

Labor Supply and the Changing Household

James A. Kahn

June 2005

Very Preliminary

Abstract

The behavior of employment and labor force participation has changed markedly over the past 40 years. Participation has drifted up, a period that has also seen substantial changes in household structure, notably more single person households and more positive assortive matching. This paper documents these facts, and then provides a model that endogenizes many of them. The benchmark model is a modest extension of the neoclassical growth model to allow for 2-person households with home as well as market production. Unbalanced technical progress results in a shift from household to market production. The result is an increase in market labor supply over time, plus changes in the dynamic response of labor supply to various kinds of shocks. The model is calibrated to data from time use surveys, and yields responses similar to those that have occurred in the data.

*Research Department, Federal Reserve Bank of New York. The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. I would like to thank Nathaniel Hilger and Ben Pugsley for valuable research assistance.

I. Introduction

The neoclassical growth model is the workhorse of modern macroeconomics. The standard “balanced growth” version is built around a number of stylized facts, initially motivated by Kaldor (1961) among others. Notable among these is the stationarity of work effort per capita. The Real Business Cycle literature since King, Plosser and Rebelo (1982) has generally assumed logarithmic preferences (or their equivalent) that build in this feature through offsetting wealth and substitution effects.

Employment per capita in the U.S. has experienced substantial low-frequency movement since W.W. II (and earlier as well, although the data are more sketchy). Understanding these movements is useful from a number of standpoints. From a public policy standpoint, changes in trend labor supply can be as troublesome as changes in productivity trends, which are widely thought to have contributed to monetary policy mistakes in the 1970s. The low-frequency behavior of hours of work has also been a focus of debate in the empirical business cycle literature, where it turns out that whether the series is treated as stationary or not has substantive effects on estimated impulse response functions (see, e.g., Francis and Ramey, 2003, Christiano, Eichenbaum, and Vigfusson, 2004).

Labor force and productivity growth are of course the two components of trend output growth. From an economic modeling standpoint, it would be useful to know whether the standard balanced growth assumptions are valid. Demographic changes, particularly in the nature of the household (overall size, number of children, etc.), have also been notable. For example, the share of the population in prime working age (25 to 54) has shifted dramatically during the postwar period—from a peak of 59 percent in the early 1950’s, down to 50 percent in the mid 1970’s, before rising again to a peak of nearly 58 percent in the late 1990’s. There have also been changes in household structure. The average number of persons per household has fallen from 3.67 just after the war, to 2.61 as of 2003. This reflects not just the decline in birth rates, and increased life expectancy (which implies more childless households), but also an increase in the

number of households with only one adult. That number has risen from 7.7 percent of all households in 1940 to 25.8% in 2000.

While college education grew dramatically over the postwar period, the growth was even more dramatic for women compared to men. According to census data, the share of women aged 25-34 with college degrees went from 0.05 in 1940, to 0.12 in 1970, to 0.29 in 2000. The same figures for men are 0.07, 0.20, and 0.26. Thus while growth in college education among men leveled off after 1980, for women it continued to the point where they have surpassed men.

There have been subtler shifts as well. Below we provide evidence of increased positive assortive matching over time. Thus even as a larger share of the population is unmatched, those who are pairing up are tending to choose partners who are more like themselves in terms of human capital. As we shall see, this affects aggregate labor supply behavior through household decision-making over and above whatever effects may be implied by

In this paper we argue that these types of demographic shifts have potentially important implications for labor supply—and not just for long-term trends, but possibly for short-term fluctuations as well. We derive relatively straightforward generalization of the growth model which endogenizes some of these shifts, and generates changes in market labor supply. The basic mechanism is unbalanced technical progress, which shifts economic activity from the household to the market, as in Greenwood et al. (2005). We make a number of extreme simplifying assumptions to keep the analysis tractable, though the goal of this research is to relax as many of them as possible. Specifically, we assume a representative household, and completely flexible hours of work (i.e. no “lumpy” participation decision, no job search or unemployment, etc.). The flexibility assumptions are reasonable for looking at long-term trends, but perhaps not for looking at short-term fluctuations (though much of the business cycle literature makes such assumptions). The justification in the latter case is that the short-run properties of the model, while not necessarily to be taken literally, provide some insight into how a model with more realistic short-run rigidities might respond.

An additional motivation for examining short-run dynamics is the accumulating evidence that the business cycle has changed dramatically in the past 20 years. We know

from a number of studies (e.g. McConnell and Perez-Quiros) that the volatility of GDP growth has declined substantially since the early 1980s. It also appears that the cyclical behavior of the labor market has changed. The last two recessions have been characterized by relatively mild declines in employment, but also very slow recoveries. The explanation for the so-called “Great Moderation” remains elusive, but at least two studies (Stock and Watson, 2003, and Ahmed, Levin, and Wilson, 2004) have argued that the source is smaller shocks rather than a change in propagation. Others (e.g. Kahn, McConnell, and Perez-Quiros, 2002) have argued the contrary, i.e. that propagation has an important role in explaining reduced volatility. Any impact that long-term shifts in the labor market have on short-run dynamics may also be another source of a change in propagation.

The first part of the paper is empirical. We examine a number of different data sets to shed light on the sources of permanent shifts in work effort per capita. We then develop a variant of the household production model, first in a version in which does not distinguish between individuals and households. Given the shortcomings of that model, we then allow individuals to form 2-person households and examine the implications for labor supply under a variety of scenarios.

II. Postwar Structural Change

A. Hours Per Capita

Measured market labor (in hours per capita) in the postwar United States has followed a U-shaped pattern over the postwar period, declining until the early 1970s, then rising again in the 80s and 90s. As emphasized by Francis and Ramey (2004), some of this is accounted for by straightforward demographic adjustment. This can be seen in Figure 1, which plots the standard notion of hours per capita based on population 16 and over versus population aged 16 to 65. Figure 2 plots hours per capita based on the decennial census (from 1950-1990) and the CPS (since 1962). This also shows evidence of a U shape, though with these data the trough occurs earlier, in the 1960s.

We control for additional demographic changes in both the census and CPS datasets and generally still find an upward trend in the 80s and 90s. Our controls are based on the notion of the household as the fundamental decision unit (see Mulligan and

Rubinstein, 2004, for more on this perspective). In the regressions we include dummy variables for numbers of adults (two, three or more), the age group of the adult head of household (under 25, 25-34, etc.), and the presence of children under the age of 6. The dependent variable is hours per adult, and we use the appropriate household weights for the regressions.

Figure 3 shows the result of fitting a cubic time trend to the March CPS data from 1962 to 2004, with and without these demographic controls. Although the 1962 start date means that much of the downward part of the “U” is absent, we still see that demographic controls have the effect of leveling out the early part of the sample, but do little to alter large post-1975 upward drift. Figure 4 presents the results from the decadal census with and without demographic controls, with similar findings for the period from 1950 to 2000. Figure 5 shows estimates of a separate trend for single-adult vs. multi-adult households (in both cases without children under 6). Here we see that most of the upward drift in hours comes from the latter group, consistent with the notion that Finally, we also compute a “fixed-weight” hours series in which we take the population shares of each demographic group (e.g. household head age 25-34, in a two-adult household with children under 6). Here for reasons of tractability we limit the age groups to cover the years 25-64, and distinguish only between one versus 2 or more adults. The resulting series, raw hours of work per capita, and the series with fixed 1986 (the midpoint of the sample) demographic weights, are displayed in Figure 6, along with their Hodrick-Prescott trends. Once again we see that controlling for demographics has only a modest impact on the upward drift since the mid-1970s.

B. Household Composition

We know that there have been substantial changes over the postwar period in the number of people per household, both adults and children. The previous results suggest that hours of work have grown notwithstanding these changes. Next we consider evidence of increased assortative matching in marriage. We examine this by looking at education levels of spouses, again in both the census and the CPS. If, say, men in a given education category are disproportionately choosing spouses in the same category, then

the share of their spouses in that category should exceed the population share. The extent to which it does so is a quantitative measure of assortive matching.

Again using both the CPS and Census, we distinguish four education groups: less than a high school degree, high school degree, some college, and college degree or higher. We use the same age groups as before, but focus on the 25-34 group, as these will have largely finished their education and will have made relatively recent marriage decisions.

Specifically, let σ_{ijt} denote the share of married women in age group i that are in education group j in year t (so that $\sum_j \sigma_{ijt} = 1$). Compare this with the same statistic (call it ω_{ijt}^h) computed over the population of women who are spouses of men in education group h . Assortive matching would imply that $\omega_{ijt}^h > \sigma_{ijt}$ when $h = j$. Random sorting would imply no systematic relationship between the $\{\sigma_{ijt}\}$ and the $\{\omega_{ijt}^h\}$, i.e. $\omega_{ijt}^h \approx \sigma_{ijt} \forall h$. Figure 6 plots $\omega_{ijt}^h - \sigma_{ijt}$ for i referring to age group 25-34, and for h referring to the two extreme education groups less than high school and college degree or higher. The results show that for both education groups, the extent of assortive matching has grown substantially over time. The implication of this for joint labor supply will be clearer below, but the essential idea is that more assortive matching leads to a smaller gain from specialization, and consequently could accelerate the increase in jointly determined market labor. This is in addition to the direct effect of the increase in women's human capital relative to men's since 1970.

C. Patterns of Time Use

Unfortunately we do not have consistent data over this period on how people use their time. Nonetheless we can potentially draw some conclusions from the cross-section documented by the recently initiated American Time Use Survey (ATUS). Some of the broad results from this survey, taken in 2003, are summarized in Table 1. It should be emphasized that an important limitation of the ATUS is that it only obtains detailed information from one household member. Thus while it is tempting (and to some degree legitimate) to draw conclusions about how time is allocated across members within a

household, firmer conclusions await further study. Nonetheless, the results suggest that in two-adult (one male, one female) households, men and women spend roughly equal time in what might be called “total labor.” At least this appears to be true for employed men and for women regardless of employment status.¹ But the division of total labor between household and market labor varies sharply across men and women depending on the presence of children. As one might expect, particularly with children under age 6, women allocate more of their time to household labor, while men allocate much of their time to market labor. See lines 5-8 and compare to lines 11-14. It would appear that with children under school age, there is some degree of specialization, with women doing substantially more household work, while men undertake more market labor. Once children reach school age, women substitute market labor for household labor, and total labor goes down for both.

D. Technical Progress in the Household Sector

As in Greenwood et al (2005), we argue that technical progress in the household sector is an important driving force. As evidence, we cite the dramatic decline in prices of household appliances relative to other goods prices (see McGrattan and Rogerson, 2004). Gordon (1990) finds a 5 percent annual decline in the relative price of household appliances over the period 1947-1984. If capital’s share in household production were on the order of 0.2 (i.e. that sector is more labor-intensive than the market sector), this would imply a one percent annual rate of relative technical progress. While this is not the precise mechanism by which technical progress occurs in this paper, the essential story is the same.

III. An Equilibrium Model of Home Production with Individual Agents

A. Model Setup

The setup is similar to that of Benhabib, Rogerson, and Wright (1991). Individuals derive utility from a standard consumption good c , a household good z that is produced from some combination of non-market work at home (“housework”) and labor

¹ Non-employed men appear to engage in less total labor, but they represent a relatively small portion of the sample and may be exceptional in other ways. For example, their status may be transitory, they may be ill or otherwise unfit for labor, or they may be spending their time in unconventional ways.

purchased on the market, and from leisure $\bar{n} - n$, where n is hours of work up to some maximum \bar{n} , which we normalize to one.

Suppose that agents have preferences given by

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t \left[\log \left[c_{it}^\rho + \alpha z_{it}^\rho \right]^{1/\rho} - \phi \frac{n_{it}^{1+1/\eta}}{1+1/\eta} \right] \right\}, \quad (1)$$

where $\rho \leq 1, \eta > 0$. The parameter η corresponds to the Frisch elasticity of labor supply in a model without home production. The production technologies are

$$c_t + i_t = A_{mt} n_{mt}^\gamma k_{mt}^{1-\gamma}, \quad 0 < \gamma < 1 \quad (2)$$

$$z_t = A_{zt} n_{ht}^\theta k_{ht}^{1-\theta}, \quad 0 < \theta < 1 \quad (3)$$

where n_h refers to household labor, n_m is market labor, and k_h and k_m are capital inputs.

The timing assumption is that k_{jt} refers to the stock at the beginning of period t . Each agent maximizes (1) subject to (2) and

$$n_t = n_{mt} + n_{ht} \quad (4)$$

$$k_t = k_{mt} + k_{ht} \quad (5)$$

$$k_t = k_{t-1}(1-\delta) + i_{t-1} \quad (6)$$

$$n_m, n_h \geq 0.$$

We can solve a planner's problem, where the value function V satisfies

$$V(k_t) = \max \log \left[c_{it}^\rho + \alpha z_{it}^\rho \right]^{1/\rho} - \phi \frac{n_{it}^{1+1/\eta}}{1+1/\eta} + \beta E_t \{ V(k_{t+1}) \} \quad (7)$$

subject to (2)-(6). The first-order conditions are (letting $\lambda_t = V'(k_t)$):

$$m_t^{-\rho} c_t^{\rho-1} = \lambda_t \quad (8)$$

$$m_t^{-\rho} c_t^{\rho-1} \gamma Y_t / n_{mt} = \phi n_t^{1/\eta} \quad (9)$$

$$\alpha m_t^{-\rho} z_t^{\rho-1} \theta z_t / n_{ht} = \phi n_t^{1/\eta} \quad (10)$$

$$\alpha z_t^{\rho-1} (1-\theta) z_t / k_{ht} = c_t^{\rho-1} (1-\gamma) Y_t / k_{mt} \quad (11)$$

$$\lambda_t = \beta E_t \left\{ \lambda_{t+1} \left(1-\delta + (1-\gamma) Y_{t+1} / k_{mt+1} \right) \right\} \quad (12)$$

where $m_k = (c_k^\rho + \alpha z_k^\rho)^{1/\rho}$ and $Y = c + i$. (8) and (12) combined is just the standard intertemporal Euler equation condition. (9) and (10) allocate labor efficiently between home and market production, while (11) does the same for capital. Note that we are

assuming that while total k_t is determined by decisions taken at date $t-1$, the allocation of capital between the home and market sectors is determined ex post at time t . This is just for simplicity and does not change the basic thrust of the results.

B. Steady State Results

Equations (2)-(12) are easily solved for steady state values of the endogenous variables. The first task is to calibrate, as best we can, the parameters of the model. Regarding the labor supply parameters, while hours are not bounded between zero and one, we will treat n as a share of non-discretionary time of about 14 hours per day devoted to work of any kind. In a model without household production, η would correspond to the Frisch elasticity of labor supply. Traditional estimates of this parameter are usually less than one (e.g. MaCurdy, 1981). Since household production potentially provides an additional margin of substitution, empirical elasticities could overstate η . Nonetheless, we will set $\eta = 0.5$ in our base case. To match the data on total work hours, we set $\phi = 5$, which results in n in the vicinity of 0.6, or 8-9 hours per day.

For technical progress in household production to result in increased market labor supply, we need a less than unit elasticity of substitution between c and z , which requires $\rho < 0$. This contrasts with the assumption in Benhabib, et al (1991), who argue for a positive value. We will take $\rho = -2$ as our benchmark. We will follow them, however, in assuming that household production is more labor intensive than market production. We will set labor's share in market production, γ , equal to 0.7, and the share in household production, θ , equal to 0.5. The discount rate β is set at 0.99, based on quarterly time periods, and the utility parameter α , which reflects the weight on the household good relative to the market good in preferences, is chosen so that, given the other parameters, roughly equal time is spent on household and market production, which turns out to be $\alpha = 2$. Finally, we set the depreciation rate equal to .05 (quarterly).

We use the simulations to illustrate the impact of relative technical progress in the household sector. Table 2a provides results for values of A_z beginning at one and increasing in steps of 10 percent, fixing $A_c = 1$. For each 10 percent increase in A_z , we

see roughly a 2.5 percent increase in market labor supply. Table 2b does some sensitivity testing, for a $\eta = 1$ and for $\rho = -3$, and also contrasts the case of “balanced” technical progress, i.e. the growth in both A_z and A_c in proportion to the sector’s labor share.

Changing the elasticity parameters has little impact on the results. Regarding the balanced technical progress case, it is worth noting that because the share parameters θ and γ differ from each other, the result is not balanced growth in the usual sense.

The single-agent household model clearly can account for an upward trend in market labor, if we accept the relative technical progress hypothesis. But we have seen that the trend has been anything but steady. Controlling for demographics, labor supply appears to have been relatively flat until the 1970s, and only then began to trend up. There is no evidence that relative technical progress followed a similar pattern. Moreover, we also know that labor force participation of men versus women has headed in opposite directions: As women’s participation has increased, men’s has declined. This is another fact that would seem irreconcilable with any single-person household model. This leads us to consider enriching the model of the household.

IV. Household Labor Supply

A. Setup

Now consider the problem of 2-person households. To provide some motivation for pairing up in this context, we assume that z is a public good within the household. We will also assume that the representative household consists of two agents of unequal skills, as quantified by effective market labor supplied per unit of time. Specifically, the first household member’s effective labor is normalized to one, while the second supplies $\mu < 1$ units of effective labor to the market per unit of time.

We assume that households maximize

$$E_t \left\{ \sum_{t=0}^{\infty} \beta^t \left[U(c_{1t}, z_t, n_{1t}) + U(c_{2t}, z_t, n_{2t}) \right] \right\}. \quad (13)$$

The production technology is as before:

$$z_t = A_{z_t} (n_{h1} + n_{h2})^\theta k_h^{1-\theta} \quad (15)$$

$$c_{1t} + c_{2t} + i_t = A_{m_t} (n_{m1} + \mu n_{m2})^\gamma k_m^{1-\gamma} \quad (16)$$

$$n_k = n_{mk} + n_{hk} \quad (k=1,2) \quad (17)$$

$$k_t = k_{mt} + k_{ht} \quad (18)$$

$$k_t = k_{t-1} (1-\delta) + i_{t-1} \quad (19)$$

$$n_{mk}, n_{hk} \geq 0$$

In contrast to the previous case, at least some of the non-negativity constraints on the various types of labor will always bind, as we shall see.

It is straightforward to see that c must be the same for both agents. We then get the following first-order conditions for this problem:

$$m_t^{-\rho} c_t^{\rho-1} = \lambda_t \quad (20)$$

$$m_t^{-\rho} c_t^{\rho-1} \gamma Y_t / n_{mt} = \phi n_{1t}^{1/\eta} \quad (21)$$

$$\mu m_t^{-\rho} c_t^{\rho-1} \gamma Y_t / n_{mt} \leq \phi n_{2t}^{1/\eta} \quad (22)$$

$$2\alpha m_t^{-\rho} z_t^{\rho-1} \theta z_t / n_{ht} \leq \phi n_{1t}^{1/\eta} \quad (23)$$

$$2\alpha m_t^{-\rho} z_t^{\rho-1} \theta z_t / n_{ht} = \phi n_{2t}^{1/\eta} \quad (24)$$

$$2\alpha z_t^{\rho-1} (1-\theta) z_t / k_{ht} = c_t^{\rho-1} (1-\gamma) Y_t / k_{mt} \quad (25)$$

$$\lambda_t = \beta E_t \left\{ \lambda_{t+1} (1-\delta + (1-\gamma) Y_{t+1} / k_{mt+1}) \right\}. \quad (26)$$

The inequalities (22) and (23) again reflect the possibility that one the partners will have either zero market labor or zero household labor. If (22) is strict, then the second (lower-skilled) agent provides zero market labor. If (23) is strict, then the first agent provides zero household labor. In general, one or both will be strict inequalities. We will refer to the case in which both are strict as “complete specialization” (CS). If only (22) is strict, we have “partial specialization--second agent” (PS2), while if only (23) is strict, we have “partial specialization--first agent” (PS1). Under the relative technical progress hypothesis, depending on initial conditions, we might expect to shift from PS2 to CS to PS1, as household production requires diminishing amounts of labor.

Of course in reality, all three cases are present in the cross-section, where the case will depend on the relative magnitudes of wages and wealth in a household, and of course the presence of (young) children. Over time technology can evolve to shift the

equilibrium for a given household, or wages (more generally, opportunities in the work force) can fluctuate. This is related to the story that Greenwood et al (2005) tell about the entry of women into the labor force. Allowing for this heterogeneity is beyond the scope of the present paper, but would be useful to explore further.

B. Steady States for the 2-person Household Equilibrium

We can once again solve for the steady state, recognizing the need to check for binding non-negativity constraints. Starting with the base case parameters used in the one-person case, we can solve for the steady states with various values of A_z and see how the different cases apply. Note that the results for the two-person household are not easily comparable to those for the one-person case, because of the public good assumption regarding z .

It turns out that for these parameters, $A_z < 0.75$ is sufficiently low for the steady state to be in the PS2 case, i.e. where both partners do household work and the first agent does market work. For $A_z \in (0.75, 1.13)$ we get CS, where the first agent does only market work and the second does only household work, and for higher values we have PS1, where the second agent joins the work force.

The more interesting finding is that whereas with the single agent household, steady relative technical progress resulted in a steady upward drift of market hours, in this case we get a sharp acceleration of hours in going from CS to PS1. Within the CS regime, each 10 percent increase in A_z results in roughly a 0.6 percent increase in market hours. Within the PS1 regime, by contrast, it produces around a 2.5% increase, similar to what the single-agent household model implied. Again, the representative household model makes this a little more stark than it would be in reality, but it seems likely that even with realistic heterogeneity, the increased presence of second wage-earners, either near the extensive margin or on the intensive margin, would get a similar acceleration in response to the forces described here.

Not surprisingly, the model also predicts that as market employment of the second agent increases, that of the “primary” wage-earner declines. This is depicted in Figure 7. Note that the primary wage-earner is not increasing his non-market work—that is still

zero. He is merely obtaining more leisure in response to the increased market work of his partner.

A second factor that could affect the trend is an increase in μ . As μ gets closer to one, the likelihood of the PS1 regime increases, i.e. the value of A_z at which the regime switches from CS to PS1 falls. Figure 8 depicts the effect of shifting to $\mu = 0.85$ from 0.8. In a setting with realistic heterogeneity, this would have the effect of increasing the share of households in the latter regime, and thus accelerating the response of hours to relative technical progress. Both the relative increase in women's education and human capital and the increased positive assortive matching in marriage would have this effect.

C. Household Dynamics

In this section we examine whether the different regimes have different implications for the dynamic response of work effort and other variables to shocks. The regimes present some difficulties for computing the dynamics of a linearized system. Starting from some steady state, assuming it is in the interior of one of the three regimes, the deviations must be small enough to stay in the same regime, or else the approximation will be invalid. That caveat aside, solving the linearized systems in each of the three cases will still provide useful insights into the impact of these structural changes on dynamics, and it is relatively easy to do, so we will adopt that approach for now. To solve the linearized model we use the technique described in Uhlig (1997).

We should note that the labor supply elasticity for n_m , denoted ε_m will be related to both η , the Frisch elasticity for total labor, and the elasticity of non-market labor with respect to the market wage, ε_h , which would generally be negative. If w denotes the wage, or marginal product of labor, we have

$$\varepsilon_m = \frac{\partial n_m}{\partial w} \frac{w}{n_m} = \eta \frac{n}{n_m} - \varepsilon_h \frac{n_h}{n_m} \quad (27)$$

which is always greater than η except in the limiting case when $n_h = 0$. What we will find is that ε_h appears to be bigger in absolute value if both partners are engaged in market labor.

For this analysis we will focus on the two “later” regimes, CS and PS1. It turns out that the answer to the question of how the dynamics are affected by the regime is, “It all depends.” We first examine the effect of a shock to A_m , the market technology. Here the response of hours to such a shock in the PS1 regime is stronger than in the CS regime (see Figure 9). This is an unambiguous signal to shift from home production to market work, and when the secondary earner is already in the region of positive market hours, the response is larger.

The second type of shock we will consider is a relative shift in the earnings power of one of the family members. The motivation for this is as follows. First, aggregate disturbances may be thought of as averages of a lot of sectoral disturbances, which are likely to hit different household members differently. This may be true even if there really is just one fundamental disturbance, to which different sectors respond to varying degrees (as in Abraham and Katz, 1986). Second, a shock of this sort highlights features of this type of model, notably the substitution of one household member’s labor for another, that appear to be important at low frequencies.

To examine this, we rename μ from the steady state section of the model μ_2 and generalize the model to include a shock to the effective labor of the primary earner, denoted by μ_1 . We then consider a shock to μ_1 . The effects are shown in Figures 10 and 11. Here, by contrast, the “later” (PS1) regime implies essentially no response of market labor to this shock, as shown in Figure 10. The reason is clear. Whatever the shock to the primary worker’s earning power, the effect is offset by the response of the secondary agent. This offsetting response is shown in Figure 11.

While it may seem implausible that the secondary earner would reduce his hours in response to a persistent but transitory increase in the primary worker’s earning power, it may be easier to think of this in the context of a negative shock. There the loss of earnings by the primary worker is offset by increased work effort by the partner.

Whether these shifts in the dynamics can account for changes in higher frequency behavior of work effort is an open question. One line of argument would be that at the same time as the economy has shifted toward a more predominantly PS1 regime, the nature of the shocks hitting the economy has changed from more purely aggregate (as in

the A_m shock) to something that has a stronger sectoral component. Some have argued (e.g. Groshen and Potter, 2003) that recent recessions have been characterized by more “reallocation,” as opposed to purely aggregate disturbances. The predominance of such sectoral disturbances would point toward a reduced cyclical response of hours.

V. Conclusions

We have developed an extension of the standard growth model that accounts for long-term changes in average work effort in the population. The first part of the paper documents those changes, after controlling for demographics. The second part develops a growth model with home and market production, first with single-agent households, which is shown to be inadequate, and then with two-agent households. The two-agent version is shown to have implications for aggregate labor market trends that are broadly consistent with the facts. We then explore the implications of that model for the dynamic response to shocks. We find that, depending on the type of shock, the model implies varying responses over time. This suggests further exploration in light the changing cyclical behavior of the labor market over past 20 years.

References

- Abraham, Katharine, and Lawrence Katz, "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?" *Journal of Political Economy* 94, No. 3, 507-22 (1986).
- Benhabib, Jess, Richard Rogerson, and Randy Wright, "Homework in Macroeconomics I: Basic Theory," C.V. Starr Working Paper 90-17 (1990).
- Benhabib, Jess, Richard Rogerson, and Randy Wright, "Homework in Macroeconomics: Aggregate Fluctuations," *Journal of Political Economy* 99, no. 6, 1166-1187 (1991).
- Christiano, L., M. Eichenbaum, and R. Vigfusson, "What Happens after a Technology Shock," manuscript, 2004.
- Francis, N., and V. Ramey, 2003, "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited," manuscript, 2003.
- Francis, N., and V. Ramey, "A New Measure of Hours Per Capita," manuscript, 2004.
- Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu, "Engines of Liberation," *Review of Economic Studies*, v. 72, n. 1, 109-133 (2005).
- Groshen, Erica, and Simon Potter, "Has Structural Change Contributed to a Jobless Recovery?" *Current Issues in Economics and Finance* 9, No. 8, Federal Reserve Bank of New York (2003).
- Hill, M., "Patterns of Time Use," in *Time, Goods, and Well-Being*, edited by F.J. Juster and F.P. Stafford, University of Michigan (1985).
- Kaldor, Nicholas, "Capital Accumulation and Economic Growth," in F. Lutz and D.C. Hague, eds., *The Theory of Capital*, New York: St. Martin's Press, pp. 177-222 (1961).
- MaCurdy, Thomas, "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy* 89, no. 6, 1059-1085 (1981).
- McGrattan, Ellen, and Richard Rogerson "Changes in Hours Worked, 1950–2000," Federal Reserve Bank of Minneapolis *Quarterly Review* 28, No. 1 (2004).

Mulligan, Casey B. and Yona Rubinstein, "Household vs. personal accounts of the U.S. labor market, 1965-2000," NBER Working Paper 10320, 2004.

Uhlig, Harald, "A Toolkit for Analyzing Nonlinear Dynamic Stochastic Models Easily." Manuscript, 1997.

Table 1: Patterns of Time Use

	Age	Sex	Children	Education	Employed? Y, N, or share	Home work	Market work	Total work	Non- discretionary	
1	25+	both	Any	High School	n/a	2.76	3.65	6.41	10.47	
2				College+		2.53	4.72	7.25	10.31	
3	18+	M	Any	Any	0.746	1.74	4.81	6.55	10.32	
4		F			0.591	3.19	2.98	6.17	10.70	
5		M	Under age 6		0.896	2.59	5.98	8.57	9.94	
6		F			0.570	5.27	2.54	7.81	10.49	
7		M			Yes	2.45	6.59	9.04	9.84	
8		F			4.42	4.34	8.76	10.29		
9		M			No	3.8	0.71*	4.51	10.89	
10		F			6.4	0.15*	6.55	10.77		
11		M			Age 6-17	0.840	1.89	5.49	7.38	10.08
12		F				0.724	3.49	3.52	7.01	10.48
13		M	Yes			1.76	6.41	8.17	9.88	
14		F	3.09			4.84	7.93	10.39		
15		M	No			2.57	0.68*	3.28	11.24	
16		F	4.54			0.06*	4.60	10.69		
17		M	None			0.667	1.49	4.29	5.78	10.48
18		F				0.550	2.45	2.91	5.36	10.86
19		M			Yes	1.2	6.26	7.46	10.13	
20		F			1.91	5.24	7.15	10.53		
21		M			No	2.07	0.35*	2.42	11.21	
22		F			3.11	0.06*	3.17	11.23		
23		25-34	M		Any	Any	1.68	5.95	7.63	10.14
24			F				3.72	3.62	7.34	10.67
25	35-44	M	2.08	5.84			7.92	10.01		
26		F	3.96	3.81			7.77	10.44		
27	45-54	M	1.72	5.93			7.65	9.98		
28		F	3.95	3.92			7.87	10.39		

Table 2a: Steady State with Technical Progress in Household Sector, One-Person Households

A_z	C	Z	n_h	n_m	Y	i	P
1	0.528	0.555	0.359	0.291	1.029	0.501	1.717
1.1	0.548	0.595	0.350	0.298	1.053	0.505	1.560
1.21	0.568	0.637	0.341	0.304	1.077	0.509	1.419
1.331	0.589	0.682	0.331	0.311	1.102	0.513	1.290

Table 2b: Steady State with Technical Progress in Household Sector, Alternative Assumptions

A_z	A_c	ρ	η	C	Z	n_h	n_m	Y	i	P
1.1	1	-2	0.5	0.548	0.595	0.350	0.298	1.053	0.505	1.560
1.04	1.03	-2	0.5	0.554	0.583	0.359	0.291	1.080	0.526	1.717
1.1	1	-3	0.5	0.555	0.590	0.347	0.300	1.061	0.506	1.419
1.1	1	-2	1	0.547	0.594	0.349	0.297	1.051	0.504	1.560

Table 3: Steady States with Two-Person Households

A_z	C	Z	n_{h1}	n_{h2}	n_{m1}	n_{m2}	Y	i	P
0.62	0.574	0.649	0.032	0.644	0.613	0	1.084	0.510	1.382
0.68	0.594	0.694	0.015	0.642	0.627	0	1.108	0.514	1.257
0.75	0.615	0.741	0	0.639	0.640	0	1.132	0.517	1.143
0.83	0.629	0.793	0	0.627	0.644	0	1.139	0.511	1.049
0.91	0.642	0.848	0	0.616	0.648	0	1.146	0.504	0.964
1.0	0.655	0.905	0	0.604	0.652	0	1.164	0.498	0.885
1.1	0.668	0.966	0	0.592	0.656	0	1.160	0.492	0.813
1.21	0.684	1.026	0	0.573	0.655	0.013	1.177	0.493	0.742
1.331	0.703	1.088	0	0.553	0.653	0.031	1.199	0.497	0.674
1.331	0.703	1.088	0	0.553	0.653	0.031	1.199	0.497	0.674

Figure 1: Hours Per Capita, Business Sector (1992=100)




Figure 2: Average Hours from Census and CPS Data

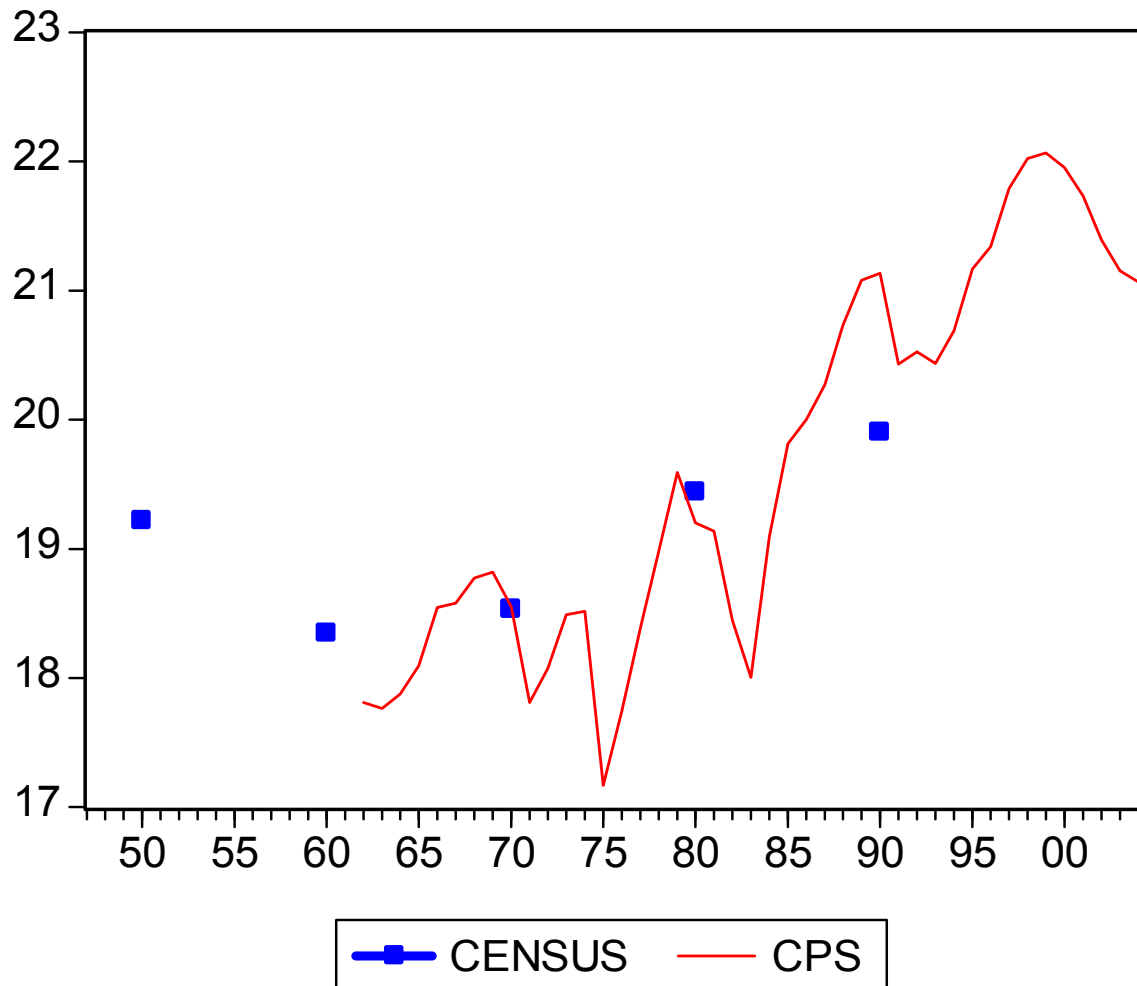


Figure 3: Hours per capita trends, CPS

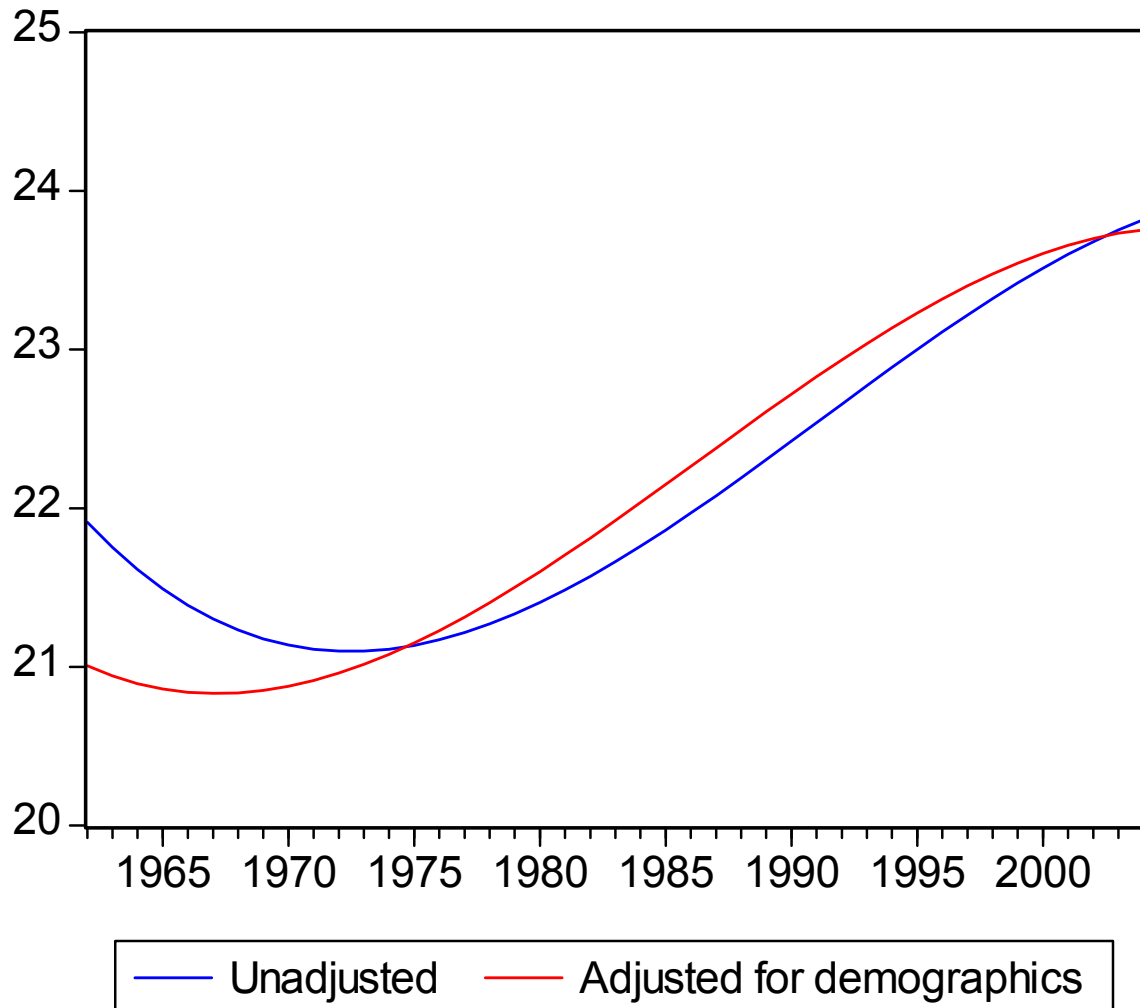


Figure 4
Census Hours, with and without Demographic Controls

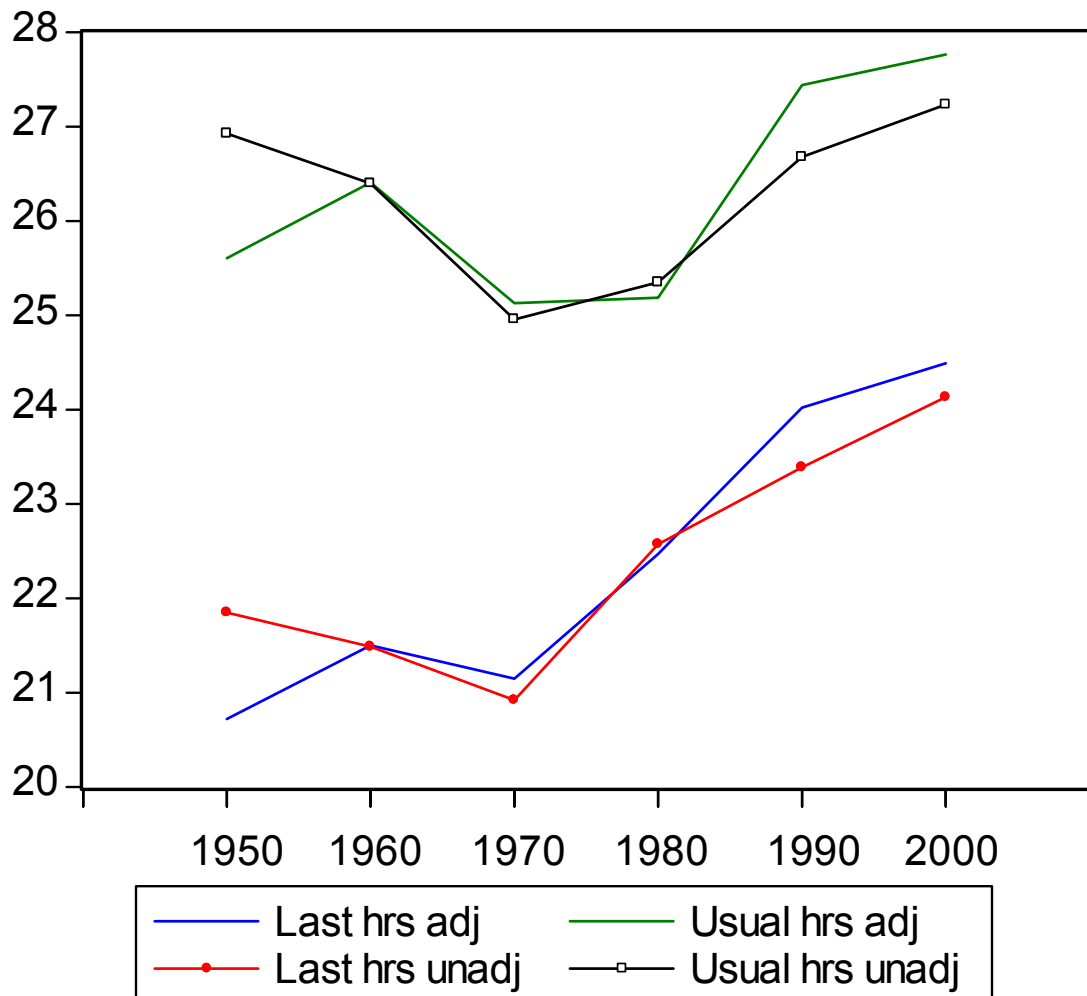


Figure 5: Effect of Fixed Demographic Weights on Hours Per Capita




Figure 6: Assortive Matching




Figure 7: Path of Total and Primary Earner's Market Hours

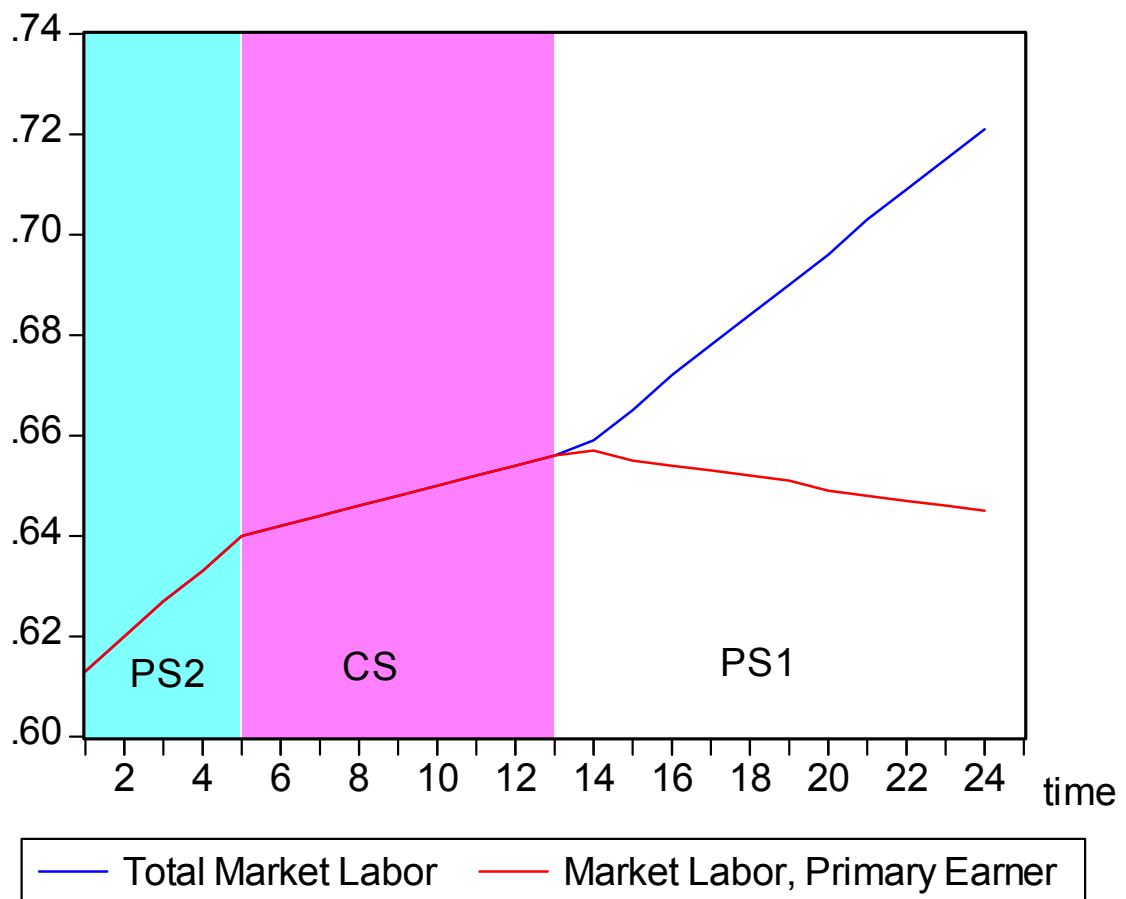


Figure 8: Effect of Relative Skill Increase of Secondary Earner on Family Market Hours




Figure 9: Impact of Ac shock (10%) on market hours

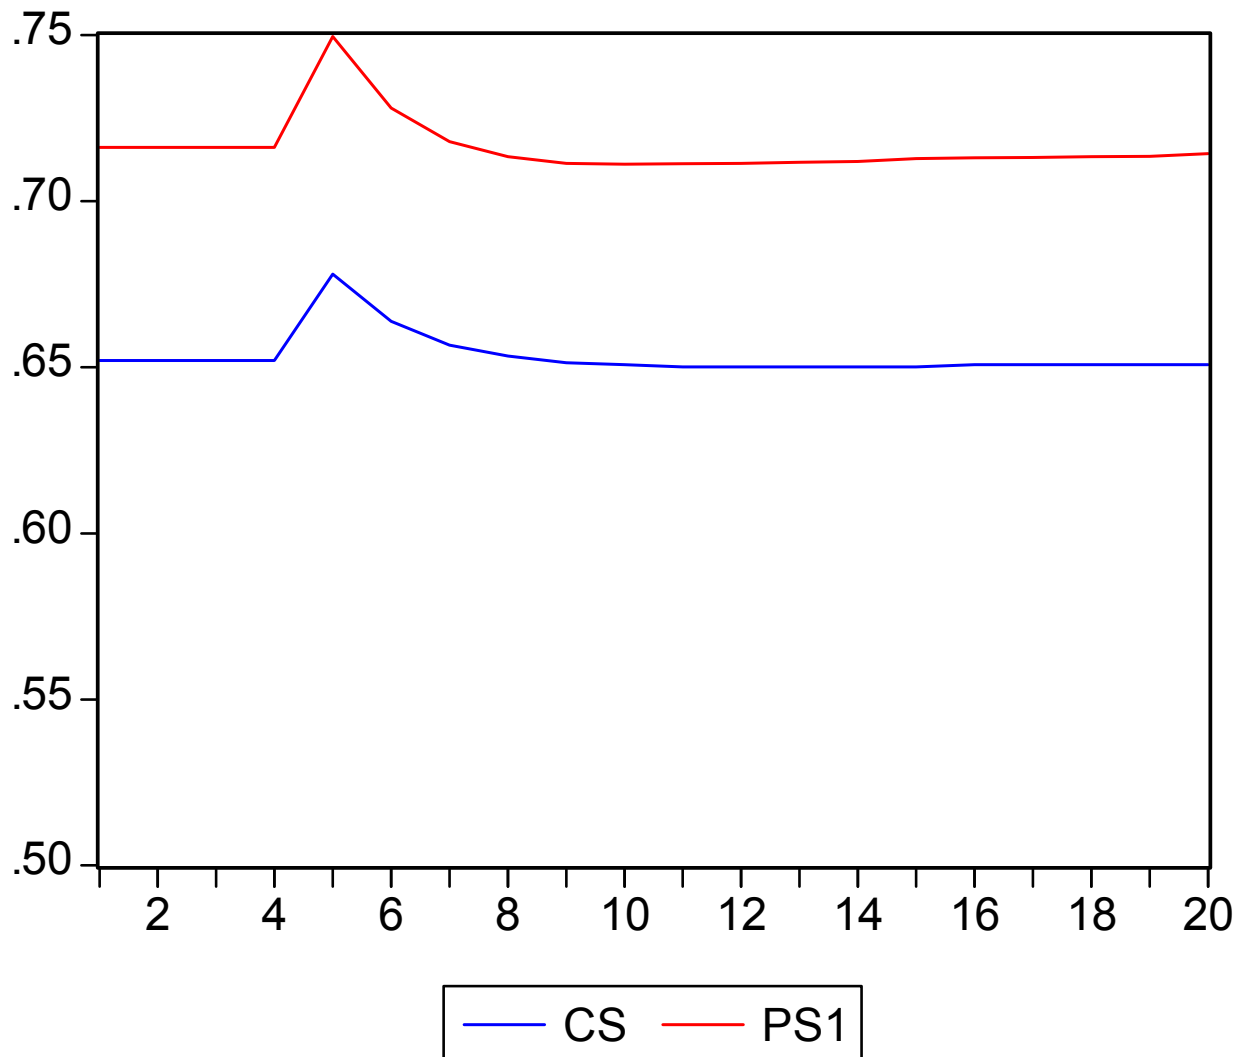



Figure 10: Impact of μ_1 shock (10%) on market hours

Figure 11:
Breakdown of market work effect of μ_1 shock in PS1 regime

