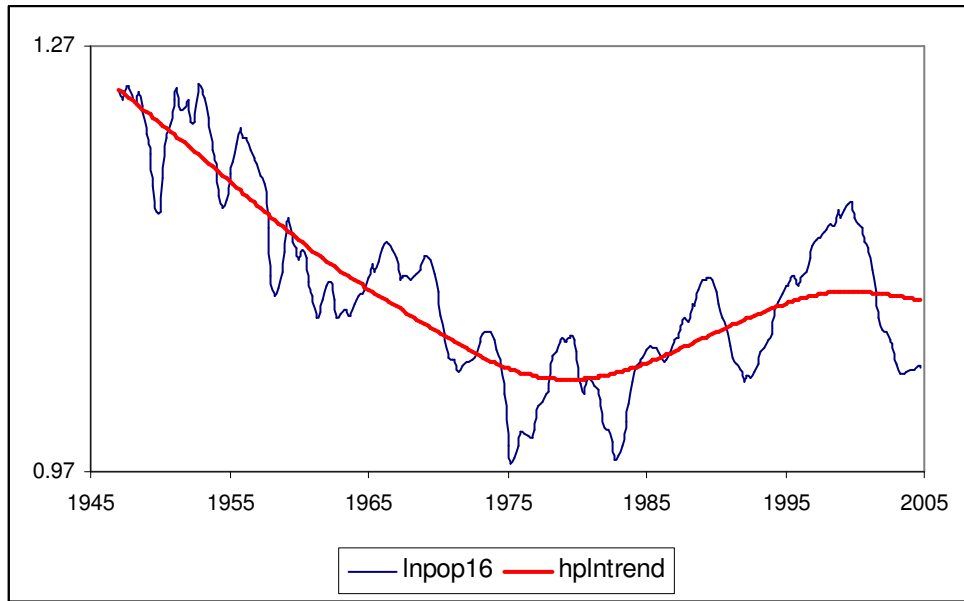



Figure 3: Plot of Hours per Capita and its HP-Filtered Trend ($\lambda = 160,000$)



**Figure 4: Impulse Responses to a Technology Shock: Quarterly 1948-2004
(HP-Filtered Hours/Civilian Population 16+, Bivariate System with 95% error bands)**

Hours in Levels



Hours in First Differences



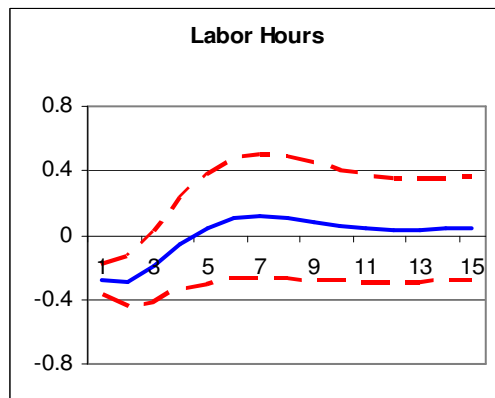
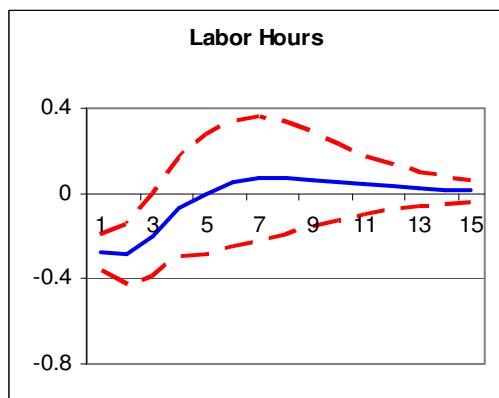



Figure 5

Figure 5A: Government Employment as a Fraction of Total Employment




Figure 5B: Number Enrolled in Grades 11 and up as a Fraction of Civilian Population 16+




Figure 5C: Fraction of the Population 65 and Older




Figure 6: Private Hours Per Capita in Post-WWII Quarterly Data: Based on New Measure of Non-institutional Population Aged 16 and Over


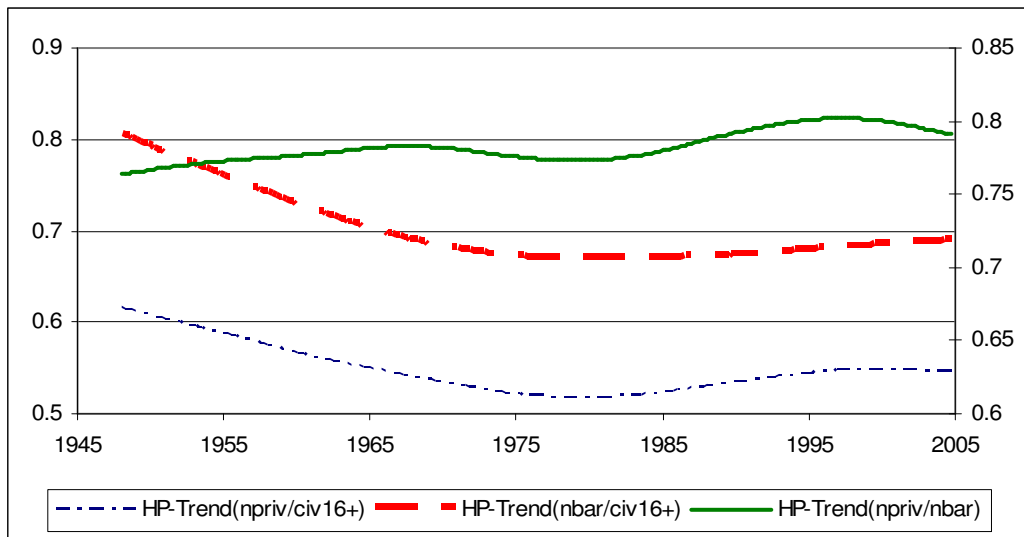


Figure 7: Plots of HP-Trends ($\lambda = 160,000$) of [Private Hours/Civilian Population 16+], [Private Hours/New Measure of Population Available] and [New Measure of Population Available to Work/Civilian Population 16+]



HP-Trend (NPRIV/NBAR) plotted on secondary axis.

Correlation HP-Trends (NPRIV/CIV16+, NBAR/CIV16+) = 0.97

Correlation HP-Trends (NPRIV/CIV16+, NPRIV/NBAR) = -0.33

Correlation HP-Trends (NPRIV/NBAR, NBAR/CIV16+) = -0.54

Keys:


NPRIV = Private Hours

CIV16+ = Civilian Population 16+

NBAR = Population Available to Work Weighting Older Workers 65+ by their Relative Participation Rates

Figure 8: Impulse Responses to a Technology Shock: Quarterly 1948-2004
 (with New Population Measure, Bivariate System with 95% standard error bands)

Hours in Levels



Hours in First Differences




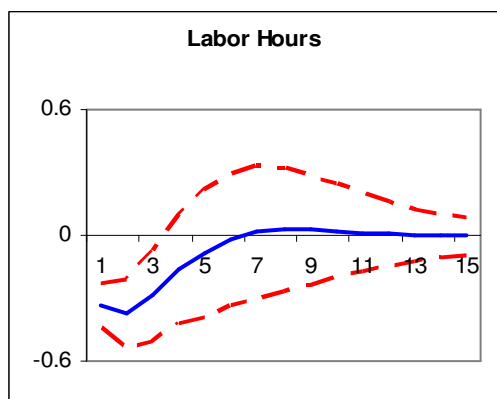



Figure 9: Impulse Responses to a Technology Shock: Quarterly 1954-2004 (New Measure, Six-Variable VAR with 95% standard error bands, Hours in Levels)

