

The Aggregate Implications of Changes in the Labour Force Composition*

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

February 8, 2018

Abstract

Labour composition by gender, age, and education has undergone dramatic changes over the last half century in the United States. Furthermore, the volatility of total market hours differs systematically between genders, age, and education groups. Reduced form exercises and a large-scale business cycle model suggest that these demographic patterns account for between 14% and 31% of the observed changes in aggregate volatility over this period of time. Additionally, these demographic changes are responsible for a significant fraction of the GDP growth observed in the considered period of time.

JEL Classification Codes: E32, J11, J22

Keywords: Demographic Transition, Business Cycle, GDP Growth

*I thank my advisor Martin Gervais for invaluable guidance. I also thank seminar participants at the the NBER 2012 SI “Consumption”, European University Institute, the Nordic Summer Symposium, EEA/ESEM 2011, RES annual conference 2012, U. of Edinburgh, Uppsala, Universitat de Barcelona, Autònoma de Barcelona, Birmingham, La Sapienza, the workshop on dynamic macroeconomics in Vigo, the Vienna Macro Café, the workshop on macroeconomic dynamics in Bologna. I also thank, Alice Schoonbroodt, John Knowles, Corrado Giuliatti, Sarolta Laczó, Ramon Marimon, Juan Carlos Conesa, Xavier Mateos-Planas, Morten Ravn, Per Krusell, John Hassler, Kjetil Storesletten, Jonathan Heathcote, Árpád Ábrahám, Salvador Ortigueira, Yang Lu, Jerome Adda, Russell Cooper, Xavier Raurich, Nils Gottfries, Roman Sustek, Ludo Visschers, Thijs van Rens, Andrea Moro, Arnau Valleredes-Esteban, Victor Rios-Rull, and especially Paul Klein for helpful comments. The usual disclaimer applies.

1 Introduction

This paper studies the role demographic change by gender, age, and education played in the prolonged period of business cycle volatility slowdown and fast economic growth that preceded the Great Recession.

In an influential paper, [Jaimovich and Siu \(2009\)](#) (J-S) found that changes in the age composition of the labour force had a large and significant effect on aggregate volatility during the postwar period.¹ Meanwhile, labour composition changes by gender and education have been equally dramatic: [Katz and Autor \(1999\)](#) and [Katz and Freeman \(1994\)](#) among others documented the well known increase in female labour supply and in the number of workers with high education. Importantly, these patterns are correlated with the increase in prime-aged population. Furthermore, changes by gender and education have the potential to affect aggregate fluctuations because total hours worked by women are less volatile than those by men, and college educated hours are less volatile than those of lower educated workers. This is a similar fact to the one that motivated the work of J-S; that the volatility of the time series of total hours worked by prime age workers is lower than that of the young and older workers.

This evidence motivates the study of the implications of all these demographic changes jointly, especially given that they are correlated. However, to add gender and education as possible driving forces of business cycle volatility is challenging because of endogeneity problems. The reduced-form techniques developed by J-S exploit the notion that changes in the age-distribution of the population, determined by birth rates lagged at least 15 years, are exogenous to changes in current economic activity. But, as this paper documents, different margins characterized the changes in the gender, age and education composition: while transitions by age and education consisted primarily of changes in population shares (mainly due to the post-War baby boom and to increasing educational attainments), changes by gender are about the number of women in employment and hours per worker. Unlike the population distribution, these more intensive margins are choice variables.

I start with simple accounting exercises to measure the contribution of these demographic changes to the Great Moderation, the large volatility decline that started

¹The 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 labour.

in the 1980s as initially documented by [Kim and Nelson \(1999\)](#) and [McConnell and Perez-Quiros \(2000\)](#). The main idea is to construct counterfactual time series for total hours, employment and aggregate labour earnings where the long run demographic trends by gender, age and education are removed. An important distinction to take into account in designing these counterfactuals is that, as mentioned, different margins characterized the changes in the gender, age and education composition. I thus construct counterfactuals through which I am able to measure the joint contribution to the Great Moderation due to trends in population shares, employment, and hours per worker. These exercises find that demographic change by gender, age and education taken jointly, account for between 14% and 29% of the change in the volatility of aggregate market hours, employment and labour earnings between the first sample (1967-1984) and the Great Moderation period (1985-2007). Furthermore, the contribution to this moderation of the changes by gender and education is more important to what was previously found by J-S and comparable to the effect of age trends.

While intuitive and informative, the paper discusses limitations of these exercises which motivate a more structural analysis. For instance, there are issues with the identification of the demographic trends. Furthermore, this type of accounting exercise takes the differences in the cyclical volatility across groups as given, i.e. it assumes that the demographic changes have not affected the individual response of market work to business cycle shocks. Should this assumption fail to hold, the effects of demographics on aggregate volatility will likely be mismeasured by this accounting exercise. As an example of how demographic change may affect the response of labour to business cycle shocks at the individual level, [Goldin \(2006\)](#) argues that female labour supply elasticity declines as the commitment to work of women increases. The theory takes an explicit stand on how the labour supply elasticity and thereby the sensitivity of market hours responds to the various demographic changes which is consistent with the empirical evidence. In particular, the evidence suggests a utility function that makes labour supply elasticity a declining function of labour input. With this specification, the model matches the empirical correlations between the evolution of hours volatility at a disaggregated level, and the demographic changes.

The framework is a business cycle model with overlapping generations similar to [Ríos-Rull \(1996\)](#) in which I introduce heterogeneity by gender and education. The demographic changes are driven by exogenous trends in birth rates, in the share of newborns with low and high education, and in the time spent in non market activities

for men and women. This way, consistently with the data, changes by gender are due to labour supply, while those by age and education involve the population margin.² Importantly, I adopt a flexible production function which, in addition to conventional Total Factor Productivity (TFP) shocks, allows for shocks to specific labor input by age, sex, and education. By modelling labour demand shocks, the model accounts for the decline in the volatility of wages which dampens the response of hours, thereby avoiding to attribute too much of the moderation in hours to demographic changes. These group specific labour demand shocks also generate changes in the wage gaps which in turn explain some of the changes in the labour composition. As a byproduct of distinguishing between aggregate and labour demand shocks, the model also matches the evolution of the labour share and provides an explanation to the employment productivity puzzle.³

A computational challenge arises from the fact that the demographic distribution is not constant and the model is solved over a large transition; this makes linearization methods around the steady state — the methodology used by [Ríos-Rull \(1996\)](#) and the typical technique used to solve large scale business cycle models — inaccurate. To address this issue, the model is solved with a new technique which consists of applying linearization methods at many points over the equilibrium path. The solution is much more accurate than that found with standard linearization methods only around the steady state and it is described in detail, and in a way that is applicable to a large class of models, in [Mennuni and Stepanchuk \(2016\)](#).

The model is used to measure the contribution that the labour compositional changes played on aggregate volatility. A counterfactual simulation where the exogenous trends are such that the labour composition fluctuates around its steady-state levels finds that output volatility would have been much lower in the period characterized by high volatility that preceded the Great Moderation (1967-1984), and slightly more volatile in the Great Moderation period when there was little volatility (1985-2007).

²A previous version — [Mennuni \(2013\)](#) — had education as a choice variable. The positive trend in educational attainments was partly due to an increasing wage premium and partly to a declining cost of education. This choice played a negligible role for the results. Intuitively what matters in this paper is that cyclical volatility differs among education groups but education decisions did not explain that. The only way the education decision was endogenous to aggregate volatility was in the extent to which the decision to pursue education depends on current business cycle conditions; but in each period this decision only pertains the newborn cohort. This had negligible aggregate effects as for all other cohorts in the labour force, education was a state variable. However, other mechanisms through which education might affect aggregate volatility are conceivable, see [Lugauer \(2013\)](#).

³The near-zero correlation between total hours and labour productivity. This fact is puzzling for Real Business Cycle models, which predict this correlation to be large. See [Hansen and Wright \(1992\)](#).

By accounting for part of the high volatility in the first sub-sample and the slowdown in the 1980s, labour reallocation accounts for 31% of the Great Moderation.⁴ Consistent with the reduced form accounting exercise, results suggest that trends in gender, and especially education, are more important than previously thought: trends in education play the most important role, followed by age and sex, in the change in output volatility. Furthermore, these results suggest that other things equal, business cycle volatility should be expected to increase somewhat in the next decades due to the ongoing ageing of the population and the fact that the counterbalancing transition by gender and education is largely completed.⁵

While these results confirm the finding of J-S that the age composition has a relevant impact on aggregate volatility, they also suggest that age is less important than previously thought. To assess whether the model understates the role of age composition, I adapt the analysis of J-S, where business cycle volatility is regressed over the age composition, to the data generated by the model. 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. So the regression overstates the role of the age composition in the model. Intuitively that changes by gender and education are correlated to those by age induces an omitted variables bias. However, adding gender and education regressors also produce biased estimates confirming that these variables are endogenous to business cycle volatility.⁶

Within the context of the model, I can also measure the contribution to growth of this demographic transition. Growth in the model comes from the productivity shocks (TFP and labour specific), initial assets below steady state, and because of

⁴Between the first and second subsample, output volatility decreased by 41.9% in the baseline simulation with demographic change. In the counterfactual simulation (where the labour composition is trendless) the moderation is only of 28.9%, so these demographic trends account for $(41.9-28.9)/41.9$ or 31% of the moderation in output volatility.

⁵But to be clear, while demographic trends explain part of business cycle volatility, this is of limited help to predict recessions such as the financial crisis: the labour composition can affect the propagation of shocks, but the severity of a recession depends fundamentally on the size of the shock. In fact, baseline and counterfactual volatility are roughly the same in the last part of the sample (2008-2013). Intuitively, the demographic trends have largely converged, so there is not much difference in the two scenarios for the last period of time.

⁶While the result suggests that the importance of age trends for the US is smaller (although still substantial) than what J-S found in their panel of countries, this result does not invalidate the analysis of J-S. They also use panel data while the regression on the model generated data has only the time series dimension specific for the U.S. A more appropriate comparison might be [Lugauer \(2012b\)](#), which does focus solely on the US. That paper estimates that the age distribution contributed 18% to the aggregate volatility decline; not as far from what is found here through the model and the reduced form accounting exercises (between 7.9 and 10.3%).

demographic change. Intuitively, the increase in female labour, and in the share of relatively more productive and hard working prime age and highly educated workers, contributed to the growth observed in the past half century. Comparing the growth path between the baseline simulation and those without demographic change, I find that 16% of the average output growth between 1967 and 2013 is due to sex trends, 9% to trends in age and 17% is due to education trends. Finally, taken together these trends account for a staggering 39% of the average output growth. It should be noted that this exercise is different from a growth accounting exercise where the growth of each input is assumed independent from the others. For instance, capital accumulation is determined by the endogenous savings of each demographic group. Furthermore, part of the labour composition is endogenous through labour supply. However, a limitation is that the long run trends in the population by age, educational attainment, and in non market hours are exogenous. In practice it is likely that these trends are jointly determined with GDP growth, especially labour participation and educational attainment.⁷ Thus these results should be taken as a first pass where the role of demographics may be overstated. Yet these numbers are conservative when compared to the literature which offers estimates for each of these dimensions taken individually.⁸ These findings suggest that the slow growth observed in recent years is likely to persist given that the demographic trends have largely converged. Thus this evidence supports the “secular stagnation” hypothesis formulated by Hansen (1939) and recently evoked by Summers (2013) among others.

The paper proceeds as follows. Section 2 documents facts and contains the reduced form exercises. Sections 3 and 4 set up and calibrate the model. Section 5 tests the model and measures the effects of labour reallocation. Section 6 concludes. Appendix A reports the model definition of a competitive equilibrium, and Appendix B includes results about the production function.

⁷Another limitation is that while the model distinguishes between population and labour per capita, the employment over population, and hours per worker margins, are not disentangled.

⁸For instance Greenwood et al. (2005) find that the rise in female labour accounts for about 19% of growth between 1900 and 1980. Feyrer (2007) finds that age labour composition changes predict almost 12% of GDP growth between 1990 and 1995. The number found here for education is slightly higher than the 15% found by Goldin and Katz (2008) for human capital: they argue that their number is likely to be an understatement.

2 Stylized facts and accounting exercises

As it is well known, aggregate output volatility in the U.S. increased during the early 1970s and declined during the 80s and 90s.⁹ Recently, volatility has increased again and there is renewed curiosity about its future unfolding. To relate these facts to the labour 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 labour data, downloaded from the Integrated Public Use Microdata Series ([King et al. \(2010\)](#), cps.ipums.org).

2.1 Total Hours

Figure 1 shows the share of paid hours worked over time by gender, age (young (15–29), prime age (30–59) and older workers) and low and high education (at least four years of college). These are shares of total hours per capita by group, i.e. hours per worker times number of workers in a group divided by the total number of hours worked by the entire working population.¹⁰

As is reported in the first two columns of Table 1, the share of hours worked by prime age workers increased relative to those by other age groups, moving from an average of 64% prior to 1984 to an average of 71% from 1985.¹¹ By contrast, the share of hours by the young and old fell. Furthermore, the volatility of hours is substantially lower for prime age workers as initially documented by [Clark and Summers \(1981\)](#). This is shown in column 3, where the standard deviation of the filtered time series of hours by age for the entire sample (1967-2013) is reported.¹²

The relative increase in prime-age hours and the fact that hours are less volatile for this age group may have contributed to the reduction in aggregate volatility. These two facts also hold by gender and education. As shown in Table 1, between the first and second sub-sample hours worked by women increased relative to those by men. It

⁹See [Kim and Nelson \(1999\)](#), [McConnell and Perez-Quiros \(2000\)](#) and [Stock and Watson \(2003\)](#).

¹⁰Aggregates for hours and labour earnings 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.

¹¹1984 is the reference year adopted by the literature as the beginning of the Great Moderation. See for instance [Stock and Watson \(2003\)](#)

¹²Like J-S, to isolate business cycle frequencies I use the Hodrick-Prescott (HP) filter on logged data with a smoothing parameter of 6.25 as suggested by [Ravn and Uhlig \(2002\)](#). Results are robust to using a parameter of 10. [Jaimovich et al. \(2013\)](#) find that these results are also robust to detrending with the band-pass filter. In a previous version, I used growth rates rather than the HP-filter throughout the paper and reached the same conclusions.

is less well known, however, that male hours are more volatile than female hours: the standard deviation of hours is 1.31% for women and 1.98% for men.¹³ Similarly, from the last two rows, the share of the highly educated increased and their hours volatility is lower than that of the less educated.

To summarize, by sex and education as well as age, there has been an increase in the hours worked by the groups with lower volatility. These are the basic facts behind the conjecture that these trends matter for aggregate volatility changes.

Next the paper decomposes total hours in population shares, employment over population, and hours per worker. This more detailed cut highlights the following: while transitions by age and education mainly involve population shares, the transition by gender involves the number of women in employment and hours per worker.

2.2 Population, employment and hours per worker

Table 2 shows population shares by gender, age and education. The population includes all individuals that are at least 15 years old. Columns one and two show that population changes are large by age and education, but minor by sex. For instance, the stable prime age population increased from 45.9 to 51.3% of the total while the volatile young decreased from 34.5 to 27.6%. Similarly, the population with high education increased from 12 to 22%. Instead, sex ratios remained roughly constant. However, sex has been characterized by a strong transition in employment over population rates reported in Table 3; employment over population for women moved from an average of 43% in the first sample to 54% in the second sample. Employment has also changed for the other groups as reported in the Table, but to a smaller degree.

Similar changes also characterized hours per worker (total market hours over employment in each group) reported in Table 4.¹⁴ Hours per worker for women increased significantly more than for men. Instead, hours per worker by age and education are roughly unchanged.

The third column in Tables 3 and 4 highlights that volatilities at these more intensive

¹³This fact has also been recognized by [Gomme et al. \(2005\)](#) and J-S; [Abraham and Shimer \(2001\)](#) and more recently [Hoynes et al. \(2012\)](#) report a similar fact for unemployment by gender. The fact is all the more surprising given that at the individual level, labour income for women is more volatile than for men. Evidently, some of the volatility at the female individual level is idiosyncratic and washes out in the aggregate. [Doepke and Tertilt \(2016\)](#) elaborates on this further.

¹⁴Variables are constructed taking care of the fact that the March CPS data on earnings and hours refer to the year preceding the interview while population and employment refers to the interview date.

margins are consistent with total hours (Table 1): standard deviations for women, prime age and the highly educated are lower than for their respective counterparts (males, other age groups and workers with lower education).

In summary, while transitions by age and education mainly involve population shares, the one by sex involves the intensive margins of female employment and hours.

2.3 Accounting exercises

To get an initial sense of the contribution of these demographic changes to the Great Moderation, let per capita aggregate hours at time t be

$$H_t = \sum_i h_{i,t} e_{i,t} p_{i,t}, \quad (1)$$

where $h_{i,t}$ is hours per worker, $e_{i,t}$ is employment over the population included in group i , and $p_{i,t}$ is the population in group i divided by total population. Counterfactual hours can be constructed by removing trends in $h_{i,t}$, $e_{i,t}$, and $p_{i,t}$.

As a complement to their regression analysis, J-S consider experiments in which they hold the population or workforce constant and then compare the volatility before and during the Great Moderation for both the baseline and the counterfactual data. Specifically, they set $\hat{p}_{i,t}$ by age constant at the average values during the pre-moderation period (1967–1984) and account for 20% of the moderation in total hours. Holding population shares by education constant accounts for 10%. They also report that holding the gender composition of the workforce constant at the premoderation levels resulted in essentially no change in the volatility of aggregate hours.

A problem with holding the workforce constant, *i.e.* setting $e_{i,t}$, or $h_{i,t}e_{i,t}$ constant, is that one essentially removes the entire business cycle as population shares show little short run variation. Then, it is difficult to say how much volatility has changed between the two samples, when volatility is virtually zero in both. Instead, I wish to keep short run fluctuations in $h_{i,t}$, $e_{i,t}$ and $p_{i,t}$ but to remove their long run trends.¹⁵ To this aim, I project gender, age, or education specific time trends on $h_{i,t}$, $e_{i,t}$ and $p_{i,t}$ disaggregated by gender, age or education. These trends are non linear and are well captured by a fourth order time trend as shown, for the case of $e_{i,t}$, in Figure 2.¹⁶

¹⁵It is worth stressing that the issue is especially severe for the gender composition where all the action involves the intensive margins rather than the population margin, which is fairly stable over the business cycle. Gender was not the focus in J-S and the accounting exercise was not their core analysis. But given the importance of these margins in this paper, I take a different approach.

¹⁶This approach to isolate low-frequency trends is also adopted by [Abraham and Shimer 2001](#)) and

To construct counterfactual employment over population \hat{e}_i , I remove these trends in the time series of e_i but keep the percent deviation from trend. This results in the dashed lines depicted in the Figure. Similarly, population trends or hours per worker are also detrended. I can then compute the volatility change in aggregate hours that would have happened had these trends been absent.

I also construct counterfactuals for aggregate employment and labour earnings per capita. Employment is $E_t = \sum_i e_{i,t} p_{i,t}$; aggregate labour earnings are $Yl_t = \sum_i w_{i,t} h_{i,t} e_{i,t} p_{i,t}$ where $w_{i,t}$ is the group specific hourly wage.¹⁷

Table 5 summarises the first takeaway of the paper, the contribution to the Great Moderation in e , H , and Yl due to the trends in population shares, employment over population and hours per worker by either gender, age, or education; for instance, counterfactual labour earnings are $\widehat{Yl}_t = \sum_i w_{i,t} \widehat{h}_{i,t} \widehat{e}_{i,t} \widehat{p}_{i,t}$, where the counterfactual hours per worker, employment and population shares are those computed above.

As done by J-S among others, the numbers in the table show the percentage of the moderation in the standard deviation explained by demographic change. For instance, removing trends by age, the fall in the standard deviation of HP-filtered aggregate employment during the Great Moderation period (1985-2007) is of 44.65 log points as opposed to 51.52.¹⁸ This explains $(51.52-44.65)/51.52$ or 13.3% of the moderation in employment, as shown in the first row, second column.¹⁹

Sahin et al. (2014). This methodology identifies trends with a much lower varying path than those identified with usual business cycle filters, e.g. HP and band pass, usually parameterized to identify trends with periodicities of 8 years or longer (Baxter and King 1999); such trends are unlikely to be due to demographics which has a longer time horizon (for instance, the baby boom produced a demographic cycle that is still ongoing). In terms of frequency, trends of order 4 can capture up to about one sinusoidal cycle (e.g. the sine function plotted over $[-\pi - \pi]$) over the entire time span 1967-2013. For a comparison, usual business cycle filters also capture trends with $(2013-1967)/8=5.75$ cycles over the time span. Changing the parameterization of the aforementioned filters can seem arbitrary but with time trends, results are entirely robust to choosing order 3 or 5 rather than 4. Furthermore, the conclusions are robust to using a band pass filter with frequency band between 2 and 32 years (I use that proposed by Christiano and Fitzgerald (2003), particularly suitable to extract low frequencies) or an HP filter with smoothing parameter 1600, typical value for quarterly data, but with annual data it isolates trends with lower frequency.

¹⁷Hourly wages are constructed as real wage income divided by hours per worker. This is a potential source of measurement error but here the unit of analysis is group averages, for which the issue is attenuated as the number of observations within cells gets large. See Angrist (1991).

¹⁸At this stage, the HP-filter with parameter 6.25 is used both on the data and counterfactuals to isolate business cycle frequencies. As mentioned, conclusions are robust to using growth rates.

¹⁹Looking at the moderation in the variance rather than in the standard deviation gives the same result: calling $x1$ and $x2$ the standard deviations in the first and second subsample, and $y1, y2$ the counterfactual ones, one can express the explained share of the moderation in the standard deviation as $(\log(x1) - \log(x2) - (\log(y1) - \log(y2)))/(\log(x1) - \log(x2))$. The latter is equal to $(\log(x1^2) - \log(x2^2) - (\log(y1^2) - \log(y2^2)))/(\log(x1^2) - \log(x2^2))$, which is the explained share of the moderation in

The table shows that education trends are very important, even more than age trends. Trends by sex are less important but not negligible.

Finally, it is possible to remove gender, age and education trends sequentially to get a measure of their joint effects. As it is common with decompositions, results are path dependent; that is, they depend on the order in which these trends are removed: gender first, then age, then education, as opposed to their permutations. So I report the largest and the smallest number. The contribution to the Great Moderation of removing all of these margins together is between 23.7 and 29.1 for per capita employment; between 24.5 and 28.0 for total hours, and between 14.3 and 15.7 for labour earnings.

2.4 Limitations

The accounting exercises above implicitly assume that the volatility of group-specific market work is not affected by demographic change: baseline and counterfactuals have different trends, but same percent deviations from the trends. Without an a priori view, this seems a reasonable starting benchmark: it helps isolating the purely mechanical role played by changing the demographic composition. But it raises an important question: is the group-specific response to business cycle shocks affected by demographic composition and would this change the results? Tables 1–4 last column document substantial changes in volatility within each group in all the margins considered: if these changes were an indirect consequence of these demographic trends, their implications for aggregate volatility should be attributed to demographics. For instance, there is evidence that labour supply elasticities declined over time, especially for females (Heim 2007 and Blau and Kahn 2007); as advocated by Goldin (2006), this elasticity decline may be related to the increase in employment and hours per worker by almost all groups and especially by females (Tables 3 and 4—first and second column). Empirically, the higher labour supply, the lower its elasticity. E.g. males versus females, prime age versus closer to retirement, women without children versus women with children; see Reichling and Whalen (2012). A causal link between labour supply and its elasticity

the variance. To get a sense of the assumption of fourth order time trends it is worth mentioning that removing age trends up to the third or the fifth order explains 13.26% or 15.04% of the moderation in employment. The number becomes 14.7% when identifying long run trends with the band pass filter with frequency 2-32 and 12.5% when using the HP filter with parameter 1600. Results are also fairly robust to changing the age bins: with 11 age bins (15–19, 20–24 and so on progressing in 5-year age groups until the last group of 65 and older), 16.6% of the moderation in employment is accounted by age changes. 3 bins are the baseline because I also consider interactions with sex and age which in the case of 11 bins would make groups too small.

implies that the mentioned increases in hours and participation played a further role on aggregate volatility by inducing the moderation within each group; so the answer of the exercise above may be wrong. Indeed, between the first and second sample, the ratio of the standard deviation of HP-filtered hours per worker or employment and that of hourly wages declined for all the groups. This is consistent with a lower sensitivity to shocks, possibly endogenous to labour demographics.

In the model, labor supply elasticity responds endogenously to changes in the labour composition. The importance of this channel in explaining aggregate volatility has also been argued recently by [Doepke and Tertilt \(2016\)](#). However, it should be noted that also the volatility of hourly wages moderated for most groups. For instance, the standard deviation of HP-filtered female hourly wages between 1985 and 2007 is 80% that between 1967 and 1984. This suggests that not all the moderation in hours should be attributed to a diminished sensitivity to shocks, but possibly also to smaller shocks, or good luck. Through the model, aggregate and groups specific shocks are identified, so the possibility that smaller shocks explain the moderation is also taken into account.

Another limitation is that, as mentioned, removing gender, age and education trends sequentially gives path dependent results. This is a general and well recognized problem with reduced form decomposition analysis. See for instance [Fortin et al. \(2011\)](#), who concludes that decompositions can be useful to uncover the main forces, but more structural work should then be used as well. Here path dependence is not strong enough to affect the conclusions: demographic trends play a non negligible role for business cycle volatility no matter the order. However, the issue highlights that to identify demographic trends is not simple. For instance: how much of the labour composition changes by education would have happened, had the female labour participation not increased? How much did the increasing presence of young female workers affect the age composition of the labour force? In the exercise above these trends are projections; it is hard to give an economic interpretation to this identification. The theoretical framework takes more meaningful and testable stands. For instance: gender composition trends will come from technological change that affects the gender wage gap as well as the housework burden. Even though agents act independently, this will have implications on the distribution by age and education as well because it affects people of different age and education differently (this is because they differ in income, labour, wealth and preferences). However, this indirect effect turns out to be modest.

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. A share $q^m \in (0, 1)$ of new born males and $q^f \in (0, 1)$ of new born females have high education.

3.1 The Problem of the Agents

Agents are distinguished by their sex $g = \{f, m\}$ for female or male; age $j \in \{1, J\}$; education $e = \{h, l\}$ for high or low; and their assets a . Each agent chooses consumption c , assets a' , and market hours ℓ , to solve the following problem:

$$V(g, j, e, a; \omega) = \max_{c, \ell, a'} \ln(c) + \chi \frac{(1 - \ell - \tilde{h})^{1-\sigma}}{1 - \sigma} + \beta \zeta_j E[V(g, j + 1, e, a'; \omega')] \quad (2)$$

subject to the constraints: $\zeta_j a' + c = a(1 + r) + w(g, j, e)\ell$,

$$a' \geq 0 \text{ if } j = J, \quad c \geq 0, \quad \ell \in [0, 1 - \tilde{h}], \quad (3)$$

where V defines expected utility discounted at factor $\beta \in [0, 1]$ times $\zeta_j \in [0, 1]$, the survival factor for the agent at age j ; it is such that people die for sure at age J , i.e. $\zeta_J = 0$.²⁰ σ regulates the Frisch elasticity of labour supply and χ is a scaling factor that helps match hours targets as detailed in Section 4. As discussed in Section 5.2.3, this utility specification makes the labour supply elasticity a function of labour input; this helps match the declining sensitivity of labour supply documented in Section 2.4. \tilde{h} is an exogenous time cost specific to each gender and education.²¹ Parameter indexation will be made explicit whenever necessary to avoid confusion; for example $\tilde{h}_{g,e}$.

r is the interest rate; $w(g, j, e)$ the wage, specific to each sex, age, and education. The period utility after death is assumed equal to 0, hence $V(g, J + 1, e, a; \omega) = 0$ and the zero debt constraint in the last period of life at age J . The expectation E is taken rationally over ω' given ω , the vector containing the aggregate state variables as explained below. Agents are born with zero assets.²²

²⁰ ζ_j in the budget constraint reflects competitive annuity markets, see Ríos-Rull (1996).

²¹In the absence of a more sophisticated theory of the household, the evolution of this parameter helps reproduce the distribution of hours by gender. Its reduction over time captures housework technology improvements and a fall in child care cost which, jointly with the reducing gender wage gap, causes increases in female market hours. See Greenwood et al. (2005) and Attanasio et al. (2008).

²²In a previous version (Mennuni 2013), individuals of opposite gender were matched stochastically based on their educational level as in Heathcote et al. (2010). Education was also an endogenous

3.2 Demographic Distribution

Denote $p : \{f, m\} \times \{1, \dots, J\} \times \{h, l\} \rightarrow \mathfrak{R}_+$ last period mass of agents by sex, age and education. $p'(g, 1, h) = p_0 q^g$ and $p'(g, 1, l) = p_0(1 - q^g)$ are the masses of newly born with high and low education respectively, with $g = f, m$. Let the mass of older households be defined recursively as $p'(g, j + 1, e) = p(g, j, e)\zeta_j$ for $j = \{1, J - 1\}$.

3.3 Firms

Competitive firms maximize profits using the following production function

$$y = A^{1/\theta} (\alpha k^\theta + (1 - \alpha)L^\theta)^{1/\theta}, \quad (4)$$

where A is total factor productivity (TFP), α is associated with the labour share of total output and θ measures the complementarity across capital and L , which is a composite of several labour groups:

$$L = \left(\sum_{i=1}^I (z_i n_i) \right). \quad (5)$$

z_i s are labour-augmenting technology shocks specific to each labour group; n_i is hours worked by all individuals categorized in group i .

z_i s capture sectorial shocks and other aspects of the production process not explicitly modeled, which move relative demands of labour inputs.

There is a mapping between firm groups i and agents: each group i is formed of agents of the same gender, age group and education level.²³ The mapping is represented by I dummy matrixes $\psi_i(g, j, e)$ which contain zeros and ones depending on whether the labour 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 ,

$$n_i = \sum_g \sum_{j=1}^J \sum_e \ell(g, j, e, a; \omega) p'(g, j, e) \psi_i(g, j, e), \quad (6)$$

choice. These two features had negligible effects on the results. Intuitively, the education decision is not that much affected by business cycle shocks because agents enter the labour market after the decision to educate. Furthermore, the return from education is spread over many years and thus only marginally affected by the business cycle conditions at the time of making the education choice.

²³Consistently with Section 2, there are three age groups: the young (1-10), the prime age (11-35) and the older agents (36-40).

where $\ell(g, j, e, a; \omega)$ is the labour input of each agent in that group.²⁴ The number of groups I is 12, i.e. the 2 genders times the 3 age groups times the 2 education levels.

The representative firm hires labour according to the following condition

$$(1 - \alpha)A^{1/\theta} (\alpha k^\theta + (1 - \alpha)L^\theta)^{1/\theta-1} L^{\theta-1} z_i = w_i, \quad (7)$$

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

$$\alpha A^{1/\theta} (\alpha k^\theta + (1 - \alpha)L^\theta)^{1/\theta-1} k^{\theta-1} = r + \delta, \quad (8)$$

where δ is the depreciation rate of capital.

The advantage of the adopted constant elasticity of substitution specification relative to a Cobb-Douglas (CD) is that with the latter, output and the marginal productivities of labour and capital move proportionally. Therefore it is in general not possible to match aggregate production given the inputs, while matching hours and wages through labour demand. Instead, with a CES with $\theta \neq 0$, it is possible, say, to increase output by increasing the productivity of capital but not that of labour, or viceversa. This is because while output, labour and capital demand move proportionally with $A^{1/\theta}$, they are not affected proportionally by z . Instead with a CD, it is impossible to disentangle labour augmenting and TFP shocks.²⁵ This is important given that I wish to match the output time series as well as hours and wages.²⁶

3.4 State Space

To make rational choices, agents need to know their type (gender, age and education) and asset position a . They also need to predict prices, which depend on the shocks and on the distribution of assets and agents across age, gender and education. The state space is defined in more detail next.

²⁴Since there are no idiosyncratic shocks within the groups, all agents of the same gender, age and education have the same asset levels, hence there is no integration over a in Equation (6).

²⁵To see this, consider the following CD: $y = Ak^\alpha(z\ell)^{(1-\alpha)}$. Let ℓ be homogeneous labour. The marginal productivities for labour and capital are $(1 - \alpha)Ak^{\alpha-1}(z\ell)^{-\alpha}z$ and $\alpha Ak^{\alpha-1}(z\ell)^{(1-\alpha)}$. It is clear that output, labour and capital demand are all proportional to $Az^{(1-\alpha)}$. Then it is not possible to increase output without increasing both labour and capital demand. This proportionality is also reflected in the fact that the capital and the labour share of income are constant and makes output and wages strongly pro-cyclical in RBC models, giving rise to the employment productivity puzzle.

²⁶A previous version also considered imperfect substitutability between labour inputs but the estimation of the complementarity was very close to the case of perfect substitution considered here.

3.4.1 Exogenous Processes

Let the logarithm of the labour productivity processes z_i , the logarithm of the TFP process A and the logarithm of the mass of new born p_0 be AR1 stochastic processes. Furthermore, the education shares q^g and housework $\tilde{h}_{g,e}$ with $g = \{f, m\}$ and $e = \{h, l\}$ are deterministic processes with an AR1 structure.²⁷ Let these variables be collected in the vector $G = [A, Z, p_0, q^g, \tilde{h}_{g,e}]$, where Z is the vector of all z_i .

3.4.2 Aggregate state

The remaining dimensions of the state space, denoted ω in Value Function (2), contain aggregate state variables that affect households' decisions through prices and expectations: the shocks, the distribution of agents by sex, age and education p , and assets across all groups of agents a_1, \dots, a_m , where $m = J \times 2 \times 2$, i.e the J generations alive, the 2 genders and the 2 education groups. The distribution of assets is a finite dimensional vector because there are no idiosyncratic shocks among people of the same age, gender and education. So $\omega = [G, p, a_1, \dots, a_m]$. Let $\Pi(\omega'|\omega)$ be the transition density for ω implied by the the exogenous processes G and the laws of motion for assets and the demographic distribution p . Appendix A reports the equilibrium definition.

4 Calibration

The model period is a year and parameter values are reported in Tables 7 and 8.

To make a sharp comparison between the model predictions and the data, the model is simulated with the shocks identified from the data. It is useful to start with the production function and the identification of the shocks.

4.1 Production

θ , the parameter governing the elasticity of substitution in the production function is calibrated to -0.25 in accordance with the literature that suggests a parameter value which induces more complementarity than the CD case. See for instance [Leon-Ledesma et al. \(2010\)](#) and [Choi and Ríos-Rull \(2009\)](#). α is such that the model predicts the average labour share found in the data: in the model, the capital share of output

²⁷Starting 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.

is $A\alpha\left(\frac{k}{y}\right)^\theta$. α is identified by normalizing total factor productivity A to be one on average. This gives $\alpha = 0.43$.²⁸ The time series for A and z_i are jointly identified through the production function and labour demands (Equations (4)—(7)) given the raw data for output, factor inputs and wages. These equations are simultaneous but conceptually, the time series for A is backed out as a Solow residual, and those for Z from the labour demand equations. With these time series in hand, I estimate the following AR1 process for each z_i and for A by Ordinary Least Squares:

$$\log(z_{i,t}) = \gamma_i + \rho_i \log(z_{i,t-1}) + u_{i,t}, \quad (9)$$

where $u_{i,t}$ are the shock innovations. The AR1 process for $\log(A)$ is:

$$\log(A_t) = \rho_a \log(A_{t-1}) + u_{a,t}. \quad (10)$$

Time trends in Equations (9) and (10) were not significant and therefore omitted.²⁹

Parameters for the productivity processes are summarized in Table 8.

4.2 Preferences, depreciation and survival probabilities

With the utility function specified in Equation (2), The Frisch elasticity of labour supply for a generic individual (e.g. a woman, with high education and a given age) is

$$\varepsilon \equiv \frac{\partial \ell}{\partial w} \frac{w}{\ell} \mid \lambda = \frac{1}{\sigma} \frac{(1 - \ell - \tilde{h})}{\ell}, \quad (11)$$

where λ is the lagrange multiplier on her budget constraint.³⁰

²⁸The 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 labour 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 .

²⁹Figure 5 shows the identified time series, the initial trends depend on the processes starting below steady state. As discussed in Section 4.4, this model accounts for much of the growth in the raw data through transitional dynamics, not through steady growth in productivity. So the model variables in levels eventually converge (very slowly given the demographic trends) to a steady state with no need to remove a balanced growth path.

³⁰The first order condition for labour of a generic individual is

$$\chi(1 - \ell - \tilde{h})^{-\sigma} = \lambda w. \quad (12)$$

Differentiating this last condition with respect to w , but holding λ_i constant, one gets

$$\frac{\partial \ell}{\partial w} \frac{w}{\ell} = \frac{\lambda}{\chi \sigma} (1 - \ell - \tilde{h})^{1+\sigma} \frac{w}{\ell}.$$

substituting out $\frac{w}{\ell}$ in the left-hand side from (12) one gets Equation (11).

I calibrate σ to have equal average Frisch elasticities of labour for men and women over the simulation.³¹ I also restrict σ to have equal average Frisch elasticities between young, prime age, and older workers.³² I let Frisch elasticities of the highly educated be lower than that of the less educated to help match the ratio between the volatilities of total market hours by education.³³

Finally, these parameters imply an average elasticity over the simulation equal to 1.1. This number lies within the range of micro estimates for women and it is also within the estimates for men, yet above the mean (which is 0.85 in the studies surveyed by [Keane \(2011\)](#)), but it is much lower than macro calibrations.³⁴ [Keane \(2011\)](#) also argues that despite the large literature, no clear consensus has emerged yet on these elasticities, and many studies for male labour supply elasticities are likely biased toward zero.

These restrictions imply that σ has to be gender, age and education specific (for instance, since from Equation (11) the Frisch elasticity of labour is decreasing in hours and men work more hours than women, men have to have a lower parameter to have the same Frisch elasticity of labour as women).

With this parametrization, the model qualitatively matches the fact that female total hours are less volatile than that of males even though men and women have the same labour Frisch elasticities; the ratio of the standard deviations of filtered hours worked by women over that of men is 0.66 in the data and 0.83 in the model.

The model also matches the fact that prime age hours are less volatile than those of the young and the old even though labour Frisch elasticities are not age specific: the standard deviation of filtered total hours of prime workers over that of the more volatile young and older workers is 0.70 in the data and 0.51 in the model simulations. This

³¹These are the averages of the Frisch elasticities for each representative agent in a given group, weighted by their mass. For instance, at each time t , the elasticity of male workers is $\frac{\sum_{j=1}^J (1/\sigma_{m,j,h} (1-\ell_t(j,m,h) - \bar{h}_{m,h,t}) / \ell_t(j,m,h)) p_t(m,j,h) + 1/\sigma_{m,j,l} ((1-\ell_t(j,m,l) - \bar{h}_{m,l,t}) / \ell_t(j,m,l)) p_t(m,j,l)}{\sum_{j=1}^J (p_t(m,j,h) + p_t(m,j,l))}$. These statistics are computed period by period in the simulation, then their averages over the time is taken.

³²The elasticity is assumed constant by age in accordance to [Jaimovich et al. \(2013\)](#) who argue that the U-shaped volatility of hours as a function of age is due to age specific sensitivity of labour demand, not labour supply.

³³As discussed below, that elasticities are education specific may reflect unmodeled labour contracts, matching frictions, and non market opportunities specific to types of worker. They could also be due to genuine differences in the characteristics of the population, which correlates with educational choices.

³⁴The business cycle literature typically chooses a high level of Frisch elasticities in order to have sufficient aggregate hours volatility, which tends to be lower than that found in the data. The lower level here makes the elasticity in Equation (11) more sensitive to changes in l , which helps match trends in hours volatility. See [Ljungqvist and Sargent \(2011\)](#), [Prescott et al. \(2009\)](#) and [Erosa et al. \(2015\)](#) among others for a discussion of how extensive margins can be incorporated into a life cycle model to reconcile micro and macro labour supply elasticities.

result supports the findings of [Jaimovich et al. \(2013\)](#) that the U-shaped age volatility of labour is not due to differences in labour supply elasticity, but to demand factors: they micro-found labour demand differences through a production function that exhibits experience-skill complementarity similarly to [Krusell et al. \(2000\)](#) and [Castro and Coen-Pirani \(2008\)](#). It is not obvious how to extend that function to the case of several labour groups with different complementarities. So here age specific labour demand differences are due to productivity shocks specific to each group identified through the data as explained below. The variance of labour productivity shocks by age implies less volatile wages for prime age workers, which is also true in the data.³⁵

While the model naturally matches the relative volatility by sex and age without having sex or age specific Frisch elasticity, this is not the case for education. The average elasticity of workers with low education is 1.33, that for workers with high education is 0.26. The reason why education specific preferences are necessary is that while hours of workers with higher education are less volatile than those of the lower educated, the opposite is true for wages: wages are more volatile for workers with high education. This pattern has also been documented in an online appendix by [Jaimovich et al. \(2013\)](#). Within a neoclassical setting where the labour supplied is always a choice, the only way to match the relative volatility of hours and wages at the same time seems to have lower labour supply elasticity for workers with higher education. Indeed, according to [Keane \(2011\)](#), it is plausible that education is related to tastes for work and [Blau and Kahn \(2007\)](#) estimate that labour supply elasticity is lower for women with higher education. [Becker and Mulligan \(1997\)](#) argue that education has explanatory power on preferences and even that education might change preferences. While this theory does not explain why elasticities differ, [Section 5.2.3](#) shows that the model predicts the evolution of relative hours and wage volatilities over time. This seems reassuring given the measurement goal of this paper.

The parameters χ in the utility function are such that the steady state matches the hours distribution by sex, age and education reported in [Table 6](#). This distribution is based on the average hours per capita by gender, age and education between 2000 and 2007, but with a caveat. In the model, workers of each age can either be of high or low education. Therefore, life starts at age 22, after college. However, [Jaimovich et al. \(2013\)](#) found that workers between 15-19 account for a non negligible share of

³⁵If anything, the model generates a too pronounced U-shaped volatility as a function of age. As explained below, sensitivity analysis suggests this has a negligible role for how business cycle volatility has evolved over time.

total hours volatility. Missing their hours share could reduce the contribution of age demographics to aggregate volatility. To avoid this, I impute to young workers (from age 22 to 29) the hours share of workers from 15 to 29. This way, hours by age groups are consistent with Section 2, where the young go from age 15 to 29. Since workers from 15 to 21 work a small number of hours, this does not majorly change the properties of labour supply over the life cycle. For instance, labour supply is hump shaped as it can be appreciated by Table 6, first 3 rows. Furthermore, this disproportion in the hours per worker of the young does not affect Frisch elasticities of labour supply as σ is age-group specific and is calibrated taking this into account.

It should also be noted that since the model does not have an employment margin, to construct hours, I multiply hours per worker by the share of people working. For instance, for prime male workers with high education, the average time spent working by those who work between 2000 and 2007 is 41.1% of their time (time is defined as 15 hours a day times 365). The working-population rate is 92.4%. So steady state hours for prime age workers with high education is 0.411 times 0.924, or 38%.

Table 6 reports $\sigma_{g,j,e}$, $\chi_{g,j,e}$, the hours targets, and the elasticities by group. Since $\chi_{g,j,e}$ and the elasticities actually varies for each age j , the table shows the average in each age bracket (young, prime and old).

Discount factor β is equal to 0.99 and capital depreciation δ is 0.06. With these values, and given the parameters of the production function, the steady state capital output ratio is 2.9, the net interest rate 4.7% and the saving rate 18%.

Survival probabilities ζ_j for $j = \{1, \dots, J\}$ come from the National Center for Health statistics Vital statistics of the US, 1992.³⁶ Since this paper focuses on active workers, attention is restricted to people from age 22 to 65. Therefore, no one can live for more than 44 periods and $\zeta_{44} = 0$. This is counterfactual, and it could affect the labour supply decision over the life-cycle. In particular the absence of retirement implies a counterfactual wealth distribution (agents die with zero assets) and it has implications for the response to shocks. A retirement decision could further increase the volatility of older workers, see Prescott et al. (2009). However, even without retirement, the model generates too much volatility of hours by the older workers. Furthermore, non reported sensitivity analysis suggests that the Frisch elasticity of old workers has negligible effects on how business cycle volatility has evolved. Intuitively, this is because the share of hours of the old is smaller and more stable over the sample than that of young

³⁶Vol II, sec. 6 life tables page 13, Washington: Public Health Service. 1996.

and prime workers as shown in Figure 1.³⁷

4.3 Trends in the composition of labour

The number of new born p_0 is modeled as an exogenous AR1 process with constant γ_{p_0} , persistence ρ_{p_0} , and variance $\sigma_{p_0}^2$, all reported in Table (7). This quite simplistic way to model birth generates variable fertility that matches changes in the labour distribution by age; see Ríos-Rull (2001) for a discussion of alternative fertility regimes.

$\tilde{h}_{g,e}$ is zero for men and it starts positive and converges to zero for women. The initial level and persistence of the deterministic AR1 process for female housework is calibrated to replicate trends in the labour composition by sex. Its initial level of 0.14 accounts for 52% of the average total (home and market work) working time for women. The gradual decay of this variable “liberates women from the home” and increases their market labour supply and educational incentives. See Greenwood et al. (2005).

q^m and q^f ; the newborn men and women with high education, are modeled as deterministic AR1 processes with parameters $\gamma_{g,e}^q$ and $\rho_{g,e}^q$ reported in Table (7). Their steady state levels are such that the model matches the share of the highly educated by gender between the young in the period 1999-2007, which is $q^f = 0.36$, $q^m = 0.29$. Initial conditions are 0.13 for men and 0.08 for women and persistence are picked to replicate trends in the labour composition by education. Since there are no shocks affecting housework and q^m and q^f , agents have perfect foresight on these trends.

4.4 Initial conditions

As mentioned, initial conditions for home production, the size of the new generation, education shares are picked to match gender, age, and education trends. Initial values for the productivity parameters are identified through the estimation of the production function. The initial population distribution reflects the initial distribution in the data.

It remains to pin down the assets’ distribution. I first take the values that solve the model for a steady state assuming that the other initial conditions were stationary.³⁸ Asset levels so determined seem a bit large: when I simulate the model over the transition with these initial conditions, agents disinvest in the first few periods. The intuition for this is that the initial conditions for gender, age and education trends

³⁷This could change in the future given the general increase in life expectancy.

³⁸This is a convenient way to pick initial assets, but to hold the initial demographic distribution stationary is inconsistent with general equilibrium. Sensitivity analysis is discussed in Section 5.3.4.

imply future growth, a positive wealth effect relative to the assumption of no growth implicit in the initial steady state; therefore agents respond by reducing investment. Therefore, I put initial assets to 85% of those in the initial steady state. There is no initial disinvestment of capital, but the business cycle statistics are qualitatively unchanged. Furthermore, with these initial conditions, the initial capital-output ratio is of 2.9, which is very close to the empirical counterpart. Importantly, the average growth rate of output over the simulation is quite close to the empirical one for output over population in working age: 1.53% in the model and 1.60% in the data.³⁹ This is striking because the model does not have exogenous productivity growth: according to this model, virtually all the growth we observed in the last 5 decades is due to the transition from initial conditions. Furthermore, this is only partly due to the fact that initial assets and productivity shocks are below the final steady state, inducing a growing transition.⁴⁰ The counterfactual exercises ran in Section 5.3 show that important factors for the observed growth are that women increased their labour supply and that there are more prime age and highly educated workers.

5 Model-based quantitative analysis

5.1 Computation

The computation of this model is challenging because the state space is quite large: 890 variables, of which 365 are state variables. Large DSGE models can be handled by perturbation methods around the steady state.⁴¹ This method is reasonably accurate when simulations remain close to the deterministic steady state. 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, a new methodology is developed which essentially consists of applying repeated local approximations over the entire transition path, between the initial conditions and the steady state.⁴²

³⁹The comparison with output over population in working age is more appropriate than the usual output per capita as the model only includes agents in working age. Output over population in working age is constructed through income data from the CPS.

⁴⁰For instance, if initial assets are 95% of those in the initial steady state, the initial capital-output ratio is 3.2 and output grows at an average of 1.50%.

⁴¹See [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.

⁴²By limiting the local approximations to the transition path, the number of grid points does not increase with the number of state variables.

The algorithm is detailed in [Mennuni and Stepanchuk \(2016\)](#); a brief summary of its logic is given here. The goal is to find the equilibrium path given initial conditions for the state vector, called 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). First, an auxiliary 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. Then, the algorithm iterates treating the new point as an initial condition and the iterations end when the initial conditions coincide with the final period of the time series of the shocks.

5.2 Assessing the model

Before carrying out the main experiments for which the model has been built, some checks are performed in order to get a sense of how satisfactory a description of the economy it provides, at least for some relevant dimensions. The facts with which the model is confronted below have direct implications for the counterfactual experiments aimed at quantifying the importance of labour reallocation to aggregate volatility.

5.2.1 Aggregate volatility and co-movements

The model is calibrated to match labour composition changes: indeed [Figure 3](#) shows a good match with the empirical shares of market hours by sex, age and education.

Since the model is restricted through demographic changes and most of the shocks are identified through disaggregated labour data, it is interesting to see whether – simulating the model with the identified shocks – it replicates the changes in volatility, as well as the typical moments used to test RBC models.

As before, the model and data are HP-filtered. The standard deviations of output over the whole sample (67-13) is 1.10 in the model and 1.48 in the data. So the model accounts for 74% of total volatility, more than what is typically accounted for by RBC models: see for instance [Prescott \(1986\)](#). The standard deviation of total hours is 36% that of output as reported in [Table 9](#): although much lower than in the data, this is an improvement to typical business cycle models given that the average elasticity is 1.1, see [Ríos-Rull et al. \(2012\)](#). A feature that contributes to this higher volatility is that the model distinguishes between TFP and labour augmenting shocks, which

induce shifts in the labour demands of each group. Furthermore, the estimation of these shocks turns out to be less persistent than the Solow residual of an aggregate CD function (see table 8). The lower the persistence of shocks is, the lower the offsetting wealth effects they induce and the greater the labour response.

The model also predicts a volatility slow down between the first and second subsamples (67 – 84 vs 85 – 07) of $(1.49 - 0.87)/1.49$ or 42%: quite close to the 50% decline documented in the literature, see for instance Arias et al. (2007).

Table 9 reports standard deviations and correlations with output of consumption, investments and total hours. Consistently with national income data (and in line with other RBC models), this model predicts that while consumption is less volatile than output, investments are much more volatile. 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).

5.2.2 Wages, labour share and the employment-productivity puzzle

As mentioned, this model reconciles the employment-productivity puzzle —the near-zero correlation between hours and labour productivity— which lies at the root of an important critique to the RBC model which predicts a high positive correlation. See Galí (1999). In this model, the correlation of HP-filtered labour productivity (output over total hours) and total hours is 0.11. As explained in Section 3.3, this result relies on the distinction between labour-specific and TFP shocks. Figure 5 plots the time series of the labour shocks z and the TFP shock defined as $A^{\frac{1}{\theta}}$.⁴³

The wedge between the marginal productivity of labour and output is also reflected in the labour share and wages: the fit with the empirical labour share, shown in Figure 4 last panel, shows a notable improvement to the constant labour share predicted with the benchmark CD technology. The Figure also shows actual versus predicted hourly wages by sex, age and education. In particular, despite the fact that female and highly educated labour supply increases over time, the model predicts a falling gender wage gap and an increasing education wage premium from 1979-1980 as documented, for instance by Card and DiNardo (2002) and Krusell et al. (2000). The reason is that the labour augmenting productivity processes specific to each group are identified through the wage equations (wage equal to marginal productivity for each group) taking into

⁴³Since θ is negative, an increase in A has a negative effect on output as is obvious from Equation (4). For this reason, TFP is defined as $A^{\frac{1}{\theta}}$ which has a positive impact on output.

account the increase in female and highly educated hours. So, to the extent that the model generates the right hours and capital time series, it also predicts the right wages.

Appendix B offers an analytical illustration of how labour marginal productivities are affected by the shocks.

5.2.3 Trends in labour elasticities

It is also interesting to see whether the model predicts how hours and wages volatilities have evolved over time. Figure 6 shows these trends in the data and in the model. Besides matching relative volatilities on average, the figure also shows that the model predicts the time path by gender, age (prime over young and old) and education for both hours and wages. The path for this ratio in volatilities is roughly stable over time for gender and age and increasing for education (*i.e.* highly educated workers are becoming more volatile). This result suggests that the model gives a good account of how group specific volatility changes as a function of demographic change. This is important because the reduced-form methodology essentially assumed that group specific volatility is independent from the demographic composition. The model offers a way to test for the robustness of the results of Section 2 to alternative assumptions.

The model's ability to replicate relative hours volatility by sex relies on the utility function adopted, which implies that as women's labour input increases over time, their Frisch elasticities of labour supply declines (Equation (11)). This feature is important to predict the essentially trend-less path in the relative volatility of female over male hours in Figure 6. This role of declining Frisch elasticities can be appreciated by observing Figure 7 which shows that an alternative utility function with constant Frisch elasticities implies that the volatility of female hours is increasing relative to that of men. Instead, declining female elasticities neutralize this upward trend in relative hours volatilities.⁴⁴

As mentioned, this decline in female labour supply elasticities is consistent with the findings of Heim (2007) and Blau and Kahn (2007). 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

⁴⁴Time varying elasticities play a negligible role for the path of relative hours volatility by age and education; intuitively this is because both in the model and in the data, the largest change in labour per worker (and thereby labour supply elasticity) characterizes women while changes for males, age and education groups are more modest.

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; that notwithstanding, the model suggests that the elasticity decline was indeed due to demographics through its effect on labour input and its implication on the labour supply Frisch elasticity.

It should be noted that in this model female labour trends come from each woman working more hours. Instead, part of the increase in female total hours is due to participation decisions. This is a weakness because the elasticity in Equation (11) only involves the intensive margin. For instance, distinguishing between an intensive and an extensive margin would affect results if they behaved differently. However, tables 3 and 4 show that at least qualitatively, the trends in employment over population by gender, age, and education evolved similarly to hours per worker. Furthermore, column 3 in these tables show that both hours per worker and employment over population by women, prime age and the highly educated are less volatile than their respective counterparts (males, other age groups and workers with lower education). Finally, the last column in the two Tables show similar patterns in the volatility decline in the Great Moderation period. As suggested by Heim (2007) and Blau and Kahn (2007), a possible intuition for a negative relationship between labour supply and its elasticity, not just at the intensive, but also at more extensive margins is that as more people participate, the number of people with reservation wage close to their market wage gets smaller and the participation elasticity will be small.

Furthermore, Section 2 finds that the most important distinction is between population margins –which characterized the changes by age and education– and the other two (employment and hours) which mattered for gender. The model merges employment and hours, but the population margin is kept separated from the other two.

While reassuring, this evidence does not mean that distinguishing between employment and hours would not matter. Indeed the drop in female total hours' elasticity between 1980 and 2000 is one quarter in the model, which is lower than the one half drop estimated by Blau and Kahn (2007): this suggests that the model gives a conservative estimate of the overall effects this elasticity channel has on aggregate volatility.

Finally, with the adopted production function, compositional changes also affect labour demand elasticities: as shown in Appendix B, the larger the labour input of one group, the lower the elasticity of the labour demand of that group to its labour shock and the more sensitive the labour demands of the other groups. Arguably, this is an

interesting channel for the sake of this study. For instance, this model abstracts from the fact that the service sector, which is less volatile than other sectors, is female intense and has grown throughout the sample. Had female labour not increased, perhaps more men would have entered that sector reducing average male hours volatility. The model captures this: had female labour supply not increased, men would make a larger share of hours and respond less to their shocks, which are the most volatile. This is as if some men would have entered less volatile sectors. This mechanism, however, has small quantitative implications. Similarly, [Blanchard and Simon \(2001\)](#) and [McConnell and Perez-Quiros \(2000\)](#) find that sectorial shifts do not explain aggregate volatility.

To summarize this section so far, the model is broadly consistent with a set of facts at the aggregate and at the group level. Even though the model inevitably lacks some important aspects of reality, it explains the evolution of the variables which are likely to matter for the questions of this paper. This is reassuring given that next I rely on the model to predict how the agents would have responded in the counterfactual experiments. This is especially interesting given that the model finds that group specific volatility is not independent from demographic change as assumed by the reduced form accounting exercises in [Section 2](#).

5.3 Counterfactual Experiments

5.3.1 Removing all trends

The counterfactual experiment consists of changing the trends in the amount of housework $\tilde{h}_{g,e}$, the birth rate dynamics p_0 and the education shares of the population— q^m and q^f — while maintaining the shocks A and Z as in the baseline.

The initial population distribution is the steady state one implied by the steady state levels of p_0 , q^m and q^f . In order to construct the counterfactual with trend-less hours' shares by sex, the time cost $\tilde{h}_{g,e}$ has to be increasing over time for women, i.e. start below steady state: equal to -0.5% of total time for women with low education and -11% for highly educated women. Evidently, increasing time costs contrast the fact that female wages are fast increasing, especially for the highly educated.

Initial individual asset holdings are constructed in the same way as in the baseline simulation as detailed in [Section 4.4](#): they are set equal to 0.85 times the assets obtained by solving for a steady state taking as given the initial demographic distribution.

[Figure 8](#) shows market hours shares by gender, age and education in the counter-

factual and original simulations: as can be seen, counterfactual hours are essentially trend-less.⁴⁵ Table 10 contains the standard deviation of HP-filtered output and market hours during the sub-samples of interest: the period before the Great Moderation 1967-1984, and the period of the Great Moderation: 1985-2007. 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: 2000-2013 (last column).

As can be deduced from the first two rows in the table, output volatility in the first turbulent 10 years of the sample is $\log(1.76) - \log(1.36)$, or approximately 23% lower in the counterfactual than in the baseline, the discrepancy is 18% until 1984. Thus, in the absence of these trends, volatility is much lower in the first sample, when volatility was very high. Instead, during the Great Moderation (85–07) counterfactual volatility is 2% higher and it converges to the baseline one in the last years of the sample (as is natural given that the labour composition also converges to the baseline toward the end of the sample). Thus in the counterfactual, output volatility is much more stable over time. The same holds for total hours as shown in the third and fourth rows.

To get a visual sense of how volatility is affected over time, Figure 9 shows the time trend of baseline and counterfactual cyclical volatility measured as a 3-year roll over standard deviation of filtered output.

To get a concise statistic that quantifies the amount of the Great Moderation explained, I proceed as in section 5.2.1, comparing the standard deviations in the first sub-sample (1967-1984) and in the second one (1985-2007). Between the two sub-samples, aggregate output volatility decreased by $(1.493 - 0.867)/1.493$ or 41.9% in the baseline simulation. Had the shares remained stable as in the counterfactual, we would have observed a reduction in volatility of $(1.243 - 0.884)/1.243$ or 28.9%. Therefore, these demographic changes account for $(41.9-28.9)/41.9$ or 31% of the moderation in output.⁴⁶ The same statistic for total hours is 12.3%.

⁴⁵Small low frequency movements are still notable in hours shares by sex. This is because hours trends by sex are fully determined by endogenous labour supply decisions as opposed to trends by education and age, which come mainly from exogenous birth rates. As a result, it is not possible to make hours by sex fully constant simply with a deterministic AR1 time cost. However, non reported simulations confirm that the results reported below are robust to small changes in these trends.

⁴⁶What not accounted for by demographics is due to a lucky draw of small innovations or “good luck”. This is similar to Arias et al. (2007) and Smets and Wouters (2007).

5.3.2 Removing trends one by one

What is the importance of each of the three factors? Table 11 includes the outcomes of counterfactual experiments where only one of the long-run trends is removed.⁴⁷ The first column reports the share explained of the Great Moderation by each demographic trend: the numbers suggest that education trends are the most important for the Great Moderation in output, followed by trends in age and sex. Trends in age are the most important for the moderation in hours.

The statistic in column 1, which is computed the same way as in the previous subsection, is concise and conventional in the literature. However, it is also instructive to see how volatility would have differed over time: columns 2–5 report the volatility ratio relative to the original simulation for each of the sub-samples considered. Trends by education and sex matter primarily because they decrease counterfactual standard deviation in the initial most volatile period: for instance, if initial education composition was as in steady state, the standard deviation of filtered output would have been $\log(1.714) - \log(1.438)$ or 17.5% lower in the most turbulent initial decade of the sample. Age trends matter by reducing volatility in the first sub-sample, but especially by increasing counterfactual volatility in the Great Moderation period (fourth column): the standard deviation of counterfactual filtered output is about 3% higher without the disproportionately high share of prime age workers between 1985 and 2007.

Overall, these results are qualitatively similar, but more conservative than those in section 2. For instance, the Great Moderation in market hours due to trends by sex is 5.4% in the reduced form exercise (Table 5) and 2.1% here. The effects by age go from 10.3% in the reduced form to 7.9% and those by education from 15.5% to 4.3%.

Results from both methodologies challenge the findings of the previous literature by arguing that that changes in the gender and especially education composition are more important compared to previous findings while changes in the age composition, while important, did not play as large a role as was previously believed. That the two methodologies are broadly consistent is reassuring given that a possible issue with the reduced-form analysis was that it assumed no change in volatility by group in response to a change in group sizes: results are robust to different assumptions on how group

⁴⁷Removing only trends by sex, hours shares by age and education are quite close to the baseline in Figure 8, while the shares by sex are close to the counterfactual ones. So there are small indirect effects. Indirect effects are also small when removing trends by age or education.

specific volatility changes as a function of demographic change.^{48, 49}

5.3.3 Growth implications

Finally, an interesting side result from this model is to measure the contribution for growth of these demographic trends.

The main factors contributing to growth in this model are the demographic trends, the path of the shocks, and the initial asset conditions. To disentangle the role of demographics for growth, I compare growth with and without demographic trends, holding the shocks and individual asset holdings as in the baseline.

In the absence of sex trends, average annual output growth between 1967 and 2013 would have been 1.28% as opposed to 1.53%. Therefore $(1.53-1.28)/1.53$ or 16.3% of the average growth rate is due to sex trends. Average output growth would have been 1.39% without age trends, and 1.27% without trends in education. Therefore 9.2% of the average growth rate is due to trends in age and 17% to education trends. Finally, removing sex, age and education trends jointly accounts for a staggering 39% of the average annual output growth. As mentioned in the introduction, the model ignores elements that could affect these results and more work should be done. However, the numbers appear conservative when compared to the literature, which provides estimates for each of these dimensions singularly: for instance [Greenwood et al. \(2005\)](#) find that the rise in female labour participation accounts for about 19% of growth between 1900 and 1980. [Feyrer \(2007\)](#) finds that changes in the labour composition by age explain almost 12% of GDP growth between 1990 and 1995. The number for education is close to the 15% found by [Goldin and Katz \(2008\)](#) for human capital: they argue that their accounting exercise is likely to be an understatement because it

⁴⁸That the endogenous response of hours volatility to demographic change does not play a primary role for the aggregate volatility changes is corroborated by simulations with the alternative utility function with constant Frisch elasticity, which reports similar results to the ones in Table 11: for instance, gender, age, and education trends account for 26% (31%) of the great moderation in output with constant (endogenous) Frisch elasticity. Of course there are calibrations of the utility function that induce larger endogenous labour supply elasticity responses and larger results. But then the model would not match trends in relative hours volatility.

⁴⁹One reason why endogenous elasticities do not change results that match is due to the behavior of wages, which contrasts the role of demographics on aggregate volatility: The change in volatility of w_i is larger in the counterfactual. Intuitively, the more elastic labour supply, the smaller the effect on wages of shifts in labour demand. With constant demographics, hours are relatively less sensitive in the first sample and more sensitive in the Great Moderation period. Then the opposite should happen to wages: more sensitive in the first sample, and less in the moderation period. This leads to larger volatility changes in the counterfactual.

misses indirect effects. See also [Acemoglu and Autor \(2012\)](#).

Sadly, to the extent that the demographic composition is converging toward a stationary distribution, we should expect less growth in future.

5.3.4 Sensitivity

These results are quite robust to sensible changes in initial assets. As mentioned, initial individual asset holdings are as in the baseline. Intuitively, this reflects the hypothesis that demographics had no bearing on initial savings; this helps isolating the role of demographics from initial conditions. However, one could argue that demographics has some bearing on initial asset holdings: without the demographic transition, the model would start closer to its steady state, so also initial assets should be closer to steady state. Results are robust to this alternative: assuming that initial individual assets are 0.95 times their initial steady state rather than 0.85, the average output growth is virtually unaffected. Initial assets have little impact also for aggregate volatility, for instance, with the latter initial conditions, sex, age and education trends account for 29.3% of the moderation in output rather than 31.0%.

Initial conditions for TFP and labour shocks are identified from the data. Nevertheless it is instructive to run a simulation with different initial conditions. I choose the steady state values implied by the AR1 processes estimated for A and Z . I keep the innovations I have identified from the estimation. Perhaps not surprisingly, the most notable fact is that output growth is now much lower: the average growth is 1.32 rather than 1.53. Interestingly, the business cycle implications of the model are essentially unchanged: for instance, the Great Moderation in output is of 42%, as with the baseline initial conditions.

Conclusions are robust to changes in the parameter values concerning the complementarity between capital and labour in the production function, the average elasticity of labour supply (where values ranging between 1 and 1.5 have been considered) and changing the age classes (moving the young-to-prime age threshold from 30 to 34 and the old threshold from 59 to 54).

5.4 Relationship with regression analysis

These results confirm the finding of J-S that the age composition has an important impact. However they suggest that age trends are not as important as previously thought but that they are largely complemented by trends in sex and education.

J-S use regression analysis, which —for the case of age trends— has 2 important advantages: 1 they convincingly argue that age trends are exogenous to the business cycle, 2 they exploit a lot of variation in the age composition, both within and across countries: this makes it less likely that the correlation between aggregate volatility and the age composition is just spurious. Sex and education labour composition trends do not have these properties. This motivates the alternative methods, which are not based on the computation of the correlations but exploit the structural mechanism of changing shares of groups of different volatility. But then a discussion of the differences in the results is warranted. In particular, even though results from the model and the reduced form methodology are in the same range, the reader may be left wondering whether the result depends on the fact that both the structural and the reduced form methodology 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.⁵⁰ The regression considered is:

$$\sigma_t = \alpha + \gamma \mathbf{share}_t + \varepsilon_t \quad (13)$$

where σ_t is 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).^{51, 52} Following J-S, results are reported for heteroskedasticity and two-period autocorrelation robust standard errors constructed using the Newey-West estimator. See Table 12.

The coefficient γ is significant and it implies a stronger role for age than the true one in the artificial economy: when the independent variable **share** moves from its first sample average (55.6%) to its Great Moderation average (48.8%), the relative change in volatility is 12.3%. This is 12.3/42 or about 29% of the overall Great Moderation in the model: quite close to the result found by J-S. Since in the model, age does not

⁵⁰However, this case does not necessarily invalidate the regression analysis in J-S: endogeneity might work differently in real life data. Furthermore, this paper uses a different data set, different methodology, and only considers the US (versus a panel of countries).

⁵¹ σ_t is constructed by taking a 5-period standard deviation of the residuals between log output and its HP-trend. Conclusions are unchanged using growth rates or a 3-period rolling window.

⁵²Reverse causality is not an issue because population shares are exogenous by model construction.

matter this much, the result suggests that the regression overstates the role of age.

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. Indeed the explanatory power the regression attributes to age trends is of the magnitude that gender, age and education trends have jointly in the model. Can one correct this possible omitted variable bias by simply adding gender and education labour composition regressors? It depends on whether they are actually exogenous or endogenous: while the age composition of the population is independent to the current business cycle, this may not be the case for the workforce by gender and education, which are choice variables that cannot be easily instrumented. Indeed, to motivate the structural methodology, it has been argued that movements in the labour composition by gender and education may be partly endogenous.

Unfortunately it is not possible to measure the endogeneity bias in the true data as one would need to know the unbiased effects of the regressors. However, it would be interesting to get a sense of the endogeneity at least in the model.⁵³ While it would be difficult, if not impossible, to identify and measure the effects of the various channels that induce endogeneity in the model, it is possible to get a sense of the overall importance of these endogenous implications by seeing if the gender and education regressors produce biased estimates. With this aim, I augment Equation (13) by including gender and education regressors (including the share of males and low education over total hours):

$$\sigma_t = \alpha + \gamma_{age}\mathbf{share}_{age_t} + \gamma_{edu}\mathbf{share}_{edut} + \gamma_{sex}\mathbf{share}_{sext} + \varepsilon_t. \quad (14)$$

The results from the estimation of equation (14) are reported in table 13.

The age coefficient is no longer significant, the one for gender is positive and significant, but it overstates the role sex has in the model, and the education coefficient is negative and insignificant: clearly a biased estimate of the role of these variables in the

⁵³There are several channels through which the labour 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 age and education are affected by aggregate risk. Labour supply is also affected by wealth, which in turn is affected by aggregate risk. Furthermore, the effect of aggregate risk on the hours composition may vary over time with the movements in labour supply elasticities. Appendix B uncovers a further impact of groups size on labour demand elasticities. Notice that these are examples of reverse causality, which is a different issue than the one affecting the analysis in Section 2: there the main concern was that groups volatility were endogenous to group size, but —contrarily to the regression— that analysis does not require that any movement in the shares is orthogonal to the business cycle. Given this distinction, it is not surprising that the two methodologies give different results.

model. So while excluding these variables generates an omitted variable bias, including these endogenous regressors also generates a bias. An argument for ignoring gender and education trends is that the correlation with trends by age and with output volatility may be spurious. This possibility cannot be ruled out with time series evidence alone (instead, for the case of age, J-S exploit data variability both over time and across countries which makes it highly unlikely that the correlation with output volatility is just spurious). However, in the model, that the correlation is spurious is rejected by construction.⁵⁴ That notwithstanding, perhaps gender and education trends affect the results at least partly through a common trend with aggregate volatility (spurious relationship). Ignoring the common trend might lead to overestimating the role of gender and education while underestimating that of age. Adding a time trend all variables become insignificant and the age coefficient takes the wrong sign.

Finally, it is worth mentioning that running the same regressions on the true data gives results that are essentially the same as those ran on model data. This could suggest that the way endogeneity unfolds in the true data may not be that dissimilar than how it does in the model.

6 Conclusion

This paper documents changes in the composition of the labour force by gender, age and education and finds that they account for a significant portion of the change in aggregate volatility over the past half century. In particular, the paper argues that the changing shares of young and prime age workers had a significant but smaller role than previously believed, while two key trends are increasing shares of college educated workers, and to a lower, yet significant degree, the rising female labour force participation. The intuition for these causal effects is that — as is the case for partitions of the labour force by age — there are systematic differences in the sensitivity to the business cycle of labour market variables partitioned by education and gender.

These conclusions are consistently reached with a reduced form methodology and with a more structural analysis. That the two methodologies are broadly consistent is reassuring given the different assumptions on the identification of the demographic trends, on the extent to which these demographic changes affect the group-specific

⁵⁴And also in real life, it would be difficult to argue that changes by age have a causal effect on aggregate volatility while changes by gender and education do not given that the main mechanism that groups with lower volatility increased in size, characterized gender, age, and education alike.

sensitivity of market work to the business cycle, and vice-versa, on the extent to which some of the compositional changes (e.g. in employment and hours) may be affected by business cycle movements. The theory takes an explicit stand on these issues with implications that are consistent with the data.

Another achievement of the model is to measure the contribution to growth of these demographic trends. Intuitively, the increase in female labour and relatively more productive and hard working prime age and highly educated workers, contributed to the growth observed in the past half century: these trends account for almost 40% of the average output growth in the last 50 years. An important caveat is that in the model the long run demographic trends are due to exogenous causes but in the data they may be partly a byproduct of growth. Yet the estimates on the growth effects of demographic change by either sex, age or education, are conservative relative to existing estimates.

One challenge has been to find a solution method for this large model which guarantees sufficient precision over the dramatic demographic transition path that has characterized the last 50 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. The method is described in detail in [Mennuni and Stepanchuk \(2016\)](#).

Finally, it is worth mentioning that while this paper highlights positive implications of demographic change, the findings may matter for a variety of public debates that affect labour supply and its composition, ranging from tax, welfare, and pension reforms, to the regulation of institutional features of labour markets. Besides the policy implications, since demographic trends are fairly predictable, the uncovered implications on volatility and growth may help improve our forecasts.

References

- Abraham, K. G., Shimer, R., 2001. Changes in Unemployment Duration and Labor-Force Attachment. Russell Sage Foundation, pp. 367–420.
- Acemoglu, D., Autor, D., June 2012. What does human capital do? A review of Goldin and Katz's *Race between Education and Technology*. *Journal of Economic Literature* 50 (2).
- Angrist, J. D., February 1991. Grouped-data estimation and testing in simple labor-supply models. *Journal of Econometrics* 47 (2-3), 243–266.

- 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).
- 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.
- Becker, G. S., Mulligan, C. B., August 1997. The Endogenous Determination of Time Preference. *The Quarterly Journal of Economics* 112 (3), 729–58.
- 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.
- Card, D., DiNardo, J. E., October 2002. Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles. *Journal of Labor Economics* 20 (4), 733–783.
- 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.
- Christiano, L. J., Fitzgerald, T. J., 05 2003. The Band Pass Filter. *International Economic Review* 44 (2), 435–465.
- Clark, K. B., Summers, L., 1981. Demographic differences in cyclical employment variation. *Journal of Human Resources* 16.
- Doepke, M., Tertilt, M., 2016. Families in macroeconomics. *Handbook of Macroeconomics* 2, 1789–1891.
- Erosa, A., Fuster, L., Kambourov, G., May 2015. Towards a micro-founded theory of aggregate labor supply. Working papers, IMDEA.

- Feyrer, J., February 2007. Demographics and Productivity. *The Review of Economics and Statistics* 89 (1), 100–109.
- Fortin, N., Lemieux, T., Firpo, S., 2011. Decomposition Methods in Economics. Vol. 4 of *Handbook of Labor Economics*. Elsevier, Ch. 1, pp. 1–102.
- Galí, J., March 1999. Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review* 89 (1), 249–271.
- Goldin, C., May 2006. The Quiet Revolution That Transformed Women’s Employment, Education, and Family. *American Economic Review* 96 (2), 1–21.
- Goldin, C., Katz, L. (Eds.), 2008. *The Race between Education and Technology*. Cambridge, MA: The Belknap Press of Harvard University Press.
- 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, A. H., 1939. Economic progress and declining population growth. *The American Economic Review* 29 (1), 1–15.
- Hansen, G. D., Wright, R., 1992. The labor market in real business cycle theory. *Quarterly Review* (Spr), 2–12.
- Heathcote, J., Storesletten, K., Violante, G. L., 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).
- Hoynes, H., Miller, D. L., Schaller, J., Summer 2012. Who Suffers during Recessions? *Journal of Economic Perspectives* 26 (3), 27–48.
- Jaimovich, N., Pruitt, S., Siu, H. E., 2013. The demand for youth: Implications for the hours volatility puzzle. *The American Economic Review* 5 (103), 3022–3044.
- 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.
- Katz, L. F., Autor, D. H., October 1999. Changes in the wage structure and earnings inequality. 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.
- Keane, M. P., 2011. Labor supply and taxes: A survey. *Journal of Economic Literature* 49 (4), 961–1075.
- 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., Rios-Rull, J.-V., Violante, G. L., September 2000. Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica* 68 (5), 1029–1054.
- Leon-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., November 2012a. Demographic Change And The Great Moderation In An Overlapping Generations Model With Matching Frictions. *Macroeconomic Dynamics* 16 (05), 706–731.
- Lugauer, S., November 2012b. Estimating the Effect of the Age Distribution on Cyclical Output Volatility Across the United States. *The Review of Economics and Statistics* 94 (4), 896–902.
- Lugauer, S., 2013. The supply of skills in the labor force and aggregate output volatility. SSRN.

- Lugauer, S., Redmond, M., 2012. The age distribution and business cycle volatility: International evidence. *Economics Letters* 117 (3), 694–696.
- 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.
- Mennuni, A., Dec. 2013. Labor force composition and aggregate fluctuations. Discussion Paper Series in economics And Econometrics 1302, University of Southampton.
- Mennuni, A., Stepanchuk, S., 2016. Dynamic perturbation. Manuscript.
- 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.
- Ríos-Rull, J.-V., Schorfheide, F., Fuentes-Albero, C., Kryshko, M., Santaeulalia-Llopis, R., 2012. Methods versus substance: Measuring the effects of technology shocks. *Journal of Monetary Economics* 59 (8), 826 – 846.
- Sahin, A., Song, J., Topa, G., Violante, G. L., November 2014. Mismatch unemployment. *American Economic Review* 104 (11), 3529–64.
- 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.
- Summers, L., 2013. Why stagnation might prove to be the new normal. *The Financial Times*.

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: Employment over population: data and trends are solid lines. De-trended series are dashed lines.




Figure 3: Data vs model share of total market hours




Figure 4: Wages per hour by group and labour share of output.




Figure 5: TFP and labour productivity by group.




Figure 6: Hours and Wages volatility ratio by sex, age and education.

Notes: Dashed lines are model simulated data. The first row shows the ratio (female over male; prime over young and old; high over low education) of the standard deviation of HP-filtered total hours. The second row shows the same statistic for wages per hour weighted by hours by group. In each t , it is plotted the ratio of the standard deviations over a period of 15 years centered at year t .




Figure 7: Comparing relative hours volatility between men and women with utility function with constant Frisch elasticity of labour supply.

Note: Time series are normalized to one at the beginning of the sample to ease comparison.




Figure 8: Shares of hours, original versus counterfactual simulation




Figure 9: Output volatility, baseline versus counterfactual simulation with trend-less hours shares.

8 Tables

Table 1: Hours by Gender, Age and Education

	Share of hours 67-84	Share of hours 85-13	St. dev	$\frac{St.dev(85-07)}{St.dev(67-84)}$
	Age			
Young	28.99	22.19	2.61	81.23
Prime age	63.50	70.70	1.49	55.17
Old	7.51	7.11	1.79	97.09
	Gender			
Women	34.28	41.99	1.31	59.01
Men	65.72	58.01	1.98	66.87
	Education			
Low	79.99	69.00	2.02	57.17
High	20.01	31.00	1.05	113.87

Notes: numbers are expressed in percentage terms. Share of hours is the ratio of hours by each group to total hours. The data contain information on weekly hours and weeks of work in a year. There are two variables for weeks of work: one starting in 1976 that contains the number of weeks for each respondent. Another, which goes back to the beginning of the sample, classifies weeks between 1 and 6. The latter is used. Since both measures exist from 1976, bins 1 to 6 are converted into weeks by taking the average number of weeks of workers in each bin.

Table 2: Population shares by Gender, Age and Education

	population shares 67-84	population shares 85-13	St. dev	$\frac{St.dev(85-07)}{St.dev(67-84)}$
Gender				
Women	52.40	51.78	0.09	88.35
Men	47.60	48.22	0.10	85.97
Age				
Young	34.51	27.63	0.44	146.74
Prime age	45.86	51.26	0.21	59.10
Old	19.63	21.12	0.39	193.26
Education				
Low	87.78	78.40	0.18	115.88
High	12.22	21.60	0.83	74.40

Table 3: Employment over population by Gender, Age and Education

	Empl / pop 67-84	Empl / pop 85-13	St. dev	$\frac{St.dev(85-07)}{St.dev(67-84)}$
Gender				
Women	43.21	54.07	1.07	56.74
Men	70.95	67.38	1.34	66.23
Age				
Young	55.92	58.55	1.90	65.97
Prime age	70.76	77.11	0.92	62.47
Old	23.75	22.20	1.28	95.77
Education				
Low	53.42	55.76	1.40	61.18
High	77.74	77.79	0.57	92.48

Table 4: Hours per worker by Gender, Age and Education

	hours 67-84	hours 85-13	St. dev	$\frac{St.dev(85-07)}{St.dev(67-84)}$
Gender				
Women	27.1	30.0	1.29	50.4
Men	34.9	35.8	1.78	56.1
Age				
Young	26.92	27.37	2.26	61.98
Prime age	35.12	35.82	1.41	49.11
Old	28.84	29.74	1.94	108.70
Education				
Low	30.59	31.60	1.86	46.50
High	37.79	37.13	1.23	93.42

Notes: numbers are percentage terms.

Table 5: % explained of the moderation

	Gender	Age	Education
Removing population, employment and hours trends			
e	1.9	13.3	16.9
H	5.4	10.3	15.5
Yl	5.0	7.1	9.0

Table 6: Hours targets and utility parameters

	male low edu	male high edu	female low edu	female high edu
Hours ($\ell_{g,j,e}$)				
young	30.6	35.9	21.7	26.8
prime	30.8	38.0	21.8	27.2
old	17.5	31.4	10.7	20.4
$\sigma_{g,j,e}$				
young	1.6	6.1	3.0	10.7
prime	1.6	5.5	3.0	10.4
old	3.9	8.9	8.2	18.2
$\chi_{g,j,e}$				
young	2.3	0.2	3.0	0.2
prime	1.8	0.2	2.1	0.1
old	0.8	0.1	0.9	0.1
Elasticities ($\varepsilon_{g,j,e}$)				
young	1.4	0.3	1.4	0.3
prime	1.5	0.3	1.3	0.3
old	1.2	0.2	1.7	0.3

Table 7: Summary of Parametrization

Parameter	Meaning	Moment to Match	Value
β	discount factor	interest rate	0.99
δ	capital depreciation	capital-output ratio	0.06
σ	curvature of leisure	hours volatility	See Table 6
χ	scaling utility of leisure	hours levels	See Table 6
ζ_j	survival rates	vital statistics	see text
$\rho_{p_0}, \gamma_{p_0}, \sigma_{p_0}$	birth rates process	share of 20-year-olds	0.92, 0.002, 0.0014
ρ_h^m, γ_h^m	education share process	male education composition	0.91, 0.026
ρ_h^m, γ_h^m	education share process	female education composition	0.91, 0.032
ρ^h	housework persistence	evolution of female market hours	0.97
α	prod. share parameter	labour share	0.43
θ	prod. elasticity parameter	capital-labour complementarity	-0.25

Table 8: Productivity processes

	m,y,l	m,y,h	f,y,l	f,y,h	m,p,l	m,p,h	f,p,l	f,p,h	m,o,l	m,o,h	f,o,l	f,o,h	A
ρ	.87	.86	.88	.88	.89	.88	.90	.89	.89	.91	.90	.88	.91
γ	.17	.24	.13	.19	.19	.25	.14	.20	.18	.18	.13	.21	0
$100\sigma(u)^2$	3.19	3.17	3.17	3.01	3.17	3.05	3.13	2.86	3.61	3.20	3.53	3.45	.09

Table 9: Standard deviations relative to output and correlations

	Standard deviations		Correlations with output	
	Data	Model	Data	Model
Output	1	1	1	1
Consumption	0.82	0.63	0.91	0.32
Investment	4.57	4.11	0.93	0.87
Hours	1.12	0.36	0.71	0.89

Table 10: Standard deviation of HP-filtered output and hours over time

	67 – 76	67 – 84	85 – 07	00 – 13
	Output			
Baseline simulation	1.71	1.49	0.87	0.92
Counterfactual simulation	1.36	1.24	0.88	0.92
	Market Hours			
Baseline simulation	0.66	0.56	0.19	0.21
Counterfactual simulation	0.55	0.49	0.20	0.22

Table 11: Standard Deviation ratios

	Great Moderation	Counterf.-Base St.Dv. 67-76	Counterf.-Base St.Dv. 67-84	Counterf.-Base St.Dv. 85-00	Counterf.-Base St.Dv. 00-13
	Output				
All	31.0	-23.1	-18.3	1.9	-0.7
Sex	3.3	-4.0	-3.4	-0.7	-0.7
Age	4.4	-0.8	-0.4	2.7	-1.8
Educ.	25.4	-17.5	-13.6	3.3	1.8
	Market Hours				
All	12.3	-17.3	-15.0	6.6	5.5
Sex	2.1	-5.4	-5.3	-1.4	-2.2
Age	7.9	-6.6	-6.3	3.1	11.4
Educ.	4.3	-7.2	-7.4	0.7	1.0

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 labour. Counterfactual-Base St.Dv. measures the percentage difference between output volatility in the baseline and counterfactual simulation.

Table 12: Regression analysis: age

	Coef.	Newey West Std. Err	$P > t $
γ	0.079	0.020	0.000
α	-0.030	0.010	0.005

Table 13: Regression analysis: age, gender and education

	Coef.	Newey West Std. Err	$P > t $
γ_{age}	0.003	0.035	0.929
γ_{sex}	0.141	0.057	0.018
γ_{edu}	-0.015	0.026	0.551
α	-0.067	0.018	0.001

Appendixes

A Equilibrium

Definition 1 *A recursive competitive equilibrium is composed of a value function for the agents $V(g, j, e, a; \omega)$, consumption, assets and labour functions $c(g, j, e, a; \omega)$, $a'(g, j, e, a; \omega)$, and $\ell(g, j, e, a; \omega)$; a agents' distribution function $p'(g, j, e; \omega)$, aggregate capital K , output Y , labour inputs n_i for all i and their aggregation L . Pricing functions $r(\omega)$ and $w(g, j, e; \omega)$ and a transition process for the aggregate state vector $\Pi(\omega'|\omega)$ such that the following conditions are satisfied:*

1. *The decision rules for consumption, assets and labour c , a' , ℓ , and the value function V solve the agents problem in 3.1.*
2. *Capital and labour inputs satisfy equations (7)–(8).*
3. *Labour markets clear. i.e. equation (6) holds for all i .*
4. *The capital market clears: $K = \sum_g \sum_{j=1}^J \sum_e a(g, j, e) p'(g, j, e; \omega)$. Where $a(g, j, e)$ is the asset position by gender, age and education.*
5. *The goods market clears: $C + K' - K(1 - \delta) = Y$, where $C = \sum_g \sum_{j=1}^J \sum_e c(g, j, e, a; \omega) p'(g, j, e; \omega)$, $K' = \sum_g \sum_{j=1}^J \sum_e a'(g, j, e, a; \omega) p'(g, j, e; \omega)$, and Y is given by the production function (4) with aggregate capital K , and L defined by Equation (5).*
6. *$\Pi(\omega'|\omega)$ is consistent with Sections 3.2 and 3.4, and with agents expectations and decision rules.*

B Labour demand sensitivity to shocks.

I show next some analytical properties of the production function referred to in Sections 5.2.2 and 5.2.3.

For illustration purposes, I only consider men and women and abstract from the other groups. The marginal productivity of female labour is

$$f_{\ell_f} = (1 - \alpha) A^{\frac{1}{\theta}} (\alpha k^\theta + (1 - \alpha)(z_m \ell_m + z_f \ell_f)^\theta)^{\frac{1}{\theta} - 1} (z_m \ell_m + z_f \ell_f)^{\theta - 1} z_f. \quad (15)$$

It is then possible to construct the elasticity of labour demand to the shocks A , z_m and z_f , holding labour and capital constant. These measure the shifts in the labour

demand curve in response to the shocks.

$$\varepsilon_{z_f} \equiv \frac{\partial f_{\ell_f}}{\partial z_f} \frac{z_f}{f_{\ell_f}} = \frac{(1-\theta)z_f \ell_f}{z_m \ell_m + z_f \ell_f} \left(\frac{(1-\alpha)(z_m \ell_m + z_f \ell_f)^\theta}{(\alpha k^\theta + (1-\alpha)(z_m \ell_m + z_f \ell_f)^\theta)} - 1 \right) + 1, \quad (16)$$

$$\varepsilon_{z_m} \equiv \frac{\partial f_{\ell_f}}{\partial z_m} \frac{z_m}{f_{\ell_f}} = \frac{(1-\theta)z_m \ell_m}{z_m \ell_m + z_f \ell_f} \left(\frac{(1-\alpha)(z_m \ell_m + z_f \ell_f)^\theta}{(\alpha k^\theta + (1-\alpha)(z_m \ell_m + z_f \ell_f)^\theta)} - 1 \right) \quad (17)$$

$$\varepsilon_A \equiv \frac{\partial f_{\ell_f}}{\partial A} \frac{A}{f_{\ell_f}} = \frac{1}{\theta} \quad (18)$$

From the last equation it is immediate that the elasticity of labour demand to A is not affected by changes in female labour l_f . The next figure shows how the elasticities to z_f and z_m , computed using Equations (16)–(18), change with l_f given the calibration in the paper and steady state shocks, capital and male labour: first, the elasticity to z_f decreases while that to z_m increases; second, the elasticity to z_f is larger and steeper in absolute value than that of z_m .




Figure 10: Elasticity of female labour demand to z_f and z_m

These elasticities imply that the response to shocks decreases with l_f . To illustrate this, it is instructive to assume perfectly correlated shocks z_f, z_m, A .¹ With constant

¹They are indeed positively correlated, I consider the empirical covariance structure later.

elasticities, the labour demand shift to an impulse that increases the shocks z_f, z_m, A by a proportion σ is

$$\frac{\Delta f_{\ell_f}}{f_{\ell_f}} = \varepsilon_{z_f} \sigma + \varepsilon_{z_m} \sigma + \varepsilon_A \sigma.^2$$

So how does $\frac{\Delta f_{\ell_f}}{f_{\ell_f}}$ change with ℓ_f ?

$$\frac{\partial \frac{\Delta f_{\ell_f}}{f_{\ell_f}}}{\partial \ell_f} = \frac{\partial \varepsilon_{z_f}}{\partial \ell_f} \sigma + \frac{\partial \varepsilon_{z_m}}{\partial \ell_f} \sigma.$$

Since $\frac{\partial \varepsilon_{z_f}}{\partial \ell_f} < 0$, $\frac{\partial \varepsilon_{z_m}}{\partial \ell_f} > 0$ and $|\frac{\partial \varepsilon_{z_f}}{\partial \ell_f}| > |\frac{\partial \varepsilon_{z_m}}{\partial \ell_f}|$, it follows that $\frac{\partial \varepsilon_{z_f}}{\partial \ell_f} + \frac{\partial \varepsilon_{z_m}}{\partial \ell_f} < 0$ so that $\frac{\Delta f_{\ell_f}}{f_{\ell_f}}$ is decreasing in ℓ_f . Similarly, it is possible to see how demand shocks to male hours respond to the increase in ℓ_f . Responses are qualitatively the opposite ($\frac{\Delta f_{\ell_m}}{f_{\ell_f}}$ increasing in ℓ_f).

More generally, it is possible to compute how the standard deviations of $\frac{\Delta f_{\ell_f}}{f_{\ell_f}}$ and $\frac{\Delta f_{\ell_m}}{f_{\ell_f}}$ move with ℓ_f , taking into account the entire covariance matrix of the shocks $\hat{z} \equiv [z, A]$. For each of the 12 groups (e.g. female, young, with low education) the standard deviation can be approximated assuming constant elasticities as

$$\text{ST.DEV} \left(\frac{\Delta f_{\ell}}{f_{\ell}} \right) = \left(\sum_{i,j} \frac{\partial f_{\ell}}{\partial \hat{z}_i} \hat{z}_i \frac{\partial f_{\ell}}{\partial \hat{z}_j} \hat{z}_j \sigma(i,j) \right)^{1/2} \quad (19)$$

where i, j index the shocks collected in \hat{z} and $\sigma(i, j)$ is the covariance between \hat{z}_i and \hat{z}_j . Equation (19) comes from the formula for the variance of a sum and stems from the fact that $\frac{\Delta f_{\ell}}{f_{\ell}} \simeq \sum_i \frac{\partial f_{\ell}}{\partial \hat{z}_i} \frac{\Delta \hat{z}_i}{f_{\ell}}$.

To get a sense of the average volatility of labour demand shifts for men and women, I compute the following average by sex

$$\frac{\sum_i \text{ST.DEV} \left(\frac{\Delta f_{\ell_i}}{f_{\ell_i}} \right) \ell_i}{\sum_i \ell_i}$$

where i either includes all groups of females or males.

With ℓ_i 's and elasticities computed using Equations (16)—(18) from the steady state where hours shares by sex are equal to the average between 67 and 84, the standard

²This expression comes from the fact that – holding capital and labour constant – $\frac{\Delta f_{\ell_f}}{f_{\ell_f}} \simeq \frac{1}{f_{\ell_f}} \left(\frac{\partial f_{\ell_f}}{\partial z_f} \Delta z_f + \frac{\partial f_{\ell_f}}{\partial z_m} \Delta z_m + \frac{\partial f_{\ell_f}}{\partial A} \Delta A \right)$ where $\Delta z_m = \sigma z_m$, $\Delta z_f = \sigma z_f$, $\Delta A = \sigma A$. Here and below it is understood that Δf_{ℓ_f} measures the shift in the marginal productivity holding capital and labour constant i.e. a labour demand shift.

deviation of female labour demand shifts over that of men is 0.503. With ℓ_i 's and elasticities computed with hours shares at the final steady state, the relative volatility is 0.492. This corresponds to a drop of 2.2%. The corresponding drop in wage volatility between the two simulations is $0.80/0.83 - 1$ or 3.6%. So these elasticities account for about 60% the change in wage volatility due to the increase in female hours.³

The described mechanism, based on labour demand, comes from the adopted production function. This property holds for θ larger or smaller than zero, the Cobb-Douglas case. As explained in the paper, in the Cobb-Douglas case it is not possible to distinguish between TFP and labour shocks, but the result holds true if one approaches $\theta = 0$ from the left or from the right.

While this analysis improves the understanding of the model, it should be noted that this mechanism has small bite in practice: the relative volatility of labour demand by sex only moves from 0.50 to 0.49, so clearly the fact that the standard deviation of female wages is lower than for males has to do with the size of the shocks. In particular, the mechanism above is unlikely to have much propagation on the response of aggregate volatility to changing shares. Several unreported robustness checks confirm that the aggregate volatility effects of changing shares are quite robust to changes in relative volatilities of the magnitude implied by the labour demand elasticities.

³The remaining part may be due to a number of factors including the endogenous movement in capital, held constant in the calculation of these elasticities and the labour supply movements.