

The Business Cycle and the Life Cycle*

Paul Gomme

The University of Iowa

and

Federal Reserve Bank of Cleveland

Richard Rogerson

Arizona State University

and

NBER

Peter Rupert

The University of Western Ontario

and

Federal Reserve Bank of Cleveland

Randall Wright

The University of Pennsylvania

and

NBER

June 7, 2004

Abstract

This paper documents the differences in variability of hours worked over the business cycle across several demographic groups and shows that these differences are large. We argue that understanding these differences should be useful in understanding the forces that account for aggregate fluctuations in hours worked. In particular, it is well known that standard models of the business cycle driven by technology shocks do not account for all of the variability in hours of work. This raises the question of to what extent the forces in this model can account for the differences across demographic groups. We explore this in the context of hours fluctuations by age groups by formulating and analyzing a stochastic overlapping generations model. Our analysis shows that the model does a good job of accounting for hours fluctuations of prime age workers, but not for young or old workers. We conclude that a key issue is to understand why fluctuations for young and old workers are so much larger.

*The authors thank Mark Gertler, Paul Klein, Eva Nagyal, Ken Rogoff, Robert Shimer and conference participants for useful comments. Rogerson and Wright thank the NSF for financial support.

1 Introduction

The representative household model is the workhorse of modern business cycle theory. One can understand this from several perspectives. First, from an empirical perspective, the business cycle is defined in terms of time series variation in the per capita values for several key aggregate variables. By construction, the representative agent model is a model of per capita values. Second, from a conceptual perspective, the process of understanding is facilitated by first analyzing economic forces in simple settings, and abstracting from heterogeneity helps to maintain simplicity in the model. Third, from a technical perspective, the appropriate theoretical framework in modern business cycle theory is dynamic stochastic general equilibrium theory, and assuming a representative agent greatly reduces the burden of such analysis, both computationally and theoretically.

These factors suggest that representative agent models are a useful starting point for analyzing the economic forces that shape aggregate fluctuations. However, the thesis of this paper is that our understanding of labor market fluctuations (in particular) will be enhanced by moving beyond the representative agent model. The essence of our argument follows from a simple empirical finding. As we document, the magnitude of business cycle fluctuations in hours of market work varies quite significantly across subgroups in the population. We believe that understanding why some groups fluctuate more than others should be relevant for understanding why the aggregate fluctuates as much as it does. Consider two scenarios. In the first, suppose that for reasonable parameterizations a given model is *unable* to account for a sizeable fraction of observed fluctuations in aggregate hours. In assessing which modifications to the theory may be most relevant, it would be important to know if the problem was that the model systematically underaccounts for fluctuations in hours across all groups, or if the problem is that it cannot account for the magnitude of fluctuations experienced by some specific groups. In the second scenario, suppose that for reasonable parameterizations a given model is *able* to account for the bulk of aggregate fluctuations in hours. While this is useful information, we would obviously be more confident that the economic forces captured in this model are indeed the relevant ones if they were also able to account for the patterns of fluctuations across various groups.

In this paper we pursue a disaggregated analysis of fluctuations in market work by considering one specific dimension of heterogeneity – age. Specifically, we document how cyclical fluctuations in hours of market work vary over the life cycle, and then assess the predictions of a life-cycle version of the growth model for the observations. Our analysis yields a simple but striking finding. The main discrepancy between the model and the data lies in the inability of the model to account for fluctuations in hours for individuals over the first half of their life cycle; it can account for most of the fluctuations for individuals aged 45 – 64 without resorting to extreme labor supply elasticities. This suggests that in looking for alternative theories to better account for aggregate labor market fluctuations, attention should be directed toward features that specifically affect individuals during the first half of their life. Although the goal of this paper is not to present alternatives to the benchmark life-cycle growth model, one is led to think about the options: e.g., to ask whether search frictions, say as opposed to sticky wage models or other candidates, may be more relevant in terms affecting workers differently at different stages of the life-cycle.¹ In this sense, our goal is to raise some issues without trying to resolve everything here.

While heterogeneity has received a lot of recent attention in macroeconomics, it is important to distinguish our emphasis from that of others. A recurring issue in many studies is whether introducing a particular type of heterogeneity, often in connection with some other feature, will influence the properties of the *aggregate* time series. In these studies, the emphasis remains on the properties of the aggregate variables, and not on the behavior of disaggregated series. One example of this is Krusell and Smith (1998), who ask whether a model with idiosyncratic income shocks and incomplete markets would produce different aggregate responses to technology shocks. Another example is Rios-Rull (1996), who studies a similar model to the one used here, but whose main objective was to see if aggregate fluctuations were different in an overlapping generations model than in the standard infinitely lived agent model. Both of these studies concluded that the properties of aggregate fluctuations were not much affected. In contrast, our goal is to ask whether allowing for heterogeneity provides more insight into the details of a particular shock and

¹For example, if one was thinking that the main shortcoming of the model was the absence of rigidities in real wages, one would have to argue that this feature is more important for young workers than for prime-age workers.

propagation mechanism by explicitly focusing on the implications of the model for fluctuations at the disaggregated level.²

Though our work is related to several papers in the literature, there are two papers that particularly relevant. The first is Clark and Summers (1981), who documented that cyclical fluctuations in employment vary across demographic groups, and the second is Rios-Rull (1996), who examined fluctuations in a life cycle economy. Our empirical work extends Clark and Summers along several important dimensions. Specifically, we analyze additional dimensions of heterogeneity, use more conventional methods to define cyclical components, examine both the intensive and extensive margins, and perform additional robustness checks. While our results for fluctuations by age are similar to theirs, we find differences along other dimensions. Our theoretical work also extends the work of Rios-Rull along several dimensions. Specifically, we consider a different class of preferences, our model allows for home production and life-cycle preference shifters, and we assume a different market structure. Most importantly, however, we carry out a detailed analysis of the role that various factors play in shaping the volatility of hours over the life cycle.

Although we do not pursue it here, we believe that the life cycle model developed and analyzed in this work is of independent interest in other contexts as well. For example, it would allow one to study how fluctuations in cohort size affect economic outcomes. Shimer (1998), for example, argued empirically that fluctuations in cohort size had a large impact on fluctuations in aggregate unemployment.

The rest of the paper is organized as follows. In Section 2 we describe a standard representative household (infinitely-lived agent) model and examine its predictions concerning fluctuations in aggregate hours. Section 3 documents the extent to which the cyclical variation in hours varies with several household characteristics. Section 4 presents and calibrates our version of the growth model populated by overlapping generations. Section 5 presents the results of the model concerning business cycle fluctuations, with a particular focus on its implications for fluctuations in hours

²Researchers have previously suggested that trying to understand fluctuations at a more disaggregate level will be useful; an example is Gertler and Gilchrist (1994) who argued that the differences in behavior of large and small manufacturing firms provided additional information about the nature and propagation of aggregate shocks.

by age. Section 6 is devoted to discussing the factors that give rise to the observed pattern of fluctuations. Section 7 presents some international evidence on fluctuations in hours worked by age, and Section 8 concludes.

2 A Representative Agent Model

For purposes of comparison, it is instructive to start with a representative agent model of the sort that serves as one of the benchmark models of business cycle analysis. Rather than formulating the model in its most general form, we restrict attention to a specification with commonly used functional forms. We add two features relative to the simplest possible specification: household production and a government sector. We include household production because previous work has shown that models with household production do a much better job of accounting for several aspects of business cycles, particularly for hours of market work.³ We include a government sector because taxes are an important element in calibrating home and market capital stocks.

2.1 Model

There is an infinitely-lived representative household with preferences

$$\sum_{t=0}^{\infty} \beta^t [\log C_t - \frac{\omega}{\gamma} H_t^\gamma],$$

where $\beta \in (0, 1)$ is the discount factor, C_t is a CES aggregator of market and home consumption in period t , and H_t is total time spent working in the market and at home in period t . That is,

$$C_t = [\psi C_{mt}^\xi + (1 - \psi) C_{ht}^\xi]^{1/\xi}$$

$$H_t = H_{mt} + H_{ht},$$

³Some of the standard references on home production in macroeconomics include Benhabib, Rogerson and Wright (1991), Greenwood and Hercowitz (1991), Greenwood, Rogerson and Wright (1995), McGrattan, Rogerson and Wright (1997), Baxter and Jermann (1999) and Gomme, Kydland and Rupert (2001).

where C_{mt} and C_{nt} are market and home consumption respectively, and H_{mt} and H_{nt} are market and home work respectively. The agent is endowed with one unit of time each period and K_0 units of capital at $t = 0$. The parameters $\gamma \geq 1$ and $\xi \leq 1$ play a key role in influencing the business cycle predictions of the model, since γ determines the substitutability in hours worked across time and $\xi \leq 1$ determines the extent of substitutability between home and market goods. As a result, these parameters dictate the amount of intertemporal and intratemporal substitution in hours of market work.

We choose the utility function $\log C - \frac{\omega}{\gamma} H_t^\gamma$ to facilitate comparison with the large literature in labor economics that tries to estimate γ . The standard life-cycle labor literature (without home production) typically assumes separability in the sense that $U(C_t, H_t) = u(C_t) + v(H_t)$. In a deterministic setting, this means the first order condition for h_t can be written

$$v'(H_t) = -w_t \lambda,$$

where λ is the Lagrange multiplier on the lifetime budget constraint and w_t is the wage in period t . Due to separability this condition does not include C_t , so one can take the equation to the data without having to observe consumption. If $v(H) = -\frac{\omega}{\gamma} H^\gamma$ then, after taking logs and rearranging, we have

$$\log(H_t) = \alpha_0 + \frac{1}{\gamma-1} \log(w_t),$$

where α_0 can incorporate a constant, a time trend and an error term, depending on assumptions. From this one can estimate the elasticity $1/(\gamma-1)$ and recover the structural parameter γ .⁴

The above analysis does not require specifying $u(C)$. It is well known in macro, however, that balanced growth requires either $U(C, H) = C^\sigma v(H)$ or $U(C, H) = \log(C) + v(H)$ for some function $v(H)$. Hence, assuming separability so that we can apply the labor supply results, we are led to

$$U(C, H) = \log(C) + v(H)$$

⁴Standard references in the labor supply literature include, for example, MaCurdy (1981), Altonji (1986), and Pencavel (1986).

for some function $v(H)$. Although in principle any function $v(H)$ satisfying the usual regularity conditions would do, we will adopt the common specification $v(H) = -\frac{\omega}{\gamma}H^\gamma$.

Incorporating home production into the analysis now merely requires reinterpreting C and H as composites of market and home consumption and of market and home work: $C = C(C_M, C_H)$ and $H = H(H_M, H_H)$. Here we follow much of the previous literature by assuming $C = [\psi C_m^\xi + (1 - \psi)C_n^\xi]^{1/\xi}$, so that we can appeal to existing estimates of the parameter ξ , and $h = h_m + h_n$, which means hours worked in the market and home are perfect substitutes. The last thing to say about preference is that although there will be a government in the model, we assume that agents derive no utility from government consumption.⁵

In terms of technology, there is a production function

$$Y_{mt} = z_t K_{mt}^\theta [(1 + g)^t H_{mt}]^{1-\theta},$$

where Y_{mt} is market output in period t , K_{mt} and H_{mt} are capital and labor services used in market production in period t , z_t is a technology shock, and g represents the constant rate of labor augmenting technological change. We assume that z_t follows the process

$$\log z_{t+1} = \rho \log z_t + \varepsilon_{t+1},$$

where ε_t is an i.i.d. random variable that is normally distributed with mean μ_m and variance σ_m^2 . The period t realization of ε is observed before any decisions are made. Market output produced in period t can be used either as market consumption C_{mt} , government consumption G_t , or investment I_t :

$$C_{mt} + G_t + I_t = Y_{mt}.$$

⁵Equivalently, we could assume that individuals do derive utility from government consumption but that it is separable with respect to the other arguments, which accords with results in the literature, like Christiano and Eichenbaum (1992) and McGrattan et al. (1997).

There is also a production function for home produced goods,

$$Y_{nt} = K_{nt}^{\eta} [(1+g)^t H_{nt}]^{1-\eta},$$

where Y_{nt} is household production in period t , K_{nt} and H_{nt} are capital and labor services used in home production in period t and g again represents the constant rate of labor augmenting technological change. We assume the same rate of technological change in the two production functions, as is required for balanced growth. Although we assume the home production function is Cobb-Douglas, following Greenwood and Hercowitz (1991) some authors have argued that departures from Cobb-Douglas are crucial for understanding certain issues, including the pattern of investments in home and market capital. Moreover, the estimates in McGrattan et al. (1997) imply the home production function is significantly different from Cobb-Douglas (the model actually allows both market and home production functions to be CES, but the estimates implied only the latter is significantly different from Cobb-Douglas). For the issues on which we focus, however, this does not matter much, so we use Cobb-Douglas for simplicity. We also abstract from shocks to the home production function, since they will not play any role in the subsequent analysis.⁶

One asymmetry between market and home production is that the only use of home produced output is as home consumption, i.e.,

$$C_{nt} = Y_{nt}.$$

That is, although capital is used in home production, it is produced only in the market sector.

Capital accumulation is given by

$$K_{mt+1} = (1 - \delta_m) K_{mt} + I_{mt}$$

$$K_{nt+1} = (1 - \delta_n) K_{nt} + I_{nt}$$

⁶For some other issues, however, shocks to the home technology are crucial; for example, Benhabib, Rogerson and Wright (1991) show how such shocks can generate a reasonable contemporaneous correlation between productivity and hours, and Hall (1997) finds that home sector shocks are actually a significant source of aggregate fluctuations.

where I_{mt} and I_{nt} are investment in market and home capital in period t , both of which are constrained to be nonnegative, while $\delta_m \in (0,1)$ and $\delta_n \in (0,1)$ are depreciation rates. Aggregate investment in period t is the sum of investment in home and market capital.⁷

$$I_t = I_{mt} + I_{nt}$$

It is well known that empirically plausible tax rates can have big effects in this model. Since we will be choosing some parameter values by calibrating to steady state values, it is important to incorporate taxes into the specification. However, given that our primary reason for doing so is to facilitate calibration, we assume constant tax rates. In particular, we assume that market labor income is taxed at the constant rate τ_h , and capital income is taxed at the constant rate τ_k . The government uses tax revenues to finance spending G_t , which we assume is a constant ratio of market output. The government faces a period-by-period budget constraint, with lump sum transfers τ_t serving to achieve budget balance.

2.2 Parameterization

Calibration of parameter values for this model is fairly standard. Because of this, and also because we will go into detail on the calibration of the overlapping generations model later in the paper, we do not provide details here, and simply report the parameter values in Table 1. Note that we set a period to be a year in this paper. While it is more common in infinitely lived agent models to use a quarter rather than a year, the basic properties of the model are not affected by this choice. We will be using an annual model once we introduce overlapping generations because the data on hours worked by age is at annual frequency. Not shown in Table 1 is that we assume government spending relative to market output is .20.

There are also four utility parameters not listed in Table 2: the elasticity parameters γ and ξ , and the coefficients giving weights on market versus home consumption and on hours versus total

⁷Alternatively, one could assume that installed home and market capital can be costlessly transformed into each other, as in Benhabib et al. (1991). For the issues on which we focus here, this would not make much of a difference.

Table 1: Parameters for Infinitely-Lived Household Calibration

β	θ	η	δ_m	δ_n	g	ρ	σ_ε	τ_h	τ_k
.954	.30	.27	.065	.057	.018	.806	.0139	.25	.50

consumption, ψ and ω . The standard procedure for determining values for these parameters is to set ω and ψ so that the steady state values of H_m and H_n are equal to some target values taken from the data, typically $H_m = 1/3$ and $H_n = 1/4$, and to set γ and ξ in accord with the empirical literature, as they cannot be easily pinned down by steady state considerations. It is well known that the values of the two elasticity parameters γ and ξ matter a lot for the cyclical properties of hours. There is also considerable controversy over these parameters, and estimates can vary a lot depending upon which group one looks at (e.g., males versus females), which features are incorporated (e.g., skill accumulation, home production), and which margins one considers (i.e., the intensive versus the extensive margin). Hence, in what follows we will present results for a wide range of values for γ and ξ , without necessarily taking a stand on any particular value.

2.3 Results

It is well known that models of this sort can mimic the broad features of cyclical fluctuations in the US economy, although the magnitude of fluctuations in market hours has received considerable attention since the model does less well on this (see Hansen and Wright 1992). In particular, we will emphasize the relative standard deviation of market hours to market output.⁸ In the data, the relevant number is .80, which is the standard deviation of market hours based on CPS data over the period 1962-2000, relative to the standard deviation of output over this same period. The standard deviation of hours is 1.79 and the standard deviation of output is 2.23, where these numbers are annual and correspond to data that has been HP-filtered with a smoothing parameter of 100. We note that our output measure excludes the service flow from housing, since we treat this as a

⁸In assessing the magnitude of fluctuations in hours, two normalizations have been used in the literature – one normalizes relative to fluctuations in output, while the other normalizes relative to fluctuations in average labor productivity. For our purposes this choice does not matter, as the benchmark specification accounts for roughly two-thirds of fluctuations by either metric. Hence we will simply report the relative volatility of hours to output.

nonmarket service.

Given the calibration strategy discussed above, Table 3 reports values for the standard deviation of market work relative to market output (both HP filtered) that come from simulating the model. We generate samples that are 39 periods in length, the same as our data, and average over 1000 runs. Table 3 reports the results for values of γ ranging from 1 to 11, which corresponds to elasticities ranging from infinite to .1, and values of ξ ranging from 0 to .8, which corresponds to elasticities ranging from 1 to 5. As can be seen, and as is fairly well known, the model can account for most of the fluctuations in market hours if (and only if) the elasticities are sufficiently large. Perhaps somewhat less well known is that it is not sufficient to know the value of the intertemporal elasticity parameter γ in order to assess the model on this dimension, since even if γ is set to 11 the model would still be able to account for the bulk of the observed fluctuations if the value of ξ were sufficiently high.⁹

Table 2: Standard Deviation of Market Hours Relative to Output

	$\gamma = 1$	$\gamma = 2$	$\gamma = 2.5$	$\gamma = 3.0$	$\gamma = 4$	$\gamma = 11$
$\xi = 0$.66	.48	.45	.43	.40	.35
$\xi = .2$.67	.50	.47	.45	.43	.38
$\xi = .4$.68	.53	.51	.49	.47	.44
$\xi = .5$.69	.56	.54	.52	.50	.47
$\xi = .6$.71	.59	.57	.56	.54	.52
$\xi = .8$.78	.71	.70	.69	.69	.68

There is considerable debate over the appropriate value of γ for the representative household, and even less is known about the parameter ξ . This notwithstanding, just to fix ideas suppose we use values in the upper part of the plausible range, in particular $\gamma = 2.5$ and $\xi = .5$, corresponding to elasticities of $2/3$ and 2 respectively. The implied relative standard deviation of hours is equal to .54, roughly two-thirds of what we observe in the data. This suggests that although the model accounts for a substantial fraction of the volatility in market hours, there is also a sizeable fraction

⁹Our choice of metric for assessing the magnitude of fluctuations in hours was made in the context of an analysis that emphasizes technology shocks. For a model driven by monetary shocks this may not be a good metric. Nonetheless, we believe the basic point—that one should consider implications for disaggregated data—to be relevant for all business cycle models.

that it misses.

One is then naturally lead to consider modifications to the model in order to better match the behavior of market hours. Many such modifications have been proposed in the literature, including alternative specifications of preferences (e.g., Kydland and Prescott (1982)), search frictions (e.g., Merz (1994) and Andolfatto (1996), Den Haan, Ramey and Watson (2000)), informational asymmetries (e.g., Gomme (1999) and Alexopoulos (2004)), restrictions on working hours (e.g., Rogerson (1988) and Hansen (1985)), alternative formulations of technology (e.g., Kydland and Prescott (1988)) and alternative wage setting mechanisms (e.g., Danthine and Donaldson (1995)). What we want to argue in the remainder of this paper is that if we are looking for ways to isolate the key empirical deficiencies of equilibrium macro models like the one described above, one ought to consider a lower level of aggregation. In the next section we document a wide range in variability of market hours over the business cycle across subgroups. In view of this, it seems interesting to ask whether the mechanism implicit in the model underaccounts for fluctuations across all subgroups, or if perhaps it does account for the fluctuations of some groups but not others. Put somewhat differently, if we are trying to understand the causes of fluctuations in hours of work over the business cycle, it seems reasonable that understanding why some groups fluctuate much more than other groups would be a key piece of information.

3 Beyond Aggregate Data

In this section we document differences in the magnitude of cyclical fluctuations in market hours across groups in the population. In particular, we disaggregate by age, education, marital status, gender, and industry of employment.

3.1 Fluctuations by Age

Using data from the March Supplement of the CPS for the period 1962-2000 we compute aggregate market hours per capita for the entire population aged 16 and above, and market hours per capita

by age for seven age groups: 16 – 19, 20 – 24, 25 – 34, 35 – 44, 45 – 54, 55 – 64, and 65+. The standard procedure for defining the business cycle component of an aggregate series is the percent deviation from a suitably defined trend, defined here using the HP filter. To extract the component of fluctuations in aggregate hours that is accounted for by each age group we use a two-step procedure described in the Appendix. The results are in Table 3.

Table 3: Relative Cyclical Fluctuations of Hours by Age Group

	16 – 19	20 – 24	25 – 34	35 – 44	45 – 54	55 – 64	65+
σ_h/σ_{Y_m}	2.23	1.23	.86	.64	.57	.59	1.26
% of H_m	4	11	26	25	20	12	2
% of σ_{H_m}	11	16	28	19	14	9	3

The first row shows fluctuations in the hours of each age group relative to output. The second row indicates the fraction of average hours worked by each age group over the entire sample period. The third row indicates the fraction of fluctuations accounted for by each age group. Several interesting patterns emerge. First, note that the pattern of fluctuations across age groups is *U-shaped*: fluctuations are highest for young and old workers, and are lowest for middle aged workers. Related to this, the second and third rows indicate the extent to which cyclical fluctuations in hours are disproportionately accounted for by fluctuations in the hours of work of younger workers. Workers aged 16 – 24 account for only 15% of total market hours, but more than 25% of fluctuations in market hours. Conversely, prime-age workers, between the ages of 35 and 54, account for 45% of total market hours but for only 33% of fluctuations in market hours.¹⁰

One may be concerned that the patterns displayed above are not due to age effects per se, but are really an artifact of a situation in which workers of different ages work at different jobs, with some jobs being more cyclically volatile than others. To explore this possibility we examine the role that the age distribution of hours worked across one digit industries may play in shaping fluctuations by age group. Table 4 shows the relative volatility of hours of work across one-digit sectors. As is well known, hours in some sectors fluctuate much more over the business cycle. In

¹⁰Using slightly different methods, this basic observation has previously been noted by Clark and Summers (1981). See also Blanchard and Diamond (1990), and Keane and Prasad (1996).

particular, goods-producing sectors display more volatility than service sectors. Table 5 indicates the distribution of hours worked by each age group across each of these eight sectors:

Table 4: Fluctuations of Sectoral Hours Relative to GDP

Sector	Agr	Min	Cons	Mfg	Trans	Wh/Re Tr	FIRE	Serv
σ_h/σ_{y_m}	1.44	1.69	2.17	1.24	.75	.80	.55	.57

Table 5: Sectoral Distribution of Hours by Age

	16 – 19	20 – 24	25 – 34	35 – 44	45 – 54	55 – 64	65+
Agr	5.17	2.99	2.91	3.20	3.81	.01	12.89
Min	.45	.92	1.08	1.06	.98	.90	.43
Cons	4.59	6.26	6.81	6.56	5.87	5.91	3.66
Mfg	15.78	22.34	23.82	24.23	24.78	25.07	11.04
Trans	2.97	5.94	7.86	8.62	8.35	7.70	3.59
Wh/Re Tr	44.30	26.10	19.52	17.98	18.30	20.43	22.47
FIRE	4.43	7.13	6.93	6.38	6.21	6.52	7.55
Ser	22.32	28.33	31.08	31.97	31.70	33.46	38.37

Table 5 reveals that there are indeed some sharp differences across age groups in how their hours of work are distributed across sectors. For example, the distribution of hours to the manufacturing sector is increasing in age up until 65. Hours of work of teenagers are heavily skewed toward the wholesale and retail trade sector. Mining, Construction, and Transportation and Public Utilities all display an inverted U shape across age groups whereas Wholesale and Retail Trade displays a U shape across age groups. Given that each age group has a distinctive pattern of hours across sectors and that fluctuations vary across sectors, we can ask what how fluctuations would vary by age if the only source of differences by age were the sectoral distribution of hours. The answer is given in Table 6, where for each group we take a weighted average of the sectoral relative volatilities using the age distribution of hours across sectors as weights.

Table 6: Relative Fluctuations Induced by Sectoral Composition

Age	16 – 19	20 – 24	25 – 34	35 – 44	45 – 54	55 – 64	65+
$\sigma_{h_m}/\sigma_{y_m}$.81	.83	.84	.84	.84	.81	.79

Two observations emerge from this exercise. First, the size of the effects induced by differences in sectoral composition across age groups is small – the range of values goes only from .79 to .84. We conclude from this finding that abstracting from sectoral composition effects is a reasonable thing to do. Second, to the extent that sectoral composition effects do matter, they actually generate an inverted U for the pattern of volatility across age, the opposite of what we see in the data. It follows that the pattern in Table 3 would be even more pronounced if we controlled for sectoral composition.

Given previous research on the importance of the intensive and extensive margins of labor hours adjustment, it is also of interest to see if there is any systematic difference across groups in the importance of the two margins. However, analysis revealed this not to be the case. We found that the percentage of total fluctuations accounted for by the extensive margin varied within the relatively narrow range of 68 – 72%.

To close this subsection, recall the discussion of the infinitely-lived representative agent model in the previous section. We suggested that an empirically plausible parameterization generates relative fluctuations in hours of work equal to .54. Recall also the first row of Table 3 which showed the relative volatility of hours by age. Looking at this row with the number .54 in mind raises a key issue. One interpretation of these findings is that the previous model is actually successful in accounting for the fluctuations of hours of prime aged individuals, and that its main shortcoming is accounting for the fluctuations in hours of younger workers. Of course, this interpretation is not warranted in the infinitely lived agent framework; we need to consider a model in which agents differ by age. We pursue exactly this in the remaining sections of this paper. However, before we do so, we think it is also of interest to examine the heterogeneity in hours fluctuations along some additional dimensions.

3.2 Fluctuations by Education

Here we repeat the previous analysis, but this time splitting the population by education. Because we do not model fluctuations by educational attainment in our theoretical analysis below, we do

not carry out as extensive an analysis of this case as we did for age. Also, due to data issues we restrict attention to the years 1974-2000. Furthermore, because measuring educational attainment for young workers is difficult, we restrict attention to individuals age 25 or greater. For each year, using the March CPS, we compute hours per person for individuals in four educational groups: 1. those with less than high school; 2. those with exactly high school; 3. those with some college but no college degree; and 4. those with at least a college degree. Again, we extract the component of fluctuations in aggregate hours accounted for by each of these groups.

Table 7: Hours Fluctuations by Education, 1974-2000

	1	2	3	4
$\sigma_{h_m}/\sigma_{Y_m}$	1.20	.84	.67	.23

The results are reported in Table 7. This table raises a similar issue to the one in the previous subsection: is it possible that the model does a good job of explaining fluctuations for those with, say, at least some college education, and that the model's shortcomings are entirely to do with the fluctuations in the other groups? Although we will not address this question explicitly in this paper, we believe that understanding the sources of these differences in volatility across education groups may also help us better understand the fluctuations in the aggregate data.

3.3 Fluctuations by Gender and Marital Status

In the model that we study below, we will continue to take the unit of analysis to be a household, and we will assume the sole dimension along which households differ is age. In particular, we will abstract from differences in household size, and we will abstract from the issue of time allocations across household members in multi-member households. However, we still think it is interesting to examine the extent to which fluctuations in market hours differ along the dimensions of marital status and gender (if for no other reason, this helps in assessing the extent to which our abstractions are warranted). Using the same procedure as above, we extract the component of aggregate hours fluctuations that is accounted for first by men and women, and then by married and unmarried

individuals. For these series the data cover the period 1962-2000.

Table 8: Hours Fluctuations by Gender, 1962-2000

	Males	Females
$\sigma_{h_m}/\sigma_{Y_m}$.87	.73

Table 8 reports relative fluctuations for men and women. Somewhat surprisingly, men display larger cyclical fluctuations than do women. One may suspect that part of this difference is accounted for by sectoral composition patterns. Table 9 shows that this is indeed a factor. This table shows that in several sectors, notably manufacturing and wholesale and retail trade, males display less volatility than females.

Table 9: Hours Fluctuations by Gender and Sector ($\sigma_{h_m}/\sigma_{Y_m}$), 1962-2000

	Agr	Min	Cons	Mfg	Tran	Wh/Re Tr	FIRE	Ser
Males	1.49	1.64	1.99	1.09	.72	.70	.61	.63
Females	1.46	2.28	1.51	1.27	.68	.83	.47	.52

It is also of interest to examine how the variability across the life-cycle varies by gender. In Table 10 we report relative variability by age, using the same procedure as before. Table 10 shows the pattern of volatility over the life-cycle is *U*-shaped for both males and females, though the timing of the trough is different across the groups – for men the volatility begins to increase in the 55 – 64 group, whereas for women it does not increase until the 65+ group. Quantitatively there are some differences – for younger workers the volatility of hours is somewhat less for females, while for individuals aged 35 – 44 the volatility is somewhat higher for females. Taken together, we interpret these findings with respect to gender as supporting our decision to abstract from the within family decision in the analysis that follows.¹¹

Table 11 reports the results disaggregated by marital status, where we note that fluctuations here are at the individual level and not the household level. Interestingly, fluctuations for single individuals are significantly larger than they are for married individuals. It is important to keep in

¹¹There are some attempts to explicitly model within-household time allocation over the business cycle; see e.g. Cho and Rogerson (1988).

Table 10: Relative Standard Deviation of Hours by Age and Gender ($\sigma_{h_m}/\sigma_{Y_m}$), 1962-2000

	16 – 19	20 – 24	25 – 34	35 – 44	45 – 54	55 – 64	65+
Male	2.60	1.52	.96	.64	.60	.73	1.16
Female	2.01	1.00	.75	.69	.55	.45	1.50

mind that among prime aged individuals the majority are married, so that for these individuals the aggregate numbers look very similar to those of the married group. However, for younger groups the reverse is true. Here again note that the basic pattern is *U*-shaped for both groups, with the exception of the observation for the 45 – 54 group that is not married. However, for this group the majority of fluctuations in hours are not correlated with movements in aggregate hours, so this number does not necessarily reflect the overall fluctuations for this group. We conclude from Table 11 that there are indeed differences across married and not married individuals. If we interpret our model as applying to married households, it follows that the target levels of volatility are somewhat lower as a result.

Table 11: Standard Deviation of Hours by Marital Status and Age, 1962-2000

	16 – 19	20 – 24	25 – 34	35 – 44	45 – 54	55 – 64	65+
Married	1.98	1.03	.79	.61	.52	.61	1.19
Not Married	2.35	1.36	1.07	.80	.87	.61	1.47

4 A Life-Cycle Model

In the remainder of the paper we focus on the nature of fluctuations by age. The obvious way to incorporate this type of heterogeneity into the model is to move from the infinitely lived agent framework to the overlapping generations framework. The goal will be to retain the basic structure of the infinitely lived household model as much as possible, while introducing life-cycle considerations. To maintain simplicity, some of these considerations will be captured in somewhat of a reduced form manner – for example, rather than explicitly modeling fertility, we will simply assume preference parameters change systematically along the life-cycle.

4.1 Model

In each period a representative T -period lived household is born. We abstract from population growth and assume that length of life is deterministic, although one could extend things on these dimensions at some cost in terms of simplicity. In our quantitative analysis we interpret a period to be a year and set $T = 55$, and think of a household beginning economic decision making at age 20 and continuing until 75. We impose exogenously that agents retire at age T_R . Agents have preferences over lifetime profiles of consumption and work. For a generic variable s we will use s_t^a to denote the value of s for an agent of age a in period t , and we use lower (upper) case letters to represent choices at the individual (aggregate) level. Hence, for an agent born in period t preferences are given by

$$\sum_{a=0}^{T-1} \beta^a U(c_{m,t+a}^{a+1}, c_{n,t+a}^{a+1}, h_{m,t+a}^{a+1}, h_{n,t+a}^{a+1}, \psi^a, \omega^a),$$

where c_m^a is market consumption, c_n^a is household consumption, h_m^a is market work, and h_n^a is home work at age a , while ψ^a and ω^a are life-cycle preference shifters.¹²

We make the same assumption regarding functional forms as earlier. In particular, the period utility function $U(c_m^a, c_n^a, h_m^a, h_n^a, \psi^a, \omega^a)$ is of the form

$$U = [\psi^a (c_m^a)^\xi + (1 - \psi^a) (c_n^a)^\xi]^{1/\xi} - \frac{\omega^a}{\gamma} (h_m^a + h_n^a)^\gamma.$$

Each household is endowed with one unit of time in each period of life, which can again be allocated among three uses: working in the market, working at home, and leisure. As in the previous model, agents derive no utility from government consumption.

An important empirical regularity is that wages exhibit significant changes over the life-cycle. We incorporate this feature by assuming that the efficiency units corresponding to a given amount of time spent working changes with age. In particular, each unit of time spent in market work

¹²The role of these shifters will become clear subsequently, but loosely speaking ψ^a allows a household's relative desire for home versus market consumption to change systematically over the life-cycle, and ω^a allows a household's value of consumption relative to leisure to change over the life-cycle.

at age a yields e^a efficiency units of market labor input.¹³ Analogously, we could assume that each unit of time spent in home work at age a yields e_n^a efficiency units of household labor input. However, given the life-cycle preference shifter ψ^a , there is really nothing to be gained by this, so we assume e_n^a does not vary with a .¹⁴

As before, we assume that home produced goods are non-traded, and can only be used as consumption. Hence, for a household of age a at time t we have

$$c_{nt}^a = (k_{nt}^a)^\eta [(1+g)^t h_{nt}^a]^{1-\eta}$$

where k_{nt}^a is the stock of home capital for a household of age a in period t , h_{nt}^a is time spent in home production by a household of age a in period t , and we assume that the rate of technological progress is the same in both production functions in order to have balanced growth.

The laws of motion for individual stocks of capital satisfy

$$\begin{aligned} k_{mt+1}^{a+1} &= (1 - \delta_m)k_{mt}^a + i_{mt}^a \\ k_{nt+1}^{a+1} &= (1 - \delta_n)k_{nt}^a + i_{nt}^a, \end{aligned}$$

while the aggregate laws of motion for the capital stocks are given by

$$\begin{aligned} K_{mt+1} &= (1 - \delta_m)K_{mt} + I_{mt} \\ K_{nt+1} &= (1 - \delta_n)K_{nt} + I_{nt}. \end{aligned}$$

We assume that each household begins economic life with an endowment of home capital k_{n0} ; that is, $k_{nt}^1 = k_{n0}$ for all t . Without this endowment agents would have no home capital in the first

¹³We are assuming the life-cycle profile of efficiency units is not affected by decisions taken by the individual, such as investment in human capital; this is similar to much of the labor supply literature, but there are exceptions, including Shaw (1989), Chang, Gomes and Schorfheide (2002), Keane and Imai (2003), and Olivetti (2001).

¹⁴The issue is that events such as having children and buying a home tend to affect the time allocated to home production. One can view these changes as affecting the efficiency of time spent in home production or as affecting one's preferences for home consumption. This choice is not likely to matter for our results. One thing that might be more interesting is to make the timing of these events endogenous, but this is beyond the scope of the current project.

period of their life, which gives rise to big differences in the time allocated to home work in the first and all other periods of life. For internal consistency, we assume that this endowment of home capital is transferred from the age T cohort to the new-born cohort each period; i.e., we require that individuals make choices such that $k_{tn}^{T+1} = k_{n0}$. Aside from this we assume no linkages between the generations.

As above, the government consumes G and finances expenditures via proportional taxes τ_h and τ_k that do not vary over time, balancing the budget every period by adjusting the lump sum transfer τ_t . Given that the total mass of households alive at any date is T , and letting w_t and r_t denote the wage per efficiency unit of labor and the rental rate of market capital, this implies

$$\tau_t = (\tau_h w_t E_{mt} + \tau_k r_t K_{mt} - G_t) / T$$

where E_{mt} is aggregate supply of labor measured in efficiency units.

Individual budget constraints are given by

$$c_{mt}^i + i_{mt}^i + i_{nt}^i = (1 - \tau_h) w_t e^a h_{mt}^a + (1 - \tau_k) r_t k_{mt}^a + \tau_t.$$

We require that k_{nt}^a is always non-negative. By contrast, although aggregate market capital cannot be negative, we assume that individuals may hold negative market capital k_{mt}^a as a way to borrow. We do not place any explicit restriction on the extent to which individuals can borrow, but we do require that everyone have zero holdings of market capital at the time of death, $k_{mt}^{T+1} = 0$. In a deterministic model these restrictions implicitly generate a maximum feasible debt at each age. In a stochastic model the situation is more complicated, but in practice we found that this issue is irrelevant in the quantitative analysis since the shocks are not that large. Hence, we impose no explicit restrictions on holdings of market capital.

The market technology is given by

$$Y_{mt} = z_t K_{mt}^\theta [(1 + g)^t E_{mt}]^{1-\theta}$$

where K_{mt} is input of market capital, E_{mt} is input of market labor measured in efficiency units, g is the constant rate of labor augmenting technological progress, and z_t is an aggregate shock that follows exactly the same process as in the representative agent model. Aggregate efficiency units of market labor are given by $E_{mt} = \sum_{a=1}^T e^a h_{mt}^a$. Market output in period t has four uses: private market consumption C_{mt} , government consumption G_t , investment in market capital I_{mt} , and investment in home capital I_{nt} . Hence feasibility requires

$$C_{mt} + G_t + I_{mt} + I_{nt} \leq Y_{mt}$$

where $C_{mt} = \sum_{a=1}^T c_{mt}^a$, $I_{mt} = \sum_{a=1}^T i_{mt}^a$, and $I_{nt} = \sum_{a=1}^T i_{nt}^a$.

We abstract from any form of public social security, and do not allow markets for risk sharing. Note that in this model all shocks are aggregate shocks, which induce changes in wages and rental rates, but all individuals face the same wage per efficiency unit of labor. In principle, once one allows for heterogeneity there is the possibility of the shocks affecting individuals differently even if they are perfectly correlated. It is possible, for example, that fluctuations in market hours differ across groups because the size of the shocks differ across groups. Our formulation implicitly rules this out, since we think this provides the most natural baseline for comparison with the standard representative agent model.

4.2 Equilibrium and Computation

Our solution concept is recursive competitive equilibrium. The aggregate state of the economy at t will be the the technology shock, plus μ_{mt} and μ_{nt} , which denote the distributions of market and home capital across agents indexed by age at t . We denote the aggregate state by S_t . For a given agent the individual state is given by their age and their two capital stocks. We denote an individual state vector by s_t . In a recursive competitive equilibrium prices at t are time invariant functions of the aggregate state variable, i.e., $w_t = W(S_t)$, and $r_t = R(S_t)$. Since our notion of equilibrium is

standard, in the interests of space we do not provide a formal definition.¹⁵

In equilibrium, individual decision rules depend on the entire state vector, including the two distributions μ_{mt} and μ_{nt} . Solving for these decision rules would be prohibitively costly in terms of computational time unless they are restricted to be linear; see Rios-Rull (1996). We therefore adopt the following procedure to solve for equilibrium numerically. We linearize the households' first-order conditions, the firms' first order conditions, and the equilibrium conditions around the model's steady state and then use a Schur algorithm to solve for the linearized decision rules; see Klein (2000) for details.

4.3 Calibration

We use the model to interpret the choices of households between ages 20 and 75. Thus, we set a period length to be one year and set $T = 55$ and $T_R = 45$, implying that agents retire at 65. We now turn to the choice of parameter values for our benchmark specification. We emphasize that these choices are only for a benchmark, and that we have carried out sensitivity analyses to explore the implications of alternative values for several parameters. Given that we have a home production model there is an important decision regarding the division of activity between the home and market. Typically, housing services are treated as a component of market consumption. This does not fit well with a model which says that market output is the result of combining market hours with market capital. We treat housing as a form of home capital, and so in our treatment of the data we subtract the flow of housing services from market activity. Because the data on market hours fluctuations by age is for the period 1962-2000 we restrict attention to this period for all of our measurement.¹⁶

As is standard we follow the procedure of requiring that parameter values are such that the model's deterministic steady state matches the time series averages for several aggregate variables.

¹⁵While we will not get into them at all, we note that there are important issues concerning the existence of recursive competitive equilibria in overlapping generations models with incomplete markets. See Kubler and Polemarchakis (2004) for a discussion and some related results.

¹⁶The one exception is in the case of determining the stochastic process for the Solow residual, since we think it is important to use as much data as possible to obtain more precise estimates of this process.

There are various specific procedures that one may adopt to carry this out. We have experimented with several of these and found that it made little difference to our conclusions. For our benchmark case we do the following. The capital share parameter for the market production function is set to $\theta = .3$ to match capital's share of market income in the data. The capital share parameter for the home production function η is set to generate the same ratio of home investment to market output as in the data. The ratio of home investment to GDP in the data is .1248 and the resulting value is $\eta = .21$. The two depreciation rates are picked on the basis of estimated depreciation by the BEA. This implies $\delta_m = .0654$ and $\delta_n = .0568$.

We use a standard procedure to pick values for the stochastic process for the market technology shock. In particular, we construct a series for the Solow residual over the period 1954-2000 using annual data and then estimate an $AR(1)$ process assuming a polynomial time trend.¹⁷ This leads us to $\rho_m = .8953$ and $\sigma_m = .0153$. The trend growth rate g is set to a value of .0184. For the government sector we choose G so that government spending on goods and services is always equal to .20 of market output, which is roughly the average ratio of government spending to output over the period 1962-2000. We set $\tau_h = .25$ and $\tau_k = .5$.

This leaves household parameters. There are three sets of such parameters: preference parameters that are constant over the life cycle (β , ξ and γ); the efficiency units profile e^a ; and the profiles for the preference shifters ψ^a and ω^a . We choose the discount factor so that households on average have investment in market capital that amounts to the share .1203 of market output. The implied value for β is .9563.¹⁸ For ξ , which determines substitutability between home and market goods, given the empirical results in Rupert et al. (1997) and McGrattan et al. (1997) a reasonable range is between .4 and .5; we set $\xi = .45$.¹⁹ For γ , which determines the degree of intertempo-

¹⁷Specifically, we use data on private GDP as the output measure, market capital as the capital input and data from the Establishment Survey on hours in private establishments as the labor input.

¹⁸Alternatively, we could have chosen β to target a particular rate of return to capital. Our choice implies an after tax rate of return of approximately 7%. Targeting a lower value would generate a much larger investment share. Ultimately there is some tension between the various statistics that we ask the model to match. Matching a lower rate of return and a reasonable investment share would require a higher capital share.

¹⁹Related evidence is contained in recent work by Aguiar and Hurst (2003), who document substantial substitution between expenditures and time in the production of food consumption in response to variation in the opportunity cost of time.

ral substitution in hours of work, there are a variety of estimates to consider, for both men and women, with the latter usually being greater. In our sensitivity analysis we explore many values for γ , but as a benchmark we set $\gamma = 2.5$, which we think is a reasonable compromise between the range of estimates for men and women. This is the estimate obtained by Rupert et al. (2001) using life-cycle data for males in a model that explicitly allowed for time spent in home work as well as market work. As they point out, this estimate is larger than those often found for males, but this is explained by the fact that neglecting home work leads to a negative bias in previous estimation procedures.

We choose the efficiency units profile for market work for the household by matching data on male wages over the life-cycle. In particular, we use cross-section data from the March CPS for the years 1975-1981 and then use the fitted values from a regression on a constant, age and age squared.²⁰ The parameters ψ^a and ω^a , which dictate the relative weights on market versus home consumption and on consumption relative to the work, are important as they influence the absolute amount of time spent working and relative amount of time spent on market and home work. Accordingly, we choose these parameters so as to match the profile of time spent on home and market work, given the other parameters. We obtain life-cycle profiles of time spent on market and home work for married couple households from the Time Use Study. In particular, we use data from the Michigan Time Use Longitudinal Panel Study for the years 1975-1981. We use data on market and home hours for married households and use the fitted values from a regression on a constant, age, age squared and age cubed. Figure 1 shows the life-cycle profile for market hours that we use in the calibration, and Figure 2 shows the life-cycle profile for home hours.

Figure 3 shows the calibrated profile for ψ over the life cycle, and Figure 4 the profile of ω over the life cycle.²¹

²⁰It would be more appropriate to use a weighted average of male and female wages over the life cycle. Given the selection issues that are more significant in estimating wages for women and the secular changes in women's wages and hours of work, we chose to use men's wages as a proxy.

²¹As a side issue we note that for a given profile of efficiency units e^a and elasticity parameters ξ and γ , one can always find values of the ψ^a and ω^a profile such that observed life-cycle hours are consistent with optimization. This should make one leery of studies that claim to identify the value of γ from life-cycle data on wages and hours worked, since one cannot make this inference without knowing the values of the preference shifters, and they are clearly unobservable.

Figure 1: Household Profile for Market Hours




Figure 2: Household Profile for Home Hours




Figure 3: Calibrated Profile for ψ





Figure 4: Calibrated Profile for ω



As can be seen, the calibrated profile for ψ has an inverted U shape, increasing during the middle part of the life cycle, whereas the profile for ω is increasing over time. This profile has a large jump at retirement age since we require the first order condition for total work to hold on either side of retirement. Perhaps not surprisingly, the calibration then requires that the disutility of working increases rather sharply upon retirement. To understand the shape of the ω profile, note that in the steady state equilibrium of this model the standard first order conditions for market hours at each point of the life cycle with positive market hours implies a relationship between total hours (market plus home) at each period relative to hours in the first period, relative efficiency units at each point relative to the first period, and the value of γ , the parameter that determines the intertemporal elasticity. Our calibrated hours series implies a series for total hours which is hump-shaped, similar to the efficiency units profile. The profile for ω is increasing over time because with $\gamma = 2.5$ the data on total hours and efficiency units can only be reconciled with an increasing profile for ω . Given the series for total hours, the series for ψ effectively justifies the split of total hours between home and market work over the life cycle.

Lastly, we need to assign a value to the endowment of home capital that a household receives when they begin economic life. We set this to .2 in our benchmark, since with this value we did not need any large departures from the profiles for the ω^a and ψ^a in order to match the life-cycle hours profiles. This condition is obviously somewhat weak; however, we found that this parameter does not matter for the model's business cycle properties. This completes the calibration. Table 12 summarizes the key parameter values for our benchmark economy.

Table 12: Parameters for Benchmark Calibration

β	θ	η	δ_m	δ_n	g	ρ	σ_m	τ_h	τ_k	γ	ξ
.967	.30	.21	.065	.057	.018	.895	.0153	.25	.50	2.5	.45

5 Results

In this section we present the results for our benchmark model. As is standard, we simulate the model for 39 years starting from the deterministic steady state and compute sample statistics from the equilibrium time series. We then repeat this 1000 times and average across the trials. Panel A of Table 13 shows the standard set of aggregate business cycle statistics for the US economy and Panel B shows the same set of statistics for our benchmark model.

Table 13: Properties of Aggregate Fluctuations

A. US Data, 1962-2000				
variable (x)	σ_x	σ_x/σ_{Y_m}	$corr(x_t, x_{t-1})$	$corr(x_t, Y_{mt})$
Y_m	2.23	1.00	.54	1.00
Y_{mp}	2.36	1.06	.52	.99
C_m	1.37	.61	.65	.91
I	5.68	2.55	.56	.89
H_{mE}	1.95	.87	.58	.86
H_{mH}	1.79	.80	.55	.75
Y_{mp}/H_{mE}	1.14	.51	.34	.47
B. Model				
variable (x)	σ_x	σ_x/σ_{Y_m}	$corr(x_t, x_{t-1})$	$corr(x_t, Y_{mt})$
Y_m	2.25	1.00	.52	1.00
C_m	1.12	.50	.56	.97
I	4.99	2.22	.51	.99
H_m	1.05	.47	.54	.98
Y_m/H_m	1.23	.55	.51	.99

A few remarks on the data are in order. The measure of output that we use in Panel A is GDP per capita, less the imputed value of owner occupied housing services. As discussed earlier, subtracting the value of owner occupied housing services is consistent with viewing this as a non-market service that derives from the stock of home capital. We also report a measure of private GDP. Although our model has a government sector, by assumption in our model the government sector fluctuates as much as the private sector and is perfectly correlated with fluctuations in the private sector. In reality, although the government sector does fluctuate about as much as the private sector, the two series are virtually uncorrelated. Our measure of consumption is spending on

consumer nondurables and services (net of the imputed service flow for owner occupied housing). Spending on consumer durables is counted as investment in home capital and hence is included in the investment category. Because our model abstracts from inventories, our investment series excludes this component. We report two hours series – one from the household series and one from the establishment series. The productivity series reported is for productivity in the private sector, and is derived from using the data on private GDP and the hours series from the establishment series.

The relationship between the model statistics and their real world counterparts is fairly typical for this literature, so we do not devote much space to it here. Note, however, that if one calibrates to annual data then Solow residuals are large enough to account for virtually all fluctuations in market output, whereas in a quarterly model the typical result is that the model accounts for roughly two-thirds of output fluctuations.²²

Our focus here is on the ability of the model to account for fluctuations in hours, and as we can see from the above tables, the model can account for only about 60% of relative fluctuations in market hours. Also, consistent with the findings of Rios-Rull (1996), note that the volatility of aggregate hours in the overlapping generations model is very similar to that of the infinitely-lived agent model. The relative standard deviation of hours here is .47, whereas it is .52 in the infinitely-lived agent model with the same values for labor supply elasticities. However, note that in the infinitely-lived agent model all labor services were equally productive, so that the variability of labor services in efficiency units was the same as the variability of labor services as measured by units of time. This is not the case in the calibrated overlapping generations model. If we instead compute the standard deviation of efficiency units of labor input we obtain a value of .50 for the relative volatility. We conclude that the ability of the two models to account for aggregate labor market fluctuations is basically the same.

²²This point is not new – Plosser (1989) found the same result in his model. Note, however, that although Rios-Rull (1996) was an annual model he also found that his model could account for roughly two-thirds of observed fluctuations in output. The reason for this apparent discrepancy is that he effectively used the Solow residuals computed using quarterly data. In particular, he computed Solow residuals using quarterly data and then aggregated the quarterly process to get an annual process.

Since these aggregate statistics have been studied extensively in this context, we do not wish to devote any additional space to them here. Rather, we wish to look more carefully at the model’s implications for market hours fluctuations by different age groups. Table 14 presents some summary statistics. The first column shows the standard deviations of hours fluctuations by age group, using the detrending procedure described earlier. As can be seen, these fluctuations exhibit a U-shaped pattern over the life-cycle, with prime aged individuals exhibiting the smallest fluctuations.

Table 14: Relative Standard Deviation of Market Hours by Age Group

Age Group	Data	Model	Model/Data
16 – 19	2.23	–	–
20 – 24	1.23	.39	.32
25 – 34	.86	.35	.41
35 – 44	.64	.35	.54
45 – 54	.57	.46	.81
55 – 64	.59	.97	1.64
65+	1.26	–	–
Aggregate	.80	.47	.59

A striking pattern emerges. In particular, the model’s ability to account for fluctuations in hours increases as we consider older age groups. Although the magnitude of fluctuations exhibits a *U* shaped profile over the life-cycle in both the data and the model, this shape is much more pronounced in the data. In the model the profile is effectively flat over the first part of the life-cycle and increasing thereafter. We also note that in the model the high variability of the age group 55 – 64 is due to the individuals in the 60 – 64 age group. If one considers the group 55 – 59 the model predicts a relative standard deviation of roughly .73, which is much closer to the actual data. A simple message emerges from Table 14. Although the various income and substitution effects present in this model are sufficient to account for only about 60% of all fluctuations in hours, the extent of the shortcoming varies dramatically across age groups. We conclude that whatever the key additional mechanisms might be to help account for hours fluctuations, these mechanisms must be very non-uniform across age groups, as evidenced by the last column in Table 14.

A key property of these results that we want to emphasize is the pattern of volatility over

the life-cycle. We will explore the economic factors behind this shape more fully in the next section. But, before doing so we want to emphasize that this pattern is very robust with regard to our calibration strategy. In particular, this pattern is basically independent of the elasticity parameters. Changing the elasticity parameters basically generates parallel shifts in the curve that shows volatility over the life-cycle.

6 Understanding the Life-Cycle Pattern of Volatility

In this section we try and shed some light on why the life-cycle profile of fluctuations takes on the shape that it does. Note that in the model all shocks are aggregate, in the sense that all individuals face exactly the same shock processes. The differing responses of individuals over the business cycle are purely the result of individuals responding differently to common shocks. There are two different aspects to heterogeneity in the model. The first is that individuals are of different ages and hence at any point in time the agents that are alive have different planning horizons. In our model this is also associated with different weights on home and market consumption, different weights on consumption and time spent working, and different productivities in market work. All of these differences represent heterogeneity in the *exogenous* component of an individual's state vector. The second source of heterogeneity is in the *endogenous* component of an individual's state vector. Optimal decision making implies that on average individuals of different ages will have accumulated different amounts of capital. Individuals with different amounts of capital will potentially respond differently to the same shock. In seeking to understand the pattern of hours volatility over the life-cycle predicted by the model it will be useful to acknowledge these two different sources of heterogeneity.

In order to learn about the role that various features play in shaping the resulting profile of hours volatility over the life-cycle, we find it instructive to compare outcomes across models in which specific model features are varied. If we do this type of analysis in the context of the full general equilibrium model that we studied earlier, a difficulty emerges, since with any change in

model features we will potentially generate different parameters from a given calibration procedure. Moreover, the equilibrium properties of the stochastic wage and rental rate processes may also vary. This makes it more difficult to assess the role of the various changes. This would also be true even if we did not redo the calibration with the new model feature present.

Because of this, in this section we have chosen to focus on a comparison of decision theoretic cases in which we take a given stochastic process for wage (or rental) rates, solve individual decision problems with different features and then compare the outcomes. We feel that this is a useful way to isolate the manner in which changes in features of the individual decision problem lead to changes in the volatility of hours for a given exogenous stochastic process for wages (or rental rates).

6.1 Case I: The Pure Effect of Time Horizon

We begin by focusing on the pure effects of differences in time horizon, i.e., we are interested in how the time horizon affects the response of an individual to a given shock holding all other factors constant, such as the stock of capital owned by the individual, or the individual's productivity in market work. We will do this in two contexts, one in which there is no retirement, and the other in which there is retirement, since it is of separate interest to understand the role of retirement. In this section we present results for the case of no retirement. For simplicity, we abstract from home production in these exercises. Hence, we consider an individual with period utility function given by

$$\log(c_t) - \frac{\omega}{\gamma} h_t^\gamma$$

We consider an individual with constant efficiency in market work and who begins life with zero assets. The individual works for T periods and then dies. We assume that the individual faces stochastic processes for wage and rental rates that approximate those in the benchmark equilibrium above, except that we assume that the mean rental rate of capital is such that the return to capital exactly offsets the effect of discounting. In each exercise we only allow for one stochastic process,

holding the other price constant. We do this analysis for various values of T . In each case we simulate the decision problem and compute the variability of hours at each stage of the life-cycle. One appealing feature of this specification is that the individual has no life-cycle motive for capital accumulation, so that in the absence of shocks the individual would not accumulate any capital – they would simply work the same amount each period and then consume their income.

All results reported below are for the case of $\gamma = 2.5$. We assume that wage and rental rate stochastic processes follow AR(1) processes with a persistence parameter of .75. Average hours of work in the no shock case are equal to .33, and average wages are equal to .96. The standard deviations in the tables below are based on the raw series and are not HP filtered, since we are not comparing these to actual data.

Comparing the volatility of, say those in years 1 – 5 of their life, across cases with different T provides a way to assess the role of time horizon in shaping the magnitude of the response. Intuitively, in this model the key mechanism through which changes in wages and capital rental rates influence hours of market work is through intertemporal substitution. The shorter the horizon, the less scope there is for intertemporal substitution. In fact, in the extreme case of a one period context the intertemporal substitution effect vanishes. Table 15 below shows the results of this exercise. As in the previous analysis, we interpret our individuals as starting life at age 20.

Table 15: Effect of Planning Horizon: Wage Shocks, No Retirement

Age Interval	$T = 5$	$T = 15$	$T = 25$	$T = 35$	$T = 45$	$T = 55$
20 – 24	.026	.115	.179	.223	.252	.272
25 – 29		.087	.161	.214	.249	.272
30 – 34		.044	.129	.193	.235	.264
35 – 39			.087	.163	.215	.249
40 – 44			.042	.128	.190	.232
45 – 49				.087	.162	.213
50 – 54				.042	.128	.190
55 – 59					.087	.162
60 – 64					.042	.128
65 – 70						.087
71 – 75						.042

Reading down the columns one sees the various life-cycle profiles of volatility. Two patterns

emerge. First, for a given life cycle, volatility decreases as the household ages. Second, holding age of the household fixed, volatility increases as we increase the number of periods remaining. The two patterns are strongly related. In fact, the table reveals that the volatility in hours is effectively determined by how many periods remain in the household's planning horizon: holding the number of periods remaining fixed it is basically irrelevant how old the household is. The decreasing pattern is consistent with the intuition that we expressed earlier – as the horizon becomes shorter there is less opportunity for intertemporal substitution. Or, put somewhat differently, as the horizon becomes shorter, the shocks appear to be more permanent, and with balanced growth preferences individuals do not change hours of market work in response to a permanent shock to wages (if they have no additional income). Recall that given our earlier comment that it is not the case that capital holdings systematically vary with age in the deterministic version of this problem.

The table also allows us to assess the quantitative significance of the time horizon effect. The table indicates that once the number of years remaining is around 30, the effect of further increasing the number of periods remaining is relatively small. However, comparing the volatility at different points in the life cycle we see that the associated effects are very large. Specifically, consider the final column of the table, which corresponds to a planning horizon of 55 years. Volatility in the first five years of working life is more than six times as large as volatility in the final five years of working life, and about one third larger than volatility during the middle five years of working life.

We have done this same exercise using the stochastic process for rental rates on capital rather than the stochastic process on wage rates. The patterns are virtually identical, though the volatility is about half as much on average. Since there is little additional information we do not present the results for this case.

6.2 The Effect of Retirement

Next we consider the same situation except that we add retirement. In particular, we consider an individual who works for 45 periods and then retires for T_R periods, where we vary T_R . The results are shown in Table 16.

Table 16: Effect of Retirement: Wage Shocks

Age	$T_R = 0$	$T_R = 10$	$T_R = 15$	$T_R = 20$
20 – 24	.252	.253	.253	.255
25 – 29	.249	.254	.256	.260
30 – 34	.235	.247	.252	.256
35 – 39	.215	.235	.242	.249
40 – 44	.190	.225	.236	.246
45 – 49	.162	.219	.235	.249
50 – 54	.128	.222	.247	.266
55 – 59	.087	.256	.291	.316
60 – 64	.042	.390	.427	.451

As before, each column depicts the life-cycle pattern of volatility for a given length of retirement. The first column in Table 16 is identical to the second last column in the Table 15 – both correspond to a case in which the worker works for 45 years and then dies. A striking new pattern appears. Once there is retirement, volatility no longer decreases monotonically over the life cycle. In fact, once T_R exceeds zero we see that the highest volatility always occurs in the final five years of working life, which is just the opposite of what we found in the case without retirement. Looking at the results more carefully, we see volatility is not monotone as we read down the columns. Loosely speaking, the pattern is for volatility to be roughly constant over the first 15 or so years of working life, then to decrease somewhat prior to increasing over the final ten or fifteen years of working life. The overall pattern is roughly U-shaped.

Why does the final five years of working life now have the highest level of volatility? The reason is once again intuitive. The presence of retirement extends the worker’s planning horizon beyond the final period in which they work. If a worker in their final year of working life realizes a positive wage shock in the no retirement case, they will increase current period consumption by the same amount by which labor income increases. And with balanced growth path preferences this results in no increase in hours of work. In contrast, a worker with a large number of periods left will not increase current consumption by the full amount by which current labor income increases, since they will save some of it to supplement consumption when wages are low sometime in the future. When we add a retirement period, a worker in the final period of working life will spread

any increased income across all retirement periods, so that current consumption increases by only a fraction of the increase in current labor income. Moreover, if the same individual were to have additional working periods after the current period they would shift less income forward, since in the face of a persistent positive shock to wages they would plan on working more not just this period but also in future periods. This lessens their incentive to work more this period and explains why the response is even larger for someone facing retirement. Put somewhat differently, if a worker is in the final period of working life prior to retirement and experiences a positive persistent shock to wages, the fact that they will retire next period makes the shock seem more transitory than it really is, and intertemporal substitution is larger in response to less persistent shocks. Of course, this effect is present not only in the final period prior to retirement but also in earlier periods, which is why we see that volatility will increase not only in the final five years of working life but also earlier.

Intuitively, the model without retirement predicts a monotone decreasing pattern for volatility over the life cycle, whereas the argument just made suggests that retirement gives rise to an increasing pattern for volatility over the life cycle. The size of these effects are not uniform over the life cycle, so when they are combined we see that one dominates over the early part of the life cycle and the other dominates over the latter part, giving rise to the rough U-shaped pattern.

It is also important to note the quantitative importance of retirement. As just remarked, the pattern of volatility over the life cycle is roughly U-shaped. However, going from the first five periods to the middle five periods the decrease in volatility is only about ten percent, whereas in going from the middle five years to the final five years the volatility of hours almost doubles. This quantitative pattern is very reminiscent of what we found in our benchmark simulations. While the model does generate a U-shaped pattern, the left hand side of the U is in fact almost flat.

The key message from this exercise is that adding retirement is likely to have a large effect both qualitatively and quantitatively on the nature of volatility over the life cycle. However, it should also be noted that once the retirement period reaches ten years the resulting profile of volatility over the life cycle is in fact relatively constant in the face of additional increases in the retirement

period.

Though we do not deal with the case of endogenous retirement, it is worth noting that in such a context one would probably expect the sharp increase in volatility just prior to retirement to be mitigated somewhat. An individual who realizes a positive wage shock at age 65 would potentially postpone retirement to take further advantage of the increased earnings opportunities rather than focusing all of the increased hours in one period. Conversely, the fact that individuals become eligible for social security benefits at age 62 could cause individuals to have much larger responses to negative shocks if a persistent negative shock leads them to opt for early retirement.

6.3 The Effect of Life-Cycle Changes in Wages

We now ask how the presence of changes in wages over the life-cycle influences the pattern of volatility over the life cycle. To better isolate the role of this factor we consider a somewhat stylized version in which wages over the life-cycle are represented by a symmetric triangle. In the benchmark case considered above, efficiency units were always equal to one. We now consider cases where the peak efficiency units are 2 and 3. Table 17 presents the results. For this exercise we assume that the worker works for 55 periods and then dies.

Table 17: Effect of Life-Cycle Earnings: Wage Shocks, No Retirement

Age	Standard Deviation of Hours		
	no peak	peak= 2	peak= 3
20 – 24	.272	.293	.303
25 – 29	.272	.252	.245
30 – 34	.264	.212	.191
35 – 39	.249	.175	.143
40 – 44	.232	.149	.112
45 – 49	.213	.146	.118
50 – 54	.190	.169	.162
55 – 59	.162	.186	.203
60 – 64	.128	.196	.243
65 – 69	.087	.211	.285
70 – 74	.042	.220	.323

As noted earlier, the life-cycle profile of volatility in the first column is decreasing. As we move

from the first column to the second column we see that the amount of volatility is decreased in the middle of the life-cycle and is increased at the two edges of the life-cycle. Note that the decrease is largest for the periods in which efficiency units are greatest. Why does this happen? We believe there is a simple intuitive explanation for this. There are two perspectives from which one can view the mechanics of intertemporal substitution in this model. One perspective is that when an individual engages in intertemporal substitution, they are effectively substituting production of income today for production of income at some future date, i.e., they are choosing to produce income when it is most efficient to do so. The other perspective is that they are trading off leisure today for leisure in the future. We argue that both of these perspectives lead us to expect volatility of hours of work to be lower during periods of high efficiency units.

We begin with the first perspective. If one is considering trading off production of income in two periods in which efficiency units differ, then the tradeoff in hours will necessarily not be one-for-one. In particular, the change in hours from the lower productivity period must be greater to compensate for the change in hours from the higher productivity period. This suggests that intertemporal substitution in this context will necessarily lead to lower changes in hours in the high productivity period and higher changes in hours in the low productivity period, as we see in Table 17. Next consider the second perspective. If leisure is lower in the period with high efficiency units, then at the margin leisure is more valuable in these periods. It follows that if the individual is trading off leisure in the different periods that it takes more leisure in the low efficiency unit periods to compensate for one unit of leisure in the high efficiency units period. Again, this suggests that hours should be less volatile in the high efficiency unit periods.

The results in Table 17 also reflect another factor which is mechanical in nature. To see this note that if we change the profile of hours over the life cycle in the absence of shocks, but keep the absolute magnitude of fluctuations in hours worked over the life cycle constant, then it would actually appear that percentage fluctuations in hours worked were lower during periods in which efficiency units are higher. To assess the magnitude of this mechanical effect, we have also computed standard deviations of the business cycle component of hours worked by age by using actual

hours rather than the log of hours. When we did this we found a U-shaped pattern of volatility that was of roughly the same quantitative magnitude as in Table 17, so we conclude that this mechanical channel is not driving the results. Alternatively, the potential size of this effect can also be gauged by noting that the variation in hours worked over the life cycle is not that large.

As with the previous factors, it is important to assess the quantitative magnitude of the effect associated with the life-cycle pattern of efficiency units. The case of peak efficiency units equal to two is of roughly the appropriate order of magnitude in terms of reality. As can be seen, this effect decreases volatility of hours in the middle of the life cycle by about one-third, and increases volatility of older workers quite substantially.

In this subsection we have focused on life cycle changes in wages. However, we note that if there were life cycle changes in the value of leisure we would get similar effects. In particular, if leisure is valued differently at different points then intertemporal tradeoffs are altered. Given the similarity to the effects just analyzed, we do not present any results for this particular specification. But it should be noted that our benchmark calibration does entail a changing value of leisure over the life cycle.

6.4 Discussion

The objective of this section was to investigate the role that various factors play in producing the life-cycle profile of volatility generated by our calibrated model. We have shown that three factors seem to be quantitatively significant. First, the finite time horizon matters. Second, the existence of a retirement period matters. And third, the variation of parameters over the life cycle to mimic life cycle patterns in wages and hours of work also matter. Based on this analysis, we feel that the basic finding regarding volatility of hours over the life cycle is robust property of the benchmark model with a reasonable parameterization. This is not to claim that our results are robust to all changes in various model features. For example, as mentioned earlier it is possible that having an endogenous retirement decision in the context of a realistic social security program may influence the nature of fluctuations for older individuals.

What is the relative importance of the three factors just described? To provide an answer to this question we redid our general equilibrium calibration exercise keeping everything the same except that we imposed no change in parameters over the life cycle. In particular, we assumed that the efficiency unit profile is constant, as are the profiles for the preference shifters. We then examined the business cycle properties of this model. The main finding is the following. Volatility of hours worked increases monotonically over the life cycle. The relative volatility of the youngest group is roughly the same as in the benchmark calibration, while the volatility of the oldest group is about two-thirds as volatile as in the benchmark calibration. The main impact of the parameters that vary over the life cycle is to depress volatility during the middle years and increase volatility in later years. It remains true in this exercise that the model's ability to account for the pattern of volatility over the life cycle is increasing in age. We conclude from this that our key quantitative finding is largely due to the finite horizon and retirement aspects.

7 International Evidence

Earlier in the paper we presented evidence pertaining to properties of labor market fluctuations in the US. It is of interest to ask to what extent these patterns are also found in other countries. This may well help us think about what factors are generating these patterns. In particular, given that labor market policies and regulations differ quite widely across economies, if these factors are playing a central role we would expect to see quite different patterns across countries.

Unfortunately, data limitations prevent us from exactly repeating our earlier analysis using data that is available from international statistical agencies such as the OECD. One would have to go directly to country level data sets in order to extract equivalent information. However, data that is available from the OECD does allow us to compare fluctuations in employment to population ratios by age groups for several countries. The time period for which this data is available does vary from country to country, but Table 18 presents summary statistics for several countries for which there was sufficient data.²³ In this table we report standard deviations relative to the age

²³The data are annual and the numbers in the table are based on standard deviations of cyclical components as

group 45 – 54.

Table 18: International Evidence

	15-24	25-34	35-44	45-54	55-64
Australia (1963-2001)	1.78	1.19	1.14	1.00	7.71
France (1968-2001)	3.06	2.36	1.21	1.00	4.20
Germany (1970-2001)	1.68	1.29	1.03	1.00	1.02
Ireland (1961-2001)	2.60	1.24	.96	1.00	.86
Norway (1972-2001)	2.20	1.33	.97	1.00	.84
Portugal (1974-2001)	3.94	1.65	1.16	1.00	1.60
Spain (1972-2001)	2.74	1.79	1.09	1.00	.98
Sweden (1963-2001)	4.24	2.22	1.54	1.00	1.53

In all countries but one we observe that volatility is highest for individuals in the 15 – 24 group, and decreases until we reach the group aged 45 – 54, though for two countries the volatility increases slightly going from the 35 – 44 group to the 45 – 54 group. For most countries relative volatility increases as we move to the oldest group, though for two countries the change is minimal and for two others there is a relatively sizeable decrease. A more complete assessment of the cross-country data is beyond the scope of this paper, but based on this first look at the data we conclude that the life-cycle pattern of volatility that we documented for the US is a robust stylized fact for a broad cross-section of countries.

8 Summary and Directions for Future Research

The motivation for this paper consisted of two simple observations. The first motivating observation is that for what many would view as reasonable parameterizations, the standard infinitely-lived representative household business cycle model cannot account for the magnitude of fluctuations in aggregate hours of market work over the cycle. According to our benchmark specification and our metric, this model can account for about 60 percent of observed fluctuations. This observation has lead many researchers to modify the model in ways to produce greater fluctuations in hours of work for the representative household. The second motivating observation is that fluctuations defined by the HP filter. Unlike in the earlier tables, we have made no additional adjustments.

in hours of market work over the business cycle vary quite dramatically across subgroups in the population. We documented this heterogeneity along two specific dimensions: age and education.

Taken together, this suggests to us a clear direction for research which has been largely ignored. If some groups experience much larger cyclical fluctuations in hours of work than do other groups, this should presumably provide substantial insight into the factors that account for these fluctuations. Or, put somewhat differently, if a model produces average fluctuations that are too small, but there is substantial heterogeneity in the magnitude of fluctuations, a natural question to ask is which groups are not fluctuating enough – is the “shortfall” of fluctuations uniform across all groups or is it concentrated in a few select groups? The answer to this question should influence the nature of modifications that researchers choose to explore.

This paper has taken a first step in this line of research. We analyzed business cycle fluctuations in a model in which household differ in age, and used it to explore the implications of standard shocks and economic forces for the pattern of fluctuations in hours by age. The finding is quite striking. As in the standard model, average hours of market work do not fluctuate enough. But significantly, the main shortfall in fluctuations is accounted for by the behavior of young individuals. Taken at face value, this suggests that whatever modifications one believes are empirically relevant for generating larger fluctuations in average hours, these modifications should be such that they interact with age in a very non-neutral manner.

It is beyond the scope of this paper to explore what these modifications might be. One plausible feature is some sort of search friction. One aspect of labor market behavior that varies with age is that younger workers are more likely to be in the process of searching for a career. Aggregate shocks may interact with this process in a distinctive manner.

Another plausible modification could involve human capital accumulation, since the nature of human capital accumulation varies over the life-cycle. While our model implies that human capital accumulation varies over the life-cycle, it implicitly assumed that human capital accumulation at any age occurred at the same rate independently of how an individual allocates their time between market work, home work and leisure. It seems very reasonable to consider modifications of the

human capital accumulation process. In this vein, the work of Imai and Keane (2003) is very relevant since they argue that allowing for endogenous human capital accumulation greatly increases the estimated labor supply elasticities.

Finally, one qualification that was mentioned earlier also bears repeating. In a model that allows for heterogeneous agents, one must also allow for the possibility that differences in volatilities might also reflect the fact that these agents face different shocks. Even if the shocks are perfectly correlated, the magnitudes of the shocks could vary.

Although our analysis focused solely on the age dimension, the empirical work that we summarized also suggests that a fuller treatment will consider age and human capital accumulation jointly in the business cycle context. It will be important to assess the key economic forces that alter the way in which individuals that differ in age and human capital respond to common shocks.²⁴ More generally, the analysis carried out suggests a research agenda in which macroeconomists take seriously the patterns of hours fluctuations at disaggregated levels in order to better assess the economy's impulse and propagation mechanisms.

²⁴An early attempt to understand differences in fluctuations in hours of market work across different skill groups is Kydland (1984).

References

- [1] Aguiar, M., and E. Hurst, "Consumption vs. Expenditure," NBER Working Paper No. 10307, 2003.
- [2] Alexopoulos, M., "Unemployment and the Business Cycle," *Journal of Monetary Economics* 51 (2004), 277-298.
- [3] Altonji, J., "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *Journal of Political Economy* 94 (1986).
- [4] Andolfatto, D., "Business Cycles and Labor Market Search," *American Economic Review* 86 (1996), 112-32.
- [5] Benhabib, J., R. Rogerson and R. Wright, "Homework in Macroeconomics: Household Production and Aggregate Fluctuations," *Journal of Political Economy* 99 (1991), 1166-87.
- [6] Baxter, M., and U. Jermann, "Household Production and the Excess Sensitivity of Consumption to Current Income," *American Economic Review* 89 (1999), 902-920.
- [7] Blanchard, O., and P. Diamond, "The Cyclical Behavior of the Gross Flows of U.S. Workers," *Brookings Papers on Economic Activity* 1990, 85-143.
- [8] Chang, Y., Gomes, J., and F. Schorfheide, "Learning-by-Doing as a Propagation Mechanism," *American Economic Review* 92 (2002), 1498-1520.
- [9] Cho, J. O., and R. Rogerson, "Family Labor Supply and Aggregate Fluctuations," *Journal of Monetary Economics* 21 (1988), 233-246.
- [10] Clark, K., and L. Summers, "Demographic Differences in Cyclical Employment Variation," *Journal of Human Resources* 16 (1981), 61-79.
- [11] Christiano, L. and M. Eichenbaum, "Current Real Business Cycle Theory and Aggregate Labor Market Fluctuations," *American Economic Review* 82 (1992), 430-450.

- [12] Danthine, J. P., and J. Donaldson, "Non-Walrasian Economies," in *Frontiers in Business Cycle Theory*, edited by T.F. Cooley, Princeton: Princeton University Press, 1995.
- [13] Den Haan, W., Ramey, G., and J. Watson, "Job Destruction and Propagation of Shocks," *American Economic Review* 90 (2000), 482-498.
- [14] Gertler, M., and S. Gilchrist, "Monetary Policy, Business Cycles and the Behavior of Small Manufacturing Firms," *Quarterly Journal of Economics* (1994).
- [15] Gomme, P., "Shirking, Unemployment and Aggregate Fluctuations," *International Economic Review* 40 (1999), 3-21.
- [16] Gomme, P., Kydland F., and P. Rupert, "Home Production Meets Time to Build," *Journal of Political Economy*, 109 (2001), 1115-1131.
- [17] Greenwood J., and Z. Hercowitz, "The Allocation of Capital and Time Over the Business Cycle," *Journal of Political Economy* 99 (1991), 1188-1214.
- [18] Greenwood, J., R. Rogerson, and R. Wright, "Household Production in Real Business Cycle Theory," in *Frontiers in Business Cycle Theory*, edited by T.F. Cooley, Princeton: Princeton University Press, 1995.
- [19] Hall, R., "Macroeconomic Fluctuations and the Allocation of Time," *Journal of Labor Economics* 15 (1997), S223-S250.
- [20] Hansen, G., "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics* 16 (1985), 309-327.
- [21] Hansen, G.D. and R. Wright, "The Labor Market Real Business Cycle Theory," *Federal Reserve Bank of Minneapolis Quarterly Review*, Sept. 1992.
- [22] Keane, M., and S. Imai, "Intertemporal Labor Supply and Human Capital Accumulation," mimeo, Yale University, 2003..

- [23] Keane, M., and E. Prasad, "The Employment and Wage Effects of Oil Price Shocks: A Sectoral Analysis," *Review of Economics and Statistics* 78 (1996), 389-400.
- [24] Klein, P., "Using the Generalized Schur Form to Solve a Multivariate Linear Rational Expectations Model," *Journal of Economic Dynamics and Control* 24 (2000), 1405-1423.
- [25] Krusell, P., and A. Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy* 106 (1998), 867-896.
- [26] Kubler, F., and H. Polemarchakis, "Stationary Markov Equilibria for Overlapping Generations," forthcoming in *Economic Theory*, 2004.
- [27] Kydland, F., "Labor Force Heterogeneity and the Business Cycle," *Carnegie-Rochester Series on Public Policy* 21 (1984), 173-208.
- [28] Kydland F., and E. Prescott, "Time to Build and Aggregate Fluctuations," *Econometrica* 50 (1982), 1345-1370.
- [29] -----, "The Workweek of Capital and Its Implications," *Journal of Monetary Economics* 21 (1988), 343-360.
- [30] MaCurdy, T., "An Empirical Model of Labor Supply in a Life Cycle Setting," *Journal of Political Economy* 89 (1981), 1059-1085.
- [31] McGrattan, E., R. Rogerson and R. Wright, "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy," *International Economic Review*, 38 (1997), 267-290.
- [32] Merz, M., "Search in the Labor Market and the Real Business Cycle," *Journal of Monetary Economics* 36 (1995), 269-300.
- [33] Mulligan, C., "Substitution Over Time: Another Look at the Evidence," in *NBER Macroeconomics Annual* 1998, edited by Ben S. Bernanke and Julio J. Rotemberg.

- [34] Olivetti, C., “Changes in Women’s Hours of Work: The Effect of Changing Returns to Experience,” mimeo, Boston University, 2001.
- [35] Pencavel, J., “Labor Supply of Men,” in *Handbook of Labor Economics*, edited by O. Ashenfelter and R. Layard, Amsterdam, North-Holland, 1986.
- [36] Plosser, C., “Understanding Real Business Cycles,” *Journal of Economic Perspectives* 3 (1989), 51-78.
- [37] Rios-Rull, V., “Life-Cycle Economies and Aggregate Fluctuations,” *Review of Economic Studies* 63 (1996), 465-489.
- [38] Rupert, P., R. Rogerson, and R. Wright, “Using Panel Data to Estimate Substitution Elasticities in Household Production Models,” *Economic Theory* 6 (1995), 179-193.
- [39] Shaw, K., “Life-Cycle labor Supply with Human Capital Accumulation,” *International Economic Review* 30 (1989), 431-456.
- [40] Shimer, R., “Why is the US Unemployment Rate So Much Lower?” in *NBER Macroeconomics Annual* 1998, edited by Ben S. Bernanke and Julio J. Rotemberg.

Appendix

In this appendix we outline in detail the procedure that we used to produce the statistics reported in Table 3. As state in the text, our data source is the CPS March Supplement for the years 1962-2000. We use the question on total number of hours of market work in the preceding week to compute average hours per person for all individuals aged 16 and over as well as for each of the seven age groups listed in the paper. We use these numbers as our estimates of hours of work per person in the aggregate and by age for each year in the sample, giving us an annual data set for each series.

We define the cyclical component of aggregate hours per person by applying the Hodrick-Prescott filter to the log of aggregate hours and applying a smoothing parameter of 100. This series has a standard deviation of 1.99.

Our basic goal is to determine how changes in aggregate hours per person are accounted for by changes in hours per person of each age group. The first step we take is to apply the Hodrick-Prescott filter to each age specific series. Doing this produces the values in Table A.1

Table A.1: Cyclical Fluctuations of Hours by Age Group

	16 – 19	20 – 24	25 – 34	35 – 44	45 – 54	55 – 64	65+
σ	6.91	3.38	2.19	1.61	1.55	1.92	5.13

These values display the *U* shaped pattern that figured prominently in the analysis in the text. However, these values are not necessarily the best measures of fluctuations by age for our purposes. There are two issues. First, in going from aggregate data to age specific data, the survey sample sizes are reduced considerably and there is the possibility of additional noise in the data. Given our detrending procedure this noise will likely show up as cyclical fluctuations. Second, there may be some non-business cycle shocks which impact on relative hours across age groups that we do not want to interpret as representing business cycle shocks. Table A.2 presents cross-correlations for the various age specific cyclical components with each other and the aggregate component.

Table A.2: Contemporaneous Correlations Across Age Groups, HP Filtered Data

	16 – 19	20 – 24	25 – 34	35 – 44	45 – 54	55 – 64	65+	Agg
16 – 19	1.00							
20 – 24	.81	1.00						
25 – 34	.81	.92	1.00					
35 – 44	.70	.84	.92	1.00				
45 – 54	.59	.77	.82	.86	1.00			
55 – 64	.44	.50	.62	.71	.79	1.00		
65+	.34	.19	.31	.41	.39	.53	1.00	
Agg	.80	.89	.96	.96	.90	.74	.43	1.00

As expected, one sees that each of the age specific series is highly correlated with the aggregate, with the exception of the over 65 group. This group also accounts for very few hours worked.

In fact, the basic pattern is that the greater the age specific hours worked, the greater is the correlation with the aggregate. All of the cross-correlations between age groups are also fairly positive. However, the basic pattern suggests that measurement error and/or some age specific shocks may be a factor.

Hence, the second step in our two step procedure is to remove the component due to measurement error and/or idiosyncratic shocks. To do this we take each age specific series for HP residuals and find the component that is correlated with changes in aggregate hours. To do this we regress each age specific series on a constant, and current and lagged aggregate hours and then use the predicted values as our measure of the cyclical component of each series. We experimented with additional lags but found that it made not difference. In fact, except for the groups aged 55 and above, the effect of adding one lag of aggregate hours was very small. Table A.3 shows the effect that this has on the measure of volatility for each group.

Table A.3: Standard Deviations by Age Group, HP Filtered Data

	16 – 19	20 – 24	25 – 34	35 – 44	45 – 54	55 – 64	65+
Raw	6.90	3.38	2.18	1.61	1.55	1.92	5.13
Adjusted	5.67	3.04	2.12	1.58	1.41	1.46	3.11

The first row repeats the standard deviations of the deviations from the HP trend, and the second row presents the values produced by our second step. As can be seen, the changes are relatively small for prime aged individuals but are sizeable for the youngest and oldest individuals. We should emphasize that for the points that we make in our analysis, the raw data would actually make our case somewhat stronger, so this process of adjustment is not to make our case stronger. As evidence that our adjustment serves its purpose, we note that if one uses the raw data and computes the weighted average of age specific standard deviations using age specific hours as weights, then the resulting series has a standard deviation that is more than 10% larger than the aggregate series, but when we do the same calculation using our adjusted series we obtain the same standard deviation as for the aggregate series. In this sense we feel that we have isolated the age specific components of the aggregate fluctuations.

There is one final adjustment that we make for all data reported in the text. This adjustment is irrelevant for purposes of comparing standard deviations across hours series, but is relevant for comparing volatility of hours with a series such as GDP. Ideally, we would have computed the annual value for aggregate average hours per person by averaging the monthly values for each year. Because we do not have the monthly values for all years we are unable to do this. Intuitively we would expect that using only the values for March rather than all months would lead to greater variance in the series. To estimate the extent of this effect we carried out a similar exercise using Establishment hours. In particular, for this series we asked how the standard deviation of the cyclical component changes when we use only the March data as our annual estimate rather than averaging over all twelve months. We find that the standard deviation is larger by 10 percent in the case in which only March is used. To retain comparability with other annual series for which we use all observations, in what follows we will make a 10% adjustment to the standard deviation of all of our hours series based on using only March data. Note that this has no impact on any comparisons of relative volatility across hours series – it is only relevant when comparing the volatility of an hours series to some other series such as GDP.