American Economic Review 2013, 103(7): 3022–3044
http://dx.doi.org/10.1257/aer.103.7.3022

The Demand for Youth: Explaining Age Differences in the Volatility of Hours[†]

By Nir Jaimovich, Seth Pruitt, and Henry E. Siu*

Labor market fluctuations over the business cycle differ greatly for individuals of different ages. Perhaps the most salient difference is between the volatility of hours worked of young workers relative to the prime aged (see Clark and Summers 1981; and Gomme et al. 2005). While this fact is well known, the literature lacks a quantitatively successful explanation.[1]

In our view, developing such an explanation is important for our understanding of business cycle and labor market dynamics. As an example, the results of Jaimovich and Siu (2009) show that work-force age composition has a strong causal impact on employment and output volatility. A theoretical explanation of this lies in understanding age differences in cyclical sensitivity. In addition, understanding age differences in the volatility of hours—and specifically, why young hours are so volatile—leads to an understanding of the volatility of aggregate hours as an important corollary. Accounting for the relative volatility of aggregate hours to output remains as one of the puzzles in real business cycle (RBC) analysis (see King and Rebelo 1999; Gomme et al. 2005).

In this article, we study this phenomenon through the lens of a neoclassical model in which households and firms optimize, taking prices as given, and interact in competitive spot markets. We view this as an important exercise given the prominence of this framework in quantitative business cycle analysis. Within this framework, age differences in the volatility of hours can arise from factors related to preferences (or succinctly, differences in labor supply), technology (labor demand), or both. We argue that the *joint* behavior of age-specific hours and wages over the cycle provides the necessary evidence to distinguish between these two channels. Specifically, we document in Section I that the volatilities of *both* hours and wages

* Jaimovich: Department of Economics, Duke University, 213 Social Sciences Building, Durham, NC 27708-0097 and NBER (e-mail: njaimo@gmail.com); Pruitt: Federal Reserve Board, Division of International Finance, 20th St. and Constitution Avenue NW, Washington, DC 20551 (e-mail: seth.j.pruitt@frb.gov); Siu: Department of Economics, University of British Columbia, #997-1873 East Mall, Vancouver, BC V6T 1Z1, Canada, and NBER (e-mail: hankman@mail.ubc.ca). We thank the anonymous referees, Manuel Amador, Paul Beaudry, Larry Christiano, Valerie Ramey, Sergio Rebelo, Victor Ríos-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. A previous version of this paper circulated under the title "The Demand for Youth: Implications for the Hours Volatility Puzzle."

[†] Go to http://dx.doi.org/10.1257/aer.103.7.3022 to visit the article page for additional materials and author disclosure statement(s).

[1] Ríos-Rull (1996), Gomme et al. (2005), and Hansen and Imrohoroglu (2009) study models with differences in hours volatility owing to life-cycle considerations. They find that such factors cannot quantitatively account for the greater volatility of young hours relative to others.

for young workers are greater than those of the prime aged over the cycle.[2] We view these as our two key labor market facts.

We show in Section II that a general class of models featuring only age-specific labor supply differences cannot reconcile these facts. Generating a higher volatility of hours of the young implies a lower volatility of wages, and vice versa.

Consequently, in Section III we present a model that rationalizes the labor market facts by allowing for cyclical differences in age-specific labor demand. Our approach represents a minimal deviation from the standard RBC model, extended to three factor inputs: capital, "young" labor, and "old" labor. We study a model where the elasticity of substitution between capital and labor can differ between young and old. The model features *capital-experience complementarity* in production, when age is equated with labor market experience.[3] We note that our model represents one particular micro-foundation for the differences in the cyclicality of labor demand. A simple alternative is to allow for shocks to young labor input in the production function that are more volatile than shocks to old labor. However, we view this approach as unappealing since it essentially assumes the desired result and lacks the discipline that our approach entails. Specifically, in Section IV, we use our model's factor demand equations and estimate the key elasticity of substitution parameters in a manner that does not target the differences in cyclical volatility of age-specific hours and wages.

We find that our capital-experience complementarity model generates relative volatilities of hours and wages across age groups that are similar to those observed in the data. As a by-product, the model also generates a relative volatility of aggregate hours to output that is essentially unity. These results are presented in Section V. Section VI concludes.

I. Age-Specific Hours and Wages

In this section, we document the empirical findings that motivate our work. We first present evidence on the large differences by age in the volatility of hours over the cycle. We then provide evidence on the cyclicality of age-specific real wages. These facts allow us to distinguish between models in the analysis that follows.

A. *Hours*

The evidence on the cyclicality of age-specific hours has been extensively addressed in Gomme et al. (2005) and Jaimovich and Siu (2009). We provide a brief summary here and refer the reader to the cited papers for greater detail.

[2] This approach takes seriously the assumption of spot labor markets in the neoclassical model. One could imagine a model where contracts insuring hours and/or wage fluctuations are provided differentially by age might also rationalize these phenomena. Nevertheless, we note that existing empirical evidence does not support the hypothesis of age differences in the extent of contracting; see McDonald and Worswick (1999).

[3] 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; and Castro and Coen-Pirani 2008). However, Mincerian wage regressions emphasize two important observable dimensions of skills: education and experience. We show that labor market experience exhibits important complementarities to capital at the business cycle frequency.

Table 1—Volatility of Hours Worked by Age Group

Age group	Filtered volatility	R^2	Cyclical volatility	Hours share	Volatility share
15–19	5.66	0.80	5.08	3.74	12.73
20–24	2.30	0.79	2.04	10.85	14.83
25–29	1.92	0.68	1.58	13.12	13.93
30–39	1.44	0.94	1.40	26.00	24.38
40–49	1.23	0.70	1.03	24.16	16.67
50–59	1.49	0.75	1.29	17.58	15.24
60–64	2.05	0.13	0.73	4.54	2.21
15–29	2.33	0.91	2.22	27.71	41.49
30–64	1.23	0.95	1.20	72.29	58.51

Notes: Data from the March CPS, 1964–2010. 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. Hours share is the sample average share of aggregate hours worked by the age group, reported in percentage. Volatility share is the age group's share of aggregate hours volatility, defined as the average of age-specific cyclical volatilities weighted by hours shares, reported in percentage.

Using data from the March supplement of the Current Population Survey (CPS) over 1964–2010 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. We extract the high frequency component of each series using the Hodrick-Prescott (HP) filter on logged data.[4]

Table 1 presents results on the time series volatility of hours worked by age. The first column 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.

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 column of Table 1 reports the R^2 from these regressions. This is high for most age groups, even for those whose hours make up a small fraction of total hours. This implies that the preponderance of high frequency fluctuations is attributable to the business cycle.[5]

The third column presents the percent standard deviation of the cyclical series. The data indicate a pattern of decreasing volatility with age. The young experience much greater cyclical volatility in hours than all others. The standard deviation of cyclical hours fluctuations for 15–19- and 20–24-year-old workers is five and two

[4] Since we are interested in fluctuations at business cycle frequencies (those higher than eight years), we use a smoothing parameter of 6.25 for annual data; Ravn and Uhlig (2002) find this to be the optimal value through analysis of the transfer function of the HP filter. One may alternatively use a smoothing parameter of ten or the bandpass filter as suggested by Baxter and King (1999) to remove fluctuations less frequent than eight years. The quantitative results are essentially identical in all cases.

[5] The exception is the 60–64 age group, where a larger fraction of fluctuations are due to noncyclical movements.

times that of 40–49-year-olds, respectively.[6] As a group, the hours of young workers aged 15 to 29 years old are about 1.85 times as volatile as for prime-aged workers, aged 30 to 64 years old.

The fourth column indicates the average share of aggregate hours worked by each age group. The fifth column 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 are disproportionately accounted for by young workers, whose share of volatility is markedly greater than their share of hours. Although those aged 15–29 make up only about one-quarter of aggregate hours worked, they account for more than two-fifths of hours' volatility. By contrast, prime-aged workers aged 30 to 64 years account for about three-quarters of hours, but a little less than three-fifths of the volatility.

Participation.—Young individuals might display greater cyclicality of hours worked relative to the prime aged 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 cyclicality 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*.[7] 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, age differences in the participation decision. If not, it would indicate that to a first-order, the primary factor generating age group differences is to be found elsewhere.

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

$$\mathrm{Var}(hpc) = \mathrm{Var}(hplf) + \mathrm{Var}(lfpr) + 2\,\mathrm{Cov}(hplf, lfpr)$$

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, using HP-filtered log data. This table shows the proportion of hours variation by age group that can be attributed to the participation margin. We focus our discussion on the proportion of cyclical variance of *hpc* owing to the participation margin, presented in panel B.[8] With covariance terms not

[6] These results corroborate the findings of Gomme et al. (2005) and Jaimovich and Siu (2009) and extend them to include data from the 2001 and 2008–2009 recessions. See also Clark and Summers (1981), Ríos-Rull (1996) and Nagypál (2007) 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 the Appendix, show essentially the same volatility pattern.

[7] 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 referenced in the macro-labor literature.

[8] Again, this is calculated as a projection on a constant, current detrended aggregate output, and current and lagged detrended aggregate hours. The interpretation of results presented in panel A on the filtered variance is essentially the same.

Table 2—Participation Margin's Share of Hours Variance

Age group	A. Filtered volatility		B. Cyclical volatility	
	Cov. not incl.	Cov. incl.	Cov. not incl.	Cov. incl.
15–19	20.29	28.87	16.59	24.91
20–24	10.03	16.71	6.02	10.75
25–29	8.13	13.98	2.95	5.57
30–39	6.88	12.10	2.73	5.19
40–49	5.72	10.27	0.61	1.20
50–59	10.06	16.74	2.88	5.45
60–64	32.52	39.41	6.75	11.90
15–29	11.14	18.21	7.51	13.06
30–64	6.25	11.12	2.76	5.23

Notes: Data from the March CPS, 1964–2010. Shown are percentage shares of total hours variation attributed to the participation margin. Total hours per age group member is the product of two variables: labor force participation per age group, and hours per labor force participant in that age group. "Cov. not incl." 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 that of labor force participation. "Cov. incl." means total variation includes covariance terms, so total variation is the sum of the variables' variances plus two times their covariance; the share attributed to the participation margin is the variance of labor force participation plus the covariance, divided by total variation. 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.

included, the participation margin explains less than one-fifth of the variation of any age group. For the 15–29-year-old age group as a whole, the participation margin accounts for only 8 percent of hours fluctuations. For teenagers this is higher at 17 percent; nonetheless, more than four-fifths of the variance of their hours worked is due to the hours margin. The bulk of all age groups' cyclical hours variation is due to variation in hours per labor force member.

The second column of panel B 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. Both columns of panel B give the same message (as indeed does panel A and its results for filtered volatilities). With the inclusion of covariance terms, the participation margin accounts for only 13 percent of hours variation for the young (15–29-year-old) age group, and only 5 percent for the old.

Hence, fluctuations in hours per labor force participant account for the bulk of hours variation for all age groups.[9] Consequently, it does not appear that explanations centered on differences in the cyclicality of participation are of first-order importance for generating greater volatility of young hours over the business cycle.

[9] Though 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.

B. *Wages*

From the March CPS, we use information on labor income and hours worked to construct annual series for hourly wages over the years 1963–2009.[10] 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 et al. (2000) and is detailed in the online Appendix.[11] We then HP-filter these series to isolate fluctuations at the business cycle frequency.

The first column in Table 3 reports the percent standard deviation of the HP-filtered hourly real wage rates by age. We see a decreasing pattern in volatility by age with an upturn in the 60–64 age group. Column 3 presents the percent standard deviation of the *cyclical* age-specific series. As in column 1, we see the familiar decreasing pattern of volatility by age, with a slight upturn at the end of the age distribution. We see that the real wage of 15–29-year-olds is about one and half times as volatile as that of the 30–64 age group.

To summarize, we identify two key findings regarding labor market differences between young and old workers. First, young workers experience hours worked volatility that is almost twice that of old workers. Second, real wage volatility of young workers is about one and a half times that of old workers. In the following sections, we explore the implications of these facts for real business cycle analysis. We refer to these two results as our *labor market facts*.

C. *Robustness Checks*

The online Appendix reports further results showing that both the age-specific wage and hours volatility patterns hold, after conditioning upon several demographic characteristics. We briefly comment on the results here.

We first investigate differences across educational attainment. Tables OA1 and OA2 represent the analogs of Tables 1 and 3 for those with less education (high school diploma and less), while Tables OA3 and OA4 report the results for those with more education (more than high school diploma). Within each education group, the young exhibit greater volatility of both hours and wages relative to the old.

Tables OA5 and OA6 represent the analogs of Tables 1 and 3 for males, while Tables OA7 and OA8 report the results for females. Within each gender group, the young exhibit greater volatility of both hours and wages relative to the old.

As such, our labor market facts are robust at these finer levels of aggregation. Moreover, this indicates that age (or equivalently, labor market experience) is not

[10] These data are taken from the March CPS questions pertaining to "last year." Hence, the surveys from 1964–2010 provide data for the years 1963–2009, which we of course take into account when deflating and constructing cyclical measures.

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

Table 3—Volatility of Real Hourly Wages by Age Group

Age group	Filtered volatility	R^2	Cyclical volatility
15–19	2.57	0.15	1.01
20–24	2.14	0.24	1.06
25–29	1.65	0.17	0.69
30–39	1.21	0.18	0.51
40–49	1.38	0.19	0.60
50–59	1.12	0.29	0.61
60–64	2.12	0.15	0.82
15–29	1.64	0.29	0.88
30–64	1.16	0.26	0.59

Notes: Data from the March CPS, 1964–2010. 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.

simply a proxy for other demographic characteristics in terms of hours and wage volatility.[12]

II. Labor Supply Channels

In this section, we demonstrate in a very general class of models that differences in labor supply characteristics alone cannot explain the two facts regarding age-specific labor market fluctuations documented above.

To focus attention on labor supply differences, throughout this section we assume that labor demand is symmetric across young and old. Specifically, consider an economy where production is summarized by an aggregate function $Y = F(A, K, H_Y, H_O)$. Here A denotes aggregate productivity/technology, K capital input, H_Y labor input of young workers, and H_O labor input for all other (i.e., old) workers. We assume that the marginal products with respect to capital and both labor inputs are positive and diminishing.

Throughout this article, we assume profit maximization and price taking by the representative firm. This implies that the wage rates for young and old labor, W_Y and W_O, respectively, are given by

$$W_Y = F_{H_Y}$$

$$W_O = F_{H_O}.$$

Denoting the log-linearized value of a variable with a circumflex, we get

$$\hat{W}_Y = \eta_{Y,A}\hat{A} + \eta_{Y,K}\hat{K} + \eta_{Y,Y}\hat{H}_Y + \eta_{Y,O}\hat{H}_O \tag{1}$$

$$\hat{W}_O = \eta_{O,A}\hat{A} + \eta_{O,K}\hat{K} + \eta_{O,Y}\hat{H}_Y + \eta_{O,O}\hat{H}_O. \tag{2}$$

[12] See also Gomme et al. (2005) who provide similar analysis with respect to hours along the dimensions of marital status and industry of employment.

Here $\eta_{x,z}$ is the elasticity of the marginal product of x with respect to z, $\eta_{x,z} = \frac{F_{x,z} \times z}{F_x}$; for the sake of exposition, we denote $\eta_{H_Y,z}$ as simply $\eta_{Y,z}$, and similarly for H_O.[13]

We define a production function, $F(\cdot)$, as being *symmetric in cyclical labor demand characteristics* if it has the following properties:

(i) $\eta_{Y,A} = \eta_{O,A}$

(ii) $\eta_{Y,K} = \eta_{O,K}$

(iii) $(\eta_{O,Y} - \eta_{Y,Y}) = (\eta_{Y,O} - \eta_{O,O}) = x$

(iv) $x \geq 0$.

The first condition says that the elasticity of the marginal product of both young and old labor input to a technology shock are equal. The second says the same thing with respect to capital. The third is a natural symmetry condition on the elasticities of the marginal products of labor. For instance, any production function that features constant returns to scale in K, H_Y, H_O and satisfies condition (ii) also trivially satisfies condition (iii).[14]

The final condition is a natural sign restriction. Given diminishing marginal products, $\eta_{Y,Y}$, $\eta_{O,O} < 0$, then young and old labor inputs being complements, $\eta_{O,Y}$, $\eta_{Y,O} > 0$, is a sufficient (but not necessary) condition for the sign restriction to be satisfied. In the case where labor inputs are substitutes, $\eta_{O,Y}$, $\eta_{Y,O} < 0$, condition (iv) states a natural requirement: that the marginal product of H_Y is more diminishing in H_Y than is the marginal product of H_O with respect to H_Y, and vice versa.[15]

In all of our propositions, we focus on results that deliver nonnegative comovement of hours and wages for young workers, as we find in the data.[16] All proofs are contained in the Appendix. Our first proposition does not require imposing any conditions on the characteristics of labor supply. Hence, Proposition 1 holds allowing for arbitrary differences in the cyclical properties of labor supply between young and old.

PROPOSITION 1: *If the production function satisfies conditions* (i)–(iv), *then for any specification of young and old labor supply, it is impossible for the response of young hours and young wages to a business cycle shock to be greater than for the old.*

[13] The characterization of log-linear responses is standard in the literature (where models are typically solved for log-linearized dynamics about steady state), and is informative regarding the cyclical properties of model for deviations of "business cycle" magnitude.

[14] See the lemma in the Appendix.

[15] Note that the production functions used in essentially all macroeconomic models satisfy these symmetry conditions. As examples, this includes the class of Cobb-Douglas functions, $Y = AK^{\alpha} H_Y^{\gamma} H_O^{\phi}$, and CES functions $Y = AK^{\alpha}[\gamma H_Y + (1 - \gamma)H_O]^{1-\phi}$, regardless of returns to scale.

[16] In the data, the correlation of cyclical hours and wages for the 15–29-year-old age group is 0.202. Note that this correlation is weaker than the correlation of age-specific hours to either aggregate hours or output. See also Bils (1989) for further evidence on the procyclicality of real wages.

The intuition for this result is straightforward. Consider the standard textbook treatment of the labor market, where labor supply (demand) is an upward (downward) sloping function of the wage. Procyclicality of hours and wages requires a business cycle shock that shifts the labor demand schedule. Symmetry in labor demand characteristics implies that the labor demand curves for young and old labor (i) have the same slope, and (ii) shift by the same amount in response to the cyclical shock.[17]

In this case, there is only one way to generate a greater response of young hours to the shock relative to old hours: the young labor supply curve must be more elastic (after both income and substitution effects are taken into account). That is, the young labor supply curve must be flatter. But this immediately implies that the response of the young wage must be smaller than the response of the old wage.

Therefore, it is not possible for a production function with symmetry in cyclical labor demand characteristics to yield a larger response of both young hours *and* wages, irrespective of how one specifies labor supply characteristics.[18] We view this as an instructive result. It is well known that standard RBC models embody weak endogenous propagation mechanisms (see, for instance, King and Rebelo 1999). As a result, the volatility properties of endogenous variables in such models are determined almost exclusively by their responsiveness to exogenous shocks. Proposition 1 indicates, for instance, that upon impact of a technology shock (i.e., a deviation in $\hat{A}$), it is impossible for $\hat{H}_Y > \hat{H}_O$ and $\hat{W}_Y > \hat{W}_O$. Moreover, the proposition indicates that this must also be true along any steady-state transition path—for instance, in response to $\hat{K}$ deviations induced by the dynamic response to a business cycle shock.

While instructive, Proposition 1 is not sufficient to characterize the variances of hours and wages. This can be seen in the proposition's proof: determining the relative magnitude of variances requires the signing of covariance terms as well. In general, this cannot be done without putting more structure on the characteristics of labor supply. Our second proposition, however, establishes a special case where this can be done.

PROPOSITION 2: *Suppose the production function satisfies conditions* (i)–(iv). *If the model is summarized by only one state variable*, *then for any specification of young and old labor supply*, *it is impossible to match the labor market facts.*

The key to Proposition 2 is that all covariances can be traced back to the variance of the single state variable, and other variables' response to that state variable.

In order to provide results for models with multiple state variables, we must make some assumptions regarding labor supply, i.e., the specification of the household side of the model. In what follows we assume that young and old workers live in

[17] It is precisely these restrictions that are removed when we move to the capital-experience complementarity model in Section III.

[18] Indeed, this is true even without symmetry conditions (3) and (4); all that is required is the weaker condition, that $(\eta_{O,Y} - \eta_{Y,Y}) \geq (\eta_{Y,O} - \eta_{O,O})$, and Proposition 1 holds.

perpetuity and belong to the same representative household.[19] The unified household construct allows us to restrict differences in the "wealth effect" on labor supply.

Our final result in this section is contained in Proposition 3. The key restriction is that the wealth effect on labor supply be equated across young and old agents.[20] Our proposition holds for any type of time-separable preferences used in the business cycle literature. These include the commonly used "balanced growth" preferences of King, Plosser, and Rebelo (1988) that feature separability between consumption and hours, as well as the nonseparable preferences of Greenwood, Hercowitz, and Huffman (1988) that feature "zero wealth effect" on labor supply.

PROPOSITION 3: *Let preferences for young and old workers be given by* $U(C_Y, H_Y)$ *and* $V(C_O, H_O)$, *respectively. Suppose* U, V *satisfy the usual regularity conditions* (*specifically*, U, V *decreasing and weakly convex in hours*) *and have identical wealth effects on labor supply. If the production function satisfies conditions* (i)–(iv), *then it is impossible to match the labor market facts.*

Hence, regardless of the age differences embodied in the utility functions U and V, it is impossible to simultaneously generate greater volatility of hours and wages of the young relative to the old when the wages and hours of the young positively covary; this is true when wealth effects are identical within the household. Thus, for a broad class of preferences, a model featuring symmetric labor demand characteristics cannot explain the labor market facts presented in Section I.

III. A Model with Age-Specific Labor Demand

Here, we present a model featuring age-specific differences in the characteristics of labor demand to rationalize the labor market facts presented in Section I. The remaining features of the model—in particular, household preferences—are specified to conform as closely as possible to the standard RBC model. This allows us to focus on the role of age differences in labor demand.

We view our model of capital-experience complementarity as speaking to complementarities 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 requires (or is embodied in) 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 young, inexperienced labor that is less tied

[19] Note that this can be viewed as a simple specification where, in every period, some young workers are "born," some young workers age and become old workers, and some old workers "die," in such a way as to maintain a constant share of young and old workers. This representation serves our purposes, since we are interested only in deriving workers' labor supply functions.

[20] In the online Appendix, we consider models that allow for *different* wealth effects on labor supply across agents. In such cases, we cannot analytically rule out the possibility of greater volatility of hours and wages for the young relative to the old. However, in a numerical exercise, we find that models allowing for such differences cannot come close, *quantitatively*, to reconciling our two labor market facts.

to these factors, while demand for old, experienced labor exhibits behavior that resembles labor hoarding.[21]

A. *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*, and let s_Y denote the share of family members that are young.[22] 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

$$E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left[s_Y U_Y(C_{Y\tau}, N_{Y\tau}) + (1 - s_Y) U_O(C_{O\tau}, N_{O\tau}) \right], \tag{3}$$

subject to

$$\begin{aligned} s_Y C_{Y\tau} + (1 - s_Y) C_{O\tau} + K_{\tau+1} = {} & (1 - \delta) K_\tau + r_\tau K_\tau + s_Y W_{Y\tau} N_{Y\tau} \\ & + (1 - s_Y) W_{O\tau} N_{O\tau}, \quad \forall \tau \geq t, \end{aligned} \tag{4}$$

with $0 < \beta < 1$, $0 \leq \delta \leq 1$. Here 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.

We specify the instantaneous utility functions 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 ψ_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 $0 \leq N_{Yt}, N_{Ot} \leq 1$.

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

$$C_{Yt} = C_{Ot} = C_t. \tag{5}$$

[21] Of course, the measurement of factors such as organizational capital or firm know-how are very difficult. 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. Finally, to the extent that information technology is embodied in physical capital, we note that one might alternatively consider an environment where capital is complementary to younger workers. We view this idea, coupled with business cycle dynamics of investment in a vintage capital model, as an interesting avenue for future research.

[22] Again, this can be viewed as a simplified framework in which young workers are born at a given rate x, young workers age and become old at rate x, and old workers die at rate x, so that the population shares of young and old workers remains constant.

The intertemporal first-order condition (FOC) is standard: $C_t^{-1} = \beta E_t\big[C_{t+1}^{-1}(r_{t+1} + 1 - \delta)\big]$. The FOCs for hours worked are given by $W_{Yt} = \psi_Y\, C_t\, N_{Yt}^{\theta_Y}$ for the young, and $W_{Ot} = \psi_O\, C_t\, N_{Ot}^{\theta_O}$ for the old. When $\theta_Y = \theta_O$, together with condition (5), the wealth and substitution effects on labor supply are equal for young and old workers.

B. *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 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.

We consider the following production function specification:

$$(6)\quad Y_t = \big[\mu(A_t H_{Yt})^{\sigma} + (1-\mu)\big[\lambda K_t^{\rho} + (1-\lambda)(A_t H_{Ot})^{\rho}\big]^{\sigma/\rho}\big]^{1/\sigma}, \qquad \sigma, \rho < 1.$$

Labor-augmenting technology follows a deterministic growth trend with stationary shocks: $A_t = \exp(gt + z_t)$ where $z_t = \phi z_{t-1} + \varepsilon_t$ and $\phi \in (0, 1)$, $E(\varepsilon) = 0$, $0 \leq \mathrm{Var}(\varepsilon) = \sigma_\varepsilon^2 < \infty$, and $g > 0$ is the deterministic growth rate of technology. Since technology augments both H_O and H_Y, and given the households' preferences, the economy exhibits balanced growth.

Profit maximization on the part of the firm entails equating factor prices with marginal revenue products. The FOCs are

$$(7)\qquad r_t = Y_t^{1-\sigma}(1-\mu)\Omega_t\, \lambda K_t^{\rho-1},$$

$$(8)\qquad W_{Ot} = Y_t^{1-\sigma}(1-\mu)\Omega_t(1-\lambda)A_t^{\rho}\, H_{Ot}^{\rho-1},$$

$$(9)\qquad W_{Yt} = Y_t^{1-\sigma}\, \mu A_t^{\sigma}\, H_{Yt}^{\sigma-1},$$

where $\Omega_t \equiv \big[\lambda K_t^{\rho} + (1-\lambda)(A_t H_{Ot})^{\rho}\big]^{(\sigma-\rho)/\rho}$.[23]

The degree of diminishing marginal product differs between young and old labor whenever $\sigma \neq \rho$. The elasticity of substitution between old workers and capital is given by $(1-\rho)^{-1}$, while the elasticity of substitution between young workers and the H_O–K composite is $(1-\sigma)^{-1}$. Adapting the terminology of Krusell et al. (2000), we define production as exhibiting *capital-experience complementarity* if $\sigma > \rho$, when we equate age with labor market experience.

To see how such a production technology can generate greater volatility of young labor input relative to the old, consider a simple example. Suppose that old labor is a perfect complement to capital (i.e., $\rho \to -\infty$), while young labor is not ($\sigma > \rho$).

[23] Krusell et al. (2000) consider a similar production specification in their analysis of trends in relative wages. They distinguish labor skill by the level of educational attainment (as opposed to experience in our model). Castro and Coen-Pirani (2008) use the same production function as Krusell et al. (2000) and consider changes in the elasticities of substitution to explain changes in the relative volatility of hours for high educated workers since 1984.

Since capital is in inelastic supply in the short run, a productivity shock generates no response in the quantity of old labor hired; the only variation is in the quantity of young labor.

C. *Equilibrium*

Equilibrium is defined as follows. Given $K_0 > 0$ and the stochastic process for technology, a *competitive equilibrium* is an allocation, $\{C_t, N_{Yt}, N_{Ot}, K_t, Y_t, H_{Yt}, H_{Ot}\}$, 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; the capital rental market clears for all t; and labor markets clear $(H_{Yt} = s_Y N_{Yt};\ H_{Ot} = (1 - s_Y) N_{Ot})$ for all t. Walras' law ensures clearing in the final goods market: $C_t + K_{t+1} = Y_t + (1 - \delta) K_t$, for all t. Finally, for the purposes of model evaluation, we define aggregate hours worked as $H_t = H_{Yt} + H_{Ot}$.

IV. Quantitative Specification

In this section, we describe the quantitative specification of our model. 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 US data. Instead, we adopt a structural, instrumental variables estimation procedure to identify these values using data from the CPS and Bureau of Economic Analysis (accessed via Haver Analytics). After describing the procedure, we discuss calibration of the remaining parameter values. Recalling the empirical results of Section I, we identify young and old workers in the model with 15–29- and 30–64-year-old age groups, respectively, in the data.[24]

A. *Estimation*

To estimate the elasticity parameters σ and ρ, consider the factor demand equations implied by our model.[25] The firm's FOC with respect to H_{Yt} (9) can be logged and first-differenced into

$$(10)\quad \Delta \log W_{Yt} = \alpha_0 + (\sigma - 1)\Delta \log (H_{Yt}) + (1 - \sigma)\Delta \log (Y_t) + \sigma u_t,$$

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

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

Hence, equation (10) represents a textbook labor demand equation.

[24] This is clearly an extreme specification, in that the distinction between young and old labor input is determined by a single age threshold. However, we have found that our estimation results are robust to five-year variation in this cut-off age.

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

Therefore our strategy for estimating the elasticity parameter σ amounts to estimating the 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 in the aggregate data. Abstracting from endogeneity issues (which we address below), σ could be estimated from a simple regression.

The age-specific wages analyzed in Section I are constructed using hours data in order to translate direct information on labor income in the CPS into measured wages. It is possible that there is error in our measurement of hours, in spite of the aggregation across individuals within each age group. This would contaminate our measurement of wages and induce unnecessary imprecision into our IV estimates. A simple fix is to estimate a variant of (10) relating hours directly to labor income:

$$\Delta \log LI_{Yt} - \Delta \log Y_t = \alpha_1 + \sigma(\Delta \log H_{Yt} - \Delta \log Y_t) + \sigma u_t, \tag{11}$$

where $LI_{Yt} \equiv W_{Yt} H_{Yt}$ denotes labor income earned by young workers. Again, direct measures of LI_{Yt} and H_{Yt} can be obtained from the CPS, while Y_t is available from the NIPA.

To estimate ρ we proceed in a similar manner. The firm's FOCs with respect to K_t and H_{Ot}, (7) and (8), 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:

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

Again, we can avoid unnecessary imprecision by estimating the following version with instrumental variables:

$$\Delta \log Q_{Ot} - \Delta \log Q_{Kt} = \alpha_2 + \rho(\Delta \log H_{Ot} - \Delta \log K_t) + \rho u_t. \tag{12}$$

Here, 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.

Importantly, our procedure does not require imposing any restrictions from the model's specification of household behavior.[26] 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, (11) and (12), 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 or not this is satisfied depends on the relation between aggregate prices and quantities observed in the data.

Endogeneneity.—Since our empirical equations are based on factor demand equations we must address the potential endogeneity of the regressors. The structural

[26] We see this as a virtue since our goal is to study the quantitative role of differences in the cyclical demand for young and old labor.

equations identify the error term as due to shocks to technology. To obtain unbiased estimates, we isolate variation in our regressors that is unrelated to shocks shifting firms' factor demand, be they technology shocks or other omitted factors from the FOCs.

We do so by adopting an instrumental variables approach using lagged birth rates. Intuitively, these instruments allow us to identify changes in current labor supply—due to changes in past fertility—that are uncorrelated to shifts in factor demand.[27] Recall that

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

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.[28]

Results.—Our theory suggests that the error terms u_t are correlated. Therefore, joint estimation is efficient. We use a system approach to estimate equations (11) and (12) by two-step GMM following Hansen (1982). Heteroskedasticity and autocorrelation robust standard errors are estimated following Andrews (1991) with optimal bandwidth chosen according to Newey and West (1994).[29] This ensures that the standard errors we report and use for hypothesis tests account for any serial correlation or heteroskedasticity that is present in u_t.

We must use instruments to consistently estimate σ and ρ and must ensure our instruments are not only valid but also *relevant* and *robust* to weak instrument issues. Hence, we test whether any possible instrument weakness leads to biased estimates or distorted hypothesis tests. Using Stock and Yogo's (2005) critical values for the first-stage F-statistic, we reject at the 5 percent level the hypothesis that weak instruments lead to bias in our estimates or distortions of our hypothesis tests' size.[30] In other words, we can reject that our instruments' weakness distorts the point estimates or hypothesis tests we report.

Table 4 presents the results of the estimation. The point estimate of $\rho = 0.201$ indicates that the elasticity of substitution is a little more than unity in K and H_O; in contrast, the estimate of $\sigma = 0.662$ indicates that the substitution elasticity is substantially larger between H_Y and the K–H_O composite. Conducting the F-test of the hypothesis that $\sigma = \rho$, we obtain an extremely low p-value, suggesting that

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

[28] See the online Appendix for further detail.

[29] In the online Appendix we show that our estimates are robust to alternative GMM specifications.

[30] In particular, we test the hypotheses that point estimates are biased by larger than 0.1 of the true value, and that hypothesis tests have size distortions in excess of 0.2. Note that in our framework with one endogenous regressor, the first-stage F-statistic is Cragg and Donald's (1993) statistic.

TABLE 4—ESTIMATION RESULTS

	Point estimate	Standard error	First-stage F-statistic	J-test p-value	$\sigma = \rho$ p-value
σ	0.662	0.048	10.829*	0.425	$<$ 0.001
ρ	0.201	0.016	13.89*		

Notes: Data from the March CPS, 1964–2010. Estimation is two-step GMM with HAC standard errors with bandwidth chosen by Newey and West (1994), using lagged birth rates as instruments. J denote's Hansen's J-test; $\sigma = \rho$ denotes the F-test of the null that the two parameters are equal. * indicates significance at the 5 percent level for Stock and Yogo's (2005) TSLS weak instrument tests based on 0.2 maximal size distortion or 0.1 maximal bias.

the difference between σ and ρ is statistically significant at the 0.1 percent level. Moreover, the difference is in the "right" direction for the interpretation of capital-experience complementarity ($\sigma > \rho$).[31] In summary, our results demonstrate strong instruments, precisely estimated parameters, robustness across a variety of specifications, and a statistically significant difference in the elasticity of substitution between young or prime-aged hours and capital.

B. *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 US 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 others. As such, we consider various values to illustrate the quantitative properties of our models.

Following Krusell et al. (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, we set μ and λ to match the 1964–2010 national income shares of $Q_K = 0.37$ and $Q_O = 0.50$.

With values for $\{\hat{\sigma}, \hat{\rho}, \mu, \lambda\}$ we back out the implied technology series $\{A_t\}$ using data on output and factor inputs. From $\{A_t\}$, we obtain a quarterly estimate of $\hat{\phi} = 0.94$ and $\hat{\sigma}_\varepsilon = 0.0064$.

V. Quantitative Evaluation

In this section, we evaluate the quantitative predictions of the capital-experience complementarity model. Specifically, we study the performance of the model with

[31] Compared to Krusell et al. (2000), who differentiate on skilled/unskilled versus experienced/inexperienced labor, our estimated σ is similar (0.66 versus their 0.40), but our ρ is different (0.2 versus their -0.5). This means that Krusell et al. (2000) find that capital and skilled labor are more complementary than we find capital and experienced labor to be. This is an interesting distinction for future research to explore.

respect to our two new empirical facts—the volatility of hours and wages of the young relative to the old.

Column 1 in Table 5 presents business cycle statistics for HP-filtered US data. As discussed in Section I, the volatility of young hours worked is greater than that of old hours worked, with a relative standard deviation of 1.85. As before, the young are defined as those aged 15 to 29, and the old as those 30–64 years old. Relative to aggregate output, young hours exhibits greater cyclical volatility, while old hours exhibits somewhat lower volatility. The volatility of aggregate hours is of a similar magnitude to that of aggregate output; in fact, in our data, which include the onset of the Great Recession, the ratio of standard deviations of aggregate hours to output exceeds unity.

The remaining rows in column 1 report volatility statistics for real wages for the two age groups. As noted in Section IB, the volatility of wages is also greater for the young than for the old. For our two age groups, the ratio of real wage volatility is 1.50.[32] Since our model embodies a standard intertemporal 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.

We begin with examination of the capital-experience complementarity model where we 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 $0.7 - 0.75$.[33]

As column 2 of Table 5 reports, the capital-experience complementarity model generates significant volatility of young hours relative to old hours. The ratio of standard deviations is 2.07, which is 12 percent greater than that observed in the US data. The model also has no difficulty in generating young hours that are more volatile than output, and old hours that are less volatile. For both statistics, the model slightly understates these relative volatilities: the standard deviation of young hours to output in the model is 1.60 compared to 1.64 in the data, and the standard deviation of old hours to output in the model is 0.77 compared to 0.89 in the data.[34]

As a byproduct of this success, the model also generates significant volatility of aggregate hours. In fact, the relative volatility of aggregate hours to output is very close to that observed in the data. The relative standard deviation is 1.04 in the model, whereas it is 1.10 in the data. In this sense, the capital-experience complementarity model represents a potential resolution to the RBC literature's inability to generate sufficient hours worked volatility.

Finally, we note that the benchmark model generates significant volatility of aggregate output. In our model, the standard deviation of output is 1.32 percent,

[32] We report the volatility for cyclical fluctuations in hours and real wages, as constructed in Section I. Given the focus on business cycle fluctuations in hours and wages, we concentrate on the variation that is due to the cycle.

[33] See, for example, Hansen (1985), Rogerson (1988), King and Rebelo (1999).

[34] Our quantitative specification has an elasticity of substitution between capital and old hours that is close to unity ($(1 - \rho)^{-1} = 1.25$), 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 above. Thus the capital-experience complementarity model generates a relative volatility of old hours to output, $\text{sd}(H_O)/\text{sd}(Y)$, similar to the relative volatility of aggregate hours to output in the standard RBC model.

TABLE 5—DATA AND MODEL MOMENTS

	Data	Model		
	(1)	(2)	(3)	(4)
$\mathrm{sd}(H_Y)/\mathrm{sd}(H_O)$	1.85	2.07	1.85	1.49
$\mathrm{sd}(H_Y)/\mathrm{sd}(Y)$	1.64	1.60	1.46	1.20
$\mathrm{sd}(H_O)/\mathrm{sd}(Y)$	0.89	0.77	0.79	0.81
$\mathrm{sd}(H)/\mathrm{sd}(Y)$	1.10	1.04	1.01	0.94
$\mathrm{sd}(W_Y)/\mathrm{sd}(W_O)$	1.50	1.00	1.17	1.50
$\mathrm{sd}(W_Y)/\mathrm{sd}(Y)$	0.47	0.26	0.31	0.41
$\mathrm{sd}(W_O)/\mathrm{sd}(Y)$	0.31	0.26	0.27	0.27
$\mathrm{sd}(Y)$	1.56	1.32	1.28	1.21
$\mathrm{sd}(Y)/\mathrm{sd}(z)$	—	1.58	1.53	1.45
θ_Y	—	0	0.04	0.14
θ_O	—	0	0	0
Target	—	—	$\frac{sd(H_Y)}{sd(H_O)}$	$\frac{sd(W_Y)}{sd(W_O)}$

Notes: Column 1 is sample moments calculated from HP filtered data from March CPS, 1964–2010. Columns 2–4 are sample moments calculated from model-simulated data. The row Target indicates what moment is targeted by the Frisch elasticity parameter.

or about 84 percent of that observed in the US data. This is a marked improvement over the indivisible-labor version of the standard RBC model and compares favorably with models that allow for variability in capital utilization (see Prescott 1986; Burnside and Eichenbaum 1996). Relaxing the assumption of unit elasticity of substitution in factor supplies allows the capital-experience complementarity model to generate significant endogenous amplification of productivity shocks. The relative volatility of output to the shock process is around 1.58, which is substantially larger than in the standard RBC model where this relative volatility is typically near unity (see Burnside and Eichenbaum 1996).

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, and the wealth effect is identical for both young and old agents. In this case, the volatility of wages is necessarily identical. This is addressed in the next two experiments.

As already discussed, the benchmark calibration (with $\theta_Y = \theta_O = 0$) overstates the volatility of young hours relative to old hours. In column 2, we consider the following modification: we change only the labor supply elasticity of young workers to match the relative hours volatility observed in the US data. This requires increasing θ_Y from 0 to 0.04, so that the Frisch labor supply elasticity of young workers is less than that of old workers.[35] Not surprisingly, this lowers the volatility of young hours to aggregate output.

The modification also lowers the volatility of aggregate hours to output; however, the model still delivers a relative volatility that is near unity (1.01) and close to

[35] We find this to be an interesting experiment since it makes clear that our results in no way rely upon assuming that young workers have greater labor supply elasticity (an assumption that, to our knowledge, has no empirical support in the literature).

that observed in the data (1.10). Lowering the labor supply elasticity of the young also lowers the volatility of output (marginally, from 1.32 to 1.28). However, the model still embodies significant amplification, as the standard deviation of output is 53 percent greater than that of the exogenous shock process. Finally, we note that this modification also allows the model to match the fact that both young workers' hours and their wages are more volatile than for old workers.

In column 3 we consider the following modification: we change only the labor supply elasticity of the young to match the observed relative wage volatility ($\text{sd}(W_Y)/\text{sd}(W_O) = 1.50$). This requires increasing θ_Y to 0.14. Moreover, this modification does very well at matching the volatility of both young and old wages, relative to aggregate output, found in the US data. Not surprisingly, the lower elasticity of young labor supply induces a fall in the volatility of young hours. The relative volatility of age-specific hours ($\text{sd}(H_Y)/\text{sd}(H_O) = 1.49$) now understates that found in the data. Finally, we note that the model still embodies significant amplification, with $\text{sd}(Y)/\text{sd}(z) = 1.45$ well above unity.

In sum, we find that the capital-experience complementarity model easily captures the fact that both hours and wages of young workers are more volatile over the business cycle than for old workers.[36] That is, modeling differences in the cyclical characteristics of labor demand quantitatively accounts for our labor market facts. As a byproduct of this success, the model generates volatility of aggregate hours that is very close to that of aggregate output. Finally, the model embodies significant internal amplification of business cycle shocks.

VI. Conclusion

We highlight two important empirical observations regarding age differences in labor market fluctuations. First, hours worked of young workers are more volatile over the business cycle than their prime-aged counterparts. Second, real wages of the young are more volatile over the business cycle than the prime-aged.

We show that a general class of models allowing for age differences in labor supply characteristics alone cannot account for these facts. Instead, a model emphasizing age differences in labor demand factors can. Our model posits a greater complementarity of prime-aged workers' labor input with capital in production than for young workers.

Our model of capital-experience complementarity represents a minor deviation from the standard RBC model and delivers factor demand equations that can be used to estimate structural elasticity parameters. We find that the data is consistent with capital-experience complementarity.

Quantitative evaluation of the model shows that it can easily reconcile the labor market facts. Moreover, the model obtains aggregate hours that have equal volatility to aggregate output and embodies strong internal amplification of business cycle shocks. Altogether, the article points to the importance of characterization of

[36] Because the model is driven by a single productivity shock, the model trivially generates a positive correlation between hours and wages, though this is counterfactually close to unity. As in the literature, this could easily be remedied by the inclusion of other shocks that would affect the correlation of hours and wages without affecting the model's relative volatility properties (see, for instance, Christiano and Eichenbaum 1992, and Benhabib, Rogerson, and Wright 1991).

age-specific differences in the demand for labor inputs in understanding business cycle fluctuations.

Appendix

LEMMA: *Condition* (ii) *implies condition* (iii) *if* $F(\cdot)$ *is constant returns to scale.*

PROOF:

Since $F(\cdot)$ is homogenous of degree one in K, H_Y, H_O,

$$0 = F_{Y,K}K + F_{Y,O}H_O + F_{Y,Y}H_Y,$$

$$\frac{-F_{Y,K}K}{F_Y} = \frac{F_{Y,O}H_O + F_{Y,Y}H_Y}{F_Y},$$

$$-\eta_{Y,K} = \eta_{Y,O} + \eta_{Y,Y}.$$

Similarly, from the marginal product of old labor,

$$-\eta_{O,K} = \eta_{O,Y} + \eta_{Y,Y}.$$

Condition (ii) then implies

$$\eta_{O,O} + \eta_{O,Y} = \eta_{Y,O} + \eta_{Y,Y},$$

$$\eta_{O,Y} - \eta_{Y,Y} = \eta_{Y,O} - \eta_{O,O},$$

which is condition (iii).

PROOF OF PROPOSITION 1:

Take the difference of equations (1) and (2). Using symmetry conditions (i) and (ii):

$$\hat{W}_Y - \hat{W}_O = (\eta_{Y,Y} - \eta_{O,Y})\hat{H}_Y - (\eta_{O,O} - \eta_{Y,O})\hat{H}_O, \tag{A1}$$

$$= x(\hat{H}_O - \hat{H}_Y),$$

where the last equality follows from condition (iii). Now impose condition (iv). If $x = 0$, then the variance of young wages is identical to that of prime-age wages; this violates our requirement of a differential response of wages, irrespective of the response of hours. Now consider the case where $x > 0$. Suppose that $\hat{W}_Y - \hat{W}_O > 0$, so that the response of young wages is greater than that of old wages. Then this implies $\hat{H}_O > \hat{H}_Y$. Given our interest in $\hat{W}_Y$ and $\hat{H}_Y$ responses that are of the same sign, this violates our requirement that the response of young hours is greater than that of old hours.

PROOF OF PROPOSITION 2:

Assume there is only one state variable, S (whether it be technology, A, capital, K, or anything else). Then the model's state space representation implies that we can express the equilibrium relationships:

$$\hat{H}_Y = B_{Y,S}\,\hat{S},$$

$$\hat{H}_O = B_{O,S}\,\hat{S}.$$

Thus $\mathrm{Var}(\hat{H}_Y) > \mathrm{Var}(\hat{H}_O)$ if and only if $B_{Y,S} > B_{O,S}$.

Given conditions (i)–(iv) characterizing symmetry in labor demand, equations (1) and (2) can be rewritten as

$$\hat{W}_Y = \eta_{YS}\hat{S} + \eta_{YY}B_{Y,S}\hat{S} + \eta_{YO}B_{O,S}\hat{S} = [\eta_{YS} + \eta_{YY}B_{Y,S} + \eta_{OO}B_{O,S} + xB_{O,S}]\,\hat{S},$$

$$\hat{W}_O = \eta_{YS}\hat{S} + \eta_{OY}B_{Y,S}\hat{S} + \eta_{OO}B_{O,S}\hat{S} = [\eta_{YS} + \eta_{YY}B_{Y,S} + \eta_{OO}B_{O,S} + xB_{Y,S}]\,\hat{S}.$$

Then $\mathrm{Var}(\hat{W}_Y) > \mathrm{Var}(\hat{W}_O)$ if and only if $B_{O,S} > B_{Y,S}$. This contradicts the necessary and sufficient condition for $\mathrm{Var}(\hat{H}_Y) > \mathrm{Var}(\hat{H}_O)$.

PROOF OF PROPOSITION 3:

The FOC with respect to hours for young workers is

$$U_{H_Y}(C_Y, H_Y) = -\Lambda W_Y,$$

where Λ is the Lagrange multiplier on the household's budget constraint. Log-linearizing the FOC obtains

$$U_{H_Y C_Y} C_Y \hat{C}_Y + U_{H_Y H_Y} H_Y \hat{H}_Y = -\Lambda W_Y(\hat{\Lambda} + \hat{W}_Y),$$

which can be rewritten using $\left(\frac{U_{H_Y H_Y} H_Y}{U_{H_Y}}\right) \equiv \mathcal{U}$ as

$$\text{(A2)} \qquad \hat{W}_Y = \left[\frac{U_{H_Y C_Y} C_Y}{U_{H_Y}}\,\hat{C}_Y - \hat{\Lambda}\right] + \mathcal{U}\hat{H}_Y.$$

Thinking of (A2) in terms of a labor supply function in $H_Y - W_Y$ space, the first term on the right-hand side represents "shifts of," while the second term represents "movements along," the labor supply curve. Thus, the first term is the wealth effect, while the second term is the substitution effect. Note that $\mathcal{U} \geq 0$, given that $U_{H_Y} < 0$ and $U_{H_Y H_Y} \leq 0$. For old workers with utility function $V(C_O, H_O)$, we use $\left(\frac{V_{H_O H_O} H_O}{V_{H_O}}\right)$ $\equiv \mathcal{V}$ and derive analogously

$$\text{(A3)} \qquad \hat{W}_O = \left[\frac{V_{H_O C_O} C_O}{V_{H_O}}\,\hat{C}_O - \hat{\Lambda}\right] + \mathcal{V}\hat{H}_O,$$

where $\mathcal{V} \geq 0$, given $V_{H_O} < 0$ and $V_{H_O H_O} \leq 0$.

Restricting the wealth effect on labor supply to be identical across agents implies that the terms in square brackets in (A2) and (A3) are equal. Using the symmetry condition (A1) implies that

$$x(\hat{H}_O - \hat{H}_Y) = \mathcal{U}\hat{H}_Y - V\hat{H}_O,$$

which can be rewritten as

$$(x + \mathcal{V})\hat{H}_O = (x + \mathcal{U})\hat{H}_Y. \tag{A4}$$

Thus, for the young to have more volatile hours requires $\mathcal{V} > \mathcal{U} \geq 0$.

Using (A4), we can alternatively express (A1) as

$$\hat{W}_Y - \hat{W}_O = x\left(\frac{x + \mathcal{U}}{x + \mathcal{V}} - 1\right)\hat{H}_Y.$$

Rearranging obtains

$$\hat{W}_O = \hat{W}_Y + r\hat{H}_Y, \tag{A5}$$

where $r \equiv \frac{x(\mathcal{V} - \mathcal{U})}{x + \mathcal{V}}$. If $x = 0$, then (A5) implies that wages have equal volatility, immediately contradicting our labor market facts. Hence, it must be that $x > 0$. This implies that $r > 0$, given that we require $\mathcal{V} > \mathcal{U}$. Taking variances we get

$$\mathrm{Var}(\hat{W}_O) = \mathrm{Var}(\hat{W}_Y) + r^2\,\mathrm{Var}(\hat{H}_Y) + 2r\,\mathrm{Cov}(\hat{W}_Y, \hat{H}_Y).$$

Since the second term on the right-hand side is positive and the third term is nonnegative, this implies that $\mathrm{Var}(\hat{W}_O) > \mathrm{Var}(\hat{W}_Y)$, a contradiction of our labor market facts. Similarly, to get $\mathrm{Var}(\hat{W}_Y) > \mathrm{Var}(\hat{W}_O)$ we need r to be negative, or $\mathcal{U} > \mathcal{V}$; but then (A4) implies that $\mathrm{Var}(\hat{H}_O) > \mathrm{Var}(\hat{H}_Y)$, again contradicting our labor market facts.

REFERENCES

Andrews, Donald W. K. 1991. “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation.” *Econometrica* 59 (3): 817–58.

Baxter, Marianne, and Robert G. King. 1999. “Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series.” *Review of Economics and Statistics* 81 (4): 575–93.

Beaudry, Paul, and David A. Green. 2003. “Wages and Employment in the United States and Germany: What Explains the Differences?” *American Economic Review* 93 (3): 573–602.

Benhabib, Jess, Richard Rogerson, and Randall Wright. 1991. “Homework in Macroeconomics: Household Production and Aggregate Fluctuations.” *Journal of Political Economy* 99 (6): 1166–87.

Bils, Mark. 1989. “Pricing in a Customer Market.” *Quarterly Journal of Economics* 104 (4): 699–718.

Bureau of Labor Statistics and US Census Bureau. 1964–2011. “Current Population Survey.” US Department of Commerce. http://www.census.gov/cps/ (accessed Feburary 9, 2012).

Burnside, Craig, and Martin Eichenbaum. 1996. “Factor-Hoarding and the Propagation of Business-Cycle Shocks.” *American Economic Review* 86 (5): 1154–74.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 1995. “Capital Utilization and Returns to Scale.” In *National Bureau of Economic Research Macroeconomics Annual 1995*. Vol. 10, edited by Ben S. Bernanke and Julio J. Rotemberg, 67–110. Cambridge, MA: MIT Press.

Castro, Rui, and Daniele Coen-Pirani. 2008. "Why Have Aggregate Skilled Hours Become So Cyclical since the Mid-1980s?" *International Economic Review* 49 (1): 135–84.
Christiano, Lawrence J., and Martin Eichenbaum. 1992. "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations." *American Economic Review* 82 (3): 430–50.
Clark, Kim B., and Lawrence H. Summers. 1981. "Demographic Differences in Cyclical Employment Variation." *Journal of Human Resources* 16 (1): 61–79.
Cragg, John G., and Stephen G. Donald. 1993. "Testing Identifiability and Specification in Instrumental Variable Models." *Econometric Theory* 9 (2): 222–40.
Gomme, Paul, Richard Rogerson, Peter Rupert, and Randall Wright. 2005. "The Business Cycle and the Life Cycle." In *National Bureau of Economic Research Macroeconomics Annual 2004*. Vol. 19, edited by Mark Gertler and Kenneth Rogoff, 415–592. Cambridge, MA: MIT Press.
Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman. 1988. "Investment, Capacity Utilization, and the Real Business Cycle." *American Economic Review* 78 (3): 402–17.
Hansen, Gary D. 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics* 16 (3): 309–27.
Hansen, Gary D., and Selahattin Imrohoroglu. 2009. "Business Cycle Fluctuations and the Life Cycle: How Important Is On-the-Job Skill Accumulation?" *Journal of Economic Theory* 144 (6): 2293–309.
Hansen, Lars Peter. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica* 50 (4): 1029–54.
Haver Analytics. 1962–2011. "U.S. Economic Statistics (USECON)." http://www.haver.com (accessed October 1, 2012).
Jaimovich, Nir, Seth Pruitt, and Henry E. Siu. 2013. "The Demand for Youth: Explaining Age Differences in the Volatility of Hours: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.103.7.3022.
Jaimovich, Nir, and Henry E. Siu. 2009. "The Young, the Old, and the Restless: Demographics and Business Cycle Volatility." *American Economic Review* 99 (3): 804–26.
Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107 (1): 35–78.
King, Robert G., and Sergio T. Rebelo. 1999. "Resuscitating Real Business Cycles." In *Handbook of Macroeconomics*. Vol. 1, edited by John B. Taylor and Michael Woodford, 927–1007. Amsterdam: Elsevier, North-Holland.
King, Robert G., Charles I. Plosser, and Sergio T. Rebelo. 1988. "Production, Growth and Business Cycles: I. The Basic Neoclassical Model." *Journal of Monetary Economics* 21 (2/3): 195–232.
Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante. 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica* 68 (5): 1029–53.
McDonald, James T., and Christopher Worswick. 1999. "Wages, Implicit Contracts, and the Business Cycle: Evidence from Canadian Micro Data." *Journal of Political Economy* 107 (4): 884–92.
Nagypál, Éva. 2007. "Learning by Doing vs. Learning about Match Quality: Can We Tell Them Apart?" *Review of Economic Studies* 74 (2): 537–66.
Newey, Whitney K., and Kenneth D. West. 1994. "Automatic Lag Selection in Covariance Matrix Estimation." *Review of Economic Studies* 61 (4): 631–53.
Prescott, Edward C. 1986. "Theory Ahead of Business-Cycle Measurement." *Carnegie-Rochester Conference Series on Public Policy* 25 (1): 11–44.
Ravn, Morten O., and Harald Uhlig. 2002. "On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations." *Review of Economics and Statistics* 84 (2): 371–76.
Ríos-Rull, José-Víctor. 1996. "Life-Cycle Economies and Aggregate Fluctuations." *Review of Economic Studies* 63 (3): 465–89.
Rogerson, Richard. 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics* 21 (1): 3–16.
Stock, James H., and Motohiro Yogo. 2005. "Testing for Weak Instruments in Linear IV Regression." In *Identification and Inference for Econometric Models: Essays in Honor of Thomas J. Rothenberg*, edited by James H. Stock and Donald W. K. Andrews. New York: Cambridge University Press.