Online Appendix – Not for Publication

for "The Demand for Youth: Explaining Age Differences in the Volatility of Hours"

by Nir Jaimovich, Seth Pruitt and Henry E. Siu

A. Hours and Wage Volatility in Further Breakdowns

Table OA1—Volatility of Hours Worked by Age Group, High School Degree and Below

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60 - 64	**15 - 29**	**30 - 64**
filtered volatility	6.07	3.02	2.84	2.09	1.78	2.02	2.90	**3.27**	**1.74**
R^2	0.82	0.74	0.69	0.72	0.62	0.64	0.18	**0.87**	**0.80**
cyclical volatility	5.48	2.60	2.36	1.78	1.40	1.61	1.23	**3.05**	**1.55**
hours share	5.90	11.18	11.62	23.54	23.56	18.89	5.30	**28.71**	**71.29**
volatility share	16.13	14.51	13.65	20.89	16.39	15.17	3.26	**44.11**	**55.89**

Note: 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 the average of age-specific cyclical volatilities weighted by hours shares, reported in percentage. For individuals whose highest educational attainment is a high school degree.

Table OA2—Volatility of Real Hourly Wages by Age Group, High School Degree and Below

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60 - 64	**15 - 29**	**30 - 64**
filtered volatility	2.78	2.24	1.88	1.48	1.38	1.04	2.52	**1.71**	**1.15**
R^2	0.15	0.14	0.03	0.16	0.12	0.19	0.03	**0.11**	**0.18**
cyclical volatility	1.09	0.84	0.33	0.59	0.47	0.45	0.42	**0.56**	**0.48**

Note: 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. For individuals whose highest educational attainment is a high school degree.

Table OA3—Volatility of Hours Worked by Age Group, More than High School

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60 - 64	**15 - 29**	**30 - 64**
filtered volatility	6.38	2.50	1.92	1.53	1.43	1.74	3.34	**1.77**	**1.06**
R^2	0.28	0.38	0.22	0.50	0.21	0.22	0.03	**0.50**	**0.51**
cyclical volatility	3.41	1.54	0.91	1.08	0.65	0.81	0.59	**1.26**	**0.75**
hours share	1.34	10.49	14.78	28.73	24.82	16.13	3.71	**26.61**	**73.39**
volatility share	4.74	16.67	13.91	32.09	16.74	13.60	2.26	**37.72**	**62.28**

Note: 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 the average of age-specific cyclical volatilities weighted by hours shares, reported in percentage. For individuals whose highest educational attainment is more than a high school degree.

Table OA4—Volatility of Real Hourly Wages by Age Group, More than High School

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60 - 64	**15 - 29**	**30 - 64**
filtered volatility	4.02	2.92	1.85	1.25	1.74	1.52	2.60	**2.03**	**1.28**
R^2	0.08	0.20	0.24	0.19	0.16	0.13	0.16	**0.36**	**0.32**
cyclical volatility	1.17	1.29	0.92	0.55	0.69	0.55	1.04	**1.21**	**0.72**

Note: 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. For individuals whose highest educational attainment is more than a high school degree.

Table OA5—Volatility of Hours Worked by Age Group, Males

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60 - 64	**15 - 29**	**30 - 64**
filtered volatility	6.34	2.95	2.12	1.63	1.41	1.68	2.75	**2.69**	**1.40**
R^2	0.81	0.69	0.69	0.91	0.72	0.76	0.14	**0.85**	**0.96**
cyclical volatility	5.70	2.46	1.76	1.55	1.20	1.46	1.04	**2.47**	**1.37**
hours share	3.40	10.01	13.26	26.85	24.12	17.65	4.71	**26.67**	**73.33**
volatility share	11.50	14.61	13.85	24.75	17.13	15.26	2.91	**39.64**	**60.36**

Note: 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 the average of age-specific cyclical volatilities weighted by hours shares, reported in percentage. For males.

Table OA6—Volatility of Real Hourly Wages by Age Group, Males

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60 - 64	**15 - 29**	**30 - 64**
filtered volatility	3.33	2.27	1.84	1.39	1.67	1.36	2.70	**1.72**	**1.36**
R^2	0.14	0.22	0.14	0.20	0.14	0.36	0.14	**0.23**	**0.26**
cyclical volatility	1.23	1.07	0.69	0.62	0.63	0.82	1.00	**0.83**	**0.69**

Note: 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. For males.

Table OA7—Volatility of Hours Worked by Age Group, Females

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60 - 64	**15 - 29**	**30 - 64**
filtered volatility	5.44	2.09	2.06	1.40	1.27	1.46	2.44	**2.16**	**1.09**
R^2	0.64	0.59	0.45	0.72	0.46	0.51	0.04	**0.77**	**0.73**
cyclical volatility	4.35	1.60	1.38	1.19	0.86	1.04	0.51	**1.90**	**0.93**
hours share	4.26	12.16	12.91	24.69	24.22	17.48	4.29	**29.32**	**70.68**
volatility share	14.68	15.39	14.06	23.24	16.45	14.45	1.74	**45.75**	**54.25**

Note: 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 the average of age-specific cyclical volatilities weighted by hours shares, reported in percentage. For females.

Table OA8—Volatility of Real Hourly Wages by Age Group, Females

	15 - 19	20 - 24	25 - 29	30 - 39	40 - 49	50 - 59	60 - 64	**15 - 29**	**30 - 64**
filtered volatility	3.17	2.47	1.79	1.63	1.72	1.46	3.04	**1.83**	**1.36**
R^2	0.06	0.21	0.19	0.19	0.23	0.09	0.14	**0.29**	**0.32**
cyclical volatility	0.75	1.13	0.78	0.71	0.83	0.45	1.12	**0.98**	**0.77**

Note: 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. For females.

B. Different Wealth Effects

Qualitatively, the twin facts are that young hours and wages are more cyclically volatile than old hours and wages; *quantitatively*, the differences are large as the standard deviation of young hours is 1.9 times that of prime-aged hours, and the standard deviation of young wages is 1.5 times that of than prime-aged wages.

The paper proves that differences in the young and old labor supply characteristics alone cannot deliver the twin facts, when wealth effects for the young and old are the same. In this section we show how when we allow for the presence of different wealth effects across young and old then we cannot rule out its ability to match the twin facts. Specifically, we begin with an example that shows how the twin facts can qualitatively be delivered *only if* the workers with a "stronger" wealth effect also have a "weaker" substitution effect. We then discuss the basic economic forces behind this result via simple graphical analysis. We conclude by confirming this intuition via numerical results and show that even in cases where the model can qualitatively match the twin facts, it cannot do it quantitatively.

Example

Let young and old have different wealth effects. In this case there is a necessary condition such that we cannot rule out the possibility of accounting for the twin facts.

Specifically, suppose preferences for the young are separable as in King, Plosser and Rebelo (1988) (KPR)

$$U_Y = \log(C_Y) - \psi_Y N_Y^{1+\theta y}$$

while the preferences for the old are as in Greenwood, Hercowitz and Huffman (1988) (GHH)

$$U_O = \log(C_O - \psi_O N_O^{1+\theta_O}).$$

Note that in this case the young have a stronger wealth effect then the old (who have zero wealth effect).

The representative household's problem is given in equations (3) and (4), leading to the following two log-linear first order conditions for the supply of hours by the young and the old

$$\begin{aligned} \theta_Y \hat{H}_Y &= \hat{W}_Y - \hat{C}_Y \\ \theta_O \hat{H}_O &= \hat{W}_O \end{aligned}$$

Substituting these two equations into equation (A1) we get

$$(x + \theta_Y)\,\hat{H}_Y + \hat{C}_Y = (x + \theta_O)\,\hat{H}_O$$

and rearranging the last equation

$$\frac{x+\theta_Y}{x+\theta_O}\hat{H}_Y + \frac{1}{(x+\theta_O)}\hat{C}_Y = \hat{H}_O \tag{OA1}$$

From equation (OA1) it follows that

$$\left[\frac{x+\theta_Y}{x+\theta_O}\right]^2 Var(\hat{H}_Y) + \left(\frac{1}{x+\theta_O}\right)^2 Var(\hat{C}_Y) + 2\frac{x+\theta_Y}{(x+\theta_O)^2} Cov(\hat{H}_Y, \hat{C}_Y) = Var(\hat{H}_O) \tag{OA2}$$

Equation (OA2) shows that as long as we restrict ourselves to the case where $Cov(\hat{C}_Y, \hat{H}_Y) \geq 0$ then a necessary condition to match the hours fact is that $\theta_O > \theta_Y$. That is, in the situation where the young have a stronger wealth effect, the old must have a stronger substitution effect. Similarly, substituting equation (OA1) to equation (A1) implies

$$\hat{W}_Y + x\left(\frac{\theta_O - \theta_Y}{x+\theta_O}\right)\hat{H}_Y = \hat{W}_O + \frac{x}{x+\theta_O}\hat{C}_Y$$

from which it follows that

$$\begin{aligned} Var(\hat{W}_Y) - Var(\hat{W}_O) = &\left(\frac{x}{x+\theta_O}\right)^2 Var(\hat{C}_Y) + 2\frac{x}{(x+\theta_O)} Cov(\hat{W}_O, \hat{C}_Y) \\ &- \left[x\left(\frac{\theta_O - \theta_Y}{x+\theta_O}\right)\right]^2 Var(\hat{H}_Y) - 2x\left(\frac{\theta_O - \theta_Y}{x+\theta_O}\right) Cov(\hat{W}_Y, \hat{H}_Y) \end{aligned} \tag{OA3}$$

Note then that the second row in (OA3) is negative. However, the presence of the terms $\left(\frac{x}{x+\theta_O}\right)^2 Var(\hat{C}_Y)$ and $2\frac{x}{x+\theta_O}Cov(\hat{W}_O, \hat{C}_Y)$ imply that for this model economy we cannot rule out its ability to match the twin facts.

To gain intuition for why this model could at least qualitatively account for our empirical labor facts it is first useful to note that in the case where the preferences for the young (old) are also as in GHH (KPR) then this model economy corresponds to proposition 3. As such, the model cannot jointly match the labor market facts. Specifically, note that in that case equation (OA2) implies that the necessary condition to match the hours fact continues to be $\theta_O > \theta_Y$ but then equation (OA3) implies that

$$Var(\hat{W}_Y) - Var(\hat{W}_O) = -\left[x\frac{\theta_O - \theta_Y}{x+\theta_O}\right]^2 Var(\hat{H}_Y) - 2x\left(\frac{\theta_O - \theta_Y}{x+\theta_O}\right) Cov(\hat{W}_Y, \hat{H}_Y) < 0$$

and the model cannot match the wage fact.

To summarize, this example shows: If two groups differ in their labor supply wealth effect, then as long as the group with a stronger wealth effect also has a weaker substitution effect it may be possible for the twin facts to be matched.

Intuition

An analytical solution of the variances of age specific hours and wages is not possible since it requires a characterization the dynamic evolution of all variables. Thus, in what follows we focus on the response of variables at the impact period of the shock. Importantly, recall that Proposition 1 established that at the impact period, irrespective of the labor supply differences, it is impossible for the response of young hours and young wages to be greater than for the old. However, via simple graphical analysis we illustrate below that with counteracting wealth and substitution effects the labor market responses are not strictly contradicted.

Specifically, assume a model economy where the old have no wealth effect (i.e. have GHH preferences) while the young do have a wealth effect. The left column of panels in Figure OA1 shows the situation where the young *also* have a stronger substitution effect (i.e. the necessary condition from the example above *does not hold*), while the right column shows the situation where the young have a weaker substitution effect (i.e. the necessary condition holds). In either case, the top row displays the model in initial steady state. The key difference between the two situations is that the young labor supply curve is flatter when the necessary condition holds.

Consider a labor demand shock. In the left column, the steepness of the young labor supply curve means that the substitution effect makes young hours less responsive than old hours, and young wages more responsive than old wages (the middle-left panel). Adding the young's wealth effect (the bottom-left panel) only exacerbates the situation because the wealth response increases the wage responsiveness and decreases the hours responsiveness. Since in this situation the substitution effect made young hours less responsive than old hours, the wealth effect mitigates the model's ability to match the hours volatility fact.

On the other hand, in the right column the young labor supply curve is flatter. Now the substitution effect implies that young hours response is greater than the old hours response (the middle-right panel) while the young wage response is less than the old wage response. Adding the wealth effect increases young wage responsiveness, counteracting the substitution effect's force (the bottom-right panel). The price paid for increasing the wage response is a decreased hours response, but it can be the case that this decrease is outpaced by the substitution effect's large hours responsiveness.

In summary, this example shows that when the wealth effect and substitution effects countervail one another then the the labor market responses are not strictly contradicted. We now turn to numerical results which support this intuition.

Numerical Results

In what follows we extend the analysis by considering more general functional forms for the young and old preferences, which necessitate numerical analysis.

Condition does not hold **Condition holds**

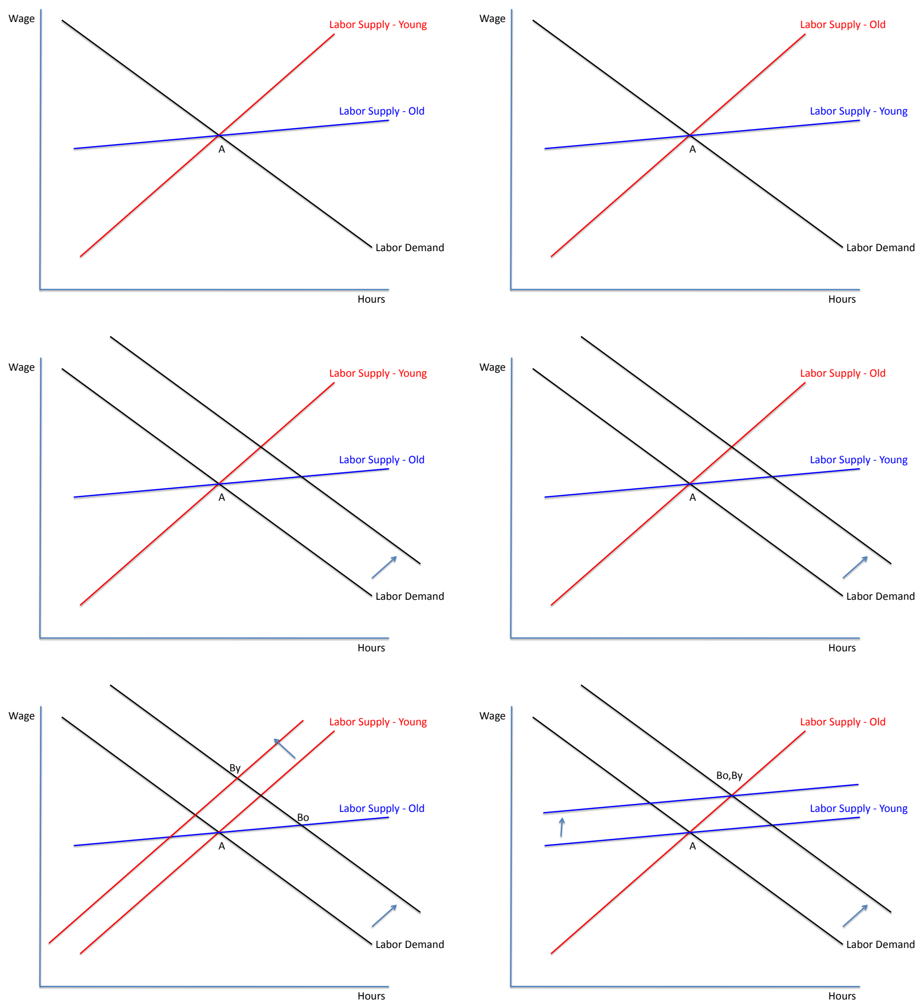

Figure OA1. Intuition for Necessary Condition on Wealth/Substitution Effects

The presentation of the household's maximization problem is identical to that of the paper.

Given the representative household framework, wealth evolves according to the same capital stock, K_t, for both young and old workers. While this appears restrictive at first, it is not so because of our specification for preferences.

Specifically, we adopt the utility representation studied in Jaimovich and Rebelo (2009). The advantage of these preferences is that they allow for an arbitrarily large range of wealth and substitution effects. The utility functions are given by:

$$\begin{aligned} U_Y(C_{Yt}, N_{Yt}) &= \log\left(C_{Yt} - \psi_Y N_{Yt}^{1+\theta_Y} X_{Yt}\right), \quad X_{Yt} = C_{Yt}^{\gamma_Y} X_{Yt-1}^{(1-\gamma_Y)}, \\ U_O(C_{Ot}, N_{Ot}) &= \log\left(C_{Ot} - \psi_O N_{Ot}^{1+\theta_O} X_{Ot}\right), \quad X_{Ot} = C_{Ot}^{\gamma_O} X_{Ot-1}^{(1-\gamma_O)}. \end{aligned}$$

As Jaimovich and Rebelo (2009) show, γ_i controls the strength of the wealth effect on labor supply, for $i = Y, O$. Thus, recalling Figure OA1, the parameter θ_i governs the slope of the labor supply curve, while γ_i governs the wealth effect's shift of the labor supply curve. Our analysis allows for arbitrarily large differences in these key properties of the labor supply curve between young and old.

Therefore, the necessary condition described above boils down to a sign restriction: $\theta_Y > \theta_O$ and $\gamma_Y < \gamma_O$, or $\theta_Y < \theta_O$ and $\gamma_Y > \gamma_O$. Our simulations will show whether, even in very general preference specifications, this sign restriction remains necessary for the model to match the twin facts.

The production function

To conduct the numerical analysis we need to specify a production function that uses capital, K_t, young labor, H_{Yt}, and old labor, H_{Ot}, as inputs. To impose *symmetry in cyclical labor demand characteristics*, we consider the following constant elasticity of substitution production function: $Y_t = A_t K_t^{\alpha} \left[\mu H_{Yt}^q + (1-\mu) H_{Ot}^q\right]^{\frac{1-\alpha}{q}}$. [1]

Business cycles in our analysis are driven by exogenous shocks to productivity, A_t. This stochastic process is identical to that considered in the paper.

The representative firm maximizes over the choice of factor inputs, taking all prices as given. In equilibrium, households and firms are optimizing and all markets clear. In particular, $H_{Yt} = s_Y N_{Yt}$ and $H_{Ot} = (1 - s_Y) N_{Ot}$.

The simulations

Here we discuss the choice of parameter values under which we evaluate the model, listed in the top panel of Table OA9.

The elasticity of substitution between young and old hours in production is given by $\frac{1}{1-q}$. At the upper end of the range, $q = 1$ corresponds to the case

[1] When $q = 1$, we obtain the standard RBC model which treats all labor input as homogenous.

Table OA9—Simulation Parameter Values

Parameter	Range of values considered
q	$\{-10, -1, -0.5, 0, 0.5, 1\}$
θ_Y	$\{0, 0.5, 1, 1.5, 2, 3\}$
θ_O	$\{0, 0.5, 1, 1.5, 2, 3\}$
γ_Y	$\{0, 0.05, 0.1, 0.15, \ldots 0.85, 0.9, 0.95, 1\}$
γ_O	$\{0, 0.05, 0.1, 0.15, \ldots 0.85, 0.9, 0.95, 1\}$
Parameter	**Value set**
ρ	0.95
α	0.36
β	0.985
δ	0.025
s_y	0.35

Note: Notes: The parameter q governs the elasticity of substitution between young and old hours in production; θ_Y and θ_O govern the Frisch elasticity of labor supply; γ_Y and γ_O govern the wealth effect on labor supply. We set the following parameters to one value for every simulation: ρ the autocorrelation of technology shocks, α the capital share of production, β the time discount, δ the depreciation rate and s_y the share of young in the household. See the text for detailed explanations.

of perfect substitutes, while $q = 0$ corresponds to the case of unit elasticity of substitution (the Cobb-Douglas case) between young and old hours. We also explore lower values of q for which H_{Yt} and H_{Ot} are more complementary than the Cobb-Douglas case.

The Frisch labor supply elasticity for young (old) workers is given by $\frac{1}{\theta_Y}(\frac{1}{\theta_O})$. A natural benchmark in the quantitative literature corresponds to "indivisible labor" case, where $\theta_Y = \theta_O = 0$. [2] However, several recent papers investigating the relationship between individual- and aggregate-level elasticities suggest values between 0.7 and 3.[3] We thus consider values of θ_Y and θ_O that cover this range.

Finally, the strength of the wealth effect on labor supply is given by γ_Y and γ_O. We consider parameterizations that span the entire range between the admissible values of γ in [0,1]. Given that the literature provides little guidance on the appropriate value of γ, we use increments of 0.05, yielding 21 values of the wealth effect parameter for the young, and 21 values for the old.

Overall, we investigate six values of q, five values of θ_Y and θ_O, and 21 values of γ_Y and γ_O. We then simulate the model 66,150 times where each solution corresponds to a different combination of parameters. The remaining parameters of the model are set at standard values, listed in in the bottom panel of Table

[2]See Rogerson (1988) and Hansen (1985).
[3]See for example Rogerson and Wallenius (2009) and references therein.

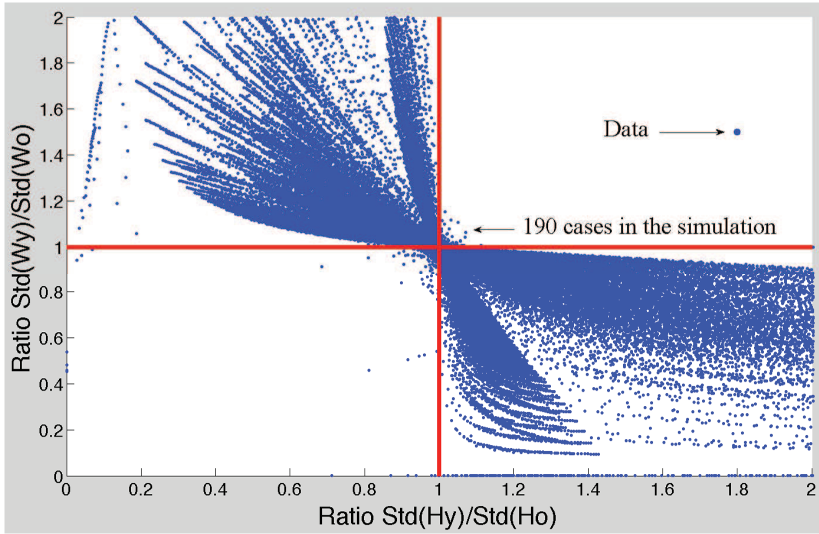

FIGURE OA2. SCATTER PLOT OF RELATIVE VOLATILITIES IN SIMULATIONS AND THE DATA

Note: The twin facts are that young hours and young wages are more volatile than old hours and old wages, respectively. The value obtained from the March CPS, 1964-2010, is marked. It lies in the northeast quadrant. Also plotted are the simulated values from the 66,150 parameterizations of the labor supply model.

OA9.

Out of all 66,150 cases, we find 190 cases where we are able to *qualitatively* reproduce the labor market facts, and indeed in all of these cases we find that the sign restriction discussed above is satisfied. Whenever the wealth effect is stronger for the young than for the old ($\gamma_Y > \gamma_O$) then the substitution effect is stronger for the old then for the young ($\theta_Y \leq \theta_O$). Vice-versa, whenever the wealth effect is stronger for the old than for the young ($\gamma_Y < \gamma_O$) then the substitution effect is stronger for the young then for the old ($\theta_O \leq \theta_Y$). That is, as discussed above, the wealth effect and substitution effect must go in "opposite" ways in order for labor supply differences alone to even *qualitatively* deliver the twin facts.

How close are the relative volatilities, $\frac{Std(H_Y)}{Std(H_O)}$ and $\frac{Std(W_Y)}{Std(W_O)}$ to the values we estimate in the data? Considering the hours volatility fact, the median, mean, and maximum value of $\frac{Std(H_Y)}{Std(H_O)}$ in the 190 cases is 1.01, 1.02, and 1.10, respectively. Meanwhile the value in the data is 1.85. Considering the wage volatility fact, the median, mean, and maximum value of $\frac{Std(W_Y)}{Std(W_O)}$ is 1.01, 1.02, and 1.15, respectively. The value in the data is 1.50. In other words, these cases are *quantitatively* unable to match the values estimated in the US data.

Figure OA2 makes the point clear. On the y-axis we plot the ratio of the volatilities of young wages to old wages $\frac{Std(W_Y)}{Std(W_O)}$, while on the x-axis we plot the ratio of the volatilities of young hours to old hours $\frac{Std(H_Y)}{Std(H_O)}$. Each dot in the scatterplot is the relative volatility of young-to-old hours and wages from one of the 66,150 simulations.[4] The northeast quadrant (where $\frac{Std(H_Y)}{Std(H_O)}$ and $\frac{Std(W_Y)}{Std(W_O)}$ are above one) is the one where the data lies.

The vast majority of cases lie in either the northwest or southeast quadrants where one of the ratio of the volatilities is less than 1. This reinforces our analytical findings that generating greater volatility of young hours results in a lower volatility of young wages (relative to the old), and vice-versa. Moreover, the cases that lie in the northeast quadrant are concentrated far from the relative volatilities observed in the data – differences in labor supply cannot *quantitatively* deliver the twin facts.

We conclude from this analysis that the necessary condition holds quite generally. Thus, differences in labor supply qualitatively deliver the twin facts by diametrically opposing the young's wealth and substitution effects relative to the old. That is, if the young have a stronger wealth effect they must also have a weaker substitution effect, or vice-versa. But then this tension leads to the fact that the model increases $\frac{Std(H_Y)}{Std(H_O)}$ by decreasing $\frac{Std(W_Y)}{Std(W_O)}$, or vice-versa, which hinders the model from generating the two volatility ratios near their values in the data. In summary, differences in labor supply seem to be an unpromising channel, at least on their own, for explaining the twin facts as estimated in U.S. data.

C. *Estimation*

We use three lagged birth rates as our instruments: The first is the birth rate 22 years prior, which represents the middle point of the young age group. The other two are birth rates 35 and 40 years earlier, which represent the old age group. Other choices of lagged birth rates lead to similar results. The results using these three instruments for both equations (11) and (12) are reported in Panel A of Table OA10. One can see that the material difference in these estimates vis-a-vis Table 4 is the strength of the instruments for equation (12). This is a sensible result. The regressor identifying ρ is the growth rate of old hours per unit of capital, $\Delta \log(H_{Ot}/K_t)$. This variable is naturally related to fluctuations in the supply of old hours which require adjustment of the capital stock. The birth rate lagged 35 years effectively captures the influx of "new" old hours (coming from "new" old workers transitioning the young group). The 22- and 40-year lagged birth rates are naturally less connected to this variable: The first captures fluctuations in young hours (not old) while the second captures fluctuations in workers who have supplied old hours for some time (and for whom, therefore, the capital stock has already been adjusted). This is exactly the empirical result we find. The 35-year lagged birth rate is significantly correlated with the growth rate of old hours per unit of capital, while the other two instruments are not. Meanwhile, since (11)'s endogenous regressor is a combination of growth rates in

[4]In some of the 66,150 cases, either $Std(H_O)$ or $Std(W_O)$ are very close to zero, generating very large ratios. These are excluded in the figure.

Table OA10—Alternative Empirical Specifications

	Point Estimate	Standard Error	First Stage F-Statistic	J-test p-value	$\sigma = \rho$ p-value
			A. All Instruments		
σ	0.656	0.022	10.829	0.592	< 0.001
ρ	0.231	0.011	5.067		
			B. One Instrument per Equation		
σ	0.600	0.225	19.656	–	0.105
ρ	0.199	0.027	13.891		
			C. Newey-West Standard Errors		
σ	0.648	0.142	10.829	0.327	0.013
ρ	0.206	0.084	13.891		
			D. Limited Residual Serial Correlation		
σ	0.626	0.217	10.829	0.098	0.289
ρ	0.266	0.256	13.891		
			E. One-Step		
σ	0.650	0.192	10.829	–	0.188
ρ	0.199	0.281	13.891		
			F. Iterated		
σ	0.741	0.055	10.829	0.457	< 0.001
ρ	0.222	0.017	13.891		

Note: Data from the March CPS, 1964-2010. J denote's Hansen's J-test; $\sigma = \rho$ denotes the F-test of the null that the two parameters are equal. Note: For several of the specifications the first stage F-statistic is indeed identical to the benchmark results reported in Table 4.

young hours and output, we reason that all three birth rates have explanatory power. Thus our main results in Table 4 use the three lagged birth rates as instruments in (11) but only the 35-year lagged birth rate as an instrument for (12).

If instead we restrict either equation to be estimated by only one instrument, we obtain the Panel B of Table OA10. Here, since the endogenous regressor identifying σ involves young hours, we use the 22-year lagged birth rate. This is significantly correlated with $\Delta \log H_{Yt} - \Delta \log Y_t$ as column four reports. There are no over-identifying restrictions in this case so that the J-statistic is undefined. The resulting standard errors are much larger which leads to the F-test of $\sigma = \rho$ being borderline. We use the overidentifying restrictions (that is, use all three instruments) in our benchmark results of Table 4 because we see no *a priori* reason why we should not and because the resulting estimates are more tightly estimated. Panel B shows that the point estimate for σ is the only one materially affected by this choice, and the effect is modest.

Our benchmark results use Andrews' (1991) quadratic-spectral kernel with Newey and West's (1994) optimal bandwidth. Panel C reports results for Newey and West's (1987) Bartlett kernel instead, again with optimal bandwidth chosen according to Newey and West (1994). Panel D reports instead results for Newey and West's (1987) kernel using only one lag.[5] Panel E reports one-step (using an identity weighting matrix) and Panel F reports iterated GMM results. Using Newey and West (1987) standard errors yields little change to any of our estimates. Limiting the residual serial correlation estimated by the HAC estimator affects the point estimates mildly, the standard errors much more, but is rejected by the test of overidentification. One-step and iterated GMM point estimates are within the conventional confidence intervals of our benchmark results.

We interpret Table OA10 as evidence that the benchmark estimates presented in Table 4 are robust.

D. Data

From the Bureau of Economic Analysis via Haver's USECON database we obtain the following aggregate series: Real GDP per Capita is LXNFA, the GDP Defaltor is LXNFIA, Aggregate Hours is either constructed from the CPS data below[6] or taken from private aggregate hours LHTPRIVA (both yield similar results for construction of cyclical measures or calculation of data moments in our quantitative exercise), the Labor Share is LXNFBL, Capital Share is one minus the Labor Share, Real Capital per Capita is given by MFPNFKH, and Birthrates are from POPBR.

[5]We use the Bartlett kernel here since the desire is to analyze the effect of restricting the serial correlation picked up by the HAC estimator. Since the Bartlett kernel involves truncation it is a natural candidate for this exercise; since the quadratic-spectral kernel does not truncate it less effectively captures the spirit of this sensitivity analysis. Nevertheless we report estimates from Andrews' (1991) HAC estimator with the smallest possible bandwidth yields $\hat{\sigma} = 0.615$ (0.227), $\hat{\rho} = 0.255$ (0.231), $J = 0.133$ and $\sigma = \rho$ F-test p-value 0.283.

[6]When constructing aggregate hours from the CPS, we use the 1962 and 1963 surveys since we do not need to form disaggregated groups based on educational attainment. We need these earliest surveys in order to provide the lag of aggregate hours onto which we project when creating cyclical measures. Forming such groups for the purposes of imputing missing values in last week hours has only minimal effects which are negligible when we do so for 1964-2010 when education data is provided.

Data on age-specific hours, employment shares, and wages is constructed from the Current Population Survey (CPS) conducted by the Census Bureau and Bureau of Labor Statistics. We use the surveys from 1964-2010 since the 1963 survey contained no education information. To obtain wage data, we use questions in the March CPS about income obtained in the previous (last) year. 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 modify Krusell, Ohanian, Rios-Rull and Violante's (2000) imputation methods 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's (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 1964-2010 our "disaggregated" groups each have at least fifty members.[7] 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:

- 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

[7] Additionally cutting by race (white/nonwhite) does not change matters much.

hours last week, number of weeks worked last year, wage and salary income last year, and CPS Person weight, respectively.[8] Then the measures of group g's "disaggregated" group average, weight, hours worked last year, and income last year are

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

$$\mu_g = \sum_{k\in g} \mu_k$$

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

$$y_g = \frac{1}{\mu_g}\left(\sum_{k\in g} y_k\mu_k\right)$$

Let γ be a set of gs: 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.[9]

$$W_\gamma = \frac{\sum_{g\in\gamma} y_g\mu_g}{\sum_{g\in\gamma} h_g f_g \mu_g}$$

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. Our data on wages come from the wage and salary income last year question, screening for self-employed and farm-working individuals. To link our labor income data with hours worked data, we use the hours worked last year in the young/old hours data used for GMM estimation of σ and ρ.

[8] In 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.

[9] Krusell, 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.

REFERENCES

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

Bureau of Labor Statistics and Census Bureau. 1964-2011. "Current Population Survey." United States Department of Labor and Department of Commerce. http://www.census.gov/cps/ (accessed Feburary 9, 2012).

Hansen, Gary D. 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics*, 16(3): 309–327.

Hansen, Lars P. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, 50(4): 1029–1054.

Haver Analytics. 1962-2011. "U.S. Economic Statistics (USECON)." Bureau of Economic Analysis, United States Department of Commerce. http://www.haver.com (accessed October 1, 2012).

Newey, Whitney K, and Kenneth D West. 1987. "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–08.

Newey, Whitney K, and Kenneth D West. 1994. "Automatic Lag Selection in Covariance Matrix Estimation." *Review of Economic Studies*, 61(4): 631–53.

Rogerson, Richard. 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics*, 21(1): 3–16.

Rogerson, Richard, and Johanna Wallenius. 2009. "Micro and Macro Elasticities in a Life Cycle Model with Taxes." *Journal of Economic Theory*, 144(6): 2277–2292.