

The Demand for Youth: Implications for the Hours Volatility Puzzle *

Nir Jaimovich[†]
Duke University
and NBER

Seth Pruitt[‡]
Federal Reserve
Board of Governors

Henry E. Siu[§]
University of
British Columbia
and NBER

June 2011

Abstract

Business cycle models cannot account for the nearly identical volatility of aggregate output and hours. Our resolution is to account for the volatility of age-specific hours, since hours volatility of the young account disproportionately for volatility in the aggregate. We show that the joint behavior of hours and real wages indicates the importance of differences in the cyclicity of labor demand. We find that a model featuring greater diminishing marginal product of prime-age labor relative to young labor input in production accounts for the volatility of age-specific hours and the volatility of aggregate hours.

Keywords: business cycle, demographics, capital-experience complementarity, labor demand

JEL codes: E00, E32

*We thank Manuel Amador, Paul Beaudry, Larry Christiano, Marty Eichenbaum, Valerie Ramey, Sergio Rebelo, Victor Rios-Rull, William Gui Woolston, and numerous workshop participants for helpful comments, and Josie Smith for superb research assistance. Siu thanks the Social Sciences and Humanities Research Council of Canada for financial support. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

[†]nj41@duke.edu

[‡]seth.j.pruitt@frb.gov

[§]hankman@mail.ubc.ca

1 Introduction

Perhaps the most salient stylized fact regarding the business cycle is the striking similarity in cyclical fluctuations of aggregate output and hours worked. Indeed, the first business cycle phenomenon discussed in Prescott (1986) is that: “*Output and hours clearly move up and down together with nearly the same amplitude.*” That is, when looking at detrended real GDP and aggregate hours worked the two series display: (i) a correlation near unity, and (ii) a relative volatility near unity.

Modern business cycle analysis has been successful in accounting for the first of these phenomena but has failed to account for the second. While the standard deviation of hours relative to that of output is near unity in the U.S. data, quantitative models generate a ratio of 0.7 – 0.75 at best (see Rogerson, 1988; Hansen, 1985; Benhabib et al., 1991). We refer to this long-standing discrepancy between data and theory as the *hours volatility puzzle*. Developing a resolution is crucial to our understanding of the mechanisms that amplify and propagate business cycle fluctuations.

In this paper, we hypothesize that a resolution can be found by modeling the cyclical behavior of disaggregated hours worked. We are motivated by the observation that the hours of young individuals fluctuate much more over the business cycle than for the prime-aged, leading to aggregate hours’ volatility being disproportionately accounted for by the young. The mechanisms embodied in standard business cycle models do a good job of accounting for hours volatility of prime-aged individuals, but not for the young. Hence, we hypothesize that explaining the volatility of aggregate hours boils down to explaining the volatility of hours for young individuals.¹ We ask whether a model that accounts for the business cycle volatility of hours worked by these different age groups can solve the hours volatility puzzle.

To maintain comparability with the real business cycle (RBC) literature, we study models that represent minimal deviations from the standard RBC model, extended to three factor inputs: capital, “young” labor, and “old” labor. Within this framework, differences across age groups can arise from factors related to preferences (or succinctly, differences in labor supply), technology (labor demand), or both.² How does one distinguish between these two channels?

The joint behavior of age-specific hours and wages over the cycle provides the necessary evidence. Specifically, any modification to the RBC framework relying on age-specific labor supply differences alone would generate either a higher relative volatility of young hours or young wages, but not both

¹See also Gomme, Rogerson, Rupert, and Wright (2004) who make a similar point.

²By RBC framework, we are referring to neoclassical models in which households and firms take all prices as given, and interact in competitive spot markets. See Nagypál (2004) for an alternative approach highlighting the interaction between age and worker-occupation match.

simultaneously (as is discussed in detail in Section 4). In Section 2 we document that *both* the volatilities of hours *and* wages of young individuals is greater than that of the prime-aged over the cycle. Hence, jointly matching the behavior of hours and wages in the RBC framework requires a role for age differences in cyclical labor demand. Since we also show that all age-specific real wages are procyclical, we restrict our analysis to models where the sole impulse is a productivity shock.

Analytically, models featuring differences in labor demand characteristics over the cycle fall into one of two categories. The first features age-specific marginal products of labor that respond differently to cyclical shocks. To study this, we construct a model that allows the variance of innovations to labor-augmenting productivity to differ across young and old labor input. In the second category are models featuring different degrees of diminishing marginal product of age-specific labor. We study this by allowing the elasticity of substitution between capital and labor to differ between young and old.

We present both economic and econometric evidence to discriminate between such models. Based on this evidence we find that the most promising explanation features different elasticities of substitution between factor inputs, so that production exhibits a greater diminishing marginal product of old labor input relative to young labor. Our analysis points to the importance of *capital-experience complementarity* in production, when age is equated with labor market experience.

Support for this approach is provided in the literature related to labor market outcomes. Previous work has emphasized the existence of complementarities between capital and skilled labor, when skill is proxied by educational attainment (see, for instance, Krusell et al. (2000)). However, as is well known, Mincerian wage regressions emphasize two important observable dimensions of skills: education and experience. In this paper we suggest that labor market experience shares similar complementarities to capital as that of education.

In our analysis, we do not simply impose values for the models' key elasticity parameters. Instead we estimate the structural parameters from the models' factor demand equations. Our estimation strategy exploits the identification that emerges from the relationship between aggregate prices and quantities observed in the data, and in no way targets the differences in cyclical volatility of age-specific hours.

Using our estimated parameters we find that a model with capital-experience complementarity generates volatilities of hours and relative wages across age groups that are very similar to those observed in the data. As a by-product, the model generates a relative volatility of aggregate hours to output that is essentially unity. That is, the capital-experience complementarity model represents

a solution to the hours volatility puzzle.

We view our theoretical mechanism as speaking to the existence of complementarity in production between experienced labor and factors that are in fixed short-run supply to the firm. These factors may include organizational capital, firm-specific capital, firm know-how, or operational/procedural knowledge that inherently require (or is embodied) in the presence of experienced labor. Since this type of knowledge or capital is hard to adjust in the short-run, it is natural that cyclical fluctuations in output result in greater variation of inputs like young or inexperienced labor that are less tied to these factors, while demand for experienced labor exhibits behavior that resembles labor hoarding. Of course, the measurement of factors such as organizational capital or firm know-how are very difficult. The challenge in this paper is to account for the observed differences in hours volatility in a quantitative manner. This motivates our modeling choice, as specifying complementarity between physical capital and experienced labor. The availability of high-quality data relating to these factor inputs allows us to discipline our analysis. Moreover, a second merit of this approach is that the model represents a minimal deviation from the standard model allowing us to clarify the exact mechanisms at work.

This is not the first paper to address age group differences in analyzing labor market fluctuations. Rios-Rull (1996) and Gomme, Rogerson, Rupert, and Wright (2004) study models with differences in hours volatility owing to life-cycle considerations (e.g., preferences for home and market production, and efficiency units of hours worked that differ exogenously by age). They show that life-cycle factors are successful at explaining volatility differences between the prime-aged and those near retirement age. However, in the data these differences are small. These same factors cannot account for the much greater volatility of young workers relative to all others. By extension, these models are unsuccessful in improving upon the standard, homogenous labor RBC model in terms of the hours volatility puzzle. Hansen and Imrohroglu (2008) consider a life-cycle model in which efficiency units of labor are accumulated while working via learning-by-doing. This mechanism generates substantial differences in volatility by age, but at the expense of dampening the volatility of hours worked for all age groups. Hence, the learning-by-doing model actually underperforms relative to the standard RBC model in matching the volatility of aggregate hours. It is also worth noting that these papers focus on age differences in the elasticity of labor supply. As a result, at a qualitative level these models are counterfactual with respect to the relative volatility of real wages over the cycle.

The paper is organized as follows. In Section 2, we document differences in the volatility of

hours worked and wages by age. Section 3 presents our class of models that allows for the cyclical properties of labor demand to differ across age. Section 4 provide analytical results on the response of age-specific hours and wages to business cycle shocks in the different models we consider. In Section 5 we present our estimation and calibration of structural parameters and discuss our ability to statistically discriminate between model specifications. In Section 6 we present results for the models' cyclical properties relative to the U.S. data. Concluding remarks are provided in Section 7.

2 The Cyclical Volatility of Age-Specific Hours and Wages

In this section, we document the empirical findings that motivate our approach. We first present evidence on the large differences by age in the volatility of hours and employment over the cycle. Within the RBC framework, these differences can arise from differences in the cyclical characteristics of labor demand or labor supply. In subsection 2.3 we provide an analysis of the cyclicity of age specific real wages. Taken together with the evidence from subsection 2.1 these findings indicate an important role for age differences in the cyclical properties of labor demand.

2.1 Hours

The evidence on the cyclicity of age-specific hours has been extensively addressed in Gomme, Rogerson, Rupert, and Wright (2004) and Jaimovich and Siu (2009). We provide a brief summary here and refer the reader to the cited papers for greater detail.

Using data from the March supplement of the CPS, 1963–2005, we construct annual series for per capita hours worked for specific age groups, as well as an aggregate series for all individuals 15 years and older.³ We extract the high frequency component of each series using the Hodrick-Prescott (HP) filter. Since we are interested in fluctuations at business cycle frequencies (those higher than 8 years), we use a smoothing parameter of 6.25 for annual data.⁴

Table 1 presents results on the time series volatility of hours worked by age. The first row presents the percent standard deviation of the detrended age-specific series. We see a decreasing relationship between the volatility of hours worked and age, with an upturn close to retirement age.

³See Appendix A.2 for detailed information on data sources used throughout the paper.

⁴Through analysis of the transfer function of the HP filter, Ravn and Uhlig (2002) find this to be the optimal value for annual data. Using a similar approach, Burnside (2000) recommends a smoothing parameter value of 6.65. Finally, see Baxter and King (1999), who recommend a value of 10 through visual inspection of the transfer function. Throughout this paper, we have repeated our analysis of annual data using the band-pass filter proposed by Baxter and King (1999), removing fluctuations less frequent than 8 years. The results are essentially identical in all cases.

Table 1: VOLATILITY OF HOURS WORKED BY AGE GROUP

| | 15 - 19 | 20 - 24 | 25 - 29 | 30 - 39 | 40 - 49 | 50 - 59 | 60 - 64 |
|-------------------------------|---------|---------|---------|---------|---------|---------|---------|
| filtered volatility | 4.351 | 2.130 | 1.471 | 1.073 | 0.790 | 0.824 | 1.309 |
| R^2 | 0.79 | 0.80 | 0.83 | 0.88 | 0.89 | 0.72 | 0.30 |
| cyclical volatility | 3.868 | 1.902 | 1.318 | 1.014 | 0.752 | 0.705 | 0.708 |
| share of hours (%) | 3.34 | 10.64 | 13.23 | 26.12 | 23.98 | 17.73 | 4.97 |
| share of hours volatility (%) | 11.62 | 18.21 | 15.70 | 23.83 | 16.23 | 11.25 | 3.17 |

Notes: Data from the March CPS, 1968-2005. Filtered volatility is the percentage standard deviation of HP-filtered log data. Cyclical volatility is the percentage standard deviation of HP-filtered log data as projected on aggregate business cycle measures, with the R^2 from this projection reported. Share of hours is the sample average share of aggregate hours worked by the age group. Share of hours volatility is the age group’s share of “aggregate hours volatility,” the average of age-specific cyclical volatilities weighted by hours shares.

We are not interested in the high frequency fluctuations in these time series per se, but rather those that are correlated with the business cycle. For each age-specific series, we identify the business cycle component as the projection on a constant, current detrended output, and on current and lagged detrended aggregate hours; we refer to these as the *cyclical* hours worked series. The second row of Table 1 reports the R^2 from these regressions. This is high for most age groups, even for those whose hours comprise a small fraction of total hours. This implies the preponderance of high frequency fluctuations are attributable to the business cycle.⁵

The third row indicates the percent standard deviation of the cyclical series. The data indicates a pattern of decreasing volatility with age. The young experience much greater cyclical volatility in hours than all others. Moreover, the age differences are large. The standard deviation of cyclical hours fluctuations for 15-19 and 20-24 year old workers is 5 and 2.5 times that of 50-59 year olds, respectively.⁶

The fourth row indicates the average share of aggregate hours worked by each age group. The fifth row indicates the share of “aggregate hours volatility” attributable to each age group. Here, aggregate hours volatility is represented by the weighted average of age-specific cyclical volatilities, with weights reflecting an age group’s share of aggregate hours. Fluctuations in aggregate hours

⁵The exception is the 60-64 age group, where a larger fraction of fluctuations are due to age-specific, non-cyclical shocks.

⁶These results corroborate the findings of Gomme, Rogerson, Rupert, and Wright (2004), and extend them to include data from the 2001 recession. See also Clark and Summers (1981), Rios-Rull (1996), and Nagypál (2004) who document differences in cyclical sensitivity across age groups. Age-specific hours aggregated by efficiency-weighting constituent groups, following the procedure discussed below and in Appendix A.2, show essentially the same volatility pattern.

are disproportionately accounted for by young workers. Although those aged 15-29 make up only about one quarter of aggregate hours worked, they account for nearly one half of aggregate hours volatility. By contrast, prime-aged workers in their 40s and 50s account for more than 40% of hours, but only about 25% of hours volatility.⁷

2.2 Participation

Young individuals might display greater cyclicity of hours worked relative to the old because of labor supply considerations; for instance, they may face different trade-offs between market work and home production, or possess a greater degree of insurance through parental ties. These possibilities indicate that if labor supply differences are of primary importance, the cyclicity of labor force participation should be more pronounced for the young. To explore this, we note that changes in per capita hours worked can be viewed as being due to changes in either hours per labor force participant, or the number of the labor force participants per capita. We refer to the former as the *hours margin*, and to the latter as the *participation margin*. It is important to note that what we refer to as the hours margin is an amalgam of both, the intensive margin (hours per worker) and extensive margin (workers per labor force participant), that are commonly referred to in the macro-labor literature. If the participation margin is the main driver of hours variation for the young, then one could argue the practical necessity of explicitly modeling labor supply differences, and specifically, differences in the participation decision between young and old. If not, it would indicate that, to a first-order approximation, the primary factor generating age group differences are to be found elsewhere.

Following Hansen (1985), the variance of hours is decomposed

$$Var(hpc) = Var(hplf) + Var(lfpr) + 2Cov(hplf, lfpr)$$

for hours per capita hpc , hours per labor force participant $hplf$, and the labor force participation rate $lfpr$. In Table 2 we present this decomposition of the variance of hours worked into these two margins.⁸ This table shows the proportion of hours variation by age group that can be attributed to the participation margin. The first row presents the ratio of the cyclical variance owing to the participation margin to the sum of the variances of the hours margin and participation margin.⁹

⁷Large differences by age remain when we undertake further breakdowns, such as gender, education, marital status, and industry of occupation. See Gomme, Rogerson, Rupert, and Wright (2004) and Jaimovich and Siu (2009) for discussion.

⁸The decomposition using filtered volatility gives similar results and is available from the authors upon request.

⁹Again, this is calculated as a projection on a constant, current detrended aggregate output, and current and lagged detrended aggregate hours.

Table 2: HOURS DECOMPOSITION, PARTICIPATION MARGIN

| | 15 - 19 | 20 - 24 | 25 - 29 | 30 - 39 | 40 - 49 | 50 - 59 | 60 - 64 |
|--------------------------------|---------|---------|---------|---------|---------|---------|---------|
| covariance not included (%) | 24.26 | 4.39 | 1.96 | 1.43 | 0.40 | 2.41 | 1.69 |
| covariance included (%) | 35.87 | 15.89 | 7.34 | 7.85 | 2.12 | 10.61 | 4.94 |

Notes: Data from the March CPS, 1968-2005. Shown are percentage shares of total hours variation attributed to the participation margin. Total hours per age group member is the product two variables: labor force participation per age group, and hours per labor force participant in that age group. “Covariance not included” means covariance terms are ignored, so total variation is just the sum of the variables’ variances and the share attributed to the participation margin is the variance of labor force participation. “Covariance included” means total variation includes covariance terms, so total variation is the sum of the variables’ variances plus two times their covariance and the share attributed to the participation margin is the variance of labor force participation plus the covariance. Cyclical volatility is the standard deviation of HP-filtered log data as projected on aggregate business cycle measures.

With covariance terms not included, the participation margin explains less than one quarter of the variation of any age group. Specifically, for 20-59 year old individuals, the participation margin accounts for no more than 5%. For teenagers, this is higher at 25%; nonetheless, nearly three quarters of the variance of their hours worked is due to the hours margin. The bulk of all age groups’ hours variation is due to variation in hours per labor force member.

The second row presents an alternative decomposition which accounts for the covariance between hours per labor force member and labor force members per capita. Specifically, the participation margin’s share is now defined as its variance plus the covariance, divided by the total variance of hours worked. Row 2 presents a similar picture to Row 1. With the inclusion of covariance terms, participation now explains at most 35% of the variation for the 15-19 year olds. But for 20-59 year old individuals, participation explains less than 15% of the variation of hours.

Hence, fluctuations in hours per labor force participant account for the bulk of hours variation for all age groups.¹⁰ Consequently, it does not appear that explanations centered on differences in the cyclical volatility of participation are of first-order importance for generating greater volatility of young hours relative to the old over the business cycle.

¹⁰Though not reported here, we also performed a decomposition exercise of the volatility of hours per labor force participant into its two components: the intensive and extensive margins. We find that the relative contributions of each margin are very similar for all age groups. Namely, between two-thirds and three-quarters of the hours margin variation is due to the extensive margin (workers per labor force participant). Hence, as has been found in the aggregate, hours variation for all age groups is accounted for largely by movements in and out of employment.

Table 3: VOLATILITY OF REAL HOURLY WAGES BY AGE GROUP

| | 15 - 19 | 20 - 24 | 25 - 29 | 35 - 39 | 45 - 49 | 55 - 59 | 60 - 64 |
|---------------------|---------|---------|---------|---------|---------|---------|---------|
| filtered volatility | 2.87 | 1.59 | 1.27 | 1.09 | 1.19 | 1.53 | 1.64 |
| R^2 | 0.35 | 0.28 | 0.23 | 0.17 | 0.14 | 0.12 | 0.14 |
| cyclical volatility | 1.69 | 0.84 | 0.61 | 0.46 | 0.44 | 0.54 | 0.61 |

Notes: Data from the March CPS, 1963-2005. Filtered volatility is the percentage standard deviation of HP-filtered log data. Cyclical volatility is the percentage standard deviation of HP-filtered log data as projected on aggregate business cycle measures, with the R^2 from this projection reported.

2.3 Age-Specific Wages

From the March CPS, we use information on labor income and hours worked to construct annual series for hourly wages, 1963–2005. These wages are then deflated by the GDP deflator to obtain real wages. Given our interest in wage cyclicality, we construct wage rates in a manner mitigating composition effects that stem from labor heterogeneity. Specifically, we classify individuals into 220 highly disaggregated demographic groups, and weight observations to derive efficiency measures of age-specific labor input. Our procedure is an extension of that used by Katz and Murphy (1992) and Krusell, Ohanian, Rios-Rull, and Violante (2000), and is detailed in Appendix A.2.¹¹ We then HP-filter these series to isolate fluctuations at the business cycle frequency.

The first row in Table 3 reports the percent standard deviation of the HP-filtered hourly real wage rates by age.¹² We see a decreasing pattern in volatility by age with an upturn beginning in the 55-59 age group. The second row reports the R^2 from projecting the age-specific series onto detrended aggregate output and hours, as done in subsection 2.1. These statistics, and the fact that all these series are positively correlated, indicate that real wages, when disaggregated by age, are indeed procyclical.

Row 3 presents the percent standard deviation of the *cyclical* age-specific series. As in Row 1, we see the familiar decreasing pattern of volatility by age, with a slight upturn at the end of the age distribution. For instance, the standard deviation of cyclical volatility for 20-24 year olds is about twice that of 45-49 year olds.

¹¹Using weekly wages, as in Katz and Murphy (1992), yields similar results to those we report here for hourly wages.

¹²We compute wage rates for 5 year age groups, as opposed to the 10 year age groups presented in the previous subsection. This is done to further minimize composition effects, eliminating heterogeneity due to the aggregation of individuals with large age differences.

Hence, the cyclical volatility of *both* hours and wages is greater for the young than for other age groups. This evidence allows us to discriminate between alternative mechanisms in our analysis. Specifically, real wages are equated to the marginal product of labor within the RBC framework. Because all real wages are procyclical, we restrict attention to models where business cycle impulses are due to productivity shocks. If age differences in labor supply were the sole factor responsible for the greater volatility of young workers' hours than the prime-aged, their wages would simultaneously be *less* volatile over the business cycle. By contrast, we find exactly the opposite.

This is not to claim that age-specific labor supply considerations are irrelevant for understanding differences in hours volatility. However, the greater volatility of wages and hours for the young indicates that, within the RBC framework, there *must* be some role played by differences in the cyclical nature of labor demand.¹³ This finding is laid out in detail in Section 4.

3 The Models

In this section, we present models featuring differences in the cyclical characteristics of labor demand for age-specific labor. The remaining features of the model – in particular, household preferences – are specified to conform as closely as possible to the standard RBC model. This specification allows us to isolate the role of age differences in labor demand in accounting for the facts presented in Section 2. It also allows us to isolate the theoretical differences implied by the two classes of labor demand models.

3.1 Households

The economy is populated by a large number of identical, infinitely-lived households. Each household is composed of a unit mass of family members. For simplicity, we assume there are only two types of family members, *young* and *old*. Let s_Y denote the share of family members that are young. Family members derive instantaneous utility from consumption C_i and disutility from hours spent working N_i , according to $U_i(C_i, N_i)$, where $i \in \{Y, O\}$ denotes either young or old.

The representative household's date t problem is to maximize

$$E_t \sum_{j=t}^{\infty} \beta^{j-t} [s_Y U_Y(C_{Yj}, N_{Yj}) + (1 - s_Y) U_O(C_{Oj}, N_{Oj})], \quad (3.1)$$

subject to

$$s_Y C_{Yj} + (1 - s_Y) C_{Oj} + \tilde{K}_{j+1} = (1 - \delta) \tilde{K}_j + r_j \tilde{K}_j + s_Y W_{Yj} N_{Yj} + (1 - s_Y) W_{Oj} N_{Oj}, \quad \forall j \geq t,$$

¹³Solon, Barsky, and Parker (1994) make the related point on the relative procyclicality of hours and wages between men and women.

with $0 < \beta < 1$, $0 \leq \delta \leq 1$. Here \tilde{K}_t denotes capital holdings at date t , r_t is the rental rate, W_{Yt} is the wage rate of young workers, and W_{Ot} is the wage rate of old workers. The household takes all prices as given. In our benchmark case, we specify the instantaneous utility function to be

$$U_Y = \log C_Y - \psi_Y N_Y^{1+\theta_Y} / (1 + \theta_Y), \quad U_O = \log C_O - \psi_O N_O^{1+\theta_O} / (1 + \theta_O).$$

The parameters $\theta_Y, \theta_O \geq 0$ govern the Frisch labor supply elasticity, while $\psi_Y, \psi_O > 0$ are used to calibrate the steady state values of N_Y and N_O . We normalize the time endowment of all family members to unity, so that $0 \leq N_{Yt}, N_{Ot} \leq 1$.¹⁴

Because of additive separability in preferences, optimality entails equating consumption across all family members:

$$C_{Yt} = C_{Ot} = C_t. \tag{3.2}$$

The first-order condition (FONC) for capital holdings is given by:

$$C_t^{-1} = \beta E_t [C_{t+1}^{-1} (r_{t+1} + 1 - \delta)].$$

The FONCs for hours worked are given by:

$$W_{Yt} = \psi_Y C_t N_{Yt}^{\theta_Y},$$

$$W_{Ot} = \psi_O C_t N_{Ot}^{\theta_O}.$$

In our benchmark calibration, we set $\theta_Y = \theta_O$ so that the substitution effect of wage changes on labor supply is equated across workers. Given this, condition (3.2) implies that the income effect of a consumption change on labor supply is equal across young and old workers. Again, the data presented in Section 2 imply that the labor demand channel is required for an RBC model to explain the stylized facts. Adopting identical income and substitution effects allows us to isolate the role of labor demand differences in generating volatility differences across young and old workers.

3.2 Firms

To study differences in demand for young and old labor over the business cycle, we relax two assumptions imposed on the standard RBC model's production technology. First, we allow hours of young, inexperienced workers and old, experienced workers to be distinct factor inputs. Second, we drop the Cobb-Douglas assumption of unit elasticity of substitution across inputs, and consider

¹⁴Francis and Ramey (2008) use a variant of this utility function to study how demographic shifts lead to low frequency movements in hours worked and productivity.

a nested CES functional form. In all of our analysis, we assume that production is constant returns to scale, and that final goods are produced by perfectly competitive firms.

With three factor inputs, there are three possible specifications for the nested CES production function:

Nestings

$$\begin{aligned} (1) & : Y_t = \Upsilon_1(H_{Yt}, \Upsilon_2(K_t, H_{Ot})), \\ (2) & : Y_t = \Upsilon_1(H_{Ot}, \Upsilon_2(K_t, H_{Yt})), \\ (3) & : Y_t = \Upsilon_1(K_t, \Upsilon_2(H_{Yt}, H_{Ot})). \end{aligned}$$

Here, Υ_1 and Υ_2 are CES aggregators, H_{Yt} is labor input of young workers, H_{Ot} is labor input of old workers, and K_t is capital services hired at date t . The representative firm's problem is to maximize profits:

$$\Pi_t \equiv Y_t - r_t K_t - W_{Yt} H_{Yt} - W_{Ot} H_{Ot},$$

taking input prices as given.

3.2.1 Differences in elasticity of substitution

Consider the following production function specification for nesting (1):

$$Y_t = \left[\mu (A_t H_{Yt})^\sigma + (1 - \mu) [\lambda K_t^\rho + (1 - \lambda) (A_t H_{Ot})^\rho]^{\sigma/\rho} \right]^{1/\sigma}, \quad \sigma, \rho < 1. \quad (3.3)$$

Labor-augmenting technology follows a deterministic growth trend with stationary shocks:

$$A_t = \exp(gt + z_t),$$

$$z_t = \phi z_{t-1} + \varepsilon_t, \quad 0 < \phi < 1,$$

where $E(\varepsilon) = 0$, $0 \leq \text{var}(\varepsilon) = \sigma_\varepsilon^2 < \infty$, and $g > 0$ is the growth rate of technology. Since technology augments both H_O and H_Y , and given the households' preferences, the economy exhibits balanced growth.

The degree of diminishing marginal product differs between young and old labor whenever $\sigma \neq \rho$.¹⁵ The elasticity of substitution between old workers and capital is given by $(1 - \rho)^{-1}$,

¹⁵ To see this, consider the extreme example when $\rho < 1$ and $\sigma = 1$. In this case, young labor and the H_O - K composite are perfect substitutes. The marginal product of old labor is diminishing in H_O , whereas the marginal product of young labor is constant (i.e., non-diminishing).

while the elasticity of substitution between young workers and the H_O - K composite is $(1 - \sigma)^{-1}$. Following Krusell, Ohanian, Rios-Rull, and Violante (2000), we define production as exhibiting *capital-experience complementarity* when $\sigma > \rho$ when we equate age with labor market experience.

Profit maximization entails equating factor prices with marginal revenue products. The FONCs are:

$$r_t = Y_t^{1-\sigma}(1 - \mu)\Omega_t\lambda K_t^{\rho-1}, \quad (3.4)$$

$$W_{Ot} = Y_t^{1-\sigma}(1 - \mu)\Omega_t(1 - \lambda)A_t^\rho H_{Ot}^{\rho-1}, \quad (3.5)$$

$$W_{Yt} = Y_t^{1-\sigma}\mu A_t^\sigma H_{Yt}^{\sigma-1}, \quad (3.6)$$

where $\Omega_t \equiv [\lambda K_t^\rho + (1 - \lambda)(A_t H_{Ot})^\rho]^{(\sigma-\rho)/\rho}$.

In nesting (2), H_{Yt} and H_{Ot} are swapped, so that capital and labor input of young workers are in the innermost nesting. Since the production function and FONCs are identical in functional form (just with H_{Yt} and W_{Yt} replaced by H_{Ot} and W_{Ot} , and vice-versa) we do not present them here. Again, the degree of diminishing marginal product of young and old labor differs whenever the elasticity parameters differ.

Krusell et al. (2000) consider a similar production specification in their analysis of low frequency trends in relative wages. They distinguish labor skill by the level of educational attainment (as opposed to age/experience in our model), and disaggregate physical capital into structures and equipment (instead of a single capital stock). Castro and Coen-Pirani (2008) use the same production function as Krusell et al. (2000), where skill is defined by education; their analysis evaluates changes in structural elasticity parameters in accounting for the increased volatility of skilled hours since 1984.

3.2.2 Differences across technology shocks to young and old labor input

In nesting (3) the functional form becomes:¹⁶

$$Y_t = \left[\tilde{\mu} K_t^{\tilde{\sigma}} + (1 - \tilde{\mu}) \left[\tilde{\lambda} (A_{Yt} H_{Yt})^{\tilde{\rho}} + (1 - \tilde{\lambda}) (A_{Ot} H_{Ot})^{\tilde{\rho}} \right]^{\tilde{\sigma}/\tilde{\rho}} \right]^{1/\tilde{\sigma}}, \quad \tilde{\sigma}, \tilde{\rho} < 1. \quad (3.7)$$

Again, the FONCs equate factor prices with marginal revenue products:

$$r_t = Y_t^{1-\tilde{\sigma}} \tilde{\mu} K_t^{\tilde{\sigma}-1}, \quad (3.8)$$

$$W_{Yt} = Y_t^{1-\tilde{\sigma}} (1 - \tilde{\mu}) \tilde{\Theta}_t \tilde{\lambda} A_{Yt}^{\tilde{\rho}} H_{Yt}^{\tilde{\rho}-1}, \quad (3.9)$$

¹⁶We note these parameters with tildes because they have the same functional interpretation as in the previous specification, but do not generate different degrees of substitutability of young and old labor with capital.

$$W_{Ot} = Y_t^{1-\bar{\sigma}}(1 - \tilde{\mu})\Theta_t(1 - \tilde{\lambda})A_{Ot}^{\tilde{\rho}}H_{Ot}^{\tilde{\rho}-1}, \quad (3.10)$$

where $\Theta_t \equiv \left[\tilde{\lambda}(A_{Yt}H_{Yt})^{\tilde{\rho}} + (1 - \tilde{\lambda})(A_{Ot}H_{Ot})^{\tilde{\rho}} \right]^{(\bar{\sigma}-\tilde{\rho})/\tilde{\rho}}$.

When $A_{Yt} \equiv A_{Ot}$, productivity shocks affect the marginal product of young and old labor in an identical fashion. As a result, the cyclical characteristics of labor demand do not differ.¹⁷ In order to generate labor demand differences in nesting (3), we must assume that shocks to labor-augmenting technology affect H_{Yt} and H_{Ot} differently. To this end we impose that:

$$A_{Yt} = \exp(\tilde{g}t + z_{Yt}), \quad A_{Ot} = \exp(\tilde{g}t + z_{Ot}),$$

$$z_{Yt} = \tilde{\phi}z_{Yt-1} + \psi\varepsilon_t, \quad \psi > 1,$$

$$z_{Ot} = \tilde{\phi}z_{Ot-1} + \varepsilon_t,$$

with $0 < \tilde{\phi} < 1$. That is, A_{Yt} and A_{Ot} are subject to stationary shocks with the same persistence and identical shock innovations. However, the impact of the innovations on A_{Yt} are larger, so that young labor input is subject to technology shocks with larger variance compared to old labor input. We define production as exhibiting *youth biased technology shocks* when $\psi > 1$.

Of course, ψ cannot be estimated or calibrated to match first moments. As a result, we must impose a value for ψ . Note that in the case of nesting (1) (or nesting (2)), no such assumption need be made; young and old labor input are subject to the same technology shock, A_t . Differences in cyclical labor demand arise naturally from the difference in elasticity of substitution with respect to capital, and moreover, need not be assumed, but are estimated from aggregate data.

3.3 Equilibrium

Equilibrium is defined as follows. Given $\tilde{K}_0 > 0$ and the stochastic process(es) for technology, a *competitive equilibrium* is an allocation, $\{C_t, N_{Yt}, N_{Ot}, \tilde{K}_{t+1}, Y_t, H_{Yt}, H_{Ot}, K_t\}$, and price system, $\{W_{Yt}, W_{Ot}, r_t\}$, such that: given prices, the allocation solves both the representative household's problem and the representative firm's problem for all t ; and factor markets clear for all t :

$$K_t = \tilde{K}_t; \quad H_{Yt} = s_Y N_{Yt}; \quad H_{Ot} = (1 - s_Y) N_{Ot}.$$

Walras' law ensures clearing in the final goods market:

$$C_t + K_{t+1} = Y_t + (1 - \delta)K_t, \quad \forall t.$$

Finally, for the purposes of model evaluation, we define aggregate hours worked as $H_t = s_Y H_{Yt} + (1 - s_Y) H_{Ot}$.

¹⁷Indeed, the standard RBC model is a special case with $A_{Yt} \equiv A_{Ot}$, $\bar{\sigma} = 0$ and $\tilde{\rho} = 1$.

4 Analytical Results

In this section, we provide analytical results illustrating how the models generate a higher relative cyclicity of young hours to old workers. We then discuss their implications for the relative cyclicity of real wages. In Section 6 we present results for quantitative versions of the models.

4.1 Youth Biased Technology Shocks

We begin with the case in which production exhibits youth biased technology shocks, nesting (3). It is easy to show that the response of young hours to a shock innovation is greater than that of the old, even when there are no differences in labor supply characteristics.

Proposition 1 *Let $\theta_Y = \theta_O \geq 0$ and $\psi > 1$. The response of hours of young workers to a business cycle shock is greater than the response of hours of old workers.*

To see this, consider the firm's FONCs with respect to labor, (3.9) and (3.10).¹⁸ In $\log W$ - $\log H$ space, these define linear labor demand curves with common slope, $(\tilde{\rho} - 1)$.¹⁹ Consider the effect of an innovation to technology. Since $\psi > 1$, this results in a larger impulse to $\log A_Y$ than $\log A_O$. That is, the vertical shift in the labor demand curve for the young is greater than that for the old. This is responsible for the result of Proposition 1.

We show this diagrammatically in Figure 1. The left panel depicts the demand curve for young labor, the right panel for old labor. The horizontal line depicts the labor supply curves derived from the household's FONCs with Rogerson-Hansen preferences and lotteries; in log-log space, both are linear with common slope, $\theta_Y = \theta_O = 0$. This restriction is made for graphical simplicity and, as indicated, the result is independent of the (common) slope of the labor supply curves.


Consider the response to a positive innovation to technology, $\varepsilon > 0$. The equilibrium income effect of this shock generates an upward shift in the labor supply curves; since our model assumes identical income effects, we abstract from these in the diagram for the sake of clarity. Using a circumflex to denote log deviations, the shock generates $\hat{Y} > 0$ and $\hat{\Theta} > 0$ in (3.9) and (3.10). Since these are identical across labor demand curves, we abstract from these as well.

Hence, the only effect that requires diagrammatic consideration is the direct effect of the innovation to the labor demand curves, and we plot these in Figure 1. The innovation generates a larger vertical shift in the demand for H_Y than for H_O : $\hat{A}_Y = \psi \hat{A}_O$ so that $\hat{A}_Y > \hat{A}_O$. The innovation

¹⁸A formal presentation of this is straightforward and available from the authors upon request.

¹⁹We are abstracting from the fact that Y_t and Θ_t are functions of H_{Yt} and H_{Ot} . Since these terms are identical across the FONCs, this is inconsequential to the analysis.

Figure 1: DIAGRAMS – YOUTH BIASED TECHNOLOGY SHOCKS MODEL
 Young Old



Notes: Red lines labeled “S” depict the labor supply curves derived from the household’s FONCs with Rogerson-Hansen preferences in log-log space with common slope $\theta_Y = \theta_O = 0$; blue lines labeled “D” depict labor demand curves. Since $\psi > 1$, $\% \Delta A_Y > \% \Delta A_O$ and therefore $\% \Delta H_Y > \% \Delta H_O$ and the innovation has a larger effect on the marginal product of young labor, the *relative shift* effect.

has a larger effect on the marginal product of young labor, and we refer to this as the *relative shift* effect. Hence, in equilibrium, $\hat{H}_Y > \hat{H}_O$.

4.2 Capital-Experience Complementarity

Now consider the case of nesting (1), when production displays capital-experience complementarity. Again, the response of young hours to a technology shock is greater than that of the old, even when there are no differences in labor supply characteristics.

Proposition 2 *Let $\theta_Y = \theta_O \geq 0$ and $\sigma > \rho$. The response of hours of young workers to a business cycle shock is greater than the response of hours of old workers.*

Here, we demonstrate this result for the special case with $\rho = 0$.²⁰ When $\rho = 0$, the H_O - K composite in equation (3.3) becomes Cobb-Douglas, and the firm’s FONCs simplify as:

$$W_{Yt} = \mu Y_t^{1-\sigma} A_t^\sigma H_{Yt}^{\sigma-1},$$

$$W_{Ot} = (1 - \mu) (1 - \lambda) K_t^{\lambda\sigma} Y_t^{1-\sigma} A_t^{(1-\lambda)\sigma} H_{Ot}^{(1-\lambda)\sigma-1}.$$

In log W - log H space, these define linear labor demand curves, with slope $(\sigma - 1)$ for young labor and slope $[(1 - \lambda)\sigma - 1]$ for old. Since $0 < \lambda < 1$ and $0 < \sigma < 1$, the demand curve for young

²⁰A detailed proof is available from the authors upon request.

labor is flatter than that of old labor. Moreover, a shock to technology (a change in $\log A$, which is common to both FONCs) generates a vertical shift in the young labor demand curve of σ , which is larger than the shift in the old labor demand curve of $(1 - \lambda)\sigma$. These two factors combine to generate the result of Proposition 2.

We show this diagrammatically in Figure 2. Again, the left panel depicts the demand curve for young labor, the right panel for old labor, and the horizontal line depicts the labor supply curves with common slope, $\theta_Y = \theta_O = 0$.

Consider a positive shock to technology, \hat{A} . As before, we abstract from the equilibrium income effect of this shock on the labor supply curves since they are identical across young and old. We also abstract from the output response, \hat{Y} , since this is identical across labor demand curves as well. Finally, since capital is a state variable, the response of capital to the shock is $\hat{K} = 0$.

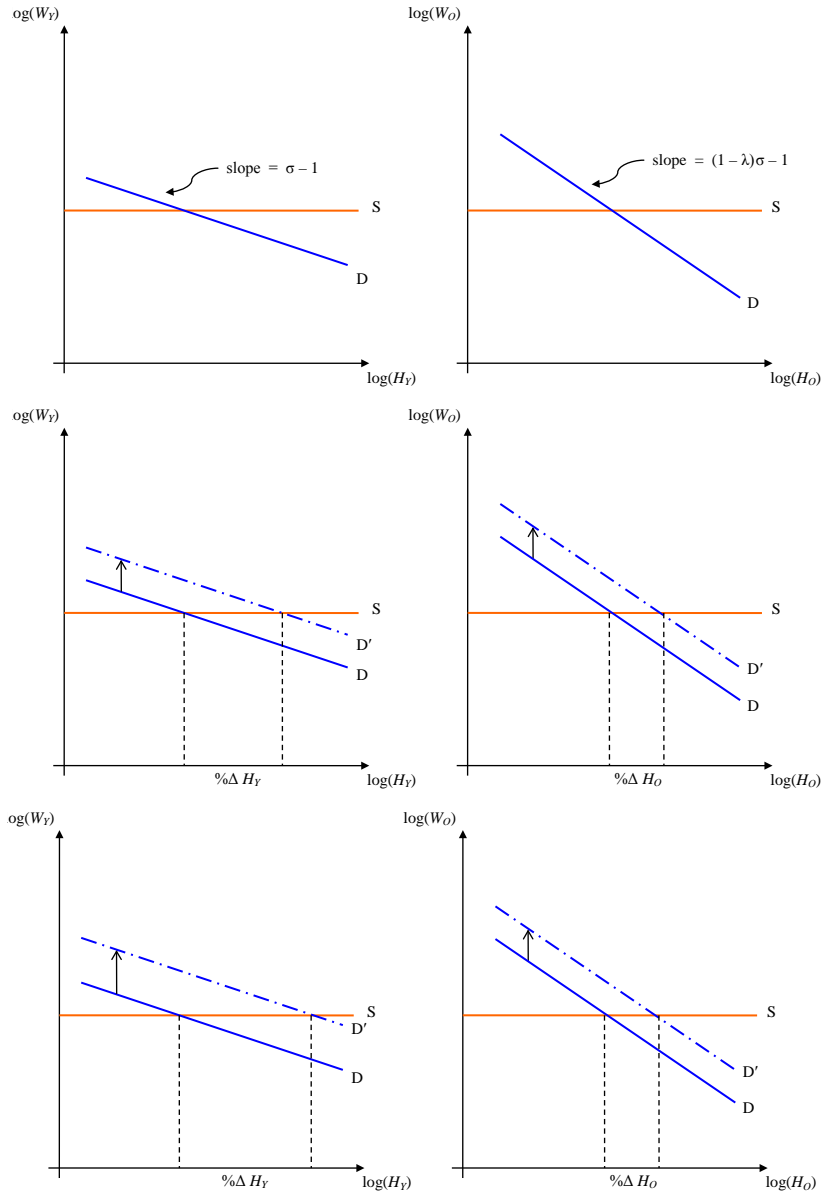
Hence, the only effect that requires diagrammatic consideration is the direct effect of the shock to the labor demand curves, and we plot these in the middle and bottom rows of Figure 2. Suppose, momentarily, that the technology shock results in identical shifts of the two demand curves: this is illustrated as the dotted lines in the middle row. This results in a larger equilibrium response of young labor input relative to the old, i.e. $\hat{H}_Y > \hat{H}_O$. This is due to the relative complementarity of old labor to capital, implying that the marginal product of labor is more sensitive to changes in H_O relative to H_Y . After a positive shift in labor demand, a smaller change in old labor is required to achieve the same change in its marginal product, and we call this the *relative slope effect*.

The positive shock also generates a larger vertical shift in the demand for H_Y than for H_O : $\sigma\hat{A} > (1 - \lambda)\sigma\hat{A}$. This is depicted by the dash-dot line in the left panel of the bottom row. The *relative shift effect* reinforces the relative slope effect. Hence, in equilibrium, $\hat{H}_Y > \hat{H}_O$.

Furthermore, the relative shift effect occurs with capital-experience complementarity, even though old and young labor demand are subject to the same technology shock, A_t . This is in contrast to the youth biased technology shock model, in which we had to assume that shock innovations had larger effects on A_{Yt} as compared to A_{Ot} .


Finally, similar results can be obtained from the production technology described as nesting (2). Given the symmetry with nesting (1), we do not provide detail here. In this case, absent labor supply differences, the response of young hours worked to a business cycle shock is greater than the response of old hours when $\sigma < \rho$.

Figure 2: DIAGRAMS – CAPITAL-EXPERIENCE COMPLEMENTARITY MODEL
 Young Old



Notes – All panels: Red lines labeled “S” depict the labor supply curves derived from the household’s FONCs with Rogerson-Hansen preferences in log-log space with common slope $\theta_Y = \theta_O = 0$; blue lines labeled “D” depict labor demand curves. **Top panels:** slope of demand curve for H_Y is flatter than the demand curve for H_O . **Middle panels:** productivity shock causes both demand curves to shift up; the *relative slope* effect is evident from $\% \Delta H_Y > \% \Delta H_O$. **Bottom panels:** the *relative shift* effect is evident from the labor demand for H_Y shifting up by more than for H_O , increasing $\% \Delta H_Y$ even more.

Figure 3: DIAGRAMS – REAL WAGE RESPONSE
 Young Old



Notes – All panels: Red lines labeled “S” depict labor supply curves; blue lines labeled “D” depict labor demand curves. **Top panels:** labor supply curves derived from the household’s FONCs with finite and positive Frisch elasticity in log-log space with common slope $\theta_Y = \theta_O > 0$, labor demand derived from the youth biased technology shock model. Since $\% \Delta A_Y > \% \Delta A_O$ we have that $\% \Delta W_Y > \% \Delta W_O$ (and as seen before, the young hours response is also larger). **Bottom panels:** young labor supply curve is perfectly elastic, old labor supply curve is less elastic, and labor demand is the same across young and old. In this case, $\% \Delta H_Y > \% \Delta H_O$ as desired, but $\% \Delta W_Y < \% \Delta W_O$.

4.3 The Response of Real Wages

Differences in the relative shift effect and/or the relative slope effect in labor demand generate greater cyclical-ity of hours worked for young workers relative to old workers. However, for the extreme case with infinite Frisch labor supply elasticities ($\theta_Y = \theta_O = 0$) displayed, the equilibrium wage response is equated across young and old labor. For the more general case of positive, but finite, Frisch elasticity ($\theta_Y = \theta_O > 0$), the response of the young wage to business cycle shocks will be greater than that of the old. This is illustrated in the top row of Figure 3 for the youth biased technology shock model. The diagram for the capital-experience complementarity model would obviously be similar and demonstrate the same result.

This can be shown analytically from the household's FONCs with respect to labor supply. Using the fact that consumption is equated across agents:

$$W_{Yt}/\psi_Y N_{Yt}^{\theta_Y} = W_{Ot}/\psi_O N_{Ot}^{\theta_O}.$$

Substituting in the labor market clearing conditions, this can be rewritten in terms of log deviations as

$$\hat{W}_Y - \hat{W}_O = \theta_Y \hat{H}_Y - \theta_O \hat{H}_O.$$

When $\theta_Y = \theta_O > 0$, $\hat{W}_Y > \hat{W}_O$ follows directly when $\hat{H}_Y > \hat{H}_O$. The reason is straightforward: when workers have identical labor supply curves, the only way to induce a greater hours response for the young is through a larger wage response.²¹

Note, however, that this condition implies a stronger result. With cyclical differences in labor demand, the wage response of young workers can be greater than that of the old, even when young labor supply is more elastic (i.e. when $\theta_Y < \theta_O$). We view this as an important result, since our point is not to claim that labor supply characteristics are identical across all agents. Relative to the old, the Frisch elasticity of young workers may be greater, or the income effect of wage changes for young workers may be smaller, or both. Models featuring labor demand differences are still capable of delivering greater cyclical volatility of hours and wages for the young relative to the old.

On the other hand, assume that there are no differences in labor demand characteristics. As discussed in Section 2, then one must assume that the Frisch labor supply elasticity of the young is higher than that of the old (and/or the income effect of wage changes is smaller) in order to match the fact that young hours are more responsive to business cycle shocks.²² However, such a model cannot match the fact that \hat{W}_Y is more responsive than \hat{W}_O . This is illustrated in the bottom row of Figure 3, where the labor supply curve of the young is depicted as being perfectly elastic. More generally, as long as labor supply is more elastic for the young relative to the old, the wage response of the young will be smaller in response to identical labor demand fluctuations. Hence, matching *both* the higher relative volatility of young hours *and* young wages requires a model where labor demand characteristics are not age neutral.

²¹This also indicates that when $\theta_Y = \theta_O = 0$, $\hat{W}_Y = \hat{W}_O$.

²²We maintain the hypothesis that business cycles are generated by shocks to technology. A model emphasizing the role of shocks to labor supply would incorrectly predict that real wages are countercyclical.

5 Quantitative Specification

In this section, we describe the quantitative specification of our models. To maintain comparability with the RBC literature, we perform a standard calibration when possible. However, the parameters governing elasticities of substitution in production cannot be calibrated to match first moments in the U.S. data. Instead, we adopt a structural estimation procedure to identify these values using data from the NIPA and CPS. After describing the procedure, we discuss calibration of the remaining parameter values.

The empirical results in Section 2, indicate that for individuals younger than 30 years, these age groups account for a greater fraction of aggregate hours volatility than for mean aggregate hours. As such, we identify young and old workers in the model with 15-29 and 30-64 year old age groups, respectively, in the data. This is obviously an extreme specification, in that the distinction between young and old labor input is determined by a single age threshold. However, we have found in our estimation results are robust to five year variation in this cut-off age.

5.1 Estimation Procedure

We estimate the elasticity parameters in production from the factor demand equations implied by our various models.²³ Across models, these equations differ in terms of factor inputs and prices, but maintain the same functional forms. As such, our strategy is identical across models, and differences occur only in the labeling of variables across estimating equations.

5.1.1 Empirical Strategy

Consider the capital-experience complementarity model, nesting (1). The firm's FONC with respect to H_{Yt} , equation (3.6) rewritten in logged, first-differenced form is:

$$\Delta \log W_{Yt} = \alpha_0 + (\sigma - 1)\Delta \log (H_{Yt}) + (1 - \sigma)\Delta \log (Y_t) + \sigma u_t, \quad (5.1)$$

where α_0 is a constant, and u_t is a function of current and lagged shock innovations,

$$u_t = \varepsilon_t - (1 - \phi) (\varepsilon_{t-1} + \phi \varepsilon_{t-2} + \phi^2 \varepsilon_{t-3} + \dots).$$

Hence, equation (5.1) represents a textbook labor demand equation (logged and first-differenced). Since $\sigma < 1$ by assumption, the firm's demand for young labor, H_Y , is a negative function of the

²³A similar approach is used in Burnside, Eichenbaum, and Rebelo (1995) and the references therein.

wage, W_Y . Moreover, a rise in the firm's scale of production, Y , results in a rise in young labor demand at any given wage.

As a result, our strategy for estimating the elasticity parameter, σ , amounts to estimating the slope or responsiveness of the young labor demand relation. Empirical identification is obtained from the response of W_Y to (exogenous) changes in H_Y and Y ; moreover, since (5.1) is in first-difference form, we are exploiting short-run, or period-to-period, variation in the aggregate data. Abstracting from endogeneity issues (which we address below), σ could be estimated from a simple restricted least-squares regression.

To implement this strategy, we must make one modification due to the nature of the data. The age-specific wage measures analyzed in Section 2 are constructed using hours data in order to translate direct information on labor income in the CPS into information on hourly wages. To avoid the obvious problems from using W_{Yt} , which is partly constructed from H_{Yt} , we estimate a variant of (5.1) for which direct data on the left-hand side variable is available. This is obtained by multiplying both sides of the FONC by H_{Yt} :

$$\Delta \log LI_{Yt} = \alpha_1 + \sigma \Delta \log H_{Yt} + (1 - \sigma) \Delta \log Y_t + \sigma u_t. \quad (5.2)$$

where $LI_{Yt} \equiv W_{Yt}H_{Yt}$ denotes labor income earned by young workers.

To estimate ρ , we proceed in a similar manner. The firm's FONCs with respect to H_{Ot} and K_t , (3.5) and (3.4), represent factor demand equations for old labor and capital. Since these are of the same functional form, they can be combined to obtain an equation that depends on the slope parameter, ρ , alone. In logged, first-differenced form:

$$\Delta \log (W_{Ot}/r_t) = \alpha_2 + (\rho - 1) \Delta \log (H_{Ot}/K_t) + \rho u_t. \quad (5.3)$$

Again, obtaining data on wages (that are independent of hours worked) and the rental rate of capital is problematic. We perform a similar manipulation as before to obtain our estimating equation:

$$\Delta \log (Q_{Ot}/Q_{Kt}) = \alpha_2 + \rho \Delta \log (H_{Ot}/K_t) + \rho u_t, \quad (5.4)$$

where Q_{Ot} denotes the share of national income earned by old labor, and Q_{Kt} the share of national income earned by capital. Identification of ρ is obtained from the response of national income shares to (exogenous) short-run variation in the factor input share.

Hence, our estimation of elasticities relies only on factor input demand relations. Importantly, our procedure does not require imposing any restrictions from the model's specification of household

behavior.²⁴ The only assumptions required to pin down σ and ρ are: (i) profit maximization on the part of firms, and (ii) that changes in factor prices reflect changes in marginal revenue products. As is obvious from our estimating equations, (5.2) and (5.4), identification does not rely upon the fact that young hours are more volatile over the cycle than old hours. Moreover, no aspect of our approach imposes that $\sigma > \rho$; whether this is satisfied depends on the relation between aggregate prices and quantities observed in the data. The empirical strategy for the other production specifications is analogous, so we do not go into detail here.²⁵

5.1.2 Exogeneity

Since our estimating equations are based on factor demand equations, we must address the endogeneity of the regressors to the error term. In (5.2) and (5.4) for instance, the structural equations identify the error term as due to shocks to technology. To obtain unbiased estimates, we must isolate variation in our regressors that is unrelated to shocks shifting firms' factor demand, be they technology shocks or other omitted factors from the FONCs.

We do so by adopting an instrumental variables (IV) approach. Intuitively, our instruments isolate variation in regressors that are due to labor supply shocks, allowing us to map out the slopes of factor demand. We use two instruments: the Ramey and Shapiro (1998) dates indicating the onset of exogenous military build-ups and lagged birth rates. In a standard RBC model like the one we consider, government spending shocks introduce exogenous shifts in labor supply due to their associated income effects (see Christiano and Eichenbaum (1992)). This results in changes in H_Y , H_O , and Y that are unrelated to shifts in factor demand.

Our second instrument is lagged birth rates. This instrument allows us again to identify changes in current labor supply – this time due to changes in past fertility – that are uncorrelated to shifts in factor demand.²⁶ Recall that:

$$u_t = \varepsilon_t - (1 - \phi) (\varepsilon_{t-1} + \phi\varepsilon_{t-2} + \phi^2\varepsilon_{t-3} + \dots).$$

²⁴We see this as a virtue since our goal is not to claim that labor supply characteristics are indeed identical across the young and old, as maintained in our benchmark calibration. Instead, our goal is to isolate the quantitative role of differences in the cyclical demand for young and old labor.

²⁵One minor difference deserves mention for the youth biased technology shock model, nesting (3). The estimate for $\tilde{\sigma}$ is obtained from the demand function for capital services as:

$$\Delta \log Q_{K_t} = \alpha_3 + \tilde{\sigma} \Delta \log (K_t/Y_t).$$

Shocks to technology do not directly enter this equation. Hence, deviations from this equation, which are observed in the data, are accounted for by the presence of classical measurement error.

²⁶See also Beaudry and Green (2003) who use exogenous demographic variation as an instrument in production function estimation.

Lagged birth rates are valid if fertility is exogenous to past technology shock innovations, $\{\varepsilon_{t-j}\}_{\forall j>0}$. If one believes that fertility decisions, say, 15 years ago might be endogenous to innovations at least 15 years ago, then some bias might be induced. However, note that in the case of the 15-year lagged birth rate, the concern is its correlation with the sum $(1 - \phi) \sum_{j=14}^{\infty} \phi^j \varepsilon_{t-j-1}$ in u_t . For standard values of shock persistence, ϕ , relevant for our analysis, this impact is almost negligible. Obviously, for birthrates of larger lag, this is even smaller. We thus conclude that, from an empirical standpoint, lagged birth rates are valid instruments.

5.2 Model Estimation and Discrimination

For the capital-experience complementarity model, we estimate equations (5.2)-(5.4) as described above. For the two other production function nestings, suitably modified versions of these equations are estimated using the same instruments. Our estimation method is the two-step version of Hansen's (1982) Generalized Method of Moments, which following Newey and McFadden (1986) is asymptotically-efficient. Heteroskedasticity and autocorrelation robust standard errors are estimated following Andrews (1991). We pursue a systems approach because our theory suggests that the error terms are correlated.²⁷ Nevertheless, estimating the system equation-by-equation does not qualitatively change our results.

Thus, each nesting involves estimation of its own pair of equations, which we list here:

$$\text{Nesting 1: } \Delta \log LI_{Yt} = a_1 + b_1 \Delta \log H_{Yt} + (1 - b_1) \Delta \log Y_t + e_t \quad (5.5)$$

$$\Delta \log(Q_{Ot}/Q_{Kt}) = a_2 + b_2 \Delta \log(H_{Ot}/K_t) + e_t \quad (5.6)$$

$$\text{Nesting 2: } \Delta \log LI_{Ot} = a_1 + b_1 \Delta \log H_{Ot} + (1 - b_1) \Delta \log Y_t + e_t \quad (5.7)$$

$$\Delta \log(Q_{Yt}/Q_{Kt}) = a_2 + b_2 \Delta \log(H_{Yt}/K_t) + e_t \quad (5.8)$$

$$\text{Nesting 3: } \Delta \log(Q_{Kt}) = a_1 + b_1 \log(K_t/Y_t) + e_t \quad (5.9)$$

$$\Delta \log(LI_{Yt}/LI_{Ot}) = a_2 + b_2 \log(H_{Yt}/H_{Ot}) + e_t \quad (5.10)$$

In nesting (1), b_1 estimates the elasticity between young labor and the old-labor/capital composite (σ) while b_2 estimates the elasticity between old labor and capital (ρ). In nesting (2), b_1 estimates the elasticity between old labor and the young-labor/capital composite while b_2 estimates the elasticity between young labor and capital. In nesting (3), b_1 estimates the elasticity between capital and the labor composite ($\bar{\sigma}$) while b_2 estimates the elasticity between old labor and young labor ($\bar{\rho}$).

²⁷This is not necessarily true for the production function of nesting (3) – see footnote 25.

Table 4: ESTIMATION RESULTS

| Nesting | Specification Test | Estimates | |
|---------|-----------------------|------------------|------------------|
| | | b_1 | b_2 |
| (1) | 0.16 | 0.72** (0.15) | 0.16 (0.17) |
| (2) | 0.00** | 0.60** (0.15) | 0.41 (0.22) |
| (3) | 0.57 | -0.34 (0.18) | 0.84** (0.06) |

Notes: Data from the March CPS, 1968-2005. Nesting (1) is the benchmark capital-experience complementarity model, estimation equations (5.5) and (5.6), b_1 estimates σ and b_2 estimates ρ . Nesting (2) is the alternate capital-experience complementarity model, estimation equations (5.7) and (5.8). Nesting (3) is the youth biased technological shock model, estimation equations (5.9) and (5.10). Estimation is two-step GMM with Andrews’s (1991) HAC standard errors (in parentheses) using lagged birth rates and the growth rate of defense spending as instruments. The specification test is the Ramsey test of null hypothesis that the linear equations are correctly specified – p -value reported. * indicates significance at the 5%-level, ** significance at the 1%-level.

Given that we are investigating three potential models which differ in terms of production technology, we pursue a simple specification test to discriminate between them. Our approach is to test for misspecification in the models’ factor demand equations to determine if any of the three alternatives are rejected by the data. To this end, we perform the Ramsey test for model misspecification. Each model’s production function delivers a pair of linear estimating equations. As discussed in Davidson and MacKinnon (2004), the Ramsey test is a straightforward and powerful method of testing the null hypothesis that the linear estimating equations are correctly specified.

The estimation results for nesting (1) are presented in the first row of Table 4. Column 1 presents the p -value for the Ramsey test. We fail to reject the null hypothesis of correct specification in equations (5.2)-(5.4); the FONCs derived from this nested CES production function are a reasonable depiction of the data.

The point estimate of $\rho = 0.16$ indicates that the elasticity of substitution is near unity in K and H_O ; in contrast, the estimate of $\sigma = 0.72$ indicates that the substitution elasticity is substantially larger between H_Y and the K - H_O composite. Conducting the F-test of the null hypothesis that $\sigma = \rho$, we obtain a p -value of 0.02 so that the difference between σ and ρ is statistically significant at the 5% level. Moreover, the difference is in the “right” direction for the interpretation of capital-

experience complementarity ($\sigma > \rho$).

Finally, we have calculated the Anderson (1951) and Cragg and Donald (1993) tests of the weak instrument null hypothesis. We reject using either statistic in both (5.2) and (5.4) with p -values below 0.05 in all cases. Thus, we conclude that weak instruments are not a cause for concern.

The results for nesting (2) are presented in the second row of Table 4. The specification test for nesting (2) is soundly rejected by the data: the Ramsey test produces a p -value of less than 0.01, indicating strong evidence of model misspecification. Therefore, we do not evaluate further the quantitative properties of this model specification.

We present results for nesting (3) in the final row of Table 4. According to the Ramsey test, the data fail to reject this specification. The point estimate of $\tilde{\sigma} = -0.34$ is statistically insignificant, while the estimate of $\tilde{\rho} = 0.84$ is close to unity.²⁸ Note that the standard Cobb-Douglas specification with perfect substitutability between young and old labor is a special case of nesting (3) with values of $\tilde{\sigma} = 0$ and $\tilde{\rho} = 1$. Finally, we note that the Anderson (1951) and Cragg and Donald (1993) tests reject the hypothesis of weak instruments in this specification.

In what follows we consider the quantitative predictions of nestings (1) and nesting (3), the capital-experience complementarity and youth biased technological shock models, respectively.

5.3 Calibration

The remaining parameters are calibrated in the standard manner. We set $\beta = 0.99$ and $\delta = 0.025$ to correspond to quarterly time periods. The values of s_Y , ψ_Y , and ψ_O are set to match the average values of the 15-29 year old population shares, and fractions of time spent in market activities by young and old individuals observed in postwar U.S. data. Since θ_Y and θ_O govern elasticities, we cannot calibrate these to match first moments. Moreover, microeconomic estimates do not necessarily correspond to the representative household's labor supply elasticity, as noted by Rogerson (1988) and many others. As such, we consider various values to illustrate the quantitative properties of our models.

Following Krusell, Ohanian, Rios-Rull, and Violante (2000), we calibrate the share parameters in production, μ and λ , to match national income shares. Specifically, given the estimated values for σ and ρ in nesting (1), and $\tilde{\sigma}$ and $\tilde{\rho}$ in nesting (3), we set μ and λ to match the 1968-2005 national income shares of $Q_K = 0.373$ and $Q_O = 0.494$.

With values for $\{\hat{\sigma}, \hat{\rho}, \mu, \lambda\}$ for nesting (1), the capital-experience complementarity model, we

²⁸The F-test for $\tilde{\sigma} = \tilde{\rho}$ is rejected at a p -value less than 0.01.

back out the implied technology series, $\{A_t\}$, using data on output and factor inputs.²⁹ Since we study log-linearized dynamics around steady state, the model’s second moment properties are invariant to the value of σ_ε^2 . As such, our primary interest is in the parameter, ϕ , governing persistence in the technology shock at a quarterly frequency. From $\{A_t\}$, we obtain an estimate of $\hat{\phi} = 0.935$. Since our estimate is not statistically different from 0.95, the value most commonly used in RBC studies, we adopt a value of $\phi = 0.95$ for our quantitative evaluation. None of the model’s implications are substantively different across values of 0.935 and 0.95.

In contrast to nesting (1), we are not able to back out series for $\{A_{Yt}\}$ and $\{A_{Ot}\}$ for nesting (3), the youth biased technology shock model. As such, we cannot identify the parameters, ϕ and ψ , governing, respectively, the persistence of shocks and the relative impact of shock innovations on technology. Given this, we consider a range of values for ψ to investigate the plausibility of the model’s predictions, and simply set $\phi = 0.95$. This last choice is made to facilitate comparison with the analysis of nesting (1) and the standard RBC model.

6 Quantitative Evaluation

Column I in Table 5 presents business cycle statistics for HP-filtered U.S. data. As is well known, the volatility of aggregate hours is almost identical to the volatility of output (the ratio of standard deviations is 0.97). The remaining rows in Column I report the relative volatility of hours and wages for the two age groups, 15-29 year olds and 30-64 year olds.³⁰ While aggregate hours worked is as volatile as output, this masks large differences across the young and the old. The hours of the young are about 50% more volatile than output, while the hours of the old are less volatile than output. As noted in Section 2.3, the volatility of real wages is also greater for the young than for the old. For our two age groups, the ratio of real wage volatility is 1.26. Since our model embodies a standard inter-temporal Euler condition, the model’s predictions for consumption and investment are unchanged relative to a standard RBC model; as such, we do not report them here.

²⁹We are calibrating a quarterly model, but our empirical results use annual data. The reason for this is that quarterly data on age-specific hours do not begin until 1976. We have constructed *semiannual* data on age-specific hours from 1968-2005 from the March CPS and the October CPS surveys. From this we see that the relevant time series display the same volatilities relative to output. Likewise, the relative volatilities of young and old hours are the same in both the annual and semiannual time series. We conclude that for these relationships the frequency of observation does not alter our results.

³⁰For proper comparison between the model and data, we compute the hours worked by 15-29 year olds in the data by efficiency weighting individuals using the procedure discussed above and in Appendix A.2. The same procedure is done for individuals within the 30-64 year olds.

Table 5: DATA AND MODEL MOMENTS

| | DATA | (1) <i>K</i> -EXP | | (3) YOUTH BIASED SHOCKS | | | | |
|-------------------|------|-------------------|---------------------------|-------------------------|-------------------------|---|-------|-------------------------|
| | (I) | (II) | (III) | (IV) | (V) | (VI) | (VII) | (VIII) |
| $sd(H)/sd(Y)$ | 0.97 | 1.03 | 0.97 | 0.56 | 0.65 | 0.63 | 0.75 | 0.81 |
| $sd(H_Y)/sd(Y)$ | 1.48 | 1.85 | 1.59 | 0.56 | 1.48 | 1.48 | 0.75 | 1.48 |
| $sd(H_O)/sd(Y)$ | 0.82 | 0.70 | 0.72 | 0.56 | 0.32 | 0.30 | 0.75 | 0.55 |
| $sd(H_Y)/sd(H_O)$ | 1.81 | 2.65 | 2.21 | 1.00 | 4.57 | 4.96 | 1.00 | 2.69 |
| $sd(W_Y)/sd(Y)$ | 0.35 | 0.28 | 0.36 | 0.31 | 0.31 | 0.40 | 0.31 | 0.31 |
| $sd(W_O)/sd(Y)$ | 0.25 | 0.28 | 0.29 | 0.31 | 0.31 | 0.31 | 0.31 | 0.31 |
| $sd(W_Y)/sd(W_O)$ | 1.26 | 1.00 | 1.26 | 1.00 | 1.00 | 1.27 | 1.00 | 1.00 |
| $sd(Y)/sd(z)$ | - | 1.51 | 1.44 | 0.97 | 1.02 | 1.11 | 1.18 | 1.18 |
| σ | - | 0.72 | 0.72 | -0.34 | -0.34 | -0.34 | 0.01 | 0.01 |
| ρ | - | 0.16 | 0.16 | 0.84 | 0.84 | 0.84 | 0.99 | 0.99 |
| θ_Y | - | 0 | 0.06 | 0 | 0 | 0.07 | 0 | 0 |
| θ_O | - | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| ψ | - | NA | NA | 1 | 1.22 | 1.36 | 1 | 1.011 |
| TARGET | - | - | $\frac{sd(W_Y)}{sd(W_O)}$ | - | $\frac{sd(H_Y)}{sd(Y)}$ | $\frac{sd(H_Y)}{sd(Y)}$ & $\frac{sd(W_Y)}{sd(W_O)}$ | - | $\frac{sd(H_Y)}{sd(Y)}$ |

Notes: Column I are sample moments calculated from HP filtered data from March CPS, 1968-2005. Columns II-VIII are sample moments calculated from model-simulated data. Columns II-III are moments simulated from the capital-experience complementarity model (nesting (1)): setting the production function parameters to those estimated in the data, for varying calibrations of the Frisch labor supply elasticities. Columns IV-VIII use the youth biased technological shock model (nesting (3)): setting the production function parameters to those estimated in the data, for varying calibrations of the labor supply elasticities and ratio (ψ) of young to old labor technology shocks. The row TARGET says what moments are targeted by setting parameters.

6.1 Capital-Experience Complementarity

We begin with examination of the capital-experience complementarity model. We initially set $\theta_Y = \theta_O = 0$, so that utility is linear in labor. This is a useful benchmark since the standard RBC model (with homogenous labor and Cobb-Douglas production function) requires very high aggregate labor supply elasticity to generate significant volatility of hours worked; the indivisible labor model (with perfectly elastic labor supply) generates a ratio of the standard deviation of hours to output of approximately 0.7 – 0.75.³¹ In this sense, the volatility of aggregate hours worked represents a puzzle to the RBC literature.

As Column II of Table 5 reports, the capital-experience complementarity model generates

³¹See for example, Hansen (1985), Rogerson (1988), King and Rebelo (1999).

volatility of total hours that is very close to that observed in the data. In fact, aggregate hours are more volatile than output; the relative standard deviation is 1.03. The next row shows that the key to this result is the model’s ability to generate hours worked by the young that fluctuates substantially more than output and old hours over the business cycle. The model generates a volatility ratio of 1.85 for young hours to output, which is greater than the value of 1.48 observed in the data. On the other hand, the model understates the volatility of old hours relative to output: the relative standard deviation is 0.70, while this is 0.82 in the data.³² As such, the model overstates the relative volatility of age-specific hours.

While the benchmark calibration is surprisingly successful along the hours dimension, it cannot account for the behavior of relative wages between the young and the old. This is expected since the Frisch labor supply elasticity is infinite for both young and old agents; as discussed in Section 4, the volatility of age-specific wages is identical in this case.

In Column III we consider the following modification: we change only the labor supply elasticity of the young to match the relative wage volatility ($sd(W_Y)/sd(W_O) = 1.26$) as observed in the U.S data. This requires a minimal change, moving θ_Y from 0 to 0.06. Under this specification the model generates a volatility of age-specific wages relative to output that is close to those in the data. Not surprisingly, the lower elasticity of young labor supply induces a fall in the volatility of young hours, and hence, aggregate hours relative to output.³³ This improves the ability of the model to match the data. The model generates values for $sd(H_Y)/sd(Y) = 1.59$ and $sd(H)/sd(Y) = 0.97$ that are close to the observed values of 1.48 and 0.97, respectively.

Because of its ability to generate volatility in hours worked, the capital-experience complementarity model embodies strong amplification of productivity shocks. Across experiments, the relative volatility of output to the shock process is around 1.4 – 1.5. This is reported in row 8 labeled $sd(Y)/sd(z)$. As we discuss below, this value is substantially larger than in the benchmark RBC model and the model with youth biased technology shocks.

In sum, we find that the capital-experience complementarity model generates a relative volatility

³²Our quantitative specification has an elasticity of substitution between capital and old hours that is close to unity ($(1 - \rho)^{-1} = 1.19$), and infinite Frisch elasticity of labor supply for the old. These are the features displayed by the homogenous labor input in the standard RBC model with indivisible labor, discussed in the preceding paragraph. Thus the capital-experience complementarity model generates a relative volatility of old hours to output, $sd(H_O)/sd(Y)$, similar to the relative volatility of aggregate hours to output in the standard RBC model.

³³Note that we are reporting the volatility for cyclical fluctuations in real wages, as constructed in Section 2.3. As previously shown, a significant portion of high frequency wage variation is not correlated with the cycle. Given the focus on business cycle fluctuations in hours and wages, we concentrate on the variation in wages that is due to the cycle. Note also that like all other RBC models driven by a single shock, the correlation between output and wages is higher than in the data.

of age-specific hours that is similar to that observed in the data. As a by-product of this success, the model generates volatility of aggregate hours that is very close to that of aggregate output. Moreover, the model accounts for the joint behavior of age-specific hours and wages.

Several recent papers study heterogenous agent models that are meant to reproduce key facts about labor supply and the wealth distribution at the micro level. In this work the intertemporal substitution elasticity of leisure is calibrated to microeconomic data, giving rise to a distinct aggregate, or ‘macro,’ labor supply elasticity.³⁴ These papers find a range for the aggregate labor supply elasticity between 0.71 and 3. Given this evidence, we conduct robustness simulations in our model. The goal is to illustrate that for any Frisch labor supply elasticity, the labor demand mechanism that we highlight provides an improvement over the standard RBC model (with Cobb-Douglas production). Specifically we consider calibrations of the labor supply elasticity ranging between 0.71 and 3. We then simulate for each case: (1) a standard RBC model, and (2) our capital-experience complementarity model.

The resulting values of the relative volatility of hours to output, $sd(H)/sd(Y)$, are summarized as follows. At the low end of the range, when the Frisch labor supply elasticity equals 0.71, the benchmark RBC model generates a value of $sd(H)/sd(Y)=0.29$ while our model gives rise to $sd(H)/sd(Y)=0.36$. At the high end, when the labor supply elasticity equals 3, the benchmark RBC model generates a value of $sd(H)/sd(Y)=0.54$ while our model gives rise to $sd(H)/sd(Y)=0.70$. Indeed, throughout the range of elasticity values considered, the capital-experience complementarity magnifies the relative volatility of aggregate hours by approximately 30% as compared to the RBC model.

A final illustration of the strength of our mechanism is as follows. Consider the case where the Frisch labor supply elasticity is 3. The capital-experience complementarity model generates $sd(H)/sd(Y)=0.70$. In order for the standard RBC model to generate the same hours volatility, it would require a Frisch elasticity of 20, which is well outside the range considered by these papers. Hence, our model is able to generate reasonable business cycle dynamics for Frisch elasticities significantly smaller than has been required in previous business cycle models.

6.2 Youth Biased Technology Shocks

Columns IV-VIII present results for the youth biased technology shock model. Again, we initially set $\theta_Y = \theta_O = 0$, so that utility is linear in labor.

³⁴See, for instance, Chang and Kim (2006), Rogerson and Wallenius (2009), Erosa, Fuster, and Kambourov (2010), and Mustre-del-Rio (2011).

In Column IV we set $\psi = 1$ for illustrative purposes. As discussed in Section 3, this specification generates identical hours volatility of young and old. Note also that this model has weak amplification which stems from the elasticity of substitution between capital and the hours composite being less than one. Even with Rogerson-Hansen preferences, the volatility of aggregate hours to output is only 0.56.

In Column V we set ψ to explicitly target the relative volatility of young hours to output, $sd(H_Y)/sd(Y) = 1.48$. This requires $\psi = 1.22$, so that technology innovations have a 22% greater impact on A_Y as compared to A_O . Again, $\theta_Y = \theta_O = 0$ to maximize the model's amplification mechanism. As is obvious, the model is grossly counterfactual regarding the relative volatility of age-specific hours. The standard deviation of young hours is almost 4.6 times that of old hours, whereas in the data, the ratio is 1.84. Moreover, the model fails to solve the hours volatility puzzle: the volatility of aggregate hours to output is only 0.65, far from the value of 0.97 in the data.

Our next experiment sets ψ to explicitly match the relative volatility of aggregate hours to output.³⁵ We do not present these results in Table 5 for brevity, but summarize them here. To obtain $sd(H)/sd(Y) = 0.97$, the youth biased technology shock model requires a value of $\psi = 2.22$. However, in doing so the model wildly overpredicts the volatility of young hours compared to output and old hours; $sd(H_Y)/sd(Y) = 8.35$ and $sd(H_Y)/sd(H_O) = 15.32$ whereas these values are 1.48 and 1.81, respectively, in the data. As such, this model does not provide a resolution to the hours volatility puzzle.

The final thing to note is the weak amplification embodied by all versions of the youth biased technology shock model. Row 8 reports that $sd(Y)/sd(z)$ is near unity, where we are displaying the standard deviation of output relative to that of the shock on old hours, z_O . This obviously maximizes the magnification statistic as much as possible, as relative to the more volatile shock, z_Y , this statistic is obviously even smaller.

Column VI repeats the exercise from Column V, this time also adjusting θ_Y to match the relative volatility of young and old wages. Of course, with lower labor supply elasticity of the young, an even larger value of ψ is required to generate $sd(H_Y)/sd(Y) = 1.48$. Moreover, the performance of the model deteriorates on the key dimensions (relative to the calibration with $\theta_Y = 0$). The overprediction of the volatility of young to old hours is even more pronounced, as is the failure to resolve the hours volatility puzzle.

³⁵Of course, this would be problematic if our interpretation was that this model resolves the hours volatility puzzle, since we are calibrating a second moment parameter (which cannot be measured) to hit the target of interest. As will be obvious, this is not the nature of our results and interpretation.

As a final robustness check, we investigate the properties of the youth biased technology shock model when we no longer use the estimated elasticities, and instead impose $\tilde{\sigma} = 0$ and $\tilde{\rho} = 1$.³⁶ This is an interesting case since it is the analog to the Cobb-Douglas assumption of unit elasticity of substitution between capital and “labor” used in the RBC literature; here, “labor” refers to the weighted sum of young and old hours worked, where the weight is time-varying and given by A_{Yt}/A_{Ot} (in the standard RBC model, the weight would be constant). Again, we set $\theta_Y = \theta_O = 0$ to maximize the model’s amplification potential and responsiveness of hours worked.

For brevity we simply present the results for: (i) $\psi = 1$, and (ii) ψ set to match the relative volatility of young hours to output. These results are displayed in Columns VII and VIII, respectively. When $\psi = 1$, this corresponds to the standard RBC model with Rogerson-Hansen preferences. The relative volatility of aggregate hours to output is 0.75. When the model matches $sd(H_Y)/sd(Y)$, the model again fails to resolve the hours volatility puzzle: $sd(H)/sd(Y) = 0.81$. This represents a minimal improvement over the Rogerson-Hansen model, and is far from the value of 0.97 found in the data.

Thus, while the youth biased technology shock model passes the econometric tests of Section 5, no variation of the model provides a plausible resolution to the hours volatility puzzle. Hence, the data favors the capital-experience complementarity model. It too passes our econometric tests, and generates a relative volatility of aggregate hours to output very close to that of the data. Moreover, the model easily accounts for the relative volatility of age-specific wages. Interestingly, this model also embodies the greatest amplification mechanism among the different specifications considered.

7 Conclusion

In this paper, we address the hours volatility puzzle by explicitly modeling age specific hours fluctuations over the business cycle. Our motivation comes from observing that aggregate hours’ fluctuations are disproportionately accounted for by the young, whose hours vary more than prime-aged workers’ hours. The fact that hours and wages of the young display greater volatility over the cycle relative to the prime-aged points to a class of models featuring labor demand differences by age.

We consider several such models which are straightforward and parsimonious extensions to the standard RBC model. This allows us to clearly identify and quantify the plausibility of the mechanism embedded in each model. We show that the most promising explanation features

³⁶To make this operational, we solve for the model’s log-linear dynamics setting $\tilde{\sigma} = 0.01$ and $\tilde{\rho} = 0.99$.

capital-experience complementarity in production which induces a greater diminishing marginal product of prime-age labor relative to the young. In our analysis we estimate the key structural parameters governing the degree of capital-experience complementarity in a manner that does not target differences in the volatility of age-specific hours.

We find that our quantitative model is able to match the relative volatility of age-specific hours to output, and as a result, also replicates the relative volatility of aggregate hours with respect to output. Moreover, the model accounts for the relative volatility of age-specific wages observed in the data. Thus, this paper demonstrates the value of understanding the volatility of hours worked by age as a potential resolution to the hours volatility puzzle, and as an important propagation mechanism in business cycle analysis.

References

- ANDERSON, T. W. (1951): “Estimating linear restrictions on regression coefficients for multivariate normal distributions,” *Annals of Mathematical Statistics*, 22, 327–351.
- ANDREWS, D. W. K. (1991): “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation,” *Econometrica*, 59(3), 817–58.
- BAXTER, M., AND R. G. KING (1999): “Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series,” *The Review of Economics and Statistics*, 81(4), 575–593.
- BEAUDRY, P., AND D. GREEN (2003): “Wages and Unemployment in the United States and Germany: What Explains the Differences?” *American Economic Review*, 93(3), 573–602.
- BENHABIB, J., R. ROGERSON, AND R. WRIGHT (1991): “Homework in Macroeconomics: Household Production and Aggregate Fluctuations,” *Journal of Political Economy*, 99(6), 1166–87.
- BURNSIDE, C. (2000): “Some facts About the HP Filter,” Discussion paper, World Bank.
- BURNSIDE, C., M. EICHENBAUM, AND S. REBELO (1995): “Capital Utilization and Returns to Scale,” vol. NBER Macroeconomics Annual, pp. 67–110. NBER.
- CASTRO, R., AND D. COEN-PIRANI (2008): “Why Have Aggregate Skilled Hours Become so Cyclical since the Mid-1980s?,” *International Economic Review*, 49(1), 135–185.
- CHANG, Y., AND S.-B. KIM (2006): “From Individual to Aggregate Labor Supply: A Quantitative Analysis based on a Heterogeneous-Agent Macroeconomy,” *International Economic Review*, 47(1), 1–27.
- CHRISTIANO, L. J., AND M. EICHENBAUM (1992): “Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations,” *American Economic Review*, 82(3), 430–50.
- CLARK, K. B., AND L. H. SUMMERS (1981): “Demographic Differences in Cyclical Employment Variation,” *The Journal of Human Resources*, XVI, 61–79.
- CRAGG, J. G., AND S. G. DONALD (1993): “Testing Identifiability and Specification in Instrumental Variable Models,” *Econometric Theory*, 9, 222–240.
- DAVIDSON, R., AND J. G. MACKINNON (2004): *Econometric Theory and Methods*. Oxford University Press.
- DEL RIO, J. M. (2011): “Aggregate Implications of Individual Labor Supply Heterogeneity,” Discussion paper.
- EROSA, A., L. FUSTER, AND G. KAMBOUROV (2010): “The Heterogeneity and Dynamics of Individual Labor Supply over the Life-Cycle: Facts and Theory,” Discussion paper.
- FRANCIS, N., AND V. A. RAMEY (2008): “Measures of Per Capita Hours and their Implications for the Technology-Hours Debate,” Discussion paper, UCSD.
- GOMME, P., R. ROGERSON, P. RUPERT, AND R. WRIGHT (2004): “The Business Cycle and the Life Cycle,” in *NBER Macroeconomics Annual*, vol. 19, pp. 415–461.
- HANSEN, G. D. (1985): “Indivisible labor and the business cycle,” *Journal of Monetary Economics*, 16(3), 309–327.

- HANSEN, G. D., AND S. IMROHOROGLU (2008): “Consumption over the Life Cycle: The Role of Annuities,” *Review of Economic Dynamics*, 11(3), 566–583.
- HANSEN, L. P. (1982): “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, 50, 1029–1054.
- JAIMOVICH, N., AND H. E. SIU (2009): “The Young, the Old, and the Restless: Demographics and Business Cycle Volatility,” *American Economic Review*, 99(3), 804–26.
- KATZ, L. F., AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963-1987: Supply and Demand Factors,” *The Quarterly Journal of Economics*, 107(1), 35–78.
- KING, R. G., AND S. T. REBELO (1999): “Resuscitating real business cycles,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 14, pp. 927–1007. Elsevier.
- KRUSELL, P., L. E. OHANIAN, J.-V. RIOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68(5), 1029–1054.
- NAGYPÁL, E. (2004): “Learning-by-Doing Versus Learning About Match Quality: Can We Tell Them Apart?,” Discussion paper, Northwestern University.
- NEWKEY, W. K., AND D. MCFADDEN (1986): “Large sample estimation and hypothesis testing,” in *Handbook of Econometrics*, ed. by R. F. Engle, and D. McFadden, vol. 4 of *Handbook of Econometrics*, chap. 36, pp. 2111–2245. Elsevier.
- PRESCOTT, E. C. (1986): “Theory ahead of business-cycle measurement,” *Carnegie-Rochester Conference Series on Public Policy*, 25(1), 11–44.
- RAMEY, V. A., AND M. D. SHAPIRO (1998): “Costly Capital Reallocation and the Effects of Government Spending,” *Carnegie-Rochester Conference Series on Public Policy*, 48, 145–194.
- RAVN, M. O., AND H. UHLIG (2002): “On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations,” *The Review of Economics and Statistics*, 84, 371–380.
- RIOS-RULL, J.-V. (1996): “Life-Cycle Economies and Aggregate Fluctuations,” *Review of Economic Studies*, 63(3), 465–89.
- ROGERSON, R. (1988): “Indivisible labor, lotteries and equilibrium,” *Journal of Monetary Economics*, 21(1), 3–16.
- ROGERSON, R., AND J. WALLENIUS (2009): “Micro and Macro Elasticities in a Life Cycle Model with Taxes,” *Journal of Economic Theory*, 144(6), 2277–2292.
- SOLON, G., R. BARSKY, AND J. A. PARKER (1994): “Measuring the Cyclicalities of Real Wages: How Important Is Composition Bias?,” *The Quarterly Journal of Economics*, 109(1), 1–25.

A Appendix

A.1 Specification Testing

To understand the Ramsey test, recall that the conditional expectation $E(Y|X)$ is a function $f(X)$. Therefore we can express the conditional expectation as a Taylor expansion of f . Let that expansion be around the linear prediction of Y , call it Xb ; a linear prediction of the left-hand side variable is what estimating equations provide. Of course, the function f itself is linear if all its higher order (second and beyond) derivatives are zero: said another way, f is *not* linear if there *is* a nonzero coefficient on a higher order expansion term. We look for evidence of higher order expansion terms by regressing the residuals on higher powers of the regression fitted values; in practice one can restrict consideration to low powers of the fitted values as suggested by Davidson and MacKinnon (2004).

For the two estimation equations involving regressors X_1 and X_2 , we run seemingly unrelated regressions of their residuals \hat{u}_1 and \hat{u}_2 on fitted values and fitted values squared³⁷

$$\hat{u}_t = \beta_i(X_i b_i) + \gamma_i(X_i b_i)^2 \quad , \quad i \in \{1, 2\}.$$

The specification test of the null hypothesis $\gamma_1 = \gamma_2 = 0$ has a $\chi^2(2)$ distribution.

A.2 Data

Data on hours, employment shares, and wages come from the Current Population Survey (CPS) conducted by the Census Bureau. To obtain wage data, we use questions in the March CPS about income obtained in the previous (last) year.³⁸ In order to turn this income data into wage data, we must know how many hours the individual worked last year. The hours for the previous year are constructed as the number of weeks worked last year multiplied by a measure of how many hours-per-week were worked by the individual last year. We follow Krusell, Ohanian, Rios-Rull, and Violante (2000) in imputing the hours-per-week from the data on how many hours the individual worked *in the previous (last) week*.

Our measure of hours-per-week is different than Krusell, Ohanian, Rios-Rull, and Violante (2000) in the following. We note whether the worker described her work last year as either full-time (FT) or part-time (PT). Her last week’s hours are imputed as the hours-per-week only if the value falls within believable values, given that her work last year was either FT or PT. If her previous week’s hours are not consistent with FT or PT work, we impute a “disaggregated” group average as the hours-per-week; by contrast, Krusell, Ohanian, Rios-Rull, and Violante (2000) impute a “disaggregated” group average only if the worker reported that she worked last year but worked zero hours last week.

Our “disaggregated” groups are formed by dividing respondents by age, education, gender, and last year’s FT/PT status. Given that there are eleven 5-year age bins (15-19, 20-24, . . . , 60-64, 65+), 5 education bins (below HS, HS, some college, college graduate, postgraduate work), 2 genders, and a FT or PT status, there are 220 possible groups. Our “disaggregated” groups combine education bins for some age-gender-FT/PT groups to ensure that for every year in 1964-2006 our “disaggregated” groups each have at least fifty members.³⁹ This is done so that the “disaggregated” group average is not overly reliant on only a few observations.

Conditional on the other characteristics we consider, we use the information on PT and FT as follows:

³⁷Higher powers of the fitted values produce similar results and the same test results.

³⁸As noted below, a specific question reporting wages only appears in the CPS survey starting in 1982.

³⁹Additionally cutting by race (white/nonwhite) does not change matters much.

- If a person claims to be PT last year and works between 1 and 34 hours last week, we impute their hours-per-week last year as their hours last week; otherwise (they worked 0 or more than 34 hours last week) they are imputed the group average
- If a person claims to be FT last year and works 35 or more hours last week, we impute their hours-per-week last year as their hours last week; otherwise (they worked 0-34 hours last week) they are imputed the group average

Let g' be the part of g with hours last week that are FT-status-fitting for imputation purposes (given the FT/PT nature of g), and g'' be those whose hours last week are not FT-status-fitting. Let h_i , m_i , y_i , and μ_i be worker i 's hours last week, number of weeks worked last year, wage and salary income last year, and CPS Person weight, respectively.⁴⁰ Then the measures of group g 's “disaggregated” group average, weight, hours worked last year, and income last year are

$$h_{g'} = \frac{1}{\sum_{i \in g'} \mu_i} \left(\sum_{i \in g'} h_i \mu_i \right) \quad (\text{A.1})$$

$$\mu_g = \sum_{k \in g} \mu_k \quad (\text{A.2})$$

$$h_g = \frac{1}{\mu_g} \left(\sum_{i \in g'} h_i \mu_i + \sum_{j \in g''} h_{g'} \mu_j \right) \quad (\text{A.3})$$

$$y_g = \frac{1}{\mu_g} \left(\sum_{k \in g} y_k \mu_k \right) \quad (\text{A.4})$$

Let γ be a set of g s: this is a larger group, such as all workers in the 15-19 age category, comprised of smaller “disaggregated” groups. Our construction of an efficiency wage measure for γ is similar to that of Krusell, Ohanian, Rios-Rull, and Violante (2000): our efficiency measurement f for each g is the average of their wage (y_g/h_g) for the years 1985-1989.⁴¹

$$W_\gamma = \frac{\sum_{g \in \gamma} y_g \mu_g}{\sum_{g \in \gamma} h_g f_g \mu_g} \quad (\text{A.5})$$

It is worth mentioning that the March CPS has a specific question “On average, how many hours per week did you work last year, when you worked?” starting in 1976. We find that making sure the hours imputation is FT-status-fitting leads to hours measures that are close to the post-1976 question when both are available. By ignoring the FT-status, one underreports the groups’ hours.

Our data on hours come directly from the hours last week question. Likewise, our labor force share data comes from a labor force status question pertaining to last week.

We have found that these last week hours have level shifts between the 1967 and 1968 survey years and therefore start our hours series at 1968. The last year information used in the wage series appears unaffected during this time, so we use data going back to the 1964 survey year (data about 1963). The statistics on wages remain virtually identical if we start the wage series at survey year 1968.

⁴⁰In the March supplement, we have both a CPS Basic Person weight, and a CPS Supplemental Person weight. Personnel at the Census Bureau have advised us to use the latter for all the data questions we are addressing, even though some of these data are not part of the March Annual Supplement.

⁴¹Krusell, Ohanian, Rios-Rull, and Violante (2000) use the wage in 1980 as the efficiency measurement. We use an average of the wage to allow for the possibility that the efficiency measure varies over the cycle. Hence, by averaging over five years we aim to smooth the efficiency measurement. The results remain the same using either efficiency measurement.