Labor Force Composition and Aggregate Fluctuations*

Alessandro Mennuni
University of Southampton
alessandro.mennuni@soton.ac.uk

December 21, 2013

Abstract

Labor composition by gender, age, and education has undergone dramatic changes over the last forty years in the United States. Furthermore, the volatility of total market hours differs systematically between genders, age, and education groups. I develop a large-scale business cycle model which suggests that demographic changes by gender and education affect labor supply elasticities at the micro level. This has important repercussions on aggregate volatility. Changes in labor composition account for 30% of the observed changes in aggregate volatility over this period of time. To solve the model over this large transition, I develop a new algorithm which extends perturbation methods to the stochastic transition path and can be applied to a broad class of DSGE models.

JEL Classification Codes: E32, J11, J22, C63
 Keywords: Business Cycle, Demographic Composition, Computational Economics

1 Introduction

This paper uses a structural model to investigate the possible significance of demographic changes on aggregate volatility. Jaimovich and Siu (2009) (J-S) exploited the exogeneity to current economic activity of the changes in the population distribution by age, to measure their impact on business cycle volatility, through regression analysis. They found that changes by age played a significant role in determining business cycle volatility over the last 40 years.[1] Meanwhile, labor composition changes by gender and education have been equally dramatic (see for instance Katz and Autor (1999) for the US, Katz and Freeman (1994) for a cross-country comparison) and appear to be correlated to the changes by age. Through a decomposition exercise, J-S show that if the response to the business cycle of hours worked by each demographic group is not affected by these compositional changes, the contribution of education is half the one of age, and the one of gender composition is negligible. This is an important benchmark but raises the question: is the group-specific response to business cycle shocks affected by demographic composition? And how much do gender, age, and education composition affect business cycle volatility once these endogenous responses are taken into account?

To address these questions I build an overlapping generations model in which males and females choose education, intrafamily time allocation, and savings. The model builds on the perfect foresight model of Heathcote et al. (2010) and extends it to an environment with aggregate uncertainty.[2]

The model is able to replicate the evolution of aggregate output volatility over time. Since most of the shocks are identified from disaggregated data, this result was not obvious and signals a good model specification.[3] It is then possible to measure the contribution that the labor reallocation played on the changes in aggregate volatility: counterfactual simulations suggest that had the labor composition remained trendless at its steady-state levels, business cycle volatility would

---

[1]The finding is robust to considering a larger pool of countries (Lugauer and Redmond (2012)), or exploiting the variation in demographic change across the United States (Lugauer (2012b)). Lugauer (2012a) reconciles the result with a search and matching model. Janiak and Monteiro (2011) find that differences in tax rates explain some of the differences in aggregate volatility across countries through their effects on the age distribution of labor.

[2]To link this model to the data, I use the National Income and Product Accounts, the March supplement of the Current Population Survey (CPS), and other US demographic information. Following Attanasio and Weber (1995), by grouping the labor force by the observables mentioned, I create a synthetic panel which has an exact counterpart in the model.

[3]The dynamic stochastic general equilibrium literature concerned with changes in business cycle volatility has mainly focused on the volatility slow-down from the 1980s, and accounts for it essentially by a reduction in the volatility of aggregate shocks (see Arias et al. (2007) and Smets and Wouters (2007)). The present framework decomposes the effects that aggregate Solow residuals have on volatility, into less aggregate productivity shocks and the labor composition.

have been 17% lower in the early 1970s and close to the observed volatility in the 1980s and 1990s (when there was little volatility). By accounting for part of the high volatility of the early 1970s and of the slowdown in the 1980s, labor reallocation also accounts for 30% of the Great Moderation, the large volatility decline in the 1980s, initially documented by Kim and Nelson (1999) and McConnell and Perez-Quiros (2000).[4] In addition, had the labor composition converged to its steady state, the business cycle would have been 5% more volatile than it was over the last decade. This last result suggests that a return to times of low business cycle volatility is not likely; in fact, aggregate volatility can be expected to increase somewhat if labor composition continues to follow current trends.[5]

The paper confirms the finding of J-S that the age composition has an impact. However it suggests that its effect is not as important. In fact, its role is smaller than the one played by educational composition and comparable to the one of gender composition. That changes by education and gender played a stronger role relies on a relationship uncovered in this paper between demographic trends by gender and education, and labor supply elasticities. Education and gender composition trends are largely due to increases in market hours per worker of the highly educated and the female population. In this model, labor supply elasticities decline with the increase in hours worked. This decline contributes to the observed moderation in aggregate volatility. Instead, demographic changes by age essentially affect a more extensive margin—the number of people in each age bracket—which does not affect labor supply elasticities.

That demographic trends induced a decline in labor supply elasticities is a prediction of the model, and relies on the utility function chosen. As such, it should be validated. First, that elasticities decrease in labor is broadly consistent with the cross-sectional evidence that groups that work more hours have less elastic labor supplies. E.g. males versus females, prime age versus closer to retirement, women without children versus women with children; see Reichling and Whalen (2012). Furthermore, with the chosen utility specification—widely adopted in the literature—the model is successful at predicting how the hours volatility ratio across groups has evolved over time, as well as the aforementioned trends in aggregate volatility. This is not the case with a specification with constant Frisch elasticities. Lastly, the prediction that female labor supply elasticities declined over time is consistent with direct estimations of married women's elasticities

---

[4]Between the two sub-samples, aggregate output volatility decreased by $\log(1.64) - \log(.99) = 50.47$ log points. Had the shares remained stable as in the counterfactual, we would have observed a reduction in volatility of $\log(1.42) - \log(1.00) = 35.07$ log points. Therefore, these demographic changes account for (50.47-35.07)/50.47 or 30.53% of the moderation in output.

[5]One intuitive reason for this result is that the fraction of prime age workers (those least sensitive to business cycle shocks compared with other age groups) declines as the baby boomers grow older.

(Heim (2007) and Blau and Kahn (2007)). These papers called for an explanation of this fact but Heim (2007) and Bargain et al. (2012) find that the decline in these elasticities is not accounted for by demographics in the sense that it is due to a decline in all female sub-groups' elasticities rather than because of the change in the demographic composition, holding sub-group elasticities constant. My model confirms their finding that elasticity declines for all women's sub-groups; however, it suggests that the decline was indeed due to demographics.

The link between labor supply elasticities and demographic composition also relates to a growing literature aimed at reconciling micro and macro estimates of labor supply elasticities by moving beyond the representative household. For instance, Dyrda et al. (2012) distinguish between stable and unstable workers and target the average relative volatility of market hours between these groups. The uncovered links between demographic trends, labor supply elasticities, trends in the relative volatility of hours, and aggregate fluctuations, may have implications for a wide array of policy questions and guide further the identification of labor elasticities; see Blundell and MaCurdy (1999) among others.

The over-all effect of gender, age and education composition on the Great Moderation (30%) is comparable to what J-S attributed to the age composition alone through regression analysis. This discrepancy in our results could either come from the fact that the model is misspecified, or from the fact that the regression analysis overstates the role of the age composition. To assess this, I adapt the analysis of J-S, where business cycle volatility is regressed over the age composition, to the data generated by the model, where age does not play a predominant role. I find that the estimated effect of age composition for aggregate volatility is significant and quantitatively similar to what was found by J-S for the true data. From this result one cannot conclude that the model understates the role of age composition. On the other hand, the exercise highlights that regression analysis overstates the importance of age composition relative to the role it has in the model.

Highlighting how demographic changes are related to the business cycle, this paper shows that heterogeneity matters for aggregate volatility. This is in contrast to the general message that has emerged from a naive reading of the literature on heterogeneous agents and, in particular, of Krusell and Smith Jr (1998) and Ríos-Rull (1996), who find no major differences between the business cycle properties of their models with heterogeneity either in income and wealth or over the life-cycle, and a representative agent model simulated around the steady state. Perhaps not surprisingly, this paper finds that large demographic changes—not just those around the steady state—do matter for aggregate volatility. Furthermore, and perhaps more interestingly, when heterogeneity changes over time as much as is observed in the data, the business cycle properties of the model are affected in a way which is consistent with how aggregate volatility has evolved over time.

A computational challenge arises from the fact that the model is solved for a period of large transition, this makes standard perturbation methods–the typical technique used to solve large scale models–an inaccurate solution method.[6] To address this issue, I develop a technique which consists of applying perturbation methods at many points over the equilibrium path. The solution is much more accurate than that found by applying perturbation methods only to the deterministic steady state, and can be applied to a broad class of dynamic stochastic general equilibrium (DSGE) models. Similarly to the Parameterized Expectation Approach (PEA), the curse of dimensionality is broken because the model is only solved over the equilibrium path, see Den Haan and Marcet (1990). However, the method proposed here does not rely on a contraction mapping of the policy functions which often leads to convergence issues with large models. Judd et al. (2012) propose a substantial improvement that allows them to solve a model with 80 variables. The model in this paper has 800 variables and is effectively solved over the transition with the proposed method, on a normal desktop computer. Matlab code is available from www.southampton.ac.uk/~alexmen/.

The paper proceeds as follows. In Section 2, a first look at the labor data will motivate the conjecture that the reallocation of labor across groups plays a role in business cycle volatility. Section 3 sets up the model and Section 4 parameterizes it. Section 5 tests the model, and measures the effects of labor reallocation on aggregate volatility. Section 6 concludes. Appendix A describes the data used, Appendix B offers a description of the original computation algorithm developed and Appendix C details the estimation of some parameters of the production function and the identification of the shocks.

2   Stylized facts

As it is well known, aggregate output volatility in the U.S. increased during the early 1970s and declined during the 80s and 90s.[7] Recently, volatility has increased again and there is renewed curiosity about its future unfolding. To relate these facts to the labor force composition, following Gomme et al. (2005) and J-S, I use data from the March supplement of the CPS to construct annual series of hours for the demographic groups considered. Figure 1 shows the share of paid hours worked over time by gender, age (young (15–29), prime age (30–55) and older workers) and low and high education (at least four years of college).

---

[6]As Caldara et al. (2012) point out, in practice perturbation methods are the only computationally feasible method for solving medium- and large-scale DSGE models, but these techniques guarantee accurate solutions only in the proximity of the steady state.

[7]See Kim and Nelson (1999), McConnell and Perez-Quiros (2000) and Stock and Watson (2003) among others.

As is shown in the first two columns of Table 1, hours worked by prime age workers increased relative to those by other age groups, moving from an average of 58% prior to 1984 to an average of 71% from 1984.[8] By contrast, hours by the young fell, while hours by the old remained more stable. Furthermore, the cyclical volatility of hours is substantially lower for prime age workers as initially documented by Clark and Summers (1981): after removing the trend from each series,[9] the standard deviation of hours is 2.37% for the young, 1.36% for the prime and 1.87% for older workers. In addition, hours volatility of the prime relative to the young and old even reduces over time, as shown in Figure 2, second panel.

The relative increase in prime-age hours and the fact that hours are less volatile for this age group may have contributed to the observed reduction in aggregate hours volatility. These two facts also characterized the distribution of hours by gender and education.

As shown in Table 1, hours worked by women increased relative to those by men between the first and second sub-sample. It is less well known however that hours worked by men are more volatile than hours worked by women: the standard deviation of hours is 1.29% for women and 1.87% for men.[10] Furthermore, the relative volatility is remarkably stable over time as shown in Figure 2, first panel.

Similarly, the share of the highly educated increased relative to that of the less educated and the cyclical volatility of the highly educated is lower. Figure 2, upper right-hand panel, shows how the relative volatility of hours over time has a kink around the mid 80s; it remains however true over the whole sample that the relative volatility between the two groups is smaller than one. Since hours may be measured with error, Figure 2, lower right-hand panel, shows the relative volatility of employment by education; apart from the period 87-92 just after the kink, the ratio is always below 1.[11]

It is important to notice, however, that hours volatility for each group decreased in the Great Moderation period as shown in the last column of Table 1.[12]

---

[8]1984 is the reference year adopted by the literature as the beginning of the great moderation. See for instance Stock and Watson (2003)

[9]Here and throughout the paper I use the Hodrick-Prescott (HP) filter with a smoothing parameter of 6.28 as suggested by Ravn and Uhlig (2002). Using 10 as suggested by Baxter and King (1999), or looking at growth rates, essentially does not affect the results.

[10]This fact is all the more surprising given that at the individual level, hours worked by women are more volatile than for men.

[11]Weights and classifications in the CPS data are such that data are essentially comparable over time. It is known however that comparability does not hold in some cases (see for instance Abraham and Shimer (2002)); the observed kink might be partly due to some reclassification that induced discrepancies over time.

[12]To highlight this volatility slowdown, the Great Moderation period is ended in year 2000 in order not to confound this period with the more turbulent past decade.

Table 1: Hours Volatility by Gender, Age and Education

	Share of hours 67-83	Share of hours 84-10	St. dev	$\frac{St.dev(84-00)}{St.dev(67-83)}$
	Age			
Young	30.11	20.39	2.37	91.82
Prime age	58.35	70.72	1.36	79.12
Old	11.54	8.89	1.87	87.08
	Gender			
Women	33.88	42.19	1.29	70.50
Men	66.12	57.81	1.87	90.38
	Education			
Low	81.69	70.36	1.16	82.13
High	18.31	29.64	1.90	99.72

Notes: numbers are expressed in percentage terms. Share of hours is the ratio of hours by each group to total hours.

2.1 A simple accounting exercise

J-S complement their analysis with a simple accounting exercise that measures the contribution of these demographic changes to the Great Moderation in total hours worked. The accounting exercise consists of constructing a counterfactual time series for total market hours using the observed group-specific hours per capita, but setting the population shares constant. They find that age changes account for 20% of the Great Moderation in total hours, while education accounts for 10%. Gender composition is not important.[13] As they acknowledge, the experiment assumes that the volatility of group-specific market work is independent of the labor composition. That is, it assumes the absence of indirect effects of changing labor composition on aggregate volatility via changes in the group-specific response to shocks. They argue that this assumption is reasonable for age groups. The structural analysis in this paper finds that indeed the labor supply elasticities are not affected by changes in the age composition. However, changes by gender and education have strong effects on these elasticities, including elasticities by age groups. This suggests that the group-specific market hours' response to shocks are endogenous to the labor composition. A theory is necessary to predict how these responses are affected by the labor composition.

---

[13]One issue with the exercise by gender is that the demographic transition involves the number of women participating in the labor market. So rather than holding population shares constant, as can be done by age and education, one can hold participation rates constant. However, this way one removes a very large portion of the business cycle, making the exercise less meaningful.

3 The Model

In each period, the economy is populated by a continuum of individuals and an equal random number p_0 of males and females are born. Following Heathcote et al. (2010), the life cycle of an individual comprises 3 sequential stages: education, matching and work. The first decision—high or low education—is made by a newly born individual before entering the marriage market. At this point, members of the opposite sex are randomly matched (no one remains single). For tractability, these two stages happen during the first period of life.[14] From the second period of life, the couple enters the working stage and jointly chooses hours of work for husband and wife, as well as consumption and savings.

3.1 Education

In each period the newly born have to make a discrete choice between college (h) or lower schooling (l). When they are born, they draw an idiosyncratic cost κ of acquiring high education from the distribution $\kappa \sim F^g(\kappa)$, with g indicating male (m) or female (f). This cost captures in reduced form the utility and financial factors that make acquiring a college degree costly and is assumed log-normal:

$$\ln(\kappa) \sim N(\bar{k}^g, \upsilon^g). \tag{1}$$

With this simple specification, it is possible to match the evolution of the education composition through changes in $\bar{k}^g$. Individuals' decisions are made by comparing their education cost with the value gain upon entering the labor market with higher education: $M^g(h; \omega) - M^g(l; \omega)$. $M^g(e; \omega)$ is the gender-specific expected lifetime utility of entering the marriage stage as a function of education $e \in \{h, l\}$ and all the other relevant state variables represented by ω.[15] They choose higher education if $M^g(h; \omega) - M^g(l; \omega) > \kappa$. For each gender, the share of the highly educated in the cohort just born is therefore

$$q^g(\omega) = F^g \left(M^g(h; \omega) - M^g(l; \omega) \right). \tag{2}$$

---

[14]The role of marriage is not predominant but is kept as in the framework of Heathcote et al. (2010) because it can improve the model capacity to predict the evolution of relative hours volatility by gender and education. Furthermore, focusing on married couples relates more closely to the bulk of empirical studies on labor supply research (see Blundell and MaCurdy (1999) and Blau and Kahn (2007) among others). However, it would be interesting to also allow for singles as the proportion of married households is declining. This would complicate the model.

[15]As is discussed in section 3.6, ω contains all the shocks, the distribution of assets across households, and the distribution of households by age and education of husband and wife.

3.2 Marriage

At this point, individuals are characterized by gender g and education e. Men and women match according to the gender-specific probability $\pi^g_{e^m,e^f}(\omega) \in [0,1]$. The expected value upon entering the matching state for a woman of education level e^f is

$$M^f(e^f;\omega) = \pi^f_{h,e^f}(\omega)V(h,e^f;\omega) + \pi^f_{l,e^f}(\omega)V(l,e^f;\omega), \tag{3}$$

where $V(e^m,e^f,\omega)$ is the expected lifetime utility of a household with education pair e^m and e^f. A similar expression can be derived for $M^m(e,\omega)$:

$$M^m(e^m;\omega) = \pi^m_{e^m,h}(\omega)V(e^m,h;\omega) + \pi^m_{e^m,l}(\omega)V(e^m,l;\omega). \tag{4}$$

Enrollment rates $q^g(\omega)$ and matching probabilities $\pi^g_{e^m,e^f}(\omega)$ jointly determine the education composition of newly formed households. For instance, the fraction of new households formed by men with high education and women with low education is

$$q^m(\omega)\pi^m_{h,l}(\omega) = (1 - q^f(\omega))\pi^f_{h,l}(\omega). \tag{5}$$

Since no individual will remain single, for any e^f

$$\pi^f_{l,e^f}(\omega) + \pi^f_{h,e^f}(\omega) = 1, \tag{6}$$

and similarly for men: $\pi^m_{e^m,l}(\omega) + \pi^m_{e^m,h}(\omega) = 1$ for any e^m. One can show that the cross-sectional Pearson correlation between education levels of husband and wife, a measure of the degree of assortative matching is

$$\varrho = \frac{q^m(\omega)\pi^m_{h,h}(\omega) - q^m(\omega)q^f(\omega)}{\sqrt{q^m(\omega)(1 - q^m(\omega))q^f(\omega)(1 - q^f(\omega))}}. \tag{7}$$

Following Heathcote et al. (2010), ϱ is treated as a parameter through which the probability function $\pi^g_{e^m,e^f}(\omega)$ is pinned down.

3.3 Work

Households are distinguished by the husband and wife's education levels e^m, e^f, their age j and their amount of assets a. They choose consumption c and assets a', and hours of work for each gender l^g, in order to solve the following problem:

$$\mathbb{V}\left(e^m, e^f, j, a; \omega\right) = \max\ u(c, 1 - l^m, 1 - l^f - l_h) + \beta\zeta_j E\left[\mathbb{V}\left(e^m, e^f, j+1, a'; \omega'\right)\right]$$

subject to the constraints: $\zeta_j a' + c = a(1+r) + w(m, j, e^m)l^m + w(f, j, e^f)l^f$,

$$a' \geq 0\ if\ j = J, \qquad a = 0\ if\ j = 1, \qquad c \geq 0, \qquad l^m, l^f + l_h \in [0,1], \tag{8}$$

where r is the interest rate and $w(g, j, e^g)$ the wage for each age, education and gender. $\zeta_j \in [0, 1]$ is the survival factor for the household at age j; it is such that people die for sure at age J, i.e. $\zeta_J = 0$.[16] When the household is alive, the period utility function u is increasing, twice continuously differentiable and strictly concave in its 3 arguments.[17] Husband and wife die together and their period utility u after death is equal to 0, hence $\mathbb{V}\left(e^m, e^f, J+1, a; \omega\right) = 0$ and the zero debt constraint in the last period of life J. l_h is an exogenous time cost specific to women.[18] The expectation is taken over ω' given ω.

New households start with zero assets. Thus, the value at the time of forming a household is equivalent to the expected lifetime utility of a formed household of age 1 and with zero assets:

$$V(e^m, e^f; \omega) = E\left(\mathbb{V}(e^m, e^f, 1, 0; \omega')\right). \tag{9}$$

3.4   Household Distribution and its Law of Motion

Denote $p_{age,edu} : \{1, ..., J\} \times \{h, l\} \times \{h, l\} \to \Re_+$ the mass of households by age and education of the couple. $p'_{age,edu}(1, h, e^f) = \pi^m_{h,e^f} q^m p_0$ is the mass of newly formed households composed of men with high education and women of education $e^f = \{h, l\}$. $p'_{age,edu}(1, l, e^f) = \pi^m_{l,e^f}(1 - q^m)p_0$ is the mass of newly formed households composed of men with low education and women with education e^f. Let the mass of older households be defined recursively as $p'_{age,edu}(j, e^m, e^f) = p_{age,edu}(j - 1, e^m, e^f)\zeta_{j-1}$.

3.5   Firms

Competitive firms maximize profits using the following production function

$$y = A^{1/\theta} \left(\alpha k^\theta + (1 - \alpha)L^\theta\right)^{1/\theta}, \tag{10}$$

where A is total factor productivity (TFP), α is associated with the labor share of total output and θ measures the complementarity across capital and L, which

---

[16]The fact that ζ_j multiplies a' in the budget constraint reflects competitive annuity markets, see Ríos-Rull (1996) for a digression.

[17]The utility function can depend on education and age. To simplify notation, u is not indexed accordingly.

[18]In the absence of a more sophisticated theory of the household, the evolution of this parameter will help reproduce the distribution of hours by gender. Its reduction over time captures housework production technology improvements and a fall in child care cost, which on top of the reducing gender wage gap help explain increases in female market hours. See among others Greenwood et al. (2005) and Attanasio et al. (2008).

is a composite of several labor groups:

$$L = \left(\sum_{i=1}^{I} (z_i n_i)^{\sigma} \right)^{1/\sigma}. \tag{11}$$

σ measures the degree of complementarity across groups.[19] z_is are labor-augmenting technology shocks specific to each labor group, n_i is hours worked by all individuals categorized in group i. This specification, distinguishing between labor-augmenting and a TFP shock makes it possible to match aggregate production given the inputs, while matching hours and wages through labor demand.[20]

There is a mapping between groups i (which have an empirical counterpart) and individuals: each group i is formed of agents of the same gender, age group and education level.[21] The mapping is represented by I dummy matrixes $\chi_i(g, e^m, e^f, j)$ which contain zeros and ones depending on whether the labor input of the agent belongs to group i. So, for instance, group 1 is formed of women, young and with low education. For a generical i,

$$n_i = \sum_{g} \sum_{e^m} \sum_{e^f} \sum_{j=1}^{J} l^g \left(e^m, e^f, j \right) p_{age,edu} \left(e^m, e^f, j \right) \chi_i(g, e^m, e^f, j). \tag{12}$$

Calling the number of age groups n_{age}, the total number of groups I is $2n_{age}2$, i.e. the two genders times the age groups times the two education levels.

The representative firm hires labor according to the following condition

$$(1 - \alpha) A y^{1-\theta} \left(\sum_{i=1}^{I} (A_i n_i)^{\sigma} \right)^{\frac{\theta}{\sigma} - 1} z_i^{\sigma} n_i^{\sigma - 1} = w_i \tag{13}$$

for every i, where w_i is the wage rate for group i; so if $\chi_i(g, e^m, e^f, j) = 1$, then $w(g, j, e^g) = w_i$. Capital is demanded according to the following condition

$$A\alpha (k/y)^{\theta - 1} = r + \delta, \tag{14}$$

where δ is the depreciation rate of capital.

---

[19]It is assumed here that all the groups have the same complementarity across them and with capital. It would be interesting to extend this function to the one introduced by Krusell et al. (2000) as done by Castro and Coen-Pirani (2008) and Jaimovich et al. (2009) to study the hours cyclicality by skill and age (see also Johnson and Keane (2007)). However, in the presence of several groups, this makes it harder to identify the shocks analytically, thereby complicating the estimation procedure. See Zoghi (2010) for a survey of possible ways to estimate a labor composition index.

[20]By Euler's theorem, the capital demand equation is also satisfied with no need for an extra shock.

[21]Consistently with the empirical section, there are three age groups: the young (1-10), the prime age (11-35) and the older agents (36-40).

3.6 State Space

To make rational choices, agents need to know their type.[22] They also need to predict prices, which depend on the shocks and on the distribution of assets and households across age and education pairs of husband and wife. The next subsections define the state space in more detail.

Exogenous Processes

Group-specific productivity levels z_i in (11) are the sum of gender, age and education specific shocks ($\varepsilon^g \in \mathcal{E}^g, \varepsilon^{age} \in \mathcal{E}^{age}, \varepsilon^{edu} \in \mathcal{E}^{edu}$) so that

$$\log(z_{t,i}) = \sum_{j=1}^{2} \varepsilon_{t,j}^g I_g(i,j) + \sum_{j=1}^{3} \varepsilon_{t,j}^{age} I_{age}(i,j) + \sum_{j=1}^{2} \varepsilon_{t,j}^{edu} I_{edu}(i,j) \qquad (15)$$

for all i, t, where $I_{edu}(i,j) = 1$ if education in labor group i is equal to j and zeros otherwise. Dummies by gender and age are defined the same way. Gender, age and education-specific shocks may be seen as capturing sectorial shocks and other aspects of the production process not explicitly modeled, which moves relative demands of labor inputs. For instance, an increase in women's productivity may be capturing an increase in the productivity of a female-intensive sector of the economy, such as the service sector.

Let the logarithm of the productivity processes $\varepsilon^g, \varepsilon^{age}, \varepsilon^{edu}$, the logarithm of TFP process $A \in \mathcal{A}$ and the logarithm of the mass of new born $p_0 \in P_0$ be AR1 stochastic processes. Furthermore, the mean of the distribution of the cost of acquiring education $\bar{k} \in \mathcal{K}$ for men and women, and women's housework $l_h \in L_h$ are deterministic processes with an AR1 structure.[23]

Let $G \equiv \mathcal{A} \times \mathcal{E}^g \times \mathcal{E}^{age} \times \mathcal{E}^{edu} \times P^0 \times \mathcal{K}^2 \times L_h$ be the state space for these variables.

Household Distribution

Since the distribution $p_{age,edu}$ is a state variable, one needs to define its set. From how $p_{age,edu}$ has been constructed in section 3.4, it follows that it depends on the series of $p_0 \in P^0$, $q^m \in [0,1]$ and $q^f \in [0,1]$ at the period of birth of each cohort which is alive.[24] $p_{age,edu}$ is therefore generated from the set $M \equiv P^{0J} \times [0,1]^{2J}$. Let $\mathcal{P}$ be the set of all admissible distributions $p_{age,edu}$ generated from the set M. The state space for each household is

$$S \equiv \{1, ..., J\} \times \{h, l\}^2 \times K \times G \times \mathcal{P} \times K^{(J-1)4}$$

---

[22]A type is gender and the idiosyncratic education shock for someone who is at the education stage, gender and education for someone who is at the marriage stage, age and education of husband and wife for households.

[23]Starting with initial values away from their steady states, these two variables will help make trends in hours by gender and education behave as in the data.

[24]q^f has not been directly used to construct $p_{age,edu}$ but it affects π_{e^m,e^f}^m.

The first dimensions of the state space, $\{1, ..J\} \times \{h, l\}^2 \times K$, contain the household's state variables: age, education of husband and wife, and asset holdings which belong to the set $K \equiv \left[\min\left(\underline{a_j}\right), \bar{a}\right]$, where $\underline{a_j}$ for $j = 1, ..., J - 1$ are age-specific lower bounds on the amount of assets implied by the borrowing constraint in equation (8). The remaining dimensions of the state space contain aggregate state variables that affect households' decisions through prices and expectations: the shocks, the distribution of households by age and education, and the distribution of assets across households grouped by age and educational composition.[25,26,27] Consistently with what was used in previous subsections, the aggregate state variables are collected in the set

$$\Omega \equiv G \times \mathcal{P} \times K^{4(J-1)}.$$

3.7 Equilibrium

Definition 1 *A recursive competitive equilibrium is composed of discounted values $M^g(e^g; \omega)$, college enrolment rates $q^g(\omega)$, matching probabilities $\pi^g_{e^m, e^f}(\omega)$ for each gender g, a value function at the time of forming a household $V(e^m, e^f; \omega)$, a value function for households $\mathbb{V}(e^m, e^f, j, a; \omega)$, consumption and assets functions $c(e^m, e^f, j, a; \omega)$ and $a'(e^m, e^f, j, a; \omega)$, labor functions $l^g(e^m, e^f, j, a; \omega')$ for each gender, a household' distribution function $p'_{age, edu}(e^m, e^f, j; \omega)$ and pricing functions $r(\omega)$ and $w(g, j, e; \omega)$ such that the following conditions are satisfied:*

1. *$q^g(\omega)$ is determined by (2). The matching probabilities $\pi^g_{e^m, e^f}(\omega)$ satisfy (5)–(6) and are consistent with the degree of assortative matching ρ in (7). Moreover, the pre-marriage discounted utilities M^f and M^m are defined by (3) and (4). The pre-labor value V is defined by (9).*

2. *The decision rules for consumption, labor and assets c, a', l^m and l^f, and the value function $\mathbb{V}$ solve the household problem in 3.3.*

3. *Capital and labor inputs satisfy equations (13)–(14).*

4. *Labor markets clear. i.e. equation (12) holds for all i.*

---

[25]Since there is no idiosyncratic risk across households belonging to the same group (defined by age and education of husband and wife), the state variable individual asset holdings is also part of the distribution of assets across all groups. While this repetition is not necessary and is avoided in the code, it is used here because it simplifies notation.

[26]The distribution of capital across groups only involves $J - 1$ age groups because at age 1 households hold zero assets.

[27]The distribution of idiosyncratic shocks κ affects enrollment rates q^g only through the mean $\bar{k}$. κ for each individual is therefore omitted from this characterization of the state space.

5. *The capital market clears:* $k = \sum_{e^m} \sum_{e^f} \sum_{j=1}^{J} a(e^m, e^f, j) p'_{age,edu}(e^m, e^f, j; \omega).$

6. *The goods market clears:* $c + k' - k(1-\delta) = y,$
   where $c = \sum_{e^m} \sum_{e^f} \sum_{j=1}^{J} c(e^m, e^f, j, \omega) p'_{age,edu}(e^m, e^f, j; \omega)$
   and $k' = \sum_{e^m} \sum_{e^f} \sum_{j=1}^{J} a'(e^m, e^f, j, \omega) p'_{age,edu}(e^m, e^f, j; \omega).$

7. *The distribution of households evolves as stated in section 3.4.*

4 Parametrization

It is useful to divide the parameters of this model into two categories: the production function parameters and all the other parameters. It is possible to calibrate the parameters belonging to the latter group by drawing on an extensive literature. On the other hand, the presence of heterogeneous groups of workers with their specific shocks makes the production side non-conventional. Therefore, some of its parameters need to be estimated. The procedure is detailed in appendix C. Parameter values are reported in Table 4.

4.1 Preferences, depreciation and survival probabilities

The utility specification for a generic household of age j, with education levels $e^m = \{h, l\}$ for men and $e^f = \{h, l\}$ for women is:

$$u_{j,e^m,e^f}(c, l^m, l^f) = \ln(c) + \gamma_{j,e^m}^{m} \frac{(1 - l^m)^{1-\sigma_e^m}}{1 - \sigma_e^m} + \gamma_{j,e^f}^{f} \frac{\left(1 - l^f - l_h\right)^{1-\sigma_e^f}}{1 - \sigma_e^f}. \tag{16}$$

σ_e^m and σ_e^f regulate the Frisch elasticity of labor supply. Heathcote et al. (2010) calibrate this parameter to be the same for both men and women and independently of education. I find that this way one cannot match how the male-female ratio of hours volatility evolves over time, an important fact for this paper. In particular, the Frisch elasticity of women has to be lower than that of men in order to make female hours less volatile than male ones.[28] Similarly, the labor supply of the highly educated has to be less elastic than that of the less educated

---

[28]This parametrization contrasts micro estimates which suggest that labor supply for women is more elastic than for men, see for instance Blundell and MaCurdy (1999). For the aim of this paper, it is however necessary to match the relative volatility between male and female total hours. The inability to reconcile micro estimates with the relative volatility between male and female total hours is a puzzle which may be interesting to study further in future research. It also highlights how the agents in this model are not "micro" agents in the sense that they are representative of their groups. Section 5.2.3 discusses the implications that this misspecification may have for the results of the paper.

in order to match hour volatility by education.[29] I calibrate the four parameters σ_e^m and σ_e^m with $e^m, e^f = \{h, l\}$ in order to match four moments:

a. The relative volatility of hours by gender,

b. The relative volatility of hours by education,

c. The relative volatility of hours by age (prime over young and older workers),

d. An average elasticity equal to 2.[30]

Table 2 shows how closely these moments are matched. To give a sense of the

Table 2: Calibration of elasticities

	Relative hours volatility			Average elasticity
	gender	age	education	
Targets	0.7	0.7	0.6	2.0
Model	0.6	0.8	0.6	2.0

Note: Relative hours volatility by age refers to the standard deviation of total hours over time of prime workers over the one of young and older workers.

elasticities implied by the chosen levels of σ_e^m and σ_e^m, table 3 reports the average labor supply elasticity by various labor groups. As can be seen in the table, with

Table 3: Average elasticities

	Total	Young	Prime	Old	High educ.	Low educ.
Total	2.0	2.2	1.9	2.2	0.70	2.5
Men	2.6	2.9	2.5	2.7	0.9	3.2
Women	1.2	1.3	1.1	1.3	0.4	1.4

Notes: Statistics are averages of the Frisch elasticities for each representative agent in a given group, weighted by their mass. For instance, total elasticity is equal to

$$\frac{\sum_{e^m} \sum_{e^f} \sum_{j=1}^{J} (1/\sigma_e^m (1-l^m(j,e^m,e^f))/l^m(j,e^m,e^f) + 1/\sigma_e^f (1-l^f(j,e^m,e^f))/l^f(j,e^m,e^f)) p_{age,edu}(e^m,e^f,j;\omega)}{2 \sum_{e^m} \sum_{e^f} \sum_{j=1}^{J} p_{age,edu}(e^m,e^f,j;\omega)}.$$

These statistics are computed period by period in the simulation, the numbers reported are the averages over the time.

the selected utility specification the model naturally matches the fact that prime

---

[29]The fact that these parameters are gender and education specific may reflect unmodeled labor contracts and matching frictions specific to sectors or types of worker.

[30]The business cycle literature typically chooses a high level of Frisch elasticities in order to have sufficient aggregate hours volatility, which tends to be significantly lower than that found in the data. The chosen level of 2, lies in between micro estimates and RBC calibrations. See Ljungqvist and Sargent (2011), Prescott et al. (2009) and Erosa et al. (2011) among others for a discussion of how extensive margins can be incorporated into a life cycle model to reconcile micro and macro labor supply elasticities.

age workers are less elastic than the young and the old.[31]

γ_{j,e^m}^m and γ_{j,e^f}^f are such that the model in steady state matches average hours per capita by gender, age and education between 2000 and 2007.

The discount factor β is equal to 0.99 and the depreciation rate of capital is 0.05. With these values, and given the parameters of the production function, the steady state capital output ratio is 2.8, the net interest rate 0.06 and the saving rate 0.14.

Survival probabilities ζ_j for $j = \{1, ..., J\}$ come from the National Center for Health statistics Vital statistics of the US, 1992. Since this paper focuses on active workers, attention is restricted to people from age 20 to 60. Therefore, no one can live for more than 40 periods and $\zeta_{40} = 0$. This is counterfactual, and it could affect the labor supply decision over the life-cycle. However, as shown in Section 5.2.3, the model is very successful at predicting relative hours volatilities over time by age group, which is crucial for the experiment of this paper. Therefore, adding retirement or more years of labor would further complicate this model without inducing an obvious improvement for what this model is constructed for.

4.2   Trends in the composition of labor

p_0, the share of new infants is modeled as an exogenous AR1 process with constant γ_{p_0}, persistence ρ_{p_0}, and variance $\sigma_{p_0}^2$. This quite simplistic way to model birth has the purpose of generating variable fertility that matches changes in the labor distribution by age; see Ríos-Rull (2001) for a discussion of alternative fertility regimes.

The initial level and the persistence ρ_h of the deterministic AR1 process for female housework l_h, is calibrated to replicate trends in the labor composition by gender and education. l_h is zero in steady state. Its initial level accounts for 36% of the average total working time for women, i.e. $l_h/(l^f + l_h) = 0.36$. The gradual decay of this variable "liberates women from the home" and increases their market labor supply and educational attainments. See Greenwood et al. (2005).

$\bar{k}^m$ and $\bar{k}^f$; the means of the cost distribution of acquiring education - equation (1) - are modeled as deterministic AR1 processes. Their steady state levels are such that the model matches the share of the highly educated by gender between age 25 and 29 in the period 1999-2007, which is $q^f = 0.36$, $q^m = 0.29$. Initial conditions and persistence are picked to replicate trends in the labor composition by gender and education between 1967 and 2010. The variance of the cost distri-

---

[31]Since these parameters are not age specific, it is remarkable that it is possible to match the relative volatility of hours by age. With these preferences, the Frisch elasticity of labor supply is $1/\sigma_e^g(1 - l^g)/l^g$. Thus, this elasticity is lower for prime age workers because they work more hours. See Jaimovich et al. (2009) for an alternative approach based on imperfect substitution in production between age groups and Dyrda et al. (2012) for a discussion of the two approaches.

bution of acquiring education for men v^m is set equal to 10, that for women, v^f is equal to 5. The lower variance for women implies more response to the wage premium, which is consistent with the increase in female enrolment rates over time.[32]

Following Heathcote et al. (2010), the degree of assortative matching in the marriage market, ϱ in equation (7), is set equal to 0.517.

Table 4 summarizes the parameters of the model, other than the ones of the productivity processes, which are set in C.

Table 4: Summary of Parametrization

Parameter	Moment to Match	Value
β	interest rate	0.99
δ	capital-output ratio	0.05
σ_l^f	relative group volatility and average labor supply elasticity	5.8
σ_h^f	relative group volatility and average labor supply elasticity	18.4
σ_l^m	relative group volatility and average labor supply elasticity	1.6
σ_h^m	relative group volatility and average labor supply elasticity	5.0
ζ_j	age-specific survival rates	see text
γ_{j,e^g}^g	average market hours by group	see text
$\rho_{p_0}, \gamma_{p_0}, \sigma_{p_0}$	share of 20-year-olds over population	0.97, 0.001, 0.002
$\bar{k}^m$	evolution of educational composition for men	see text
$\bar{k}^f$	evolution of educational composition for women	see text
l_h, ρ_h	evolution of female market hours	0.05, 0.975
v^m	sensitivity of educational composition to wage premium for men	10
v^f	sensitivity of educational composition to wage premium for women	5
ϱ	intra-family correlation of education at ages 25-35	0.517
σ	see appendix C	0.91
α	labor share	0.41
θ	see appendix C	-0.25

To run simulations, as initial conditions for the remaining state variables I take the values that solve the model for a steady state in which the level for the shocks is the average in the first 5 years of the sample (1967-1971).

Figure 3 shows actual versus predicted shares of hours by sex, age and education; the model does quite well at replicating these trends. However, shares by age cannot be matched perfectly because of the presence of migration, absent in the model, and because the initial age distribution does not perfectly match the data.[33]

---

[32]These variances are quite large in the sense that they imply a low elasticity of the enrolment rates to wage premia. This feature is convenient because it permits matching the evolution of enrolment rates with the evolution of the other parameter of the education cost distribution: the average cost of acquiring education $\bar{k}^g$, which can be easily changed in the counterfactual experiments aimed at removing trends in the labor distribution by education.

[33]The initial age distribution at the beginning of the sample is not stationary. One could

5 Quantitative analysis

5.1 Computation

The computation of this model is challenging because the state space is quite large: 805 variables, of which 324 are state variables. Large DSGE models can be handled by perturbation methods around the steady state.[34] This method generates two sources of inaccuracy: one due to the assumption of certainty equivalence and the other one due to the fact that policy functions are evaluated at points different than the deterministic steady state. When simulations remain fairly close to the deterministic steady state, the solution method is quite accurate even for models that induce large Jensen inequalities (Caldara et al. (2012)). This is not the case here because the model is simulated from starting conditions which are quite far from the steady state. To resolve this, I propose a new methodology which essentially consists of applying repeated local approximations over the entire transition path, between the initial conditions and the steady state.[35]

The algorithm is detailed in B and a brief summary of its logic is given here. The goal is to find the equilibrium path given initial conditions for the state vector, call x_0, and time series for the shocks (the actual time series of shocks is known to the computer programmer, but is unknown to the agents in the model). A path from x_0 to the steady state is drawn through the policy functions obtained by perturbation around the steady state. Then, new perturbations are computed backward along this path: from the proximity to the steady state back to the initial conditions. The policies approximated at the initial conditions are used to compute the next point, x_1. To compute x_2, a new path is drawn through the steady state policies, treating x_1 as initial conditions. New perturbations are computed along this path from the proximity to the steady state till x_1. The policies approximated at x_1 are used to compute the next point x_2. Then, the algorithm iterates until initial conditions coincide with the final period of the time series of the shocks. In practice, for all the applications tried so far, the maximum error from the equilibrium conditions is several orders of magnitude smaller with this method than with 2nd order perturbation of the steady state. Figure 19 compares the solution computed with these two methods applied to a version of the RBC model with full depreciation, for which the true solution has analytical form.[36]

---

impose it but it would be inconsistent with the other initial conditions found by solving for an initial steady state.

[34]See Ríos-Rull (1996) for an application of linear quadratic methods in a model that shares a similar OLG structure to the one in this paper.

[35]By limiting the local approximations to the transition path, the number of grid points does not increase with the number of state variables.

[36]Since the true solution of this model is linear in the logs of the variables, Taylor expanding on

18

5.2 Testing the model

Before carrying out the main experiments for which the model has been built, some tests are performed in order to get a sense of how satisfactory a description of the economy it provides, at least for the dimensions that are relevant for this study. The tests proposed below have direct implications for the counterfactual experiments aimed at quantifying the importance of labor reallocation to aggregate volatility. Therefore, the degree of success over these dimensions will help assess how reliable the outcomes of these experiments are.

5.2.1 Aggregate volatility trend

Since this model is restricted to match the observed changes in the labor composition and shocks come from disaggregated data, it is interesting to see whether it matches the typical aggregate statistics analyzed by the RBC literature initiated by Kydland and Prescott (1982) and whether it replicates the observed trends in output volatility.

Table 5, column 1 contains data and model standard deviations for the whole sample (67-10). It can be inferred that the model accounts for 1.29/1.45, or about 89 % of total volatility, more than the 60-85% typically accounted for by RBC models.[37] Two features of the model that contribute to the higher volatility compared to other RBC models are that it distinguishes between TFP and labor augmenting shocks, and that these shocks are less persistent than the Solow residual of an aggregate Cobb-Douglas production function (see table 13). The lower the persistence of shocks is, the lower are the offsetting wealth effects they induce and the greater is the labor response.

To quantify and compare the volatility slowdown of the mid 80s, columns 2 and 3 show the standard deviation of output over a period of high volatility (67-83) and over the period of moderate volatility (84-00). The model predicts the volatility slow down between the first and second sub-samples; the slowdown predicted by the model is actually larger than the one in the data, where volatility reduced by 38 log points $(\log(1.75) - \log(1.20))$, as opposed to the 51 log points predicted by the model. This prediction is actually in line with measurements of the moderation computed with aggregate data.[38] Strikingly, between the periods $84 - 00$ and $00 - 10$ the model predicts a $\log(1.22) - \log(.99)$ or 21 % increase

the logs of the variables gives the exact solution with both methods. To introduce approximation error, the Taylor expansion is computed on the variables in levels. This way the computed policy functions do not coincide with the true ones, but are only tangent around the point where the Taylor expansion has been taken.

[37]See for instance Prescott (1986). Volatility is measured as the standard deviation of log output minus its HP-trend with parameter 6.28.

[38]For instance, Arias et al. (2007) computes a moderation of about 50% using quarterly NIPA data. The statistic changes very little when using annual NIPA data. Another difference between

in volatility, which is the same size as that in the data $(\log(1.48) - \log(1.20)$, equal to 0.21). An alternative way to appreciate the extent to which the model can replicate aggregate volatility over time is offered by Figure 4, which shows the trend over time of aggregate output volatility.[39] This Figure shows that the model is not only successful at predicting the business cycle volatility decrease from the 80s mentioned, but it can also replicate the initial volatility increase from 67 to the early 70s. The Figure, however, highlights a misalignment in the volatility trend from the 90s.

Table 5: Standard deviation of output

	$67 - 10$	$67 - 83$	$84 - 00$	$00 - 10$
Data	1.45	1.75	1.20	1.48
Model	1.29	1.64	0.99	1.22

Notes: Statistics are computed after having HP-filtered the logs of the annual data with parameter 6.28.

5.2.2 Aggregate business cycle statistics

Table 6 reports standard deviations and correlations with output of consumption, investments and total hours. Consistently with the data, the model predicts that while consumption is less volatile than output, investments are much more volatile. Similarly to other RBC models, the model matches correlations rather well but under-predicts the volatility of hours. These statistics remained fairly stable over the whole sample and cannot be held responsible for the changes in aggregate volatility. See for instance Arias et al. (2007).

An interesting result of this model is that it reconciles the employment-productivity puzzle,[40] which lies at the root of an important critique to the RBC model. See Galí (1999). In this model, the correlation of labor productivity (output over total hours) with output is 0.003 and the correlation of labor productivity with hours is -0.078. This result relies on the fact that distinguishing between labor-specific

---

NIPA and CPS data is that using the latter dataset, output volatility started declining in the late 70s rather than in the 80s. Indeed, Blanchard and Simon (2001) suggest that the sharp and sudden volatility decline of the 80s masks a trend decline which started earlier. With CPS data, when the turbulent sub-sample is restricted to the period 70-76, the volatility slowdown is of 50%.

[39]Output volatility is measured as the standard deviation over three consecutive periods. This statistic is computed period by period to construct a time series. To highlight its trend the Figure plots the HP trend with smoothing parameter 6.28

[40]The near zero and often negative correlation between total hours and labor productivity found in the data. This fact is puzzling for RBC models, which predict this correlation to be very high. See Hansen and Wright (1992).

Table 6: Standard deviations relative to output and correlations

	Standard deviations		Correlations with output	
	Data	Model	Data	Model
Output	1	1	1	1
Consumption	0.8	0.6	0.90	0.61
Investment	4.6	4.4	0.94	0.87
Hours	0.9	0.5	0.85	0.91

Notes: Statistics for Consumption and Investment in the Data column are computed using annual NIPA data in logs and HP-filtered with parameter 6.28.

and TFP shocks, upward shifts in the labor demand schedules do not necessarily imply an increase in output over hours as with a simple Cobb-Douglas technology and homogeneous hours. See also Ballern and van Rens (2011). The discrepancy between labor productivity and output is also reflected in the labor share: the fit with the empirical labor share shown in Figure 5 is remarkable if compared with the constant labor share predicted with the benchmark Cobb-Douglas production function. Figure 6 shows actual versus predicted wages by sex, age and education.

5.2.3 Trends in Hours Volatility Ratios

Labor supply elasticities have been calibrated to match the average hours volatility ratios across gender, age and education. But does the model predict how these ratios have evolved over time? Figure 7, shows these trends in the data and in the model. Apart from a misalignment in the 70s, the model predicts the trendless path by gender, the decreasing ones by age (prime over young and old) and the increasing one by education.

This result is best understood in the light of the Frisch labor supply elasticity by gender, which, with the adopted utility function is $1/\sigma_e^g(1-l^g)/l^g$, decreasing in l^g.[41] The fact that hours volatility ratio by age is smaller than one is consistent with the fact that elasticity is lower for prime age workers, who work more hours. The fact that in the model hours volatility ratio by age is decreasing over time depends upon the fact that prime age hours increase over time.

The increasing hours volatility ratio by education is mainly due to an increase in the volatility of highly skilled wages. However, this effect is mitigated by the contrasting effect that the increase in white collar hours has on Frisch elasticities.[42]

---

[41]That elasticities decrease in labor is broadly consistent with the fact that groups that work more hours have less elastic labor supplies. E.g. males versus females, prime age versus closer to retirement, women without children versus women with children; see Reichling and Whalen (2012).

[42]This last effect on Frisch elasticities is important to match the trend in Figure 7. This

The model replicates the roughly constant trend in relative hours volatility by gender. Marriage plays a role in this: as female hours and education levels increase, men respond more to shocks because their offsetting wealth effect decreases as the wealth they earn becomes a smaller fraction of the whole family wealth. The contrary is true for wives: the more women work and the higher their education level, the less responsive they are to the income of men. This aspect is fostered by the increasing share of families with highly educated women who command a high share of family income.

The model also predicts that as women's labor input increases over time, their Frisch elasticities decline. This decline is consistent with the findings of Heim (2007) and Blau and Kahn (2007) on married women. These papers called for an explanation of this fact. Through a decomposition exercise, Heim (2007) and Bargain et al. (2012) find that the decline is not accounted for by demographics in the sense that the decline in elasticities occurred for all female sub-groups rather than because of the change in the demographic composition, holding sub-group elasticities constant. My model confirms their finding that elasticity declines for all women's sub-groups as shown in Figure 8; however, it suggests that the decline was indeed due to demographics.

It is worth mentioning that in this model female labor trends come from each woman working more hours. Instead, part of the increase in female total hours is due to participation decisions. This may imply an overstatement of the repercussion of female labor trends on their Frisch elasticities. However, it is likely that distinguishing between the intensive and the extensive margin would in fact enhance the results because participation elasticities decrease as participation rates increase: for example, if a large number of women have a market wage close to their reservation wage, then the participation elasticity will be high: a small increase in wages will lead to a number of women's market wages exceeding their reservation wage, and hence they will participate. As more women participate, the number of such women near the margin gets smaller and the participation elasticity will be small (Heim (2007)). Indeed, the drop in female total hours' elasticity between 1980 and 2000 is one third in the model, which is lower than the one half drop estimated by Blau and Kahn (2007). Thus, the contribution to

---

observation can be appreciated by observing Figure 9, produced with the following alternative utility function with constant Frisch elasticities:

$$u_{j,e^m,e^f}(c, l^m, l^f) = \ln(c) - \left(\gamma_{j,e^m}^m \frac{(l^m)^{1+\sigma_e^m}}{1+\sigma_e^m} - \gamma_{j,e^f}^f \frac{(l^f + l_h)^{1+\sigma_e^f}}{1+\sigma_e^f} \right). \qquad (17)$$

In this case, the relative volatility by education has a much steeper and counterfactual trend. This exercise also shows how matching trends in relative hours volatilities is not obvious and it is a useful exercise to choose among alternative utility specifications.

the aggregate volatility slowdown due to the decline in these elasticities may have been even larger than what is found in this paper.

These results are reassuring given the important role played by these statistics for the questions of this paper.[43] Indeed, part of the merit of the structural approach adopted is being able to explain how these statistics evolve over time, so that then one can rely on the model to predict how the agents would have responded in the counterfactual experiments.

5.3 Counterfactual Experiments

The counterfactual experiments consist of changing the exogenous low frequency trends in the amount of female housework l_h, the birth rate dynamics p_0 and the cost of acquiring education–$\bar{k}^m$ and $\bar{k}^f$–so that shares remain closer to their steady state levels, while maintaining all the shocks as in the original simulation.[44]

5.3.1 Removing all trends

Figure 10 shows hours shares by gender, age and education in the counterfactual and original simulations. As can be seen from comparing with the original simulation, most of the trends in these shares have been removed. Table 7 contains the standard deviation of output during the sub-samples of interest: the period before the Great Moderation 1967-1983, and the period of the Great Moderation: 1984-2000 (until before the 2001 recession). It is also instructive to focus on the initial 10 periods of especially high volatility, 1967-1976 (column 1), and the last part of the sample: 2000-2010 (last column).[45] As can be deduced from the table, aggregate volatility in the counterfactual is $\log(1.53) - \log(1.81)$, or approximately 17% lower than in the original simulation in the 70s, essentially unchanged in the 80s and 90s and 5% more volatile in the last decade. To get a visual sense of how volatility is affected over time, Figure 11, first panel, shows the trend of actual and counterfactual cyclical volatility measured as a 3-year roll over standard de-

---

[43]In particular, it is reassuring that the model is consistent with these trends, given that the average relative volatilities have been matched through a calibration that contrasts micro estimates. Presumably, additional structure missing in this model should be included to reconcile relative volatilities with micro estimates on labor supply. However, this missing structure does not seem to prevent this model from predicting the endogenous evolution of these trends.

[44]Alternatively, one could try to keep share of hours at their initial levels. However, this would imply a different calibration of the parameters of the exogenous fertility rates, educational costs and female housework, thereby implying a different steady state for the model. Then, the differences between the original and counterfactual simulation would depend not only on a different transition, but also on a different calibration, making the comparison between the original and counterfactual simulation less transparent.

[45]This last sample starts in year 2000, rather than 2001, to capture the 2001 recession: starting in 2001, the fact that output contracted would not be apparent.

Table 7: Standard deviation of output over time

	$67-76$	$67-83$	$84-00$	$00-10$
Original simulation	1.81	1.64	0.99	1.22
Counterfactual simulation	1.53	1.42	1.00	1.28

Notes: statistics are computed on the deviation from the HP-trend of the logs.

viation of output. The fact that counterfactual volatility is slightly higher in the 2000s (and from the late 90s as indicated in Figure 11) seems interesting: since in the counterfactual, share of hours are roughly constant at their steady state levels, counterfactual volatility in the last decade can be taken as a prediction of the future volatility that we should expect as share of hours by gender, age and education are converging to a steady state. From this prediction one may hazard that a return to the Great Moderation is unlikely; in fact, given the size of the shocks, volatility should slightly increase rather than decline.[46]

To get a concise statistic that quantifies the amount of the Great Moderation explained, I follow J-S and proceed as in section 5.2.1, comparing the standard deviations in the first sub-sample (1967-1983) and in the second one (1984-2000). Between the two sub-samples, aggregate output volatility decreased by $\log(1.64) - \log(.99) = 50.47$ log points. Had the shares remained stable as in the counterfactual, we would have observed a reduction in volatility of $\log(1.42) - \log(1.00) = 35.07$ log points. Therefore, these demographic changes account for (50.47-35.07)/50.47 or 30.53% of the moderation in output.

Figure 11, second panel, shows the volatility of total hours; as can be seen, labor reallocation accounts for a sizable part of the high volatility in the 70s.[47] This pattern can be traced back to the elasticities of labor supply shown in Figure 12. As can be observed, in the counterfactual elasticities were lower at the beginning of the sample. Instead, they become higher than in the original simulation from the late 90s for men, the young, and workers with low education. Thus, these groups are likely to be the ones responsible for the fact that counterfactual volatility becomes higher from the late 90s. One issue with this exercise is that removing female housework, directly affects female Frisch elasticities (see Section 5.2.3). Thus, an alternative exercise is also carried out where, instead of removing female housework, the parameter γ^f that multiplies female leisure in the utility

---

[46]It should be noted that this prediction abstracts from the aging of the population on top of the retirement of the baby boomers as in the model life expectancy is constant.

[47]In the original (counterfactual) simulation the standard deviation of the percentage deviation of total market hours from the trend is 0.90 (0.80) between 67 and 76, 0.84 (0.75) between 67 and 83, 0.44 (0.44) between 84 and 2000 and 0.48 (0.49) between 2000 and 2010. Therefore, hours volatility is 12 % lower between 67 and 76, unaffected in the period of the Great Moderation and slightly (2 %) higher between 2000-2010.

function (16) is rescaled. Other things equal, this parameter does not affect Frisch elasticities. Indeed labor supply elasticity is slightly higher at the beginning but the business cycle results are not that much affected: hours volatility between 67 and 76 is now 11% rather than 12% lower. The model accounts for 25% rather than 30% of the Great Moderation in output.

Lastly, Figure 13 shows how labor reallocation played an important role in output levels: in the counterfactual, output is much higher in the early part of the sample: this depends on the fact that in the counterfactual the labor distribution is roughly constant at the end-of-sample levels, with a higher level of education and prime age workers, and more hours worked by women.[48]

5.3.2 Removing trends one by one

What is the importance of each of the three factors? Table 8, lines 2, 3 and 4 report the outcomes of a counterfactual experiment where only one of the long-run trends is removed. Table 9, column 1, reports the various contributions to the

Table 8: Standard deviation of output over time

	$67 - 76$	$67 - 83$	$84 - 00$	$00 - 10$
Original simulation	1.81	1.64	0.99	1.22
Gender	1.77	1.61	1.01	1.25
Age	1.75	1.59	0.98	1.25
Education	1.62	1.50	1.00	1.27

moderation; the other columns report the evolution of the volatility ratio relative to the original simulation.

These numbers suggest that the major contributors to changes in volatility were changes in the labor composition by education, followed by gender and age.

The results do not change qualitatively and all the conclusions remain true with different parameter values concerning the complementarity of labor groups in the production function (where values ranging between 0.7 and 1 have been considered), the average elasticity of labor supply (where values ranging between 1.8 and 3 have been considered) and changing the age classes, moving the young-to-prime age threshold to 34 and the old threshold to 50.

---

[48]See Marimon and Zilibotti (1998) for an analysis of the importance of reallocation to growth: they find that sectoral effects account for more than 80% of the long-run differentials across countries and industries in employment growth.

Table 9: Standard Deviation ratios

	Great Moderation	Counterf.-Actual St.Dv. 67-76	Counterf.-Actual St.Dv. 84-00	Counterf.-Actual St.Dv. 00-10
All	30.5	-15.4.	0.7	4.6
Gender	8.7	-2.3	2.3	2.3
Age	6.2	-3.4	-0.3	2.4
Education	19.9	-10.7	1.2	1.2

Notes: numbers are expressed in percentage terms. Great Moderation is a measure of the size of the volatility reduction that is accounted for by changes in the composition of labor. Counterfactual-Actual St.Dv. measures the percentage difference between output and counterfactual output volatility.

5.3.3 Comparison with Jaimovich and Siu (2009)

The result confirms the finding of J-S that the age composition has an impact. However it suggests that its effect is not as important. A similar answer is found when only looking at hours: counterfactual volatility by age is 4.4% lower in the first decade, it is also 4% lower in the 80s and 1% higher from the mid 90s. This is nowhere near to explaining the 20% of the Great Moderation in hours as found by J-S in their accounting exercise. In fact, the role of age is smaller than the one played by educational composition and comparable to the one of gender composition.

There is a simple intuition for why changes by education and gender play a stronger role, likely not to be captured by reduced-form accounting exercises such as the one in Section 2.1: educational and gender trends have important repercussions on labor supply elasticities; age trends do not. This is shown in Figures 14 and 15: it is evident how in the gender and educational counterfactuals, labor elasticities are much lower in the first sub-sample, thereby lessening the effect of shocks when turbulence is at its peak. This is not the case for demographic trends by age (Figure 16). The reason for this is that education and gender composition induce large changes in the intensive margin, increasing labor supply of women and of agents with high education, thereby reducing labor supply elasticities as explained in Section 5.2.3. Instead, changes in the age composition are essentially changes in the extensive margin (number of people in each age bracket).

The effect of age is also lessened for the following reason: removing age trends implies that there are fewer young in the first part of the sample, which should imply less volatility. This effect is mitigated by the fact that the fewer young, the higher the relative volatility by gender and education.[49] This result is consistent

---

[49]Indeed, in the absence of age trends, the average relative volatility by gender and education is 0.64 and 0.61; higher than those in the original simulation, 0.63 and 0.56.

with the fact that relative volatility differences by gender are more pronounced for the young as pointed out by Gomme et al. (2005), in Table 10.[50]

Notwithstanding the intuitive explanation above, the reader may be left wondering whether the result depends on the fact that the model generates counterfactual data that understate the role of age relative to the true data. To see if this is the case, I adapt the panel regression analysis of J-S to the single-country data generated by the model. If the regression on the simulated data correctly detects a minor role for the age composition, then it is likely that the model understates the role of the age composition. If instead the regression overstates the role of age relative to what it is in the model, so that the data generated by the model are consistent with the results found by J-S, then one may not conclude that the model understates the role of the age composition.[51]

The regression considered is:

$$\sigma_t = \alpha + \gamma \mathbf{share}_t + \varepsilon_t \tag{18}$$

where σ_t is a measure of output volatility at year t generated by the model and $\mathbf{share}_t$ is the fraction of the population share which is not in its prime age (young and old).[52, 53] Following J-S, results are reported for heteroskedasticity and two-period autocorrelation robust standard errors constructed using the Newey-West estimator. See Table 10. The coefficient γ is significant and it implies a stronger

Table 10: Regression analysis: age

	Coef.	Newey West Std. Err	$P > |t|$
γ	.038	.015	.019
α	-.002	.005	.75

role for age than the true one in the artificial economy: when the independent variable **share** moves from its first sample average (41%) to its second sample average (32%), the relative change in predicted volatility is 23%, which is quite close to the results found by J-S. This result suggests that the regression overstates the role of age.

---

[50]Furthermore, McGrattan and Rogerson (2004) find that the decline in hours worked by older individuals is primarily accounted for by a decline in the hours worked by males, whereas the increase in hours of prime-age workers comes primarily from females.

[51]This case would not strictly imply that the analysis of J-S is biased as they are also exploiting the cross-country panel dimension, which here is ignored.

[52]σ_t is constructed by taking a 5-period standard deviation of the residual between log output and its HP-filter with parameter 6.28.

[53]Reverse causality is not an issue because population shares are exogenous by model construction.

A possible explanation is that the regression may be capturing the overall effects of demographic changes, including those in gender and education, which are correlated with the changes by age. Is it possible to verify this conjecture by simply including the labor composition by gender and education as regressors? It depends on whether they are actually exogenous or endogenous. Indeed, to motivate the structural methodology, it has been mentioned that movements in the labor composition by gender and education are partly endogenous, thereby invalidating the results of a regression analysis. It is possible to assess the importance of these endogenous implications in the model and see whether in practice a regression would be biased.[54] With this aim, I augment equation (18) by including regressors for the labor composition by gender and education (including the share of males and low education over total hours):

$$\sigma_t = \alpha + \gamma_{age}\textbf{share}_{age_t} + \gamma_{edu}\textbf{share}_{edu_t} + \gamma_{gend}\textbf{share}_{gend_t} + \varepsilon_t. \qquad (19)$$

The results from the estimation of equation (19) are reported in table 11. The

Table 11: Regression analysis: age, gender and education

	Coef.	Newey West Std. Err	$P > \lvert t \rvert$
γ_{age}	-.022	.026	.399
γ_{edu}	.002	.065	.976
γ_{gend}	.094	.068	.178
α	-.039	.013	.070

coefficient for age is no longer positive, and all the coefficients tend to neutralize each other as none of them are significant. This exercise suggests that endogeneity plays a role and motivates further the structural methodology adopted in this paper.

6   Conclusion

This paper has documented that while the composition of the labor force by gender, age and education has changed substantially over the last 40 years, the relative volatility of market hours between these groups differs systematically.

---

[54]There are several channels through which the labor force composition is affected by aggregate risk: female income is less sensitive to aggregate risk than male income, so the gender composition of hours can be affected by the level of aggregate risk. For a similar argument, the hours composition by education and educational choices are affected by aggregate risk. Furthermore, the effect of aggregate risk on the hours composition may vary over time with the movements in labor supply elasticities that this paper has uncovered, which affect within-group variances.

These facts lead to the conjecture of this study: that changes in the composition of labor have a causal impact on the evolution of aggregate volatility over time.

To take into account the possible change in behavior induced by these demographic changes and the fact that they might be affected by aggregate fluctuations, a general equilibrium model of the business cycle with overlapping generations, educational choices, intrafamily time allocation, and savings, has been developed.

The model accounts for the initial increase, subsequent slowdown and recent surge in aggregate volatility and is consistent with several cross-sectional facts. The demographic changes considered play a non-negligible role in aggregate volatility. Changes in education levels have the greatest impact of all the factors considered. The role of age composition changes in business cycle fluctuations is relevant but greatly curtailed compared with that previously found through reduced-form analysis by J-S. Key to this discrepancy is the endogenous effect that demographic trends have on labor supply elasticities: they decline with the increase in the labor supply of highly educated workers and females. While this paper highlights the business cycle implications of this newly found relationship, it may prove fruitful to assessing its implications for a variety of public debates, ranging from tax and welfare reforms to the regulation of institutional features of labor markets.

One challenge has been to find a solution method for this large model of more than 800 variables which guarantees sufficient precision over the dramatic demographic transition path that has characterized the last 40 years. This has been done by developing a technique that can be applied to a wide range of dynamic stochastic general equilibrium models, which essentially consists of applying perturbation methods at many points along the equilibrium path.

This computational innovation may enable researchers to extend this transitional analysis to a wide array of aspects neglected here, which are evolving in modern societies. To name a few, marital status, number of children and retirement decisions affect labor supply elasticities. Furthermore, the model abstracts from labor supply participation decisions, sectors and occupations.

References

Abraham, K., Shimer, R., 2002. Changes in unemployment duration and labor force attachment. In: Krueger, A., Solow, R. (Eds.), The Roaring Nineties: Can full employment be sustained? New York: Russell Sage Foundation, pp. 367–420.

Arellano, M., Bond, S., April 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. Review of Economic Studies 58 (2), 277–297.

Arias, A., Hansen, G. D., Ohanian, L. E., July 2007. Why have business cycle fluctuations become less volatile? Economic Theory 32 (1), 43–58.

Attanasio, O., Low, H., Sanchez-Marcos, V., 2008. Explaining changes in female labor supply in a life-cycle model. American Economic Review 98 (4).

Attanasio, O., Weber, G., December 1995. Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey. Journal of Political Economy 103 (6), 1121–57.

Ballern, A., van Rens, T., June 2011. Skill-Biased Technological Change and the Business Cycle. CEPR Discussion Papers (8410).

Bargain, O., Orsini, K., Peichl, A., Jul. 2012. Comparing labor supply elasticities in europe and the us: New results. IZA Discussion Papers 6735, Institute for the Study of Labor (IZA).

Baxter, M., King, R., November 1999. Measuring business cycles: Approximate band-pass filters for economic time series. Review of Economics and Statistics 81 (4), 575–93.

Blanchard, O., Kahn, C., July 1980. The solution of linear difference models under rational expectations. Econometrica 48 (5), 1305–11.

Blanchard, O., Simon, J., 2001. The long and large decline in U.S. output volatility. Brookings Papers on Economic Activity 32 (1), 135–174.

Blau, F. D., Kahn, L. M., 2007. Changes in the labor supply behavior of married women: 1980–2000. Journal of Labor Economics 25, 393–438.

Blundell, R., MaCurdy, T., October 1999. Labor supply: A review of alternative approaches. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics. Vol. 3 of Handbook of Labor Economics. Elsevier, Ch. 27, pp. 1559–1695.

Caldara, D., Fernandez-Villaverde, J., Rubio-Ramirez, J., Yao, W., April 2012. Computing DSGE models with recursive preferences and stochastic volatility. Review of Economic Dynamics 15 (2), 188–206.

Castro, R., Coen-Pirani, D., 02 2008. Why have aggregate skilled hours become so cyclical since the mid-1980s? International Economic Review 49 (1), 135–185.

Choi, S., Ríos-Rull, J.-V., 2009. Understanding the dynamics of the labor share: the role of non-competitive factor prices. Annales d'Economie et de Statistique (95-96), 251–277.

Clark, K. B., Summers, L., 1981. Demographic differences in cyclical employment variation. Journal of Human Resources 16.

Den Haan, W. J., Marcet, A., January 1990. Solving the stochastic growth model by parameterizing expectations. Journal of Business & Economic Statistics 8 (1), 31–34.

Dyrda, S., Kaplan, G., Ríos-Rull, J.-V., Mar. 2012. Business cycles and household formation: The micro vs the macro labor elasticity. NBER Working Papers 17880, National Bureau of Economic Research, Inc.

Erosa, A., Fuster, L., Kambourov, G., July 2011. Towards a micro-founded theory of aggregate labor supply. Working Papers 2011-13, Instituto Madrile?o de Estudios Avanzados (IMDEA) Ciencias Sociales.

Galí, J., March 1999. Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? American Economic Review 89 (1), 249–271.

Gomme, P., Klein, P., April 2011. Second-order approximation of dynamic models without the use of tensors. Journal of Economic Dynamics and Control 35 (4), 604–615.

Gomme, P., Rogerson, R., Rupert, P., Wright, R., 2005. The business cycle and the life cycle. In: NBER Macroeconomics Annual 2004, Volume 19. NBER Chapters. National Bureau of Economic Research, Inc, pp. 415–592.

Greenwood, J., Seshadri, A., Yorukoglu, M., 01 2005. Engines of liberation. Review of Economic Studies 72 (1), 109–133.

Hansen, G. D., Wright, R., 1992. The labor market in real business cycle theory. Quarterly Review (Spr), 2–12.

Heathcote, J., Storesletten, K., G.L.Violante, 08 2010. The macroeconomic implications of rising wage inequality in the United States. Journal of Political Economy 118 (4), 681–722.

Heim, B. T., 2007. The incredible shrinking elasticities: Married female labor supply, 1978–2002. Journal of Human Resources 42 (4).

Jaimovich, N., Pruitt, S., Siu, H. E., Jan. 2009. The demand for youth: Implications for the hours volatility puzzle. NBER Working Papers 14697, National Bureau of Economic Research, Inc.

Jaimovich, N., Siu, H. E., June 2009. The Young, the Old, and the Restless: Demographics and Business Cycle Volatility. American Economic Review 99 (3), 804–26.

Janiak, A., Monteiro, P. S., 2011. Towards a quantitative theory of automatic stabilizers: the role of demographics. Tech. rep.

Johnson, M., Keane, M., 2007. A dynamic equilibrium model of the U.S. wage structure, 1968–1996, mimeo.

Judd, K., Maliar, L., Maliar, S., Aug. 2009. Numerically stable stochastic simulation approaches for solving dynamic economic models. NBER Working Papers 15296, National Bureau of Economic Research, Inc.

Judd, K. L., Maliar, L., Maliar, S., Nov. 2012. Merging simulation and projection approaches to solve high-dimensional problems. NBER Working Papers 18501, National Bureau of Economic Research, Inc.

Katz, L. F., Autor, D. H., October 1999. Changes in the wage structure and earnings inequality. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics. Vol. 3 of Handbook of Labor Economics. Elsevier, Ch. 26, pp. 1463–1555.

Katz, L. F., Freeman, R., 1994. Rising wage inequality: The United States vs. Other Advanced Countries. Russell Sage Foundation for NBER, NY, pp. 29–62.

Kim, C.-J., Nelson, C., November 1999. Has the U.S. economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle. The Review of Economics and Statistics 81 (4), 608–616.

King, M., Ruggles, S., Alexander, J., Flood, S., Genadek, K., Schroeder, M., Trampe, B., Vick, R., 2010. Integrated public use microdata series, current population survey: Version 3.0. Minneapolis: University of Minnesota.

Krusell, P., Ohanian, L. E., R?os-Rull, J.-V., Violante, G. L., September 2000. Capital-skill complementarity and inequality: A macroeconomic analysis. Econometrica 68 (5), 1029–1054.

Krusell, P., Smith Jr, A.-A., October 1998. Income and wealth heterogeneity in the macroeconomy. Journal of Political Economy 106 (5), 867–896.

Kydland, F. E., Prescott, E. C., November 1982. Time to build and aggregate fluctuations. Econometrica 50 (6), 1345–70.

Leandoacute;n-Ledesma, M. A., McAdam, P., Willman, A., September 2010. Identifying the elasticity of substitution with biased technical change. American Economic Review 100 (4), 1330–57.

Ljungqvist, L., Sargent, T., May 2011. A labor supply elasticity accord? American Economic Review Papers and Proceedings 101 (3), 487–491.

Lugauer, S., 2012a. Demographic change and the great moderation in an overlapping generations model with matching frictions. Macroeconomic Dynamics forthcoming.

Lugauer, S., 2012b. Estimating the effect of the age distribution on cyclical output volatility across the united states. The Review of Economics and Statistics forthcoming.

Lugauer, S., Redmond, M., 2012. The age distribution and business cycle volatility: International evidence. Economics Letters forthcoming.

Marcet, A., Lorenzoni, G., 1999. The parameterized expectation approach: some practical issues. In: Marimon, R., Scott, A. (Eds.), Computational Methods for the Study of Dynamic Economies. pp. 143–171.

Marimon, R., Zilibotti, F., January 1998. 'Actual' versus 'virtual' employment in Europe is Spain different? European Economic Review 42 (1), 123–153.

McConnell, M., Perez-Quiros, G., December 2000. Output fluctuations in the United States: What has changed since the early 1980?s? American Economic Review 90 (5), 1464–1476.

McGrattan, E., Rogerson, R., 2004. Changes in hours worked, 1950–2000. Quarterly Review (Jul), 14–33.

Prescott, E. C., 1986. Theory ahead of business cycle measurement. Quarterly Review (Fall), 9–22.

Prescott, E. C., Rogerson, R., Wallenius, J., January 2009. Lifetime aggregate labor supply with endogenous workweek length. Review of Economic Dynamics 12 (1), 23–36.

Ravn, M. O., Uhlig, H., 2002. On adjusting the Hodrick-Prescott filter for the frequency of observations. The Review of Economics and Statistics 84 (2), 371–375.

Reichling, F., Whalen, C., Oct. 2012. Review of estimates of the frisch elasticity of labor supply. CBO Working Papers 18501, Congressional Budget Office.

Ríos-Rull, J.-V., July 1996. Life-Cycle Economies and Aggregate Fluctuations. Review of Economic Studies 63 (3), 465–89.

Ríos-Rull, J.-V., 2001. Population changes and capital accumulation: The aging of the baby boom. The B.E. Journal of Macroeconomics (1), 7.

Schmitt-Grohe, S., Uribe, M., January 2004. Solving dynamic general equilibrium models using a second-order approximation to the policy function. Journal of Economic Dynamics and Control 28 (4), 755–775.

Smets, F., Wouters, R., June 2007. Shocks and frictions in U.S. business cycles: A Bayesian DSGE approach. American Economic Review 97 (3), 586–606.

Stock, J. H., Watson, M. W., 2003. Has the business cycle changed and why? In: NBER Macroeconomics Annual 2002, Volume 17. NBER Chapters. National Bureau of Economic Research, Inc, pp. 159–230.

Zoghi, C., 2010. Measuring labor composition: A comparison of alternate methodologies. In: Labor in the New Economy. NBER Chapters. National Bureau of Economic Research, Inc, pp. 457–485.

7   Figures
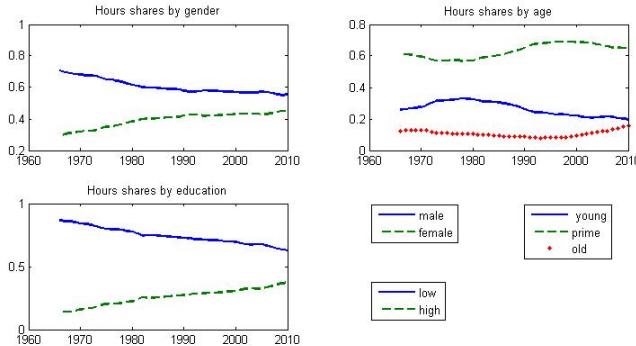
Figure 1: Share of hours by gender, age and education
Notes: young workers range from 15 to 29 years old. Prime age ranges from 30 to 55. The old are those 56 and above. By high education is meant at least four years of college.
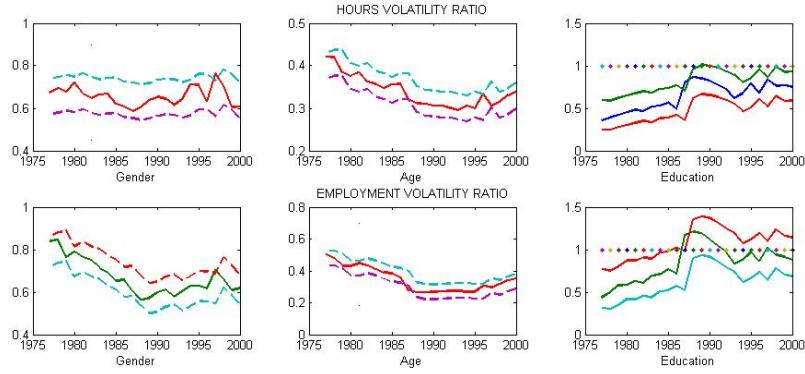
Figure 2: Hours and employment volatility ratio by gender, age and education.
Notes: In each period t, the figure plots the ratio of the standard deviation of hours or employment between genders, education groups and age groups (prime versus young and old), over a period of 19 years centered at year t. Confidence intervals are calculated assuming that the time series follow an AR 1 process.
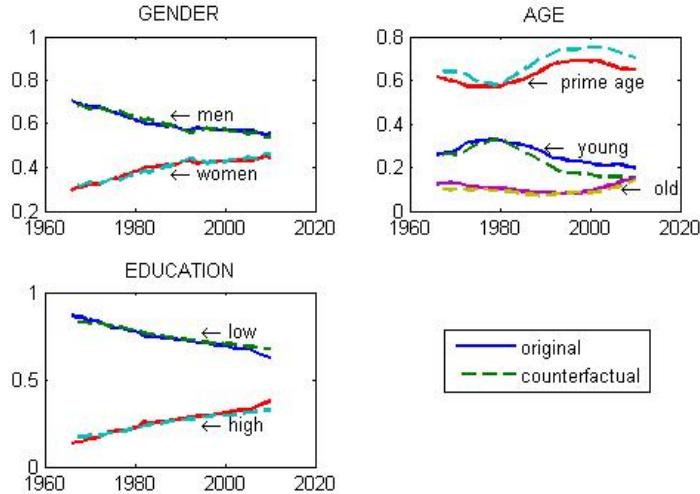

Figure 3: Data vs model share of hours
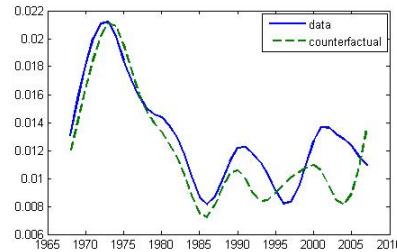
Figure 4: Output volatility over time, data versus simulation

35
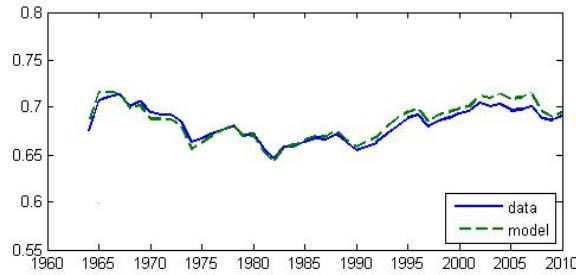
Figure 5: Labor share model vs data
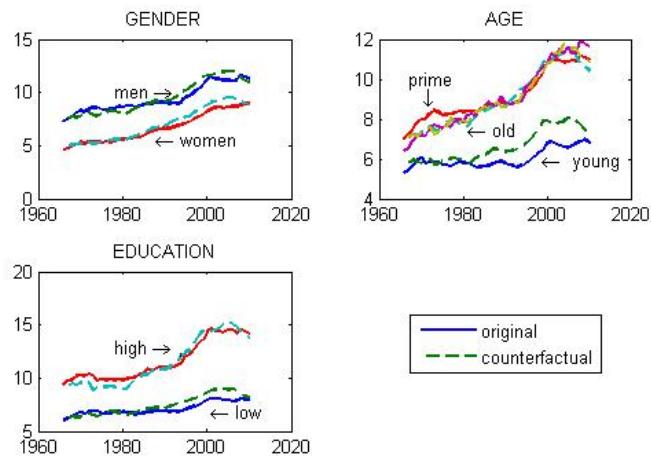
Figure 6: Wages model vs data
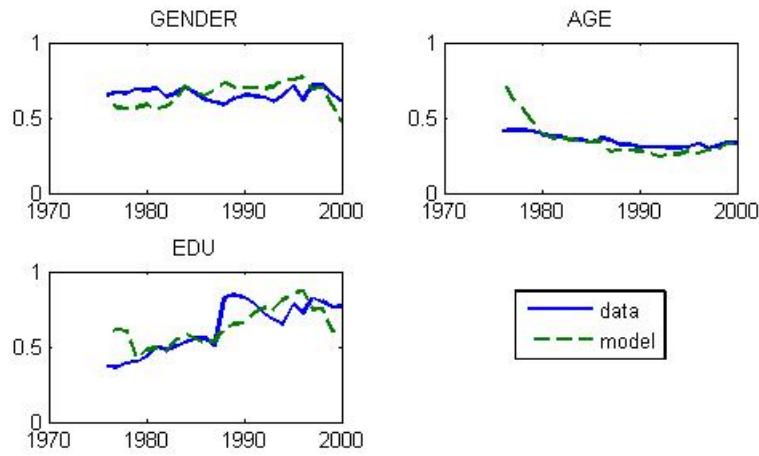
Figure 7: Data vs model relative hours volatility
Note: In each period t, the figure plots the ratio of the standard deviation of hours over a period of 19 years centered at year t.
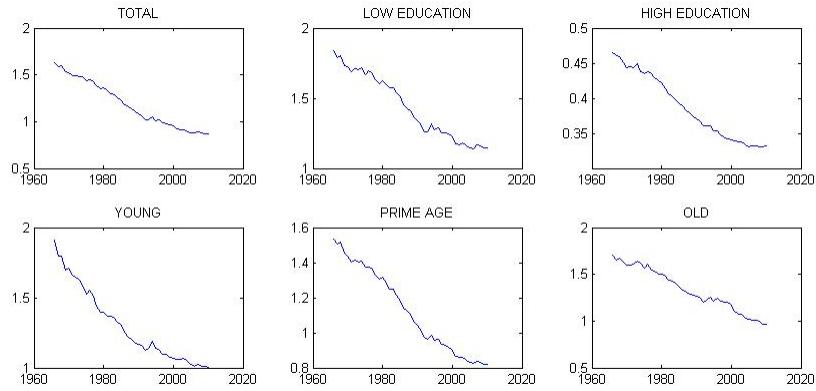
Figure 8: Women's elasticities,

Note: statistics are averages of the elasticities for each representative agent in a given group, weighted by their mass. For instance, total elasticity is equal to $\frac{\sum_{e^m}\sum_{e^f}\sum_{j=1}^{J}(1/\sigma_e^f(1-l^f(j,e^m,e^f))/l^f(j,e^m,e^f))p_{age,edu}(e^m,e^f,j;\omega)}{\sum_{e^m}\sum_{e^f}\sum_{j=1}^{J}p_{age,edu}(e^m,e^f,j;\omega)}$.
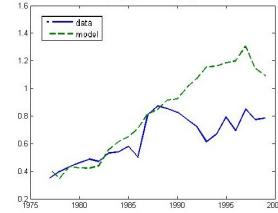

Figure 9: Data vs model relative hours volatility (high over low education) with alternative preferences
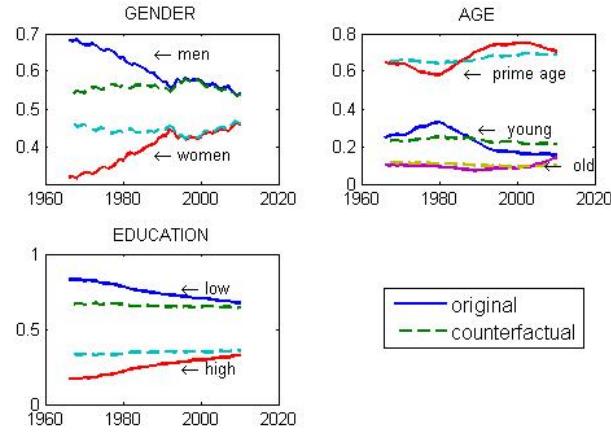

Figure 10: Shares of hours, original versus counterfactual simulation
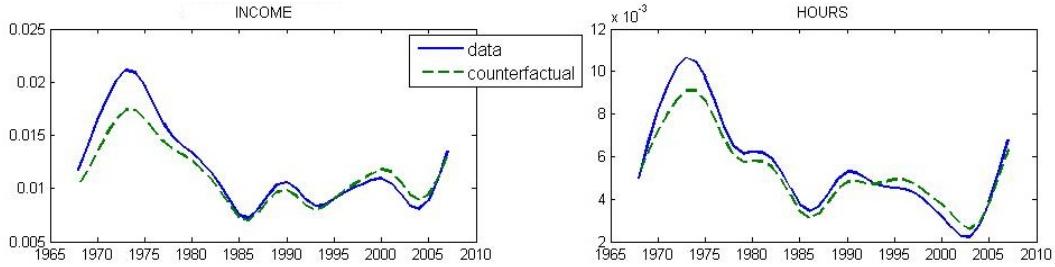
Figure 11: Aggregate income and hours volatility, original versus counterfactual simulation
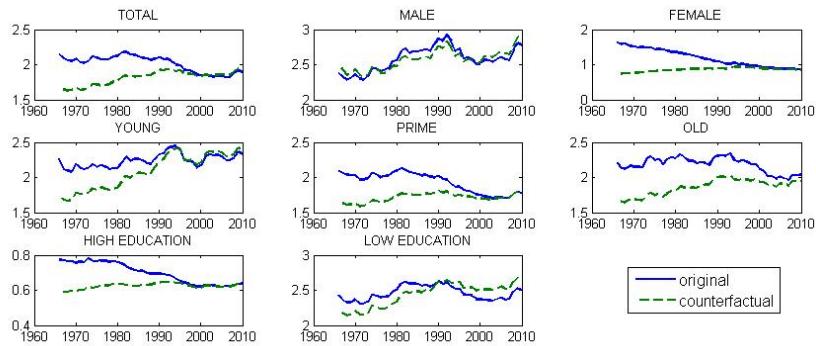

Figure 12: Average elasticities, original versus counterfactual simulation removing all trends

Note: statistics are averages of the elasticities for each representative agent in a given group, weighted by their mass. For instance, total elasticity is equal to

$$\frac{\sum_{e^m}\sum_{e^f}\sum_{j=1}^{J}(1/\sigma_e^m(1-l^m(j,e^m,e^f))/l^m(j,e^m,e^f)+1/\sigma_e^f(1-l^f(j,e^m,e^f))/l^f(j,e^m,e^f))p_{age,edu}(e^m,e^f,j;\omega)}{2\sum_{e^m}\sum_{e^f}\sum_{j=1}^{J}p_{age,edu}(e^m,e^f,j;\omega)}.$$
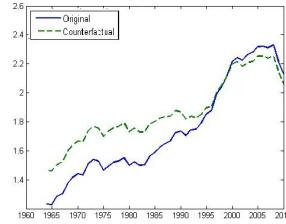

Figure 13: Output, original versus counterfactual simulation

38
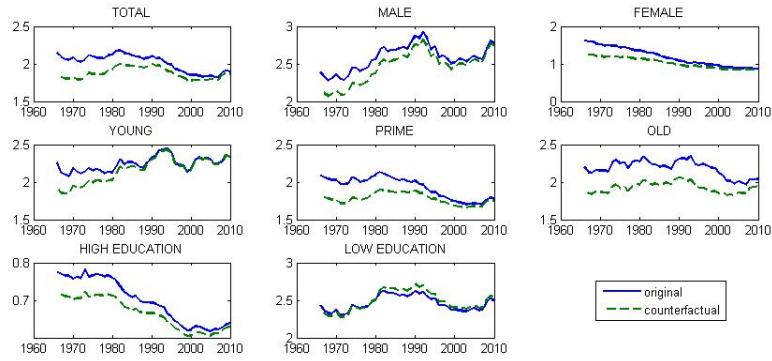
Figure 14: Average elasticities, original versus counterfactual simulation removing educational trends.
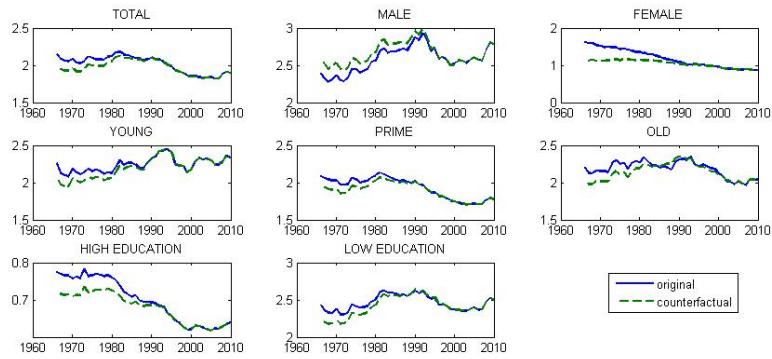
Figure 15: Average elasticities, original versus counterfactual simulation removing gender trends.
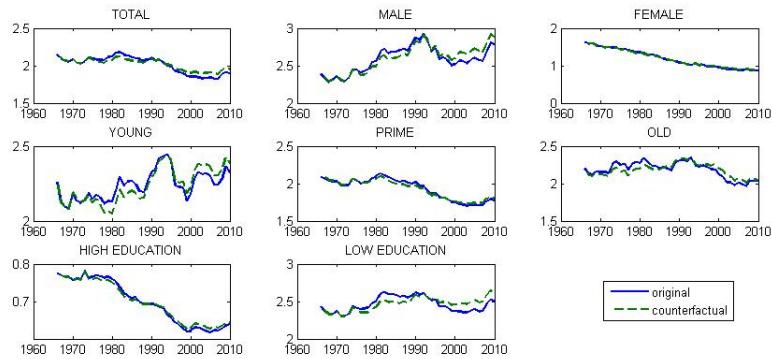

Figure 16: Average elasticities, original versus counterfactual simulation removing age trends .
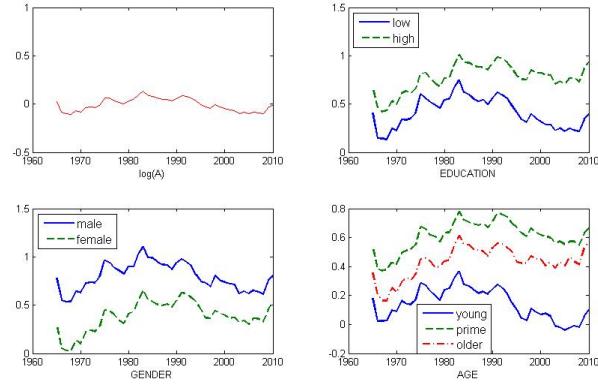
Figure 17: TFP and labor productivities
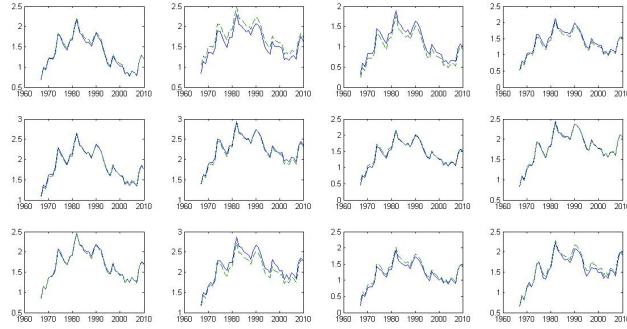
Figure 18: Group factors versus labor-specific shocks z.
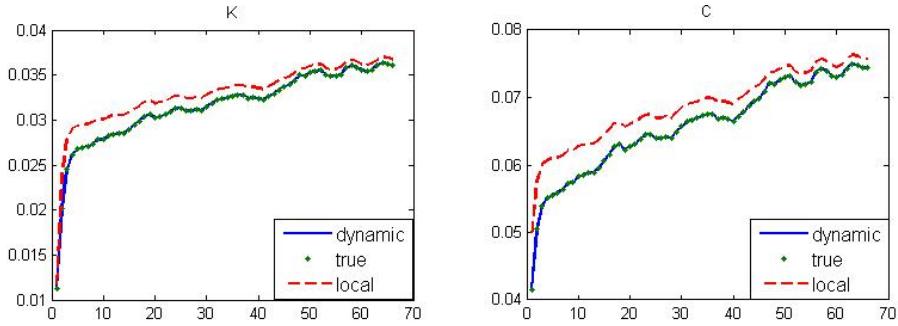
Figure 19: Capital and consumption from the analytical example of Appendix B compared with the solution computed with local and dynamic perturbation.

A   Data

The data series for aggregate consumption are from the Bureau of Economic Analysis (BEA). Aggregate capital is constructed by multiplying the capital output ratio by aggregate income constructed using CPS data.

The capital output ratio is constructed by dividing non-residential fixed assets with GDP from BEA.

Survival probabilities are from the National Center for Health statistics Vital statistics of the US 1992, Vol II, sec. 6 life tables page 13, Washington: Public Health Service. 1996.

The remaining data are from the March supplement of the CPS, downloaded from the Integrated Public Use Microdata Series (King et al. (2010), cps.ipums.org). Data for hours and wages are constructed by including individuals of at least 15 years old who reported their gender and education level and declared they worked a positive amount of weeks and for a positive wage.

B   Dynamic Perturbation

Following Schmitt-Grohe and Uribe (2004) and Gomme and Klein (2011), the model can be expressed as

$$E_t[f\left(x_{t+1}, y_{t+1}, x_t, y_t\right)] = 0, \tag{20}$$

where E_t is the expectation given information at time t, x_t is a vector of state variables sorted so that all shocks enter in the lower part of the vector, y_t is a vector of all the other variables of the model.[1] A recursive representation of the

---

[1]To familiarize with the notation, consider the following simple model:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \log(c_t)$$

subject to feasibility and to the productivity process:

$$k_{t+1} + c_t = k_t(1 - \delta) + e^{A_t} k_t^{\theta}, \tag{21}$$

$$A_t = \rho A_{t-1} + \sigma u_t, u_t \sim d(0, var_u). \tag{22}$$

The equilibrium conditions are the last two equations 21-22 and

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}} (1 + A_{t+1} \alpha k_{t+1}^{\alpha-1} - \delta) \right). \tag{23}$$

With $x_t = [k_t, A_t]$ and $y_t = c_t$, the three equilibrium conditions 21-23 are easily casted into equation 20.

solution to Equation 20 which satisfies transversality conditions takes the form

$$y_t = g(x_t, \sigma) \tag{24}$$

and

$$x_{t+1} = h(x_t, \sigma) + \sigma \eta_{t+1}, \tag{25}$$

where, following Schmitt-Grohe and Uribe (2004), σ is a scalar that scales the variance of $\eta_t = [0, u_t]$, where $\{u_t\}$ is an i.i.d. sequence of innovations with zero mean and variance matrix Σ. An approximated solution can be found by Taylor expanding the deterministic version of equation 20 around its steady state, where $x_t = x_{t+1} = x_{ss}$, $y_t = y_{t+1} = y_{ss}$ such that

$$f(x_{t+1}, y_{t+1}, x_t, y_t) = 0. \tag{26}$$

This can be done, for instance, by applying the algorithm of Gomme and Klein (2011). Taylor expansions of Equation 26 are done around the deterministic steady state because this point is typically the easiest point to find where Equation 26 holds and equations 24-25 hold with $\sigma = 0$.[2] However, if there is another point $\hat{x}$ in the state space where one knows the values $\hat{x}_1 = g(\hat{x}, 0)$, $\hat{y} = h(\hat{x}, 0)$ and $\hat{y}_1 = h(\hat{x}_1, 0)$ so that

$$f(\hat{x}_1, \hat{y}_1, \hat{x}, \hat{y}) = 0,$$

one could take the Taylor expansion there and have a solution approximated around that point. The algorithm that I am about to introduce seeks to find such points and use them to derive better approximated policy functions with which to run a simulation.

Call

$$F(x, \sigma, h, g) \equiv E_x \left(f[h(x, \sigma) + \sigma \eta, g(h(x, \sigma) + \sigma \eta, \sigma), x, g(x, \sigma)] \right), \tag{27}$$

where E_x is the expectation over η given the state variables x. Pick $\hat{\varepsilon} > 0$. The algorithm seeks to find points $\{x_t\}_1^T$ given an initial condition x_0 and for a given sequence of shocks' innovations $\{u_t\}_1^T$, with precision

$$|F(x_t, 0, h_{x_t}, g_{x_t})| < \hat{\varepsilon} \quad for \quad t = 0, ..., T-1, \tag{28}$$

where h_{x_t}, g_{x_t} are policy functions approximated around x_t.[3]

---

[2] A point x_1, y_1, x, y that satisfies Equation 26 but not 24-25 is on a path that violates transversality conditions.

[3] Note that in equation (28) $\sigma = 0$. This is because the expectation operator in equation 27 is replaced by the certainty equivalence assumption of zero innovations.

1. Put $t_0 = 1$.

2. Taylor expand $f\left(x_{ss}, y_{ss}, x_{ss}, y_{ss}\right)$, where x_{ss}, y_{ss} is the deterministic steady state of the model, and obtain the policy functions $h_{x_{ss}}(x, \sigma)$, $g_{x_{ss}}(x, \sigma)$. If these are stable, then go to step 3 (for stability, see for instance Blanchard and Kahn (1980)).[4]

3. Put $\widetilde{h}(\cdot) = h_{x_{ss}}(\cdot)$ and $\widetilde{g}(\cdot) = g_{x_{ss}}(\cdot)$

4. Simulate from x_{t_0-1} with the policy function $\widetilde{h}(\cdot)$ and with $\sigma = 0$, generating a time series $\{\widetilde{x}_t\}_{t_0}^{\bar{T}}$ with $\bar{T} > T$.[5] If this time series does not converge to the steady state, increase $\bar{T}$ and go back to step 4.

5. Put $t = \bar{T}$.

6. Pick a point $\hat{x} = \alpha\widetilde{x}_{t-1} + (1 - \alpha)\widetilde{x}_t$ with $\alpha \in (0, 1]$ such that $|F(\hat{x}, 0, \widetilde{h}, \widetilde{g})| < \hat{\varepsilon}$.[6]

7. If $|F(\hat{x}, 0, \widetilde{h}, \widetilde{g})| > 0$, find $\hat{x}_1$ and $\hat{y}$ such that $f\left(\hat{x}_1, \hat{y}_1, \hat{x}, \hat{y}\right) = 0$ where $\hat{y}_1 = \widetilde{g}(\hat{x}_1, 0)$.[7]

8. Derive the functions $h_{\hat{x}}(\cdot)$, $g_{\hat{x}}(\cdot)$ Taylor expanding $f\left(x_1, y_1, x, y\right) = 0$ around the latter point $(\hat{x}_1, \hat{y}_1, \hat{x}, \hat{y})$.

9. Put $\widetilde{h}(\cdot) = h_{\hat{x}}(\cdot)$ and $\widetilde{g}(\cdot) = g_{\hat{x}}(\cdot)$.

   - If $\alpha < 1$, increase it to a number smaller or equal to 1 and such that $|F(\hat{x}, 0, \widetilde{h}, \widetilde{g})| < \hat{\varepsilon}$ where $\hat{x}$ has been updated accordingly: $\hat{x} = \alpha\widetilde{x}_{t-1} + (1 - \alpha)\widetilde{x}_t$. Go back to step 7.
   - If $\alpha = 1$ and $t > t_0$, put $t = t - 1$ and go back to step 6.
   - If $\alpha = 1$ and $t = t_0$, policy functions at x_{t_0-1} have been found. Store $x_{t_0} = \widetilde{h}(\hat{x}, \sigma) + \sigma\eta$, $y_{t_0-1} = \widetilde{g}(\hat{x}, \sigma)$ and go to the next step.

10. If $t_0 = T$ the whole solution has been found! Otherwise, put $t_0 = t_0 + 1$ and go back to step 3.

---

[4]This algorithm is described for models that are stable around the steady state. In fact, it could be extended to models that do not have a steady state provided that a point $(\hat{x}_1, \hat{y}_1, \hat{x}, \hat{y})$ such that $f\left(\hat{x}_1, \hat{y}_1, \hat{x}, \hat{y}\right) = 0$ is known. Indeed, it has worked for models that are locally unstable in some regions of the state space.

[5]Since $\sigma = 0$, this simulation is independent of any time series for the innovations to the shocks. It provides a path along which to move backward from the steady state.

[6]One can start with $\alpha = 1$, thereby moving from $\widetilde{x}_t$ to $\widetilde{x}_{t-1}$. However, sometimes this movement is too large so that $|F(\hat{x}, 0, \widetilde{h}, \widetilde{g})| > \hat{\varepsilon}$; in this case it is advisable to pick a smaller α.

[7]This step is similar to a step in the policy function iteration algorithm. Here, this step makes sure that the function $f\left(\hat{x}_1, \hat{y}_1, \hat{x}, \hat{y}\right) = 0$ holds and hence a Taylor expansion is admissible.

Variations of this algorithm can be conceived; for instance, to increase speed one could avoid going backward through all the points on the equilibrium path, but make larger jumps from the steady state until x_0.

The iteration procedure over the equilibrium path is reminiscent of the Parametrized Expectation Approach (PEA. See Den Haan and Marcet (1990) and Marcet and Lorenzoni (1999)): both algorithms break the curse of dimensionality by only approximating the global policy function over the equilibrium path rather than over the entire state space. In practice however, the PEA may show convergence problems that make its implementation hard for high-dimensional applications.[8] An important advantage of the proposed method is that, unlike the PEA, it does not iterate on the equilibrium path and hence does not rely on a contraction mapping, which explains why convergence problems do not arise. On the other hand, the invertibility conditions necessary to derive policy functions through perturbation methods need to be satisfied at all the points where the perturbation is applied; these invertibility conditions have not been violated in the models solved so far.[9]

To evaluate this algorithm, I test it on the model in note 1, equations 21-23, and with full depreciation, for which the analytical solution is known. I then compare the true equilibrium path $\{x_t^*, y_t^*\}_0^T$ with the one generated by this algorithm, $\{x_t^{**}, y_t^{**}\}_0^T$, and with the one generated by a second-order expansion around the steady state $\{x_t^{***}, y_t^{***}\}_0^T$. For an initial condition quite far from the steady state, $x_0 = [.2k_{ss}, -.5]$, with variance of the shock equal to 0.007, [10] the maximum error

$$\max_t[\max\left(|x_t^* - x_t^{***}|, |y_t^* - y_t^{***}|\right)] \tag{29}$$

using second-order approximation around the steady state is 0.0077. Using the proposed algorithm, the maximum error

$$\max_t[\max\left(|x_t^* - x_t^{**}|, |y_t^* - y_t^{**}|\right)]$$

is 2.2610^{-5}, which is 340 times smaller than taking the expansion only around the steady state. The simulation computed with the two methods is compared with the true solution in Figure 19. I conclude that this method makes a notable

---

[8]See in this respect the improvements made by Judd et al. (2009) (typically, this approach is also less accurate because it interpolates across the points).

[9]In addition, this method is only valid on the specific equilibrium path generated by the initial conditions and the time series of innovations. However, it is possible to simulate for a long time horizon and then interpolate over the solution to obtain policy functions valid over the whole ergodic set.

[10]This is the typical calibration of a TFP shock in the RBC model. The other parameters are $\theta = .33$, $\rho = .99$ and $\beta = .99$.

improvement in terms of accuracy versus perturbation around the steady state.[11]

In this example the code takes 27 seconds to run a simulation of 60 periods on a laptop. Solving the main model of the paper in section 3 takes about 11 hours and 20 minutes. As a measure of accuracy, I compute the error

$$|f\left(E_t(x_{t+1}), E_t(y_{t+1}), x_t, y_t\right)| \tag{30}$$

for all t. Abstracting from Jensen's inequality, a solution to the model is such that the error (30) is equal to 0 for all t. Hence, the size of this error gives a sense of the accuracy of the approximated solution. The maximum error with local approximation of the steady state is 0.29; with this algorithm it is 1.710^{-12}.[12]

C Production function estimation

C.1 Estimating the complementarity across labor groups

First divide (13) for each i by the same equation for group 1.

$$\left(\frac{z_i}{z_1}\right)^\sigma \left(\frac{n_i}{n_1}\right)^{\sigma-1} = \frac{w_i}{w_1}. \tag{31}$$

Multiplying by $\frac{n_i}{n_1}$ and taking logs one gets [13]

$$\log(z_i) - \log(z_1) + \log\left(\frac{n_i}{n_1}\right) = 1/\sigma \log\left(\frac{w_i n_i}{w_1 n_1}\right), \tag{32}$$

which gives $I-1$ linear equations from which σ can be estimated directly, without knowing the other parameters of the production function.[14] In order to facilitate notation I define $\epsilon_i \equiv \log(z_1) - \log(z_i)$, $\lambda \equiv \log\left(\frac{n_i}{n_1}\right)$ and $x \equiv \log\left(\frac{w_i n_i}{w_1 n_1}\right)$. I assume that the vector of ϵ_is follows an $AR1$ process:

$$\epsilon_t = \bar{\nu} + \rho_\lambda \epsilon_{t-1} + \nu_t, \tag{33}$$

---

[11]Furthermore, the accuracy seems robust to initial errors. In fact, using a small $\bar{T}$ such that the initial path does not converge to the steady state and the backward procedure starts with an initial error, has a negligible effect on accuracy. Intuitively, step 7 corrects for such errors.

[12]Although this result is quite reassuring, it should be noticed that an explosive solution, or a solution which alternates explosive to implosive patterns, can be consistent with this result. However, from a graphical inspection, the solution does not present an oscillating path and all variables converge to a steady state.

[13]The multiplication by $\frac{n_i}{n_1}$ is done to work with wage income rather than wage rates, thereby attenuating measurement error.

[14]It is equivalent to dividing (13) for each i by the same equation for a group j different to group 1 as done in equation (31). This expression can be obtained as the ratio of (31) for groups i and j and therefore does not convey any further information.

where $\bar{\nu}$ is a vector and ρ_λ is a scalar. I rewrite (32) with the new notation

$$\lambda_t = 1/\sigma x_t + \epsilon_t + \mu_t, \tag{34}$$

where μ_t is an $I-1$ vector of measurement errors. Since λ_t affects x_t, the latter is not orthogonal to $\epsilon_t + \mu_t$. However, through (33)–(34) one can derive the following:

$$\lambda_t = 1/\sigma x_t + \bar{\nu} + \rho_\lambda(\lambda_{t-1} - 1/\sigma x_{t-1}) + \nu_t + \mu_t - \rho_\lambda\mu_{t-1}. \tag{35}$$

Similarly to Arellano and Bond (1991), I pick the parameters in order to match the following moment conditions: $E[x'_{t-2-i}(\Delta\eta_t)]$, $E[\lambda'_{t-2-i}(\Delta\eta_t)]$, with $i = 1, .., 5$, where $\eta_t = \nu_t + \mu_t - \rho_\lambda\mu_{t-1}$ and Δ stands for first difference.[15] I use these $10(I-1)$ moments to estimate ρ_λ and σ.[16] σ is estimated to be 0.91, which suggests that there is not very much complementarity across these groups.[17] As can be inferred from equation (31), the fact that σ is close to one depends upon the fact that changes in the input of labor in one group have little impact on relative wages.

A virtue of this specification is that the estimation above is independent of the complementarity between labor and capital, θ: this parameter is calibrated to -0.25 in accordance with the literature that suggests a parameter value which induces more complementarity than the Cobb-Douglas case. See for instance Leandoacute;n-Ledesma et al. (2010) and Choi and Ríos-Rull (2009). α is such that the model predicts the average labor share found in the data: in the model, the capital share of output is $A\alpha\left(\frac{y}{k}\right)^\theta$. α is identified by normalizing total factor productivity A to be one on average, its value is 0.4.[18]

Table 12 summarizes the key parameters estimated.

Table 12: Estimation Results

	σ	ρ_λ	α	θ
Value	0.909	0.527	0.409	-0.250
St. Error	0.021	0.056	-	-

---

[15]μ_{t-2}, which is present in $\Delta\eta_t$, could influence λ_{t-2} and x_{t-2}. Only λ_{t-2-i} and x_{t-2-i}, with $i \geq 1$, are uncorrelated with μ_{t-2}.

[16]$\bar{\nu}$ has been removed by differentiating η_t

[17]Including more or fewer lags does not change the result that groups are quite substitutable into the production process. This result persists when the estimation is carried out with residuals in levels, considering x and λ in first differences, or adding a trend to equation (33).

[18]The average level of A is not pinned down and can be normalized: for any level of A, it is possible to find a value of α and the shocks z_i such that output is preserved, as well as the marginal productivity of capital and of labor in each group. Since α is a constant, changes in A and z_i are identified independently of the normalization on the average level of A.

C.2 Identifying the productivity shocks

From the labor demand equation for group one, one can derive:

$$(1-\alpha)Ay^{1-\theta}\left(\sum_{i=1}^{I}(z_in_i)^\sigma\right)^{\theta/\sigma} = \frac{w_i\left(\sum_{i=1}^{I}(z_in_i)^\sigma\right)}{z_1^\sigma n_1^{\sigma-1}}. \tag{36}$$

Since z_i/z_1 has been identified through (32), it is convenient to rewrite the expression above as follows

$$(1-\alpha)Ay^{1-\theta}\left(\sum_{i=1}^{I}(z_in_i)^\sigma\right)^{\theta/\sigma} = \frac{w_i\left(\sum_{i=1}^{I}(z_in_i/z_1)^\sigma\right)}{n_1^{\sigma-1}}. \tag{37}$$

Substituting this into the production function and solving for A gives

$$A = \left(y^\theta - \frac{w_i\left(\sum_{i=1}^{I}(z_in_i/z_1)^\sigma\right)}{n_1^{\sigma-1}}\right)/(\alpha k^\theta). \tag{38}$$

As a result, z_1 can be backed out from (37) and finally, z_i for every i can be derived through (32). $a \equiv \log(A)$ is assumed to be an AR1 process:

$$a_t = \rho_a a_{t-1} + u_{a,t}. \tag{39}$$

A time trend is found not to be significant.[19] Figure 17 panel A, shows the time series for A. The initial trend depends on the process starting below steady state.

C.3 Decomposition of labor productivity shocks

Consistently with the theoretical model, group-specific shocks z_i in (11) are decomposed into gender, age and education-specific shocks $(\varepsilon^g, \varepsilon^{age}, \varepsilon^{edu})$ so that

$$\log(z_{t,i}) = \sum_{j=1}^{2}\varepsilon_{t,j}^g I_g(i,j) + \sum_{j=1}^{3}\varepsilon_{t,j}^{age} I_{age}(i,j) + \sum_{j=1}^{2}\varepsilon_{t,j}^{edu} I_{edu}(i,j) + \nu_{t,i} \tag{40}$$

for all i,t, where $I_{edu}(i,j) = 1$ if education in group i is equal to j and zeros otherwise. Dummies by gender and age are defined the same way. $\nu_{t,i}$ is a residual

---

[19]A time trend on the neutral shock would also be inconsistent with a balanced growth path.

capturing what cannot be accounted for by a combination of the other shocks.[20,][21] The problem can be written in vectoral form:

$$\log(z_t) = \Lambda \varepsilon_t + \nu_t, \tag{41}$$

where z_t and ν_t are respectively the vector of I labor specific shocks and residuals at time t, ε_t is the vector of the seven group-specific shocks by gender, age and education. Λ is a $I \times 7$ matrix that collects the group dummies introduced in equation (40). ε_t are identified by minimizing the sum of squares of residuals ν_t: $\min_{\varepsilon_t} \nu_t' \nu_t$. This problem can be considered a factor model with factors ε_t and where the factor loading matrix Λ is given by the theoretical structure.

Figure 18 shows how closely $\Lambda \varepsilon_t$ can replicate labor-specific shocks z_t. As is evident from the figure, the the difference is negligible and hence in the simulations I will only include $\Lambda \varepsilon_t$ and abstract from ν_t. I assume an AR1 process for ε_t:

$$\varepsilon_t = \gamma_\varepsilon + \rho \varepsilon_{t-1} + u_t, \tag{42}$$

where ρ is assumed to be a diagonal matrix.[22] Figure 17 panels B,C and D plots these shocks. Table 13 contains parameter values estimated with OLS and the covariance matrix for the innovation components.

Table 13: Estimation Results

	ε_f	ε_m	ε_l	ε_h	ε_{young}	ε_{prime}	ε_{old}	A
ρ	.89	.87	.89	.89	.90	.89	.90	.89
γ	.05	.10	.05	.09	.01	.07	.05	0
Innovation covariances								
ε_f	0.0048	0.0047	0.0033	0.0031	0.0032	0.0050	0.0045	0.0021
ε_m		0.0050	0.0032	0.0032	0.0033	0.0050	0.0047	0.0022
ε_l			0.0024	0.0021	0.0020	0.0035	0.0030	0.0015
ε_h				0.0021	0.0020	0.0033	0.0030	0.0014
ε_{young}					0.0025	0.0033	0.0032	0.0014
ε_{prime}						0.0054	0.0046	0.0022
ε_{old}							0.0046	0.0020
A								0.0010

---

[20]Residuals $\nu_{t,i}$ come about from the fact that labor-specific shocks z_i for all the groups cannot be accounted for by only 7 shocks: 2 by gender, 2 by education and 3 by age.

[21]This exercise being a mere decomposition does not affect the estimated complementarity across groups.

[22]A time trend was found not to be significant and therefore it is omitted. Over the sample, the model predicts growth that essentially comes from the transition of the shocks, capital and demographics to the steady state.