University of Southampton

Faculty of Law, Arts & Social Sciences

School of Social Sciences

Aggregate and Disaggregated Fluctuations

by

Alessandro Mennuni

Thesis for the degree of Doctor of Philosophy

June 2011

To Monica

Abstract

The thesis comprises three essays, broadly concerned with the sources and consequences of business cycle fluctuations. The first essay, in the tradition of the Real Business Cycle literature started by Kydland and Prescott (1982), focuses on the propagation mechanism of aggregate technology shocks. This paper extends the original framework, allowing for curvature in the transformation frontier between consumption and investments, which has been assumed linear by the preceding literature. In the second paper I take a less aggregate approach and investigate how income fluctuations to relatively small groups of the population, can lead to aggregate fluctuations. In this attempt to open the black box represented by aggregate technology shocks, I ask whether changes in aggregate fluctuations and in the labour composition, can be related and studied together within a unified structural framework. The third essay studies optimal fiscal policy with a particular timing of events that makes capital elastic in the short run. This feature is found to have important policy implications.

Contents

List of Figures

List of Tables

Declaration Of Authorship

I, Alessandro Mennuni, declare that the thesis entitled Aggregate and Disaggregated Business Cycles and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;

Signed: ....

Date

Acknowledgements

This work was undertaken with financial support of the Economic and Social Research Council (PTA-031-2006-00383). This thesis would not have been possible without the support of many people. I would like to express my sincere gratitude to:

- My supervisors, Xavier Mateos-Planas and Martin Gervais. I am grateful to Xavier for his advice and for his support in important moments of my carrier. I thank Martin for his constant advice, assistance and support throughout my studies. His contribution to this thesis cannot be overstated.

- My colleges: many of them, over the years and in different places, gave me practical and moral support. In some cases contributed substantially to the development of my work. A special mention goes to Andrea Waddle and Ronald Wolthoff when I was visiting the University of Pennsylvania. Sarolta Lazco, while we both were at the European University Institute, and Corrado Giulietti, here in Southampton.

Chapter 1

Introduction

My doctoral dissertation, which comprises three separate essays, is broadly concerned with the sources and consequences of business cycle fluctuations.

In the first paper in order of time, in the tradition of the Real Business Cycle literature started by Kydland and Prescott (1982), I focus on the propagation mechanism of aggregate technology shocks. The literature typically distinguishes between neutral shocks (that affect all sectors of the economy) and investment specific technology shocks, which only affect the investments sectors. This paper extends the original framework, allowing for curvature in the transformation frontier, which has been assumed linear by the preceding literature. It studies how this affects the relative price of investment goods, which is key for the identification of investment shocks and for the co-movement of consumption and investment goods.

In the second paper I take a less aggregate approach and investigate how income fluctuations to relatively small groups of the population, can lead to aggregate fluctuations. In this attempt to open the black box represented by aggregate technology shocks, I ask whether changes in aggregate fluctuations and in the labor composition, can be related and studied together within a unified structural framework. How successful can this framework be in explaining at the same time cross-sectional changes and aggregate volatility changes, seems a major test for this economy. And it may enable to derive targeted policy implications, more effective and less expensive than aggregate fiscal and monetary policies.

The third paper has a normative connotation and studies optimal fiscal policy. With Martin Gervais, we revisit the age-old question of whether and how governments should respond to changes in economic activity. Perhaps the best way to understand the work is to think about the reasons why previous attempts to address this question within the neoclassical framework produced counterfactual results when capital taxation is a policy instrument. Their main prescription involves running budget *surpluses* during recessions, and budget *deficits* during expansions.[1] The paper shows how the particular timing assumed in the neoclassical framework implies that capital is completely inelastic in the short run, thereby leading to the mentioned policy prescription, which is clearly at odds with what we observe. The paper considers an alternative timing of events that makes capital elastic in the short run and studies the implication for optimal fiscal policy.

The rest of the thesis is organized as follows: chapters 1, 2 and 3 contain the three mentioned papers. The last chapter derives some broad conclusions and possible extensions.

[1] See for instance Chari et al. (1994).

Chapter 2

The role of curvature in the transformation frontier for measuring technology shocks

Abstract

In the usual version of the neoclassical growth model used to identify neutral (N-Shock) and investment shocks (I-Shock), a linear transformation frontier between consumption and investment goods is assumed. This paper extends the original framework, allowing for curvature in the transformation frontier, and studies how this affects the relative price of investment goods and hence the identification of investment shocks. A concave frontier allows a substantial improvement in the prediction of the saving rate. Furthermore, a concave frontier induces short-run aggregate effects of relative demand shifts, thereby fostering the propagation of the shocks under consideration, which overall account for 86% of the aggregate fluctuations. When I identify shocks with curvature, the N-shock appears to be stationary while the I-shock is a unit root. This leads the N-shock to play a major role: 91% of the fluctuations explained are due to the N- shock.

2.1 Introduction

Recent work in the field of measuring the importance of technology shocks for the business cycle is based on the neoclassical growth framework. Greenwood et al. (1988) showed that shocks to the productivity of investment goods (I-shock) are an important source of fluctuations, together with *neutral* productivity shocks that hit all sectors of the economy (N-shock). This recognition engendered several studies of this mechanism, including Greenwood et al. (2000), Cummins and Violante (2002) and Fisher (2006).

To identify the I-shock, these papers use the fact that, in the framework considered, the relative price of investment goods only moves with I-shocks. More precisely, I-shocks (V) are identified from the relative price equation between consumption and investment goods, which, under the assumptions of this model - linear transformation frontier between consumption and investments - is simply $p = 1/V$. This is a key identification assumption because, without identifying investment shocks from the price equation, Justiniano et al. (2008) find that the I-shock should be 4 times more volatile in order to match business cycle fluctuations. The sharp contrast that comes from this price equation calls for further investigation of its specification.

This paper investigates whether this oversimplified specification may bias the measuring and the propagation of the shocks. Two signs of a potential misspecification further motivate this investigation: i) the original model fails to replicate the saving rate and ii) the two shocks identified through that framework are negatively correlated. It will be argued that these two observations suggest that the transformation frontier should be concave.

The first sign of misspecification (fact i) is that, with the preferences commonly used, the model's prediction of the saving rate is very poor:[1] the fit of the predicted saving rate

[1] To increase co-movement Greenwood et al. (1988) have to rely on very low short-run wealth effects in the labor supply, as recently emphasized by Jaimovich and Rebelo (2009)

on the actual time series, given the shocks identified, gives a very low R^2. This suggests a possible way to pin down the curvature in the transformation frontier: one that maximizes the fit of the saving rate. To motivate this choice, it is important to notice that curvature in the transformation frontier makes the relative price sensitive to any changes in the relative demand of the two types of goods, which are summarized by changes in the saving rate. Then, a better fit between the saving rate and the relative price should imply a better price equation. Under a linear transformation frontier, the effect of relative demand changes in the relative price is neglected and this may have important quantitative implications.

Of course, how well the model fits the saving rate would be an ideal way to test this model. However, the fact that the moment is used to pin down the parameters does not allow it to be used to test the model. Fortunately, fact ii leads to an alternative way to pin down the degree of curvature. Furthermore, fact ii - that the two shocks identified assuming a linear frontier are negatively correlated - is a direct sign of concavity in the transformation frontier for the following reason: a concave frontier implies a positive relation between the price and the N-shock. This is because an increase (decrease) in the N-shock, will increase (decrease) the saving rate (i.e. the relative demand between investments and consumption) because of households' desire to smooth consumption; firms, due to the concave transformation frontier, will be induced to meet demand through an increase (decrease) in the relative price.

Neglecting this channel, one would have to wrongly attribute the increase in the relative price that comes after a positive N-shock to a decrease in the I-shock: every time p increases as a consequence of a positive N-shock the researcher armed with the simplified price equation would impute the increase in p to a decrease in V. This would make the two shocks appear negatively correlated. Indeed, identifying the shocks in the usual way leads to negatively correlated shocks: after removing unit roots from the shocks, I find a significant correlation between the two shocks of -19%. According to the reasoning above, this negative correlation is the sign of a concave transformation frontier. It follows that one possible way to estimate the curvature is to pick the curvature that makes the two shocks appear independent, in order to avoid capturing as an I-shock the increase in the price due to the N-shock.

Strikingly, the curvature under the two strategies is very close and leads to the same implications. In particular, the fit of the saving rate under the second moment condition is very close to the one obtained under the first strategy, where the parameters were picked to maximize the fit of the predicted saving rate to that of the data.

The model is enhanced with curvature in the transformation frontier by adding only one parameter to the original framework. This is convenient in that it allows the one-sector characterization of the original framework to be preserved, and aggregate data to be used to fully calibrate the model. Importantly, this specification does not affect the balanced growth path predictions of the original framework. This is a virtue of this specification because the original framework is capable of reconciling the downward trend observed in the relative price and the increase in the relative production of investments goods, and make these two facts consistent with a Balanced Growth Path, as shown by Greenwood et al. (1997).[2]

To focus on the role of curvature, the model is kept as simple as possible. The real frictions usually included by the recent literature, although important to improve the fit to the data, are not considered in this paper. Indeed, adding capital adjustment costs and capital utilization would not change the message of the paper as long as they do not affect the relative price equation.[3] This is indeed the case for the usual way in which

[2]Another advantage of the proposed specification is that, although making the price equation function of both the shocks, it allows the shocks to be backed out analytically, given the parameters. This permits them to be backed out without the use of a filter or by using a simulated method, thereby increasing precision and saving on computing time.

[3]These frictions introduce *inter-temporal* adjustment costs. Instead, concavity in the transformation frontier is a concept that is closer to the *intra-temporal* adjustment costs considered by Huffman and Wynne 1999. See Guerrieri et al. (2009) for an interpretation of I-shocks in a fully- fledged multi-sector model.

these frictions are modeled, for example Schmitt-Grohe and Uribe (2008), Justiniano et al. (2008) and Justiniano et al. (2009) consider medium-scale models with several frictions that do not affect the investment price equation. In fact, abandoning a linear frontier is a necessary condition for affecting the price equation. Hence, while allowing for intertemporal adjustment costs permits improving the saving rate prediction of the model, it does not affect the identification of the I-shock, which this paper argues has been mis-measured.

One finding of this paper is that the N-shock identified appears to be stationary, while the investment shock is a unit root. In most of the previous studies the two shocks were either considered both stationary, as in Greenwood et al. (2000), or both unit roots as in Fisher (2006). This difference has important implications for the relative importance of the two shocks. When the N-shock is stationary and the I-shock a unit root, the first plays a major role in explaining aggregate fluctuations. This overthrows previous findings, where the I-shock played the main role.[4] There is a simple intuition for this result: when there is a permanent shock, productivity grows but so does expected wealth. Therefore, the expected marginal utility of consumption decreases, lowering the boost in the saving rate and in the labour supply. This implies that the households' reactions to a permanent shock are weaker than the reactions to a transitory shock. This explains why transitory shocks play a stronger role in accounting for the business cycle.

Whether the Business Cycle is about stationary fluctuations around a deterministic trend, or is due to a stochastic trend has been debated since the paper by Nelson and Plosser (1982). The present finding may reconcile the two views in that both things happen; this paper suggests that there is a stochastic trend due to the I-shock and stationary fluctuations around it through the N-shock.

[4]This result is in line with the findings of Schmitt-Grohe and Uribe (2008), where investment shocks play little role in aggregate fluctuations.

With these slight changes to the original framework, this simple model accounts for 86% of the Business Cycle - substantially more than what is predicted by the usual framework. For a comparison, the point estimation is above the 95% confidence interval extremum considered by Fisher (2006).

The paper is organized as follows, the next section identifies and discusses the misspecification, Section 3 modifies the framework in order to allow for curvature in the transformation frontier, Section 4 reports the calibration, Section 5 reports the findings and section 6 concludes. Some technical details are relegated to the Appendix at the end of this chapter.

2.2 Identifying the Misspecification

Below follows a description of the standard growth model with investment-specific technological change like, for instance, the one adopted in Fisher (2006).
The representative household solves the following problem, taking prices as given:

$$\max_{\{ct,k_{t+1},n_t\}} E_0\left[\sum_{t=0}^{\infty}\beta^t\left(\log(c_t)-\chi\frac{n^{t\,1+1/\nu}}{1+1/\nu}\right)\right]$$
$$s.t. \quad c_t+p_tk_{t+1}=w_tn_t+p_tk_t(1+r_t-\delta).$$

These preferences are adopted for instance by Fuentes-Albero et al. (2009), as they point out ν is the Frisch elasticity of labour. Capital evolves according to the law of motion

$$k_{t+1}-k_t(1-\delta)=V_tA_tk_t^{\theta}n_t^{1-\theta}(1-s_t),$$

while non-durable consumption is

$$c_t=s_tA_tk_t^{\theta}n_t^{1-\theta} \tag{2.1}$$

where s_t is the fraction of physical production allocated to consumption. V_t is the investment shock, which only hits the production devoted to increasing the capital stock. It follows the following process:

$$V_t = V_0 \gamma_v^t \exp(u_{vt}) \tag{2.2}$$

$$u_{vt} = \rho_v u_{vt-1} + \varepsilon_{vt}. \tag{2.3}$$

A_t is a neutral shock that hits both sectors in the same way and evolves according to the following process:

$$A_t = A_0 \gamma_a^t \exp(u_{at}) \tag{2.4}$$

$$u_{at} = \rho_a u_{at-1} + \varepsilon_{at}. \tag{2.5}$$

Firms are competitive: given prices they solve the following static problem:

$$\max_{k_t, n_t, s_t} y_t - w_t n_t - p_t k_t r_t$$

$$s.t.$$

$$y_t = s_t A_t k_t^{\theta} n_t^{1-\theta} + p_t V_t (1 - s_t) A_t k_t^{\theta} n_t^{1-\theta}. \tag{2.6}$$

The first order conditions for the firm are as below:

$$\theta A_t k_t^{\theta-1} n_t^{1-\theta}(s_t + p_t V_t (1 - s_t)) = p_t r_t$$

$$(1-\theta) A_t k_t^{\theta} n_t^{-\theta}(s_t + p_t V_t (1 - s_t)) = w_t$$

and

$$p_t = 1/V_t. \tag{2.7}$$

The price equation (2.7) reflects the fact that a firm can choose where to allocate its inputs

with no costs. Hence, it will be indifferent between producing consumption or investment goods if and only if (2.7) holds. This strong implication of the model is what is disputed in the present paper. This assumption is innocuous for the growth analysis of the model as in Greenwood et al. (1997) [5] for which the model was originally built, but it matters for the

[5] The modification introduced in section 2.3 leaves the balanced growth path unchanged and therefore it maintains the same growth implications of the original framework as shown in Appendix 2, section 2.7.2.

business cycle analysis.

From (2.1), (2.7) and (2.6) $s_t = \frac{c_t}{y_t}$ holds. Therefore aggregate production is

$$y_t = A_t k_t^{\theta} n_t^{1-\theta}.$$

From this and (2.7), time series for A and V are identified as follows

$$A_t = \frac{y_t}{k_t^a n_t^{1-a}} \quad (2.8)$$

$$V_t = 1/p_t. \quad (2.9)$$

2.2.1 Correlation Between Shocks

To identify the neutral shock A, α is assumed equal to 1/3 and the results of this section are robust to changes in this parameter. Data on the relative price of investment goods are those constructed by Fisher (2006), who extended the analysis by Gordon (1990), and successively by Cummins and Violante (2002). These data are available from 1947 IV to 2006 IV. [6], [7] In the first years of the sample, however, the shocks seem to exhibit properties opposite to what the data shows later on. In the early years, the two variables appear to be positively correlated, although the economy does not yet seem to be on a balanced growth path: as Figure 2.1 shows, the capital output ratio exhibits an increasing trend before stabilizing when the period considered starts. A positive correlation of the shocks never occurs later in the sample. That initial period is therefore omitted from the analysis, because it presumably captures mis-measurement due to the transition to the BGP, where the model, calibrated trough the BGP properties of the data, fails to properly measure the shocks. The transition and structural breaks after the Korean war mean that much of the macro analysis starts after the Korea war.

[6] I thank Maxym Kryshko for giving me the relative price time series. The rest of the time series comes from the Bureau of Economic Analysis (c, I) and from the Bureau of Labor Statistics (n, pop).

[7] Given the emphasis on technological shocks, the fact that the data do not include the last years may not be considered a disadvantage in that the last recession may be due to sources not considered in this framework, thereby increasing the misspecification of the model and hence biasing the findings.

The analysis is therefore restricted to the sample that goes from 1957 IV to 2006 IV [8]. ADF and Phillips-Perron tests accept the hypothesis of a unit root for $\ln(A)$ and for $\ln(V)$. I therefore estimate the regression

$$d\ln(A_t) = \underset{.06}{.0025} - \underset{.045}{.32} d\ln(V_t) + \varepsilon_t \tag{2.10}$$

$$corr[d\ln(A), d\ln(V)] = -.19. \tag{2.11}$$

The two time series in the sample that goes from 1956 to 2006 seem to be strongly negatively correlated.

Considering sub-samples of this sample gives the same result. Extending the sample a little more, considering years before 1956 does not change the results. For example, starting from 1952 one still gets a significant negative correlation of the relative change of the shocks of -15%. This however does not hold in the first years of the sample, where the two time series appear to be highly positively correlated. This is in such a sharp contrast with what happens in the rest of the sample (and in sub-samples of it), that it casts doubts on the reliability of the model to measure the shocks over that early period of time[9].

I conclude that the two time series for the shocks identified through the usual framework appear to be negatively correlated.

2.2.2 Saving Rate

The other dimension where the misspecification is notable is that the model predicts counter-factual savings rates. With a conventional calibration (that of section 2.4, but with the parameter that governs curvature in the transformation frontier $\rho = 0$, which implies a linear frontier), and the shocks identified from the data, let $\widehat{s}$ be the time series of s predicted by the model, and s the time series realized. At least with the utility function considered,

[8] 1956 is when the capital output ratio reached a level from which continued over time on a trend-less path.

[9] Fisher (2006) argues that the quality bias in the NIPA data is stronger in the earlier part of the sample.

the two time series are so different from one another, that the $R^2 = 1 - var(\widehat{s} - s)/var(s)$ is even negative. This implies that the saving rate predicted by the model is countercyclical with respect to the actual one, and therefore even the simple mean performs better, leaving a smaller residual variance.

The two facts above - he tnegative correlation between the shocks and the counterfactual saving rate - suggest that there is a misspecification in the model.
Section 2.3 describes a modification of the framework where a different specification for the two sectors is modelled. Before that, an attempt is made to interpret the negative correlation in (2.11) and the bad fit in the saving rate as being suggestive of curvature in the transformation frontier.

2.2.3 The Case for Curvature in the Transformation Frontier

Since the model can be expressed in recursive form with state variables A, V, K, assume that the true price equation is of the form

$$p = f(A, V, K) \tag{2.12}$$

and let the total production measured in consumption units be

$$y = y(A, V, K). \tag{2.13}$$

Considering instead the price equation (2.7) and the aggregate resource constraint (2.6) would wrongly impute all the increase(decrease) in the relative price to a decrease (increase) in V, and all the variation in production not explained by k to A. If instead $\frac{\partial p}{\partial A} > 0$, increases in p may be due to increases in A, and when this happens, also y increases through A. With the misspecified policy functions, the increase in the price would be attributed to a decrease in V, while instead only an increase in A occurred. This leads to the negative correlation between A and V, which is not a pure negative correlation between the two shocks, but is due to the misspecification of the model.

The misspecification also leads to counter-factual saving rates: when there is an increase in A, according to the true policy function (2.12) p grows. When this happens, the original model identifies a decrease in V. Because the productivity of investments decreased, the saving rate predicted by the model decreases. If, on the contrary, no I-Shock occurred, the increase in A would imply an increase in the saving rate. Therefore, the counter-factual saving rate predicted by the original model is a consequence of the misspecified price equation.

It follows from these considerations that a model used to measure these shocks for the business cycle should be specified in a way such that the time series for the shocks that it predicts, appear to be independent and match as closely as possible the saving rate time series. These are the two facts that will be targeted in the calibration of the model presented in the next section.

2.3 The Modified Framework

Consider the following modification to the model: a generic firm produces

$$y_t = A_t k_t^a n_t^{1-a} s_t^{1-\rho} + p_t V_t A_t k_t^a n_t^{1-a} (1-s_t)^{1-\rho}, \tag{2.14}$$

where $\rho \, \epsilon \, [0,1)$. s_t measures the share of inputs allocated to the production of consumption goods.

Therefore,

$$\begin{aligned} c_t &= A_t k_t^a n_t^{1-a} s_t^{1-\rho} \\ k_{t+1} - k_t(1-\delta) &= V_t A_t k_t^a n_t^{1-a} (1-s_t)^{1-\rho}. \end{aligned}$$

The firm can produce for both sectors, but the marginal productivity of producing for one sector is decreasing. This makes the firm willing to produce for both the sectors, even if $p_t V_t \neq 1$.

The firm's problem is

$$\max_{k,n,s} A_t k_t^a n_t^{1-a} s_t^{1-\rho} + p_t V_t A_t k_t^a n_t^{1-a} (1 - s_t)^{1-\rho} - wn - rpk. \tag{2.15}$$

When $\rho > 0$ the marginal productivity of consumption and of investment goods is decreasing[10], and therefore the firm will always choose to produce both types of goods. This technology has constant returns to scale, and therefore the Euler theorem holds, so the problem is consistent with perfect competition where the size and number of firms does not matter and firms take prices as given.

The equilibrium conditions are reported in appendix 2.7.2 and are essentially unchanged with respect to the usual framework, except for the resource constraints above and for the price equation, which comes from the optimal choice of s_t:

$$p_t V_t = \frac{(s_t)^{-\rho}}{(1 - s_t)^{-\rho}}. \tag{2.16}$$

It can be noticed how when $\rho = 0$ the price equation (and the whole model) boils down to the usual framework. This price equation shows that the change in the relative price is not only due to a change in V_t, but it also depends on the change in s_t, i.e. on the change in the relative demand for the two goods. This in turns depends on both the shocks and on capital. The relative importance of one shock with respect to the other depends on the parameter ρ. How the two shocks affect p is what in the preceding section was indicated to be crucial for the two shocks to be uncorrelated. Since this depends on ρ, this parameter will be pinned down to make the two shocks uncorrelated, and to improve the fit of the saving rate.

[10] In this sense, ρ can be interpreted as an intra-temporal adjustment cost as in Huffman and Wynne (1999)

2.4 Calibration

The parameters of the model are $\beta, \alpha, \delta, A_0, \gamma_a, V_0, \gamma_v, \sigma\varepsilon_a, \sigma\varepsilon_v, \rho, \chi, \nu, \rho_a, \rho_v$.

The crucial parameter choice is ρ. As mentioned in the introduction, two strategies are employed. The first is to pick ρ such that the shocks identified are uncorrelated.

In this model the shocks can be identified from the price equation (2.16) and (2.14) as

$$V = \left(\frac{s}{1-s}\right)^{-\rho} \frac{1}{p} \tag{2.17}$$

$$A = \frac{Y}{k_t^{\alpha} n_t^{1-\alpha} \left[s^{1-\rho} + pV(1-s)^{1-\rho}\right]}. \tag{2.18}$$

As becomes clear from observing the two equations above, to identify the shocks, it is first necessary to identify s. From (2.16) and from the fact that

$$\frac{s^{1-\rho}}{pV(1-s)^{1-\rho}} = \frac{c}{y-c}, \tag{2.19}$$

which comes from the two resource constraints, one gets the convenient fact that

$$\frac{s}{1-s} = \frac{c}{y-c}. \tag{2.20}$$

This highlights the close relation of s with the consumption rate and it is used to identify s from the data.

The strategy to calibrate ρ is the following: a guess for ρ very close to zero (as in the usual framework) is made. V and A are identified through the above equations. The relevant correlation between the two shocks is estimated as follows: [11], if it is other than zero, ρ is increased, otherwise ρ has been found. In order to implement this procedure, a value for α has to be fixed in order to back out A. In this model $1-\alpha$ is equal to the labour share and hence α is calibrated equal to 1/3.

[11]'relevant' meaning that if the shocks are unit roots, the correlation of the first differences is run; if they are stationary, the correlation of the levels is run; if one is stationary and the otherone is non-stationary, only the non-stationary one is differentiated.

This procedure leads to $\rho = .028$. This means that the marginal productivity of the two goods is decreasing and this makes the price equation 2.17 deviate from the one used in the main framework. As already noted, if $\rho = 0$, the price equation and the whole model boils down to the original one.

Given these shocks, $A_0, \gamma_a, V_0, \gamma_v, \sigma\varepsilon_a, \sigma\varepsilon_v, \rho_a, \rho_v$ are picked by running an OLS regression on the log of the shocks as identified above. The parameter values are $A_0 = 5.46$, $V_0 = 1$, $\gamma_v = 1.0074$, $\sigma\varepsilon_A = .00918$, $\sigma\varepsilon_V = .00498$.
.00

γ_a is not significantly different from 1, which means zero growth in the neutral shock. Therefore it is set equal to 1. This implies a growth rate of the I-shock of 0.74 %.

δ is picked to match the average investment capital ratio, as inferred from the law of motion for capital: assuming that the economy fluctuates around a balanced growth path

$$\delta = \frac{pI}{pK} - (\gamma_a \gamma_v)^{\frac{1}{1-\theta}} - 1 \tag{2.21}$$

this gives a value of $\delta = .0177$.

From the Euler Equation for consumption, on a BGP one gets [12]

$$\beta = E\left(\frac{\frac{c_{t+1}}{c_t}\frac{p_t}{p_{t+1}}}{1 + \theta y/(pK) - \delta}\right) = 0.98. \tag{2.22}$$

It remains to calibrate the parameters of the supply of labour: the critical one is ν, which represents the Frisch elasticity. As Prescott mentioned during his Nobel Prize Lecture, how much of the business cycle can be explained by technology shocks depends crucially on this parameter. The problem is that there is not a clear way to calibrate it; micro studies suggest $\nu = .2$ but they may understate adjustments to the extensive margin. The quasi-linear preferences of Hansen, where all the adjustment is on the extensive margin, imply $\nu = \infty$. Fuentes-Albero et al. (2009) estimate with Bayesian methods $\nu = .3$ with 95%

[12] (2.22) implies $\beta = \frac{(\gamma_a \gamma v)^{\frac{1}{1-\theta}} \gamma_V}{1+\theta E(y/(pK)) - \delta}$ if γ_a is not restricted to be 1 and therefore $E(\frac{c_{t+1}}{c_t}) = (\gamma_a \gamma v)^{\frac{1}{1-\theta}}$.

confidence interval [0.05 0.53][13]. The same choice is made here, and given the importance of this parameter and the weak arguments to motivate a particular choice, some sensitivity analysis will be carried out in the next section. Finally, χ is chosen to match the observed average labour supply $n = .3$.

The second strategy to pin down ρ is to maximize R^2 of the saving rate predicted by the model given the shocks identified. This is done by setting a greed on ρ, and for each value of ρ, doing the following: 1. given the other parameters, back out the two shocks time series through 2.17 and 2.18; 2. Estimate the parameters of the shocks' processes. 3. Solve the model. 4. Simulate given the shocks identified, and compute the R^2.

The value of ρ that gives the highest R^2 is 0.024. Strikingly, this value is very close to that obtained with the other procedure. R^2 of 0.37 is a substantial increase in the portion of variance of the saving rate explained by this model compared to the original framework[14]. Figure 2.2 reports the actual and predicted saving rate time series. With this value of ρ, the parameters of the shock process are essentially unchanged. The following tables summarize the parameter values.

Table 2.1: Curvature Parameter ρ

1^{st} strategy	0.28
2^{nd} strategy	0.24

Table 2.2: Other Parameter Values

β	α	δ	γ_A	γ_V	$\sigma\varepsilon_A$	$\sigma\varepsilon_V$	ρ_a	ρ_v	ν	χ
.98	.33	.018	1	1.0074	.00918	.00498	.95	1	.3	180

[13]The modification made to the original framework does not have remarkable effects on the labour reaction to the shocks. This makes their estimation valid even for the present framework.
[14]In the original framework the variance explained is essentially zero.

2.5 Results

This section describes the properties of the shocks identified. The main finding is that the N-Shock is trend stationary, while the I-Shock is trend stochastic. The quality of the shocks identified is tested by seeing how well the policy function for the relative price can predict the actual time series of the relative price when the shocks identified are used as inputs in the policy function. The model, if misspecified, could provide a wrong prediction, given that the shocks have not been identified through that policy function. Successively, a qualitative analysis of the propagation mechanism is carried out. In the last subsection the model is confronted with the main business cycle facts, and compared with the performance of the baseline framework. Under the preferred calibration, the model can predict 86% of the aggregate fluctuations. 91% of the fluctuations explained are due to the N-shock.

2.5.1 The Shocks Identified

A relevant result is that while the investment shock is a unit root, the neutral one is trend stationary. Under all the calibrations, ADF and Philip Perron tests, with various lags, reject the hypothesis of a unit root for the neutral shock, with p-values that range between 1% and 9%. Under the preferred calibration ($\nu = .3$ and $\rho = .024$) the autoregressive parameter is $\rho_a = .95$. As with the baseline framework, the process for the N-Shock does not show a significant trend: all the growth is captured by the I-Shock.

This overthrows the previous findings, that the I-shock plays the main role in accounting for the Business Cycle. Here the N-shock accounts for 91% of the total fluctuation explained. As explained in the introduction, agents react more strongly to a transitory shock than to a permanent one. Intuitively, agents realize that they have plenty of time to benefit from a permanent shock, while they have to extract the potential benefits that arises from a positive transitory shock more quickly. This explains why transitory shocks play a stronger role in accounting for the business cycle, and having the N-shock being transitory and the I-shock permanent makes the N-Shock the most important one, overthrowing the results found in the previous literature. It turns out that considering both the shocks, the model accounts for 86% of the Business Cycle, much more than generally predicted in previous studies: in Greenwood, Hercowitz, and Krusell (2000) the variance explained is around 70%; in Fisher (2006) it is between 40% and 60%. To the overall higher output variance found here, it contributes the new propagation mechanism implied by the presence of curvature in the transformation frontier, which makes shifts in the relative demand for goods cause a change in aggregate production. The mechanism is explained in the next subsection.

The stationarity of the N-shock is quite a relevant result for that branch of the quantitative literature that uses long-run restriction based on the unit root assumption to identify the shocks in a VAR model, such as Fisher (2006). According to the present model, only the investment shock can be identified through its long run properties, since the neutral shock is trend stationary. Whether the business cycle is about stationary fluctuations around a deterministic trend or stochastic changes in the trend has been debated since the influential paper by Nelson and Plosser (1982). The fact that one shock is trend stochastic and the other one is trend stationary may reconcile the two views: there is a stochastic trend, and also stationary fluctuations around that trend. Appendix 2.7.3 derives the equivalent stationary conditions when there is a trend stochastic shock and a stationary one. The transformed stationary model proves the existence of a Balanced Growth Path and allows for a recursive formulation. In 2.7.2, the model is also detrended under the assumption that both the shocks are trend stationary. From this it becomes clear that the model has the same long-run implications as the original framework: the expected growth rates of all the variables are unchanged.

A problem with the method, that uses fully specified theoretical models to identify the shocks, is that there is no proof that the shocks are close to the "true" ones. It is argued in this paper that the fit of the saving rate to the actual saving rate time series is a good measure of the quality of the specification, and allowing for curvature in the transformation frontier induces a clear improvement over this dimension. Another specification check can be made by exploiting the fact that the identification of V has not been done through the policy function for the relative price. Therefore, this policy function can reproduce prices that differ from the real ones when filled with the shock time series identified. Comparing the prediction with the actual price series is a robustness check that has been carried out as follows. The shocks identified are plugged into the computed policy function $p = p(k, A, V)$ and the predicted $\hat{p}$ is compared with the actual time series. If the model predicts a policy function for the relative price that is incorrect, then this fitting exercise will suggest that the model is misspecified. A good fit may be reassuring that the model is well-suited for the question at hand: to quantify the two technology shocks.

In the usual framework, this checking exercise cannot be performed because the investment shock is identified using the policy function for p. However, by tending ρ to zero this model boils down to the original framework and the policy function would become $p = 1/V$ and therefore the fit would be total. Given that ρ is quite close to zero, a good fit should not surprise. A good result therefore cannot be used to claim success. However a bad fit would be a clear sign of a wrong policy function. It is interesting to compute the R^2 on the deflated variables and for the ones in levels. The model performs surprisingly well: the predicted prices in levels are so close to the actual ones that $\mathrm{R}^2 = .9997$ Some of the good fit depends on the trend; removing it, the fit remains substantially high: $\mathrm{R}^2 = .79$. Figures 2.3 and 2.4 show the predicted and actual time series in levels and in growth rates. From figure 2.3 it is evident how most of the variance of the price time series is due to the trend. That is why the R^2 in levels is so high and therefore less informative than the one in growth rates.

2.5.2 Qualitative Results

The main differences of this model compared to the usual framework are due to the fact that the price also depends on the neutral shock. This is at the heart of the identification of different time series for the shocks and it has three important implications.

1. Perhaps surprisingly, the presence of curvature in the transformation frontier increases the effects of the shocks.

   This is explained as follows: after a neutral or I- shock, the consumption share s decreases because agents find it optimal to increase investments. The decrease in s implies that the marginal productivity of consumption goods $(1-\rho)A_t k_t^{\theta} n_t^{1-\theta} s_t^{-\rho}$ increases. The marginal productivity of investments, measured in consumption goods $p_t V_t (1-\rho) A_t k_t^{\theta} n_t^{1-\theta} (1-s_t)^{-\rho}$ also has to increase, since the two marginal productivities must be equal in equilibrium. This calls for an increase in $p * V$. Compared with the original framework, the price reacts less to a change in the investment shock, making the product $p * V$ procycliclal. Unlike what happens in the original framework, the fact that $p*V$ increases even after an investment shock, makes aggregate productivity increase. Therefore, the proposed mechanism increases the propagation of the shocks, increasing the proportion of the business cycle explained by productivity shocks.

2. After a positive neutral shock A, households want to increase the investment rate; because of the curvature in the transformation frontier. Firms, however, are reluctant and the price has to increase to induce them to adjust supply and meet demand.

   This highlights the fact that the change in the relative price of goods is not all due to the I-shock and how it could be misleading to identify the investment shock in the usual way.

   The fact that p increases after an N-shock implies that consumption is more volatile to a change in A with respect to the framework without curvature in the transformation frontier; the increase in the relative price induces agents to increase consumption with respect to what they would do in the usual framework, where the investment price does not depend on the N-shock. This is a feature typical of two-goods models with imperfect input reallocation, which turned out to imply a high equity premium as has been shown by Boldrin et al. (2001).

3. Unlike in the usual framework, here consumption increases after an I-shock. This is because, due to the smaller (compared to the usual framework) decrease in the price that follows the investment shock, pV increases, increasing GDP measured in consumption goods. Given this increase in GDP, households, although increasing the saving rate, can also increase consumption a little. This helps the matching of the positive correlation of consumption and total production, which is hard to match in the usual framework.

2.5.3 Accounting for the Business Cycle and Labour Supply Elasiticity

This section studies quantitatively the business cycle implications of this model. Since results may depend on the labour supply elasticity, the exercise is carried out for $\rho = .28$ but for two different values of ν. [15]

The model with $\nu = 0.3$

With this parametrization, the model accounts for 86% of aggregate fluctuations.
Labour fluctuations in the model account for 17% of the actual labour fluctuations. Investment fluctuations are substantially what they are in the actual data: the variance is 96% of the data. These two facts imply that part of the aggregate fluctuations are due to the change in the relative price and the saving rate as explained in paragraph 2.5.2: the original propagation mechanism implies oscillations not accounted for by changes in labour,

[15] Putting $\rho = .24$ essentially does not affect the results.

capital and TFP. What remains unexplained of the Business Cycle should be due to the low volatility of hours. Consumption fluctuates 1.9 times more than in the data . To this, it may contribute the low relative risk aversion parameter used. This choice was necessary to have a balanced growth path with the adopted utility function.

The model matches the usual correlations observed to confront the model with the data reasonably well: the correlation of consumption and GDP is .99 in the data and .98 in the model. The correlation of Investments and GDP is .93 in the data and .95 in the model. The model under-predicts the correlation of labour: .45 in the data and .25 in the model. The following two tables summarize these results:

Table 2.3: Standard Deviation Explained by the Model

Output	Consumption	Investments	Hours
86%	190%	96%	17%

Notes: Numbers are espressed in percentage terms of the actual data.

Table 2.4: Correlations with Output

	Output	Consumption	Investments	Hours
Data	1	.99	.93	.45
Model	1	.98	.95	.25

Although usually not considered, the correlations of the growth rates are reported; they highlight dimensions in which this model (and the original one) performs poorly. This may be useful for future research in that it addresses weaknesses of the framework.

The fact that consumption fluctuates too much and investments fluctuate essentially what they should, implies a lower correlation of consumption and investment goods. This becomes evident when one observes the growth rates'correlation. -0.763, compared to the observed one 0.25. This is a problem that this model shares with the usual framework. The decreasing marginal productivity of the two goods should help mitigate the problem with respect to the usual framework, but in practice, given the very small curvature in the transformation frontier, the improvement with respect to the usual framework is very weak. In principle, the decreasing marginal productivities in the production of the two goods induce the sectors to covary to a certain extent -in order to maintain the consumption share, and therefore the ratio of the two goods, as smooth as possible. However, the small curvature of the transformation frontier implied by $\rho = .28$ is not enough for this purpose. Nevertheless, it helps to increase the correlation between GDP and output, which is .48 in this model, .66 in the data and -.36 in the usual framework.

The fact that labour fluctuations are small suggests that increasing the Frisch elasticity may improve the results. This is done in the next section, where a very elastic labour supply is considered. As in Hansen (1985), this calibration captures the case of adjustments on the extensive margin.

The model with $\nu = 100$

With such an elastic labour supply, the model predicts a variance of GDP 4 times higher than the actual one. This is due to the labour volatility, which now varies 16 times more than in the actual data. Also, the price prediction with the policy function is not as good as before: the R^2 of the detrended data is .3. The only good news with this parametrization is that now the correlation between consumption and investments is around 0, an improvement with respect to before, especially if one compares this with the usual framework, where the correlation is around -.9.

The last experiment suggests that labour supply cannot be so elastic. Considering a value of ν somewhere in between the two values considered and a higher curvature in the transformation frontier, $\rho > .28$ may help match the data. However, this would reduce the fit that the model has in predicting the relative price, which is quite an important dimension for the purpose of the paper: measuring technology shocks and their importance for the Business Cycle. Nonetheless, this parametrization highlights the fact that other sources of fluctuations and propagation mechanisms may also be relevant in explaining the Business Cycle. The choice of a parametrization that shows the strengths and weaknesses of this model characterizes the calibration method with respect to other estimation techniques.

2.6 Conclusions

Counter-factual saving rates and the presence of a negative correlation between the Neutral and Investment-Specific Technology Shocks identified through the neoclassical Growth model as implemented at first by Greenwood et al. (1988), and recently by several authors, casts doubts on the specification of that model for the question at hand - to measure the N-shock and the I-shock and to quantify their importance for the Business Cycle.

It is claimed that the absence of curvature in the transformation frontier between consumption and investment goods is the cause of the two observed facts. The model is enhanced with the above feature in a way that is convenient, in that it allows us to use aggregate data, to fully calibrate a two-sector model, and not to alter the balanced growth path prediction of the original framework, which is able to reconcile the decline in the relative price of investments with the relative increase in the production of investment goods.

The distinctive prediction of this model is that the relative price is now a function of both the shocks, and not only of the investment one as in the original framework. This depends on the curvature of the transformation frontier, which makes the relative price depend on the relative demand for the two goods, i.e. the saving rate, which in turn depends on both the shocks. The degree of curvature in the transformation frontier is accounted for by only one parameter, which is calibrated in two different ways. The first strategy is to maximize the fit of the predicted time series of the saving rate with the observed one. Given that in the present model the price depends on the relative demand for the two goods, i.e. the saving rate, it is judged to be important to have a good prediction of the saving rate. The second calibration strategy is to pick the parameter that generates uncorrelated time series for the two shocks. The two strategies deliver very close parameters and business cycle predictions, suggesting that the model is well-specified for the question at hand.

While the model shares the growth implications of the original framework, it has different predictions of Business Cycle frequencies. It is relevant that, when identifying the shocks through this framework, the neutral shock is stationary, and the investment one has a stochastic trend. This fact implies that the N-shock plays the most important role in accounting for aggregate fluctuations. This is due to the fact that a transitory shock induces a stronger and more sudden reaction than a permanent one would.

A good feature of the proposed identification technique for the investment shock is that the prediction can be tested using the policy function for the relative price to predict a relative price time series, given the identified shocks, that then can be compared to the actual price time series. Unlike in the original framework, the predicted prices can be wrong[16], since the shocks have not been identified through that equation. A good fit like the one obtained ($R^2 = .99$ for the price in levels and .79 in growth rates) is comforting as evidence of a well-specified price equation.

In this model technology shocks together account for a higher share of aggregate fluctuations than predicted by the original framework: in the preferred calibration, 86% of aggregate fluctuations are explained by the two shocks considered. The presence of curva-

[16] As they are, for instance, when the model is calibrated with an excessively elastic labor supply.

ture in the transformation frontier makes any shift in the relative demand of goods cause a change in aggregate production and in the relative price. This further propagation mechanism, not present in the usual framework, is responsible for the larger portion of fluctuations explained by the aggregate shocks. The fact that not all price variations are captured by I-shocks, contributes to making this shock less important for the Business Cycle.[17].

The presence of curvature in the production frontier allows an improvement in the prediction of the correlation between the growth rates of consumption and aggregate production, a dimension where the original framework performs poorly. However, due to the small degree of curvature, this correlation is still too low compared with the actual data. This is also highlighted by the counter-factual negative correlation of the growth rates of investment and consumption, where, due to the chosen parametrization, the improvement with respect to the original framework is very small. A higher curvature would improve over this dimension, but the aim of identifying uncorrelated shocks and a good prediction of the saving rate, imposes discipline on the calibration, and the good fit of the price equation suggests that the proposed parametrization is well suited to the question at hand, to measure technology shocks. The good and bad features highlighted by the calibration chosen, suggest that future research should be carried out aiming to improve the present mechanism in a way that would allow for a greater curvature in the production frontier, which would allow consumption and investments to co-vary, while at the same time predicting the right saving rate time series and essentially uncorrelated technology shocks. On this dimension, the introduction of other frictions, such as inter-temporal adjustment costs, may be found to be complementary to the present one and improve the fit of the model with the data.

[17]The mechanism is explained in more detail in section 2.5.2

2.7 Appendix

2.7.1 Data

The relative price time series is as in Fisher (2006).

Non-farm hours of work and population come from The Bureau of Labor Statistics.

Non-durable consumption, investments, capital stock comes from the Bureau of Economic Analysis.

2.7.2 Balanced Growth Path with trend-stationary shocks

The equilibrium conditions are

$$\frac{c_{t+1}}{c_t} = \beta E\left\{\left[1+\theta A_{t+1}k_{t+1}^{\theta-1}n_{t+1}^{1-\theta}\left[s_{t+1}^{1-\rho}+p_{t+1}V_{t+1}(1-s_{t+1})^{1-\rho}\right]/p_{t+1}-\delta\right]\frac{p_{t+1}}{p_t}\right\} \quad (2.23)$$

$$\frac{(1-\theta)A_t k_t^{\theta} n_t^{-\theta}\left[s_t^{1-\rho}+p_t V_t(1-s_t)^{1-\rho}\right]}{c_t} = \xi n_t^{\nu} \quad (2.24)$$

$$A_t k_t^{\theta} n_t^{1-\theta} s_t^{1-\rho} = c_t \quad (2.25)$$

$$V_t A_t k_t^{\theta} n_t^{1-\theta}(1-s_t)^{1-\rho} = k_{t+1} - k_t(1-\delta) \quad (2.26)$$

$$p_t V_t = \frac{(s_t)^{-\rho}}{(1+s_t)^{-\rho}} \quad (2.27)$$

$$V_t = V_0\gamma_v^t \exp(u_{vt}) \quad (2.28)$$

$$u_{vt} = \rho_v u_{vt-1} + \varepsilon_{vt} \quad (2.29)$$

$$A_t = A_0\gamma_a^t \exp(u_{at}) \quad (2.30)$$

$$u_{at} = \rho_a u_{at-1} + \varepsilon_{at} \quad (2.31)$$

In a stationary environment ($\rho_v < 1$, $\rho_a < 1$) the model oscillates around the deterministic growth path. In this case Equation 2.23 is consistent with a B.G.P if

$$(1+\mu_a)(1+\mu_k)^{\theta-1} = (1+\mu_p) \quad (2.32)$$

$$\mu_s = 0, \quad (1+\mu_v)(1+\mu_p) = 1. \quad (2.33)$$

Equation 2.24 is consistent with a B.G.P if

$$(1+\mu_a)(1+\mu_k)^{\theta} \;=\; (1+\mu_c) \tag{2.34}$$

$$\mu_s \;=\; 0, \quad (1+\mu_v)(1+\mu_p)=1. \tag{2.35}$$

Conditions $(1+\mu_a)(1+\mu_k)^{\theta-1}=(1+\mu_p)$ $\&(1+\mu_a)(1+\mu_k)^{\theta}=(1+\mu_c)$ imply

$$(1+\mu_k)=\frac{(1+\mu_c)}{(1+\mu_p)}. \tag{2.36}$$

Equation 2.26 is consistent with a B.G.P if

$$(1+\mu_v)(1+\mu_a)(1+\mu_k)^{\theta-1}=1. \tag{2.37}$$

This condition is already implied by $(1+\mu_a)(1+\mu_k)^{\theta-1}=(1+\mu_p)$ & $(1+\mu_v)(1+\mu_p)=1$.
Equation 2.25 is consistent with a B.G.P if 2.34 holds.
Finally, from 2.27 if $(1+\mu_v)(1+\mu_p)=1$ then $\mu_s=0$, consistently with the Euler equation for consumption.
The above conditions summarize to

$$\mu_s \;=\; 0, \quad (1+\mu_p)=1/(1+\mu_v), \tag{2.38}$$

$$(1+\mu_k) \;=\; \left[(1+\mu_a)(1+\mu_v)\right]^{\frac{1}{1-\theta}}, \tag{2.39}$$

$$(1+\mu_c) \;=\; (1+\mu_a)^{\frac{1}{1-\theta}}(1+\mu_v)^{\frac{\theta}{1-\theta}} \tag{2.40}$$

$$1+\mu_a \;=\; \gamma_a,\; 1+\mu_v=\gamma_v\;. \tag{2.41}$$

Define the variable $\hat{x}_t=\frac{x_t}{(1+\mu_x)^t}$ with $x=\{A,V,k,c,p\}$ and rewriting the model with these variables one gets a model with a globally stable steady state, which corresponds to the balanced growth path for the original economy. The model has the same implication for growth as the usual framework.
Detrending the variables in the way mentioned, the model is equivalent to the following:

$$\frac{\hat{c}'}{\hat{c}}=\hat{\beta}\left\{1+\theta\hat{A}'\hat{k}'^{\theta-1}n'^{1-\theta}\left[s'^{1-\rho}+\hat{p}'\hat{V}'(1-s')^{1-\rho}\right]/\hat{p}'-\delta\right\}\frac{\hat{p}'}{\hat{p}} \tag{2.42}$$

$$\frac{(1-\theta)\hat{A}\hat{k}^{\theta} n^{-\theta}\left[s^{1-\rho}+\hat{p}\hat{V}(1-s)^{1-\rho}\right]}{\hat{c}}=\xi n^{\nu} \tag{2.43}$$

$$\hat{A}\hat{k}^{\theta} n^{1-\theta}s^{1-\rho}=\hat{c} \tag{2.44}$$

$$\hat{V}\hat{A}\hat{k}^{\theta} n^{1-\theta}(1-s)^{1-\rho}=\hat{k}'\gamma_v\gamma_a-\hat{k}(1-\delta) \tag{2.45}$$

$$s=\frac{\left(\hat{p}\hat{V}\right)^{-\frac{1}{\rho}}}{1+\left(\hat{p}\hat{V}\right)^{-\frac{1}{\rho}}} \tag{2.46}$$

$$\begin{aligned}
\hat{V} &= V_0\exp(u_v) && (2.47)\\
u_v' &= \rho_v u_v+\varepsilon_v' && (2.48)\\
\hat{A} &= A_0\exp(u_a) && (2.49)\\
u_a' &= \rho_a u_a+\varepsilon_a' && (2.50)
\end{aligned}$$

where

$$\hat{\beta}=\beta\left(\gamma_v\gamma_a\right)^{\frac{1}{1-\theta}}$$

2.7.3 Balanced Growth Path with a trend-stationary and a trend-stochastic shock

Identifying the shocks through this framework, the neutral shock appears to be stationary ($\rho_a = .95$), while the investment one has a stochastic trend ($\rho_v = 1$). The equivalent stationary model when there is a trend-stochastic shock and a stationary one is derived below. Consider the variables

$$\begin{aligned}
z_t &= A_0^{\frac{1}{1-\theta}}\gamma_a^{\frac{(t)}{1-\theta}}V_t^{\frac{\theta}{1-\theta}},\ \tilde{c}_t=\frac{c_t}{z_t},\ \tilde{k}_t=\frac{k_t}{z_tV_t}, && (2.51)\\
\tilde{p}_t &= p_tV_t,\ \tilde{A}_t=\frac{A_t}{A_0\gamma_a^t}. && (2.52)
\end{aligned}$$

Substituting these variables into 2.23-2.31 one gets the following:

$$\frac{1}{\tilde{c}_t} = \beta E \frac{1}{\tilde{c}_{t+1}} \left(\gamma_a \gamma_v exp(\varepsilon_{v,t+1})\right)^{\frac{1}{\theta-1}} \left\{ 1 + \theta \tilde{A}_{t+1} \tilde{k}_{t+1}^{\theta-1} n_{t+1}^{1-\theta} \left[s_{t+1}^{1-\rho} + \tilde{p}_{t+1} (1 - s_{t+1})^{1-\rho} \right] / \tilde{p}_{t+1} - \delta \right\} \frac{\tilde{p}_{t+1}}{\tilde{p}_t} \tag{2.53}$$

$$\frac{(1-\theta) \tilde{A}_t \tilde{k}_t^{\theta} n_t^{-\theta} \left[s_t^{1-\rho} + \tilde{p}_t (1-s_t)^{1-\rho} \right]}{\tilde{c}_t} = \xi n_t^{\nu}$$

$$\tilde{A}_t \tilde{k}_t^{\theta} n_t^{1-\theta} s_t^{1-\rho} = \tilde{c}_t$$

$$\tilde{A}_t \tilde{k}_t^{\theta} n_t^{1-\theta} (1-s_t)^{1-\rho} = \tilde{k}_{t+1} \left(\gamma_a \gamma_v exp(\varepsilon_{v,t+1}\right)^{\frac{1}{1-\theta}} - \tilde{k}_t (1-\delta)$$

$$s_t = \frac{\tilde{p}_t^{-\frac{1}{\rho}}}{1 + \left(\tilde{p}_t\right)^{-\frac{1}{\rho}}}$$

2.8 Figures

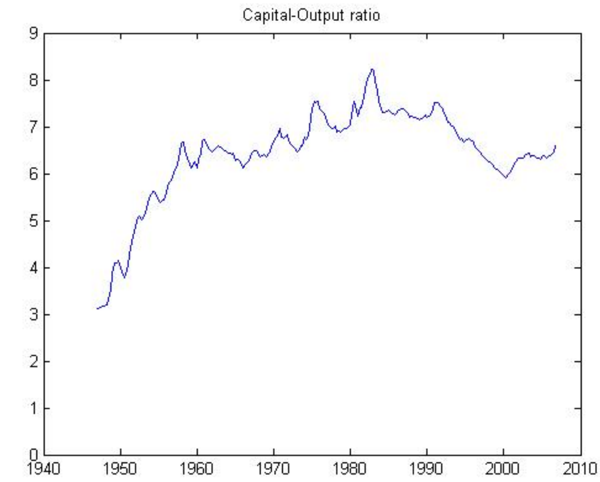

Figure 2.1: Capital Output Ratio

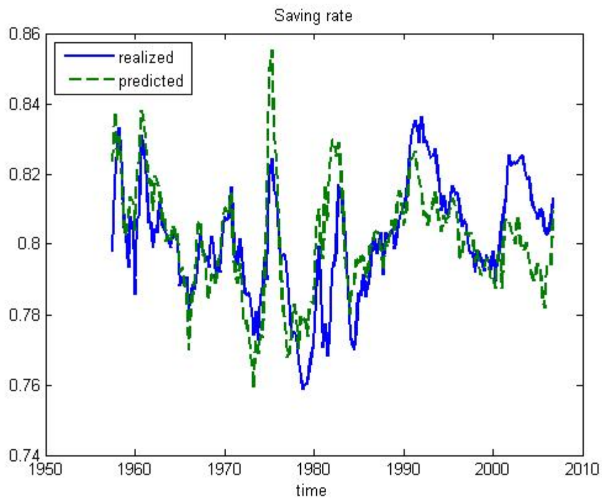

Figure 2.2: Saving Rate Fit

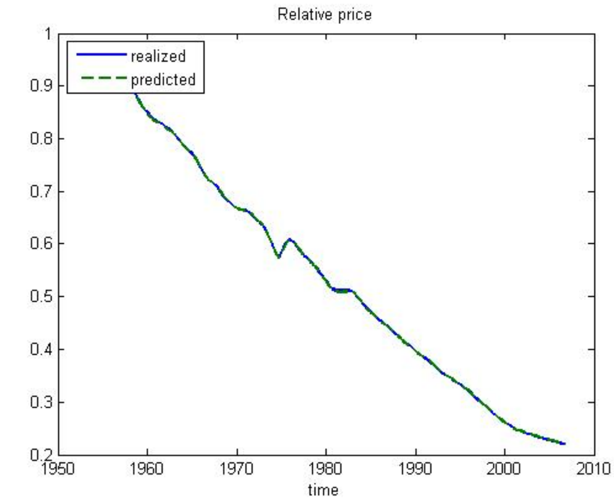

Figure 2.3: Relative Price Fit

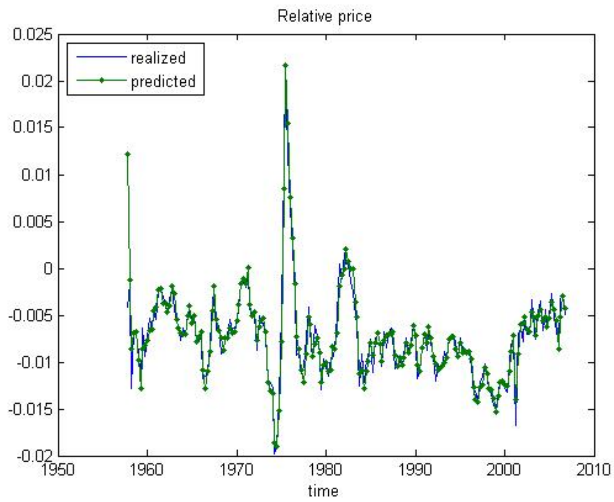

Figure 2.4: Change in Relative Price Fit

Chapter 3

Labour Force Composition and Aggregate Fluctuations

Abstract

I document that while the composition of the labour force within groups defined by gender, age and education has changed substantially over the last 40 years, the relative volatility of hours within these groups has remained remarkably stable. Together, these facts suggest that trends in the composition of the labour force may be reflected in aggregate volatility. To investigate this conjecture, I develop a large-scale real business cycle model where agents are distinguished by gender, age and education, which reproduces the aforementioned facts. To solve the model, I develop a computational technique which consists of applying perturbation methods at several points over the transition path. This methodology breaks the curse of dimensionality and enables to solve large scale models with accuracy over all the relevant part of the state space. The model is then used to show, through counterfactual simulations, that changes in the labour composition explain part of the high aggregate volatility during the 70's as well as the subsequent moderation.

3.1 Introduction

Aggregate output volatility increased during the 1970s, decreased during the 1980s and 1990s and it appears to be increasing in the most recent years. Meanwhile, several changes characterized the composition of the labour force since the late 1960s, with substantial increases in women's participation rate, the fraction of highly educated workers and the fraction of prime-age workers. I complement this empirical evidence by documenting stable differences in aggregate hours volatility by gender, age and education. These volatility differences and the changes in the labour composition imply a redistribution of workers across groups of different volatility. The paper investigates whether this transformation may have contributed to the mentioned changes in aggregate fluctuations.

The possible presence of omitted variables and simultaneous causality between these groups and aggregate volatility calls for a structural approach aimed at isolating the effects of changes in the labour composition on aggregate volatility.[1] An example of reverse causality is the fact that education tends to be countercyclical: poor labour market conditions can lead to substitute labour with education. In turbulent times, when recessions are frequent, one could therefore observe an increase in education. However, this positive relation is due to the effect of aggregate volatility on educational choices, and not vice versa; using regression analysis one may wrongly conclude that education tends to increase aggregate volatility. Similarly, female labour force participation may be affected by aggregate volatility.

These considerations motivate the following methodology. Following Attanasio and Weber (1995), I create a syntectic panel by grouping the labour force by the mentioned observables from the March supplement of the Current Population Survey (CPS). I docu-

[1] Jaimovich and Siu (2009) use panel data techniques to relate aggregate volatility changes to changes in the population by age. While it seems convincing to assume that the age composition cannot respond to the current business cycle, it may be incautious to extend their methodology to the other mentioned factors.

ment the changes in the share of each group and their cyclical volatility. I then build an overlapping generations model with educational and marriage choices, similar to Heathcote et al. (2010), but with stochastic productivity shocks.

One challenge was to find a solution method for this large model[2] which guarantees sufficient precision over all the transition path that characterized the last 40 years. To this aim, I have developed a technique which essentially consists of applying perturbation methods at many points over the equilibrium path.

The model, calibrated to match the evolution of the labour composition and the relative volatility across groups, is able to account for the mentioned changes in aggregate volatility.

Since the model replicates the mentioned facts, it is possible to run counterfactual simulations aimed at isolating the role of demographic changes on aggregate volatility. If the labour force composition remained trendless at the levels of the most recent years, volatility would have been 5% lower in the early 1970s and 6.5% more volatile in the late 80s and 90s. By accounting for part of the high volatility of the early seventies and of the slowdown in the 1980s, labour group reallocation accounts for about a quarter of the great moderation.

The paper proceeds as follows. Section 3.2 offers a review of the literature that relates aggregate volatility and the considered structural changes. In section 3.3, a first look at the labour data will motivate the conjecture made in this paper, that the reallocation of labour across groups plays an important role for cyclical volatility of hours, and thereby of aggregate income. An accounting exercise to get a preliminary quantification of the effects of labour allocation on aggregate volatility is conducted in section 3.3.4. Section 3.4 describes the methodology, section 3.5 sets up the model and section 3.6 parameterizes it. Section 3.7 offers a brief description of the original computation technique designed for this

[2] 805 variables, of which 324 state variables. The large number of variables comes for the fact that in this model there are several groups of agents, distinguished by gender, age and education.

very large model. Section 3.8 tests the model, section 3.9 quantifies the effects of labour reallocation on aggregate volatility and section 3.11 concludes.

3.2 Literature review

The large and sudden volatility decline from the mid 1980s, initially documented by Kim and Nelson (1999) and McConnell and Perez-Quiros (2000), stimulated a debate on its causes. Researchers were mainly divided between three explanations: some structural change, improved macroeconomic policies and good luck. At present, in the aftermath of the crisis, it becomes policy relevant to readdress this question and in general to understand whether there is a reason for aggregate volatility changes beyond good luck: the answer has implications for the amount of aggregate volatility to expect in the future, which should be relevant for the regulation of the financial sector. In the words of Peter A. Diamond, "Among the issues being debated now is how much we should increase capital requirements for banks. Selecting the proper size of the increase requires a balance between reducing the risk of a future crisis and ensuring the effective functioning of financial firms in ordinary times ", see Diamond (2011). The trade-off highlighted by Peter A. Diamond depends on the actual amount of aggregate risk and thereby on the risk of future crisis.

The present paper contributes to the structural change hypothesis, but it also confirms what suggested by Bernanke (2004), that these explanations are jointly needed to explain all of the changes in volatility over time. Of course, the less one needs to rely on shocks identified as residuals, the more one can feel confident about his or her understanding of changes in aggregate volatility.

The dynamic stochastic general equilibrium literature emphasized how the great moderation is accounted for mainly by a reduction in the volatility of aggregate shocks, see Arias et al. (2007) and Smets and Wouters (2007). The present framework allows to investigate whether the observed redistribution of labour, from highly volatile to more stable groups, can be held responsible for both the reduction in hours volatility and of measured aggregate productivity shocks. [3]

During the considered period of time, several studies, for example Katz and Autor (1999) for the U.S, documented several ongoing changes in the labour composition (see Freeman and F.Katz (1994) for a cross-country comparison). Davis and Kahn (2008) conjectured that changes in volatility by sector, and the increasing share of services, may have contributed to the great moderation. Similarly, Owyang et al. (2008) find that across states, changes in volatility were related to the initial share of the durable sector and Carvalho and Gabaix (2010) point at changes in the sectorial composition of output. Interestingly, the latter paper accounts for the mentioned changes in aggregate output volatility through firm level data and it may be seen as the dual of the present paper which focuses on household level data. The two studies, which have similar qualitative results, may help deriving different and perhaps complementary policy implications: one pointing at the implications of the changes in the composition of firms, and the present one, at changes by demographic characteristics, observable through household level data. As already mentioned, Jaimovich and Siu (2009) relate changes in aggregate volatility to the fact that, as a consequence of the baby boom, prime age workers, less sensitive to the business cycle than young and old workers, increased in their relative size.

For a more extended review that also focuses on other explanations of aggregate volatility changes see Blanchard and Simon (2001), Stock and Watson (2002) and Stock and Watson (2005).

[3] Aggregate labour supply and labour productivity are in fact decomposed into the contribution of the considered labour groups.

3.3 Stylized facts

This section documents the trend and cyclical volatility in hours by gender, age and education. The major empirical contribution is to show that while the composition of the labour force changed substantially during the period considered, from 1967 to 2009, the relative volatility across groups is different than one and remarkably stable over time. [4] This evidence leads to conjecture that at least part of the high volatility in the early 1970s and the subsequent reduction in cyclical volatility of hours worked can be accounted by the reallocation of workers over the labour force as detailed below.

3.3.1 Gender

Figure 3.1 shows hours worked by women and men from 1967 to 2009. Data come from the Current Population Survey, March supplement. As it is well known and shown in the first two columns of table 3.1, employment of women increased relative to men: women employment over total population moved from an average of 13.89 % between 1968 and 1983, to an average of 19.48 % in the second sub-sample.[5]. It is less well known however that hours worked by men are more volatile than hours worked by women; after removing the trend from each series using the Hodrick-Prescott (HP) filter, [6] the standard deviation of hours is 1.44 % for women and 1.86 % for men. Furthermore, the relative volatility is remarkably stable over time: it is 0.77 between 1967 and 2009; starting from a later date well inside the second sub-sample, for example 1992, the relative volatility is even lower: 0.71. Figure 3.2 shows the evolution of the relative volatility of hours by women and men; the figure confirms that this statistic remains stable over time.

[4] The sample starts in 1967 because before that date, CPS data do not seem comparable over time.

[5] 1984 is the reference year adopted by the literature as the beginning of the great moderation, see for instance Stock and Watson (2002)

[6] I use a smoothing parameter of 6.28 as suggested by Ravn and Uhlig (2002). Using 10 as suggested by Baxter and King (1999) does essentially not affect the results.

The relative increase of women hours and the fact that they are less volatile than the one for men may have contributed to the observed reduction in aggregate hours volatility. It is important to notice, however, that hours volatility for the two genders decreased substantially in the great moderation period: the standard deviation was approximately two times higher in the pre-moderation period (until 1983) relative to the period 1984-2007. The part of the reduction in hours volatility that happened within the genders, cannot be accounted for simply by the gender reallocation of the labour force. [7] In light of this last consideration, it becomes perhaps even more remarkable how, although the volatility of total hours by gender moves over time, the relative volatility remains stable. The next table summarizes these facts.[8]

Table 3.1: Hours Volatility Gender

	workers' share 68-83	workers' share 84-07	st. dev
Women	13.89	19.48	1.44
Men	20.89	22.30	1.86

Notes: numbers are expressed in percent terms. Workers' share is the ratio of workers over total population. St.deviation measures the standard deviation of the HP filtered hours with parameter 6.28

3.3.2 Age

Figure 3.4 shows hours worked by the young, prime age and older workers. The young are individuals from 15 years old to 29, the prime age range is from 30 to 50, while the old are aged 50 to 70. Jaimovich and Siu (2009) documented the increase in the relative size of the prime age population, mainly as a consequence of the baby boom. As shown in table 2, in

[7] Hours volatility changes for each gender could however be explained by the reallocation of the labour force by other observables: changes by education, age and possibly sector in favor of stable groups may induce a moderation in hours by gender.

[8] Figure 3.3 plots employment over population by gender for the remaining G7 countries, where the volatility of aggregate income had a similar pattern to the one of the U.S. Data are from OECD.stat, at annual frequencies. As it is evident from the figure, similarly to what has been documented for the U.S., female employment grew more than for men.

the first sub-sample, 15% of the population were prime age workers, they increase to 22% in the second sub-sample. Jaimovich and Siu (2009) relate this fact to the great moderation because hours worked by prime age are less volatile. Indeed I find that the cyclical volatility of hours is substantially lower for prime age workers and its volatility relative to the young it even reduces over time over time as shown in figure 3.5. Hence the increase in their relative size should contribute to the reduction in aggregate hours volatility. Furthermore, the decrease in the share of prime age workers in the last decade could be responsible for part of the observed volatility increase. It is however interesting to notice that, compared to young and old, prime age hours had the largest reduction in volatility in the period from 1984 onward (see the last column of table 2). The fact that some of the moderation happens within the prime age group implies that the redistribution of workers by age, although able to account for a substantial part of the overall moderation in hours as shown by Jaimovich and Siu (2009), cannot explain why hours volatility diminished within the prime age group. This reduction could be explained by the redistribution within each age group, in terms of the other considered observables. I refer to Jaimovich and Siu (2009) for a cross country

Table 3.2: Hours Volatility Age

	workers' share 68-83	workers' share 84-07	st. dev	$\frac{st.dev(84-07)}{st.dev(68-83)}$
Young	12.02	11.68	2.32	74.51
Prime age	15.23	21.82	1.36	61.87
Old	7.54	8.28	1.33	.69

Notes: numbers are expressed in percent terms. Workers' share is the ratio of workers over total population. St.deviation measures the standard deviation of the HP filtered hours with parameter 6.28

comparison of the age share dynamics.

3.3.3 Education

Figure 3.6 plots per capita hours worked by low and high education (at least four years of college). The share of highly educated increased more rapidly than the one of less educated, moving from an average of 5.8 % in the first sub-sample, to 11.4 % in the second one. The cyclical volatility of highly educated is 1.16 %, lower than the one of low educated, 1.91 %. The effect on aggregate hours volatility of this change in the composition are mitigated by the fact that the cyclical volatility of hours worked by low education workers reduced in the second sub-sample relatively more than the ones of high education workers. This can be seen in the last column of table 3.3 and from figure 3.7 which shows how the relative volatility of hours over time has a kink in the second sub-sample, it remains however true over the whole sample that the relative volatility between the two groups is smaller than one. [9] Freeman and F.Katz (1994) confirm that the employment-population ratio of the less educated fell relative to that of more educated workers in many countries during the 1980s.

Table 3.3: Hours Volatility Education

	workers' share 68-83	workers' share 84-07	st. dev	$\frac{st.dev(84-07)}{st.dev(68-83)}$
Low	28.99	30.31	1.16	60.29
High	5.80	11.39	1.91	98.59

Notes: numbers are expressed in percent terms. Workers' share is the ratio of workers over total population. St.deviation measures the standard deviation of the HP filtered hours with parameter 6.28

[9] Weights and classifications in the CPS data are such that data are essentially comparable over time. It is known however that comparability does not hold in some cases; the observed kink may be partly due to some reclassification that induced discrepancies in the way data are constructed over time.

3.3.4 An accounting exercise

In light of what found by Jaimovich and Siu (2009), on the importance of demographic changes by age on aggregate volatility, and given the other facts introduced in the previous section, it becomes natural to ask how much overall these compositional shifts matter for aggregate volatility changes. As mentioned in the introduction, the extension of the analysis of Jaimovich and Siu (2009) to other partitions of the labour force brings some methodological difficulties, mainly due to the potential endogeneity between the labor composition and aggregate volatility. Therefore, before dueling with these methodological challenges it would be nice to get a sense of how much changes by gender and education may contribute on top of changes by age, already considered by Jaimovich and Siu (2009).

To this aim, I design an initial partial equilibrium exercise to account for the importance of the mentioned factors. The exercise attempts at answering the following question: what would have aggregate hours volatility be, had we not observed the mentioned compositional shifts? I take the following approach: I assume that hours of each group moved as they did, but around a counterfactual linear trend such that the average labour composition remains unchanged over the entire sample. This creates some counterfactual hours for each group, where their average composition remains constant, while their higher frequencies are left as observed.

The advantage of this procedure is that it enables to maintain the observed reaction of labour to the shocks. This isolates the role of labour reallocation, while maintaining unchanged the short run elasticities of each labour group.

The next subsection details the procedure to estimate counterfactual labour and quantifies the role of this counterfactual redistribution for volatility changes in aggregate hours.

Counterfactual hours

To generate counterfactual hours I take a procedure which is very close to the one adopted by Jaimovich and Siu (2009) to quantify the role of the change in the distribution of the population by age. Aggregate hours per capita can be written as follows

$$H_t = \sum_i h_{i,t} p_{i,t}, \tag{3.1}$$

where $h_{i,t}$ is hours per employees in group i at time t, and $p_{i,t}$ is employees of group i over total population.

Counterfactual per capita hours are instead

$$\widehat{H}_t = \sum_i h_{i,t} \widehat{p}_{i,t} \tag{3.2}$$

where $\widehat{p}_{i,t}$ is counterfactual employees over population for group i at time t. To construct $\widehat{p}$ I remove a linear trend in the time series of p.[10] Since aggregate labour participation has increased over time for each group, I do not want to remove this positive long run trend in aggregate hours per capita. Therefore I add to the de-trended data the same aggregate trend. This creates some counterfactual hours for each group, where their average composition remains constant, while their high frequencies are left as observed. I choose to keep as the average composition level, the one from 1976-1978. This is convenient in that it precedes what is typically considered to be the beginning of the great moderation, 1983-1984. A second advantage is that this date comes after a period of volatility acceleration, the early seventies, which include two official National Bureau of Economic Research recessions, in 1969-1970 and 1973-1975. By isolating this period of time, it is also possible to measure the role played by labour composition in this earlier period of relatively high volatility.

[10] I also tried to hp-filter the data with smoothing parameter equal to 6.25, which is the one suggested for annual data by Ravn and Uhlig (2002). However I found the results quite sensitive to the smoothing parameter. Although under some parameter values I could explain a larger portion of the observed volatility reduction than what I do with linear trends (I account for a larger portion of the great moderation this way), the fact that the results depend on the parameter choice made me opt for the linear trend, which gives weaker results.

One clarification is necessary at this point: overall I have 12 groups, however I am not de-trending them one by one because this would remove all trends, even the ones that are not due to gender, age or education. The aim is instead to identify the role of each of these variables separately.

The effect of changes in labour composition

In what follows I quantify the role of redistribution by gender, age, education one by one.

Figure 3.22 shows realized and counterfactual hours worked by women and men. As can be seen, the exercise consists of shifting upward the trend for men, and reducing the trend for women; this leaves unchanged the average shares between the first and second sub-sample. Figures 3.24, 3.26 show realized and counterfactual employment by age and education.

Figures 3.23, 3.25 and 3.27 show observed and counterfactual aggregate hours by gender, age and education. It can be noticed how for the last periods of the sample total hours are larger in the counterfactual by gender, where men command a larger share of employment. This is due to the fact that men work more hours per worker than women. Similarly, once age trends are removed, counterfactual hours are lower in the last part of the sample. This happens because the counterfactual reduces the share of prime age workers, who work more hours per worker.

Following Jaimovich and Siu (2009), comparing the standard deviations of the growth rates in the first sub-sample (1968-1983) and in the second one (1984-2007), [11] one gets a measure of the moderation in aggregate hours: between the two sub-samples, aggregate hours volatility decreased by 46.21 log points.[12] Had the shares by gender remained on aver-

[11] I end the second sub-sample before the last recession.

[12] $46.21 = \log(st.dev(hoursgrowth_{67-83})) - \log(st.dev(hoursgrowth_{84-00}))$

age what they were in the first sub-sample, we would have observed a reduction in volatility of 45.23 log points. Therefore, changes by gender account for $(46.21 - 45.23)/46.21$ or approximatively 2% of the moderation in aggregate hours. Changes by age, education and sector account respectively for 5%, [13] and 10% of the moderation in aggregate hours. Table 3.4 summarizes these results. From this simple exercise one may hazard that changes

Table 3.4: % explained of the moderation

Gender	Age	Education
2.1	5.2	9.7

in education are the most important contributors to the great moderation, followed by changes in age and gender composition. Some words of caution are however necessary to state that this numbers do not quantify the relative importance of each demographic change considered. Changes in education, for instance, may reflect the optimal educational choices of women once their participation rates started increasing. Changes by education would then be a consequence of changes by gender and the exercise above would understate the importance of gender composition for aggregate volatility. Similar examples may be constructed to over or understate the importance of each partition. Put differently, omitted interactions may bias the results of the pervious accounting exercise. This said, the result

[13]The number for age is lower than the one obtained by Jaimovich and Siu (2009), who account for about the 20 % of the moderation in hours. Two factors explain the difference. First, they split the population in finer age groups. Considering bins of 5 years as they did, I get a slight increase in the portion of the moderation that age can account for. Given the small difference, I choose not to consider finer bins because later I will consider the interactions with the other partitions, and I need to keep the number of groups small in order to maintain the samples representative. The second source of the discrepancy is that while they consider counterfactual population shares, I focus on counterfactual employment shares. It turns out that changes in the population shares are larger than changes in the employment shares. This is due to the fact that the reduction in the young portion of the population is mitigated by an increase in their working rate, and this depends mainly on the increased labour participation of women. It is debatable whether it is preferable to focus on population or on employment changes; I focus on employment changes because this allows me to analyze jointly changes in age and education with the ones by gender and sector, where the relevant changes happened in labour, rather than in population shares. In defence of my approach it can be argued that by focusing on employment instead of population shares, one also takes into account the endogenous changes in participation rates, which seem to be important.

above seems to be sufficient to conjecture that compositional changes by gender and education may also be important on top of the ones by age, thereby justifying the investigation beyond the one considered by Jaimovich and Siu (2009). As already mentioned, however, this extension brings some methodological difficulties analyzed in the next section, which outlines the methodology.

3.4 The Methodology

Jaimovich and Siu (2009) used regression analysis based on an orthogonality assumption between the age distribution and cyclical volatility. While it seems reasonable to assume that the age composition is orthogonal to the business cycle because, abstracting from migration, is a reflection of fertility decisions made several years before workers enter the labour force, this condition would be hard to justify when studying the other partitions of the labour force considered in this study. One needs a methodology that isolates the role of changes in the labour composition on aggregate volatility, from reverse causality and from other factors that may affect both the labour composition and aggregate volatility. To this aim I follow a structural approach which consists of estimating a model of the business cycle with the heterogenous labour groups considered. And then use the model as a laboratory to run counterfactual experiments aimed at quantifying the importance of the labour composition on aggregate volatility. An advantage of the model relative to the accounting exercise of the section above is that it explicitly models the endogenous choices of the several labour groups. To take the results seriously, however, the model has to be considered reliable for the question at hand. This is done by testing the extent to which the model can match the several facts that this model can predict and that concern this study.

3.5 The Model

In each period the economy is populated by a continuum of individuals, equally many males and females and a new cohort of random size $2 * p_0$ is born. The process for the size of the new generation is recursive and such that the total amount of the population is a stochastic stationary process.

Following Heathcote et al. (2010), the life cycle of an individual comprises 3 sequential stages: education, matching and work. The first decision -high or low education- is made by a newly born individual man or woman before entering the marriage market. At this point, members of the opposite sex are randomly matched (no-one remains single). For tractability, these two stages happen during the first period of life.

From the second period of life, the couple enters the working stage and jointly choose hours of work of husband and wife, as well as consumption and savings, until they die.

In the next sub-sections, these three stages are described in more detail.

3.5.1 Education

In each period the newly born have to make a discrete choice between college (h) or lower schooling (l). When they are born, they draw an idiosyncratic cost κ of acquiring high education from the distribution

$$\kappa \sim F^g(\kappa) = \ln(\kappa) \sim N(\bar{k}^g, \upsilon^g) \tag{3.3}$$

with g equal men (m) or female (f). This cost captures in reduced form the utility and financial factors that make acquiring a college degree costly. Individuals decisions are made by comparing their education cost with the value gain upon entering the labour market with high education: $M^g(h;\omega) - M^g(l;\omega)$. $M^g(e;\omega)$ is the gender specific expected life time utility of entering the marriage stage as a function of education $e \in \{h, l\}$ and all the other relevant state variables represented by ω. [14] They choose high education if

$$M^g(h;\omega) - M^g(l;\omega) > \kappa. \tag{3.4}$$

For each gender, the share of highly educated in the cohort just born is therefore

$$q^g(\omega) = F^g\left(M^g(h;\omega) - M^g(l;\omega)\right). \tag{3.5}$$

3.5.2 Marriage

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

$$M^f(h;\omega) = \pi^f_{h,h}(\omega)V(h,h;\omega) + \pi^f_{l,h}(\omega)V(l,h;\omega), \tag{3.6}$$

$$M^f(l;\omega) = \pi^g_{h,l}(\omega)V(h,l;\omega) + \pi^f_{l,l}(\omega)V(l,l;\omega), \tag{3.7}$$

where $V(e^m, e^f, \omega)$ is the expected future life time utility of a household formed by a man with education e^m and a woman with education e^f. Similar expressions can be derived for $M^m(e,\omega)$:

$$M^m(h;\omega) = \pi^m_{h,h}(\omega)V(h,h;\omega) + \pi^m_{h,l}(\omega)V(h,l;\omega), \tag{3.8}$$

$$M^m(l;\omega) = \pi^g_{l,h}(\omega)V(l,h;\omega) + \pi^f_{l,l}(\omega)V(l,l;\omega). \tag{3.9}$$

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

$$q^m(\omega)\pi^m_{h,l}(\omega) = (1 - q^f(\omega))\pi^f_{h,l}(\omega). \tag{3.10}$$

Since no individual will remain single

$$\pi^m_{e^m,l}(\omega) + \pi^m_{e^m,h}(\omega) = 1$$

[14] As it will be discussed in section 3.5.6, ω contains all the shocks, the distribution of assets across households, and the distribution of households by age and education of husband and wife.

for any e^m and similarly for women:

$$\pi^f_{l,e^f}(\omega) + \pi^f_{h,e^f}(\omega) = 1 \tag{3.11}$$

for any e^f.

One can show that the cross-sectional Pearson correlation between education levels of husband and wife, a measure of the degree of assortative matching is

$$\varrho = \frac{q^m(\omega)\pi^m_{h,h}(\omega) - q^m(\omega)q^f(\omega)}{\sqrt{q^m(\omega)(1-q^m(\omega))q^f(\omega)(1-q^f(\omega))}}. \tag{3.12}$$

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

3.5.3 Work

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

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

Subject to the budget constraint

$$\zeta^j a' + c = a(1+r) + w(m, j, e^m)l^m + w(f, j, e^m)l^f.$$

where r is the interest rate and $w(g, j, e^g)$ the wage for each age, education and gender. Age specific borrowing limits are also imposed to the households' problem: [15]

$$a' > \underline{a_j}.$$

$\zeta^j \in [0,1]$ is the survival factor at age j, it will be parameterized so that people die for sure at age J, i.e. $\zeta^J = 0$. l_h is an exogenous time cost specific to women. [16] The expectation

[15] Age specific borrowing limits are imposed to avoid Ponzi schemes in the presence of age specific survival factors.

[16] Absent a more sophisticated theory of the household, the evolution of this parameter will help reproducing the distribution of hours across gender. Its reduction over time captures in reduced form housework

is taken over ω' given ω. The Value function V defines expected discounted utility as a function of the household's state variables. The value at the time of forming a household is equivalent to the expected life time utility of a formed household of age 1 and with zero assets

$$V(e^m, e^f; \omega) = E\left(V(e^m, e^f, 1, 0; \omega')\right), \tag{3.13}$$

where the expectation is taken over ω' given ω.

3.5.4 Households Distribution and its Law of Motion

Denote $p_{age,edu} : \{1, ..., J\} \times \{h, l\} \times \{h, l\} \to \Re_+$ the mass of households by age and education of husband and wife.

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

3.5.5 Firms

Competitive firms maximize profits using the following production function

$$y = A^{1/\theta}\left(\alpha k^\theta + (1-\alpha)L\right)^{1/\theta}, \tag{3.14}$$

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

$$L = \left(\sum_{i=1}^{I}(z_i n_i)^\sigma\right)^{1/\sigma}. \tag{3.15}$$

production technology improvements and fall in child care costs which on top of the reducing gender wage gap help explaining increases in women participation rates. See among others Greenwood et al. (2005) and Attanasio et al. (2008)

σ measures the degree of complementarity across groups.[17] z_i's are labour augmenting techology shocks specific to each labour group, n_i is hours worked by all individuals categorized in group i. There is a mapping between groups i and individuals: each group i is formed by agents of the same gender, age group and education level. The mapping is represented by I dummy matrixes $\chi_i(g, e^m, e^f, j)$ which contain zeros and ones depending whether the labour input of the agent belongs to group i. So for instance, group 1 is formed by women, young and with low education. For a generical i,

$$n_i = \sum_g \sum_{e^m} \sum_{e^f} \sum_{j=1}^{J} l^g \left(e^m, e^f, j\right) p_{age,edu} \left(e^m, e^f, j\right) \chi_i(g, e^m, e^f, j). \tag{3.16}$$

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

The representative firm hires labour according to the following first order condition

$$(1-\alpha)Ay^{1-\theta} \left(\sum_{i=1}^{I} (A_i n_i)^{\sigma} \right)^{\theta/\sigma - 1} z_i^{\sigma} n_i^{\sigma-1} = w_i \tag{3.17}$$

for every i, where w_i is the wage rate for group i.

Capital is demanded according to the following condition

$$A\alpha \left(\frac{k}{y} \right)^{\theta - 1} = r. \tag{3.18}$$

Where r is the gross interest rate of capital.

3.5.6 State Space

To make rational decisions agents need to know their type, [18] and need to predict prices, which depend on the shocks and on the distribution of assets and households across age and

[17] It is assumed here that all the groups have the same complementarity across them and with capital. It would be interesting to extend this function to the one introduced by Krusell et al. (2000) as done by Castro and Coen-Pirani (2008) to study hours cyclicality by skill and by Jaimovich et al. (2009) to study the hours volatility by age (see also Johnson and Keane (2007)). However, this would make it harder to identify the shocks analytically, thereby complicating the estimation procedure, and is left for future research. See section 3.6 for a further discussion of this production function.

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

education pair of husband and wife. The next sub-sections define the state space in more detail.

Exogenous Processes

Let the logarithm of the productivity processes z_i, the logarithm of TFP process A, the logarithm of the mass of new born p_0 be AR1 stochastic processes. Furthermore, the cost of acquiring education $\kappa \in \mathcal{K}$ and women housework $l_h \in L_h$ are deterministic processes with an AR1 structure. Starting with initial values away from their steady states, these two variables will help making hours by gender and education enrollment rates behave as in the data.

Let $G \equiv \mathcal{A} \times Z_1 \times ...Z_I \times P^0 \times \mathcal{K} \times L_h$ be the state space for these variables.

Households Distribution

Since the distribution $p_{age,edu}$ is a state variable, one needs to define its set. From how $p_{age,edu}$ has been constructed in section 3.5.4, it follows that it depends on the series p_0^0, $q^m \in [0,1]$ and $q^f \in [0,1]$ at the period of birth of each cohort which is alive. [19] $p_{age,edu}$ is therefore generated from the set $M = {P^0}^J \times [0,1]^{2J}$. Let $\mathcal{P}$ be the set of all admissible distributions $p_{age,edu}$ generated from the set M. The state space for this economy is

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

The first three dimensions of the state space, $\{1, ..J\} \times \{h, l\}^2 \times K$ contain the household's state variables : age, education of husband and wife, and asset holdings which belong to the set $K \equiv \left[min\left(\underline{a_j}\right), \overline{a}\right]$. The second part of the state-space contains aggregate state variables that affect households decisions through prices and expectations: the shocks, the

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

distribution of households by age and education, and the distribution of assets across groups. [20], [21], [22]

Collecting the aggregate state variables in the set

$$\Omega = G \times \mathcal{P} \times K^{(J-1)*4},$$

the state space can be rewritten as

$$S = \{1, ..., J\} \times \{h, l\}^2 \times K \times \Omega,$$

which maps with what used in previous subsections, where I distinguished between individual state variables and ω, element of Ω.

3.5.7 Equilibrium

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

1. *The education decision rules $e^g(\kappa; \omega)$ solve the individual problem* (3.4) *and $q^g(\omega)$ is determined by* (3.5).

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

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

[22]The distribution of idiosyncratic shocks κ affects enrollment rates q^g through its parameters, which are exogenous and constant. κ is therefore omitted from this characterization of the state space, κ is however considered when dealing with individual educational choices.

2. *The matching probabilities* $\pi^g_{e^m,e^f}(\omega)$ *satisfy* (3.10)–(3.11) *and are consistent with the degree of assortative matching* ρ^* *in* (3.6.2)*. Moreover the pre-marriage discounted utilities* M^f *and* M^m *are defined by* (3.6)–(3.7) *and* (3.8)–(3.9)*. The pre labour values* V *are defined by* (3.13)*.*

3. *The decision rules for consumption, labour and assets* c*,* a'*,* l^m *and* l^f*, and the value function* V *solve the household problem in 3.5.3.*

4. *Capital and labour inputs are allocated optimally; that is, they solve the firm problem satisfying equations* (3.17)–(3.18)*.*

5. *Labour markets clear. i.e. equation* (3.16) *holds.*

6. *The capital market clears:*

$$k = \sum_{g}\sum_{e^m}\sum_{e^f}\sum_{j=1}^{J} a(e^m, e^f, j). \tag{3.19}$$

7. *The goods market clears:*

$$c + k' - k(1-\delta) = y, \tag{3.20}$$

where

$$c = \sum_{e^m}\sum_{e^f}\sum_{j=1}^{J}\sum_{j_s=1}^{n_s} c(e^m, e^f, j, \omega) \tag{3.21}$$

and

$$k' = \sum_{e^m}\sum_{e^f}\sum_{j=1}^{J} a'(e^m, e^f, j, \omega). \tag{3.22}$$

8. *The distribution of households evolves as stated in section 3.5.4.*

3.6 Parametrization

It is useful to divide the parameters of this model into two categories: the production function parameters and all the other parameters. It is possible to calibrate the parameters belonging to the latter group for which an extensive literature is available. On the other hand, the presence of heterogeneous groups of workers makes the production side of this economy non conventional. Furthermore, since this study is interested in the implications of labour reallocation on aggregate volatility, it makes sense to allow for group specific shocks; there would be very little hope in replicating the observed trends in the composition of labour without having shocks that move wage premia across groups as observed in the data.

In choosing a technology for production, I make the smallest possible deviation from the Cobb-Douglas production function, typically assumed in the real business cycle literature, and consider the constant elasticity of substitution production function defined by equations (3.14)-(3.15) . This specification distinguishes between a labour augmenting (LA) and a total factor productivity (TFP) shock. This will allow to match aggregate production given capital and labour inputs, while identifying LA shocks through labour demand.[23] Since this technology is non conventional, some of its parameter values cannot be found in the literature and need to be estimated, namely the complementarity across labour groups. This is done in the next subsections.

3.6.1 Heterogeneous labour aggregator

The productivity shocks z_i are identified through the labour demand equations for every i (equation 3.17). In principle one could identify the I group shocks z_i and the neutral shock A by solving (3.14)–(3.17), then construct a likelihood function of these shocks and maximize it with respect to the parameters. At least three difficulties arise: i. these equations are non linear simultaneous functions of all the shocks, therefore it is not obvious how to back them out; ii. the shocks may not be exogenous to the right hand side variables; and iii. there may be measurement error.

[23]By Euler theorem, also the capital demand equation will be satisfied with no need of an extra shock.

Although one could use a simulated method to overcome point i., one can instead derive analytical closed form solutions for the shocks as shown below. This will facilitate computation and will help dealing with endogeneity and measurement error.

Estimating the complementarity across labour groups σ

First divide (3.17) for each i by the same equation for group 1.

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

Multiplying by $\frac{n_i}{n_1}$ and taking logs one gets [24]

$$log(z_i)-log(z_1)+log\left(\frac{n_i}{n_1}\right)=1/\sigma log\left(\frac{w_i n_i}{w_1 n_1}\right), \tag{3.23}$$

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

$$\epsilon_t=\bar{\nu}+\gamma t+\rho\epsilon_{t-1}+\nu_t, \tag{3.24}$$

where $\bar{\nu}$ and γ are vectors and ρ is a scalar. I rewrite (3.23) with the new notation

$$z_t=1/\sigma x_t+\epsilon_t+me_t, \tag{3.25}$$

where me_t is an $I-1$ vector of measurement errors. Since z_t affects x_t,the latter is not orthogonal to ϵ_t+me_t. However, through (3.24)–(3.25) one can derive the following expression:

$$z_t=1/\sigma x_t++\bar{\nu}+\gamma t+\rho(z_{t-1}-1/\sigma x_{t-1})+\nu_t+me_t-\rho me_{t-1}. \tag{3.26}$$

[24]The multiplication by $\frac{n_i}{n_1}$ is done to work with wage income rather than wage rates. Since wage income is directly measured by the CPS, this should attenuate measurement error.

Because the shocks contain effects that may be correlated with the right hand side variables, I pick the parameters in order to match the following $2(I-1)$ moment conditions: $E[x'_{t-2}(\Delta\eta_t)]$, $E[z'_{t-2}(\Delta\eta_t)]$ where $\eta_t = \nu_t + me_t - \rho me_{t-1}$ and Δ stands for first difference.[25] I use those to estimate ρ and σ through the generalized method of moments, for each value of ρ and σ, the vectors $\bar{\nu}$ and γ are estimated with OLS.[26] The estimation above is independent of the complementarity between labour and capital, θ, this parameter is calibrated to $-.25$ in accordance to the literature that suggests a parameter value which induces more collinearity than the Cobb-Douglas case, see for instance Choi and Rios-Rull (2009). The next table summarizes the key estimated parameters. Sensitivity analysis over σ does not

Table 3.5: Estimation Results

	σ	ρ	α	θ
Value	0.85	0.80	0.18	-.25
St. Error	0.48	0.59	-	-

seem to affect the results. Testing whether the over-identifying restrictions hold gives a p-value of .86, this model is therefore not rejected by this test.

Identifying the labour productivity shocks

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

$$(1-\alpha)Ay^{1-\theta}\left(\sum_{i=1}^{I}(z_i n_i)^{\sigma}\right)^{\theta/\sigma} = \frac{w_i\left(\sum_{i=1}^{I}(z_i n_i)^{\sigma}\right)}{z_1^{\sigma} n_1^{\sigma-1}}. \tag{3.27}$$

[25] This solves endogeneity of x, but strictly speaking, measurement error may still be biasing the estimation: the second moment conditions are true under the assumption that me has autocorrelation equal to ρ, otherwise me_{t-1} could influence z_{t-1}. To deal with this I could instrument z_{t-2}.

[26] Alternatively, one could pick σ to match the mentioned moments and use Arellano and Bond to solve (3.24) for $\bar{\nu}$, γ and ρ. Either way, the estimator for σ is consistent.

Since we have identified z_i/z_1 through (3.23), it is convenient to rewrite the expression above as follows

$$(1-\alpha)Ay^{1-\theta}\left(\sum_{i=1}^{I}(z_i n_i)^{\sigma}\right)^{\theta/\sigma} = \frac{w_i\left(\sum_{i=1}^{I}(z_i n_i/z_1)^{\sigma}\right)}{n_1^{\sigma-1}}. \tag{3.28}$$

Substituting this into the production function and solving for A gives

$$A = \left(y^{\theta} - \frac{w_i\left(\sum_{i=1}^{I}(z_i n_i/z_1)^{\sigma}\right)}{n_1^{\sigma-1}}\right)/(\alpha k^{\theta}). \tag{3.29}$$

As a result, z_1 can be backed out from (3.28) and finally, z_i for every i can be derived through (3.23). With this information, it is now possible to run the counterfactual experiment mentioned in the introduction.

3.6.2 Calibration of the other parameters of the model

The calibration strategy is to match data between 1999 and 2007 in steady state. The extent to which the model can also replicate data from 1967 onward provides an interesting challenge for this model which is used as a test for this framework.

Preferences

I choose the following utility specification for a generic household:

$$u(c, l^m, l^f) = \ln(c) - \left(\gamma^m \frac{l^{m1+\sigma^m}}{1+\sigma^m} + \gamma^f \frac{l^{f1+\sigma^f}}{1+\sigma^f}\right). \tag{3.30}$$

These preferences are consistent with a balanced growth path. $1/\sigma^g$ for each gender g, is the constant Frisch elasticity of labour supply. This choice is suggested by the fact that the relative volatility between groups is stable over time. [27] This parameters are calibrated to match average relative volatility between groups. To summarize, mean Frisch elasticity is

[27] It also makes the computation somewhat easier: non separable specifications either between labour and consumption, would make the choice of labour a complex simultaneous system of equations. This way, instead, the interdependence of each labour decision in households' first order conditions is limited to the joint presence of consumption.

.88, which is within the range of micro-estimates. However, to match the fact that total men hours are more volatile than the ones for women, men elasticity has to be lower than the one of women. This parametrization contrasts micro estimates, which suggest that labour supply for women is less elastic than for men. For the aim of this paper, is however necessary to match the relative volatility between men and women hours. The inability to reconcile micro estimates with the relative volatility between men and women hours is an interesting puzzle which may be interesting to study further in future research.

The discount factor β is equal to .99 and the depreciation rate of capital is .07. With these values, and given the parameters of the production function, the average capital output ratio predicted by the model is 2.26, the interest rate 0.04 and the saving rate 0.14.

υ^g, $\bar{k}^g$; the parameters of the cost distribution of acquiring education - equation (3.3)- are set to match the gender specific elasticity of the enrollment rate to the wage premium and the steady state share of highly educated by gender between age 25 and 29 in the period 1999-2007, which is $q^f = 0.36$, $q^m = 0.29$. [28]. Following Heathcote et al. (2010), the degree of assortative matching in the marriage market, in equation , is set equal to .517.

Shocks

It is convenient to decompose the group specific shocks z_i in (3.1) into gender, age and education specific shocks so that

$$\log(z_{t,i}) = \sum_{j=1}^{2} \varepsilon_{t,i}^{g} I_g(i,j) + \sum_{j=1}^{3} \varepsilon_{t,i}^{age} I_{age}(i,j) + \sum_{j=1}^{2} \varepsilon_{t,i}^{edu} I_{edu}(i,j) + \nu_{t,i} \tag{3.31}$$

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

[28]See also Heathcote et al. (2010)

cannot be accounted by a combination of the other shocks. [29], [30] The problem can be written in vectorial form:

$$\log(z_t) = \Lambda\varepsilon_t + \nu_t \tag{3.32}$$

where z_t and ν_t are respectively the vector of I labour specific shocks and residuals at time t, ε_t is the vector of the 7 group specific shocks by gender, age and education. Λ is a $I \times 7$ matrix that collects the group dummies introduced in equation (3.31). ε_t are identified by minimizing the sum of squares of residuals ν_t:

$$Min_{\varepsilon_t} \nu_t' \nu_t. \tag{3.33}$$

This problem can be considered a factor model with factors ε_t and where the factor loading matrix Λ is given.

Figure 3.11 shows how closely $\Lambda\varepsilon_t$ can replicate labour specific shocks z_t. As it is evident from the figure, the two are almost identical and hence in the simulations I will only include $\Lambda\varepsilon_t$ and abstract from ν_t. I assume an AR1 process for ε_t:

$$\varepsilon_t = \gamma_\varepsilon + \rho \times \varepsilon_{t-1} + u_t \tag{3.34}$$

Where ρ and u_t are assumed to be diagonal matrixes.

p_0, the share of new born by gender is simply modeled as an exogenous AR1 process with parameters ρ_{p_0}, σ_{p_0}, the process is estimated over CPS data from the March supplement for the proportion of people of age 20 over people from 20 to 60. This quite simplistic way to model fertility has the purpose of generating variable fertility in order to match changes in the population distribution by age.

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

[30] Being this exercise a mere decomposition, it does not affect the estimation of the complementarity across the labour groups.

Survival probabilities ζ^j for $j = \{1, ..., 40\}$ come from the National Center for Health statistics Vital statistics of the US, 1992. Since nowone can leave for more than 40 periods, $J^{40} = 0$.

Labour trends

The fertility rate time series is such that the model replicates the evolution of population size over time. Initial conditions for the deterministic AR1 processes for women housework l_h, and the cost of acquiring education κ are calibrated so that the model can replicate trends in the labour composition over the last 45 years. To run simulations, as initial conditions for the remaining state variables I take the values that solve the model for a steady state in which the level for the shocks is the average in the first 5 years of the sample (1967-1971).

Figure 3.12 shows actual versus predicted labour shares by education; shares by education are partly driven by the exogenous time series for the cost of acquiring education and partly by the increased wage premium, which comes from labour shocks. Furthermore also the marriage market plays a role: with positive assortative matching, the value of education depends on the probability of matching with an educated partner, this component increases over time as the share of educated increase .

Figure 3.13 shows actual versus predicted labour shares by gender; even in this case the model does quite well at replicating the increase in the share of hours by women. This pattern is mainly driven by the exogenous time series of women housework and partly by the reduction in the gender wage gap.

Hours by age groups is plotted in figure 3.12. It is not to surprising that the model is somewhat successful as the fertility rate time series is such that the model replicates the age distribution observed in the data.

3.7 Computation

The computation of this model presents some challenges that come from the fact that the state space is quite large: 805 variables, of which 324 are state variables. Furthermore, the model has to be simulated over a period of time characterized by transition. For this reason, it is desirable to use a solution method that remains accurate over all the relevant part of the state space. Large dynamic stochastic general equilibrium models can be effectively handled by perturbation methods around the steady state when simulations remain fairly 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 this aim I propose a new methodology which essentially consists of applying repeated local approximations over the entire transition path over the state space, from the initial conditions to the steady state. In practice, since the transition path is unknown, it is first approximated through the policy functions obtained by perturbation around the steady state. Then a new perturbation is conducted in the proximity of the steady state and the new policy functions are used to simulate the transition path again. This step is iterated until the approximation is made in the proximity of the initial condition. This method is detailed in the appendix 3.12.1. In practice, the maximum error with this method is several times smaller than the one with 2nd order perturbation.

3.8 Testing the model

Before carrying out the main experiment of the paper aimed at quantifying the importance of labour composition shifts for aggregate fluctuations, some tests of the model are performed in order to get a sense of how this model provides a satisfactory description of the economy, at least for the dimensions that are relevant for this study.

Being this framework a theory of aggregate volatility and having calibrated the model so that it can match the observed changes in the labour composition, the obvious thing to check is whether the model can replicate the observed trends in aggregate output volatility, as well as the typical statistics analyzed by the real business cycle literature. [31]

The degree of success over this dimensions will help assess how reliable are the outcomes of the counterfactual experiment aimed at quantifying the importance of labour reallocation for aggregate volatility.

3.8.1 Aggregate volatility trend

Table 3.6 contains data and model standard deviations for the whole sample (67-09), the pre-moderation period 67-83 and for the moderation period 84-00. Looking at the standard deviation for the whole sample, it can be inferred that the model accounts for 1.09/1.69 or about 65 % of total volatility. Furthermore, columns 2 and 3 show that the model is quite successful at replicating the volatility slow down between the first and second sub-samples: in the data, volatility reduced by 31.1 log points ($log(2.07) - log(1.51)$), in the counterfactual the reduction is of 29.4 log points. An alternative way to appreciate the extent to which the model can replicate aggregate volatility over time is offered by figure 3.15 which shows the trend over time of aggregate output volatility. [32] This figure shows how the model successfully predicts the initial increase (until around 1974) and subsequent decrease in aggregate volatility. The model under-predicts the volatility rise of the last 10 years, perhaps justifying the effort made by the profession to find theories alternative to the real business cycle to explain the most recent past episodes.

[31] Being able to replicate aggregate volatility changes with a neoclassical framework and with a reasonable parametrizations is not obvious: Arias et al. (2007) show that a Real Business Cycle model with indivisible labour (see Hansen (1985) for details) driven by productivity shocks can successfully account for the decline in cyclical volatility of output.

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

Table 3.6: Standard deviation of output

	67 – 09	67 – 83	84 – 00
Data	1.69	2.07	1.51
Model	1.09	1.38	1.03

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

3.8.2 Aggregate business cycle statistics

Tables 3.7 and 3.8 report respectively standard deviations and correlations with output of consumption, investments and total hours. [33] Consistently with the data, the model predicts that while consumption is less volatile than output, investments are way more volatile. The model under-predicts the volatility of hours. The model matches correlations rather well.

Table 3.7: Standard deviation relative to output

	Data	Model
Output	1	1
Consumption	0.82	0.68
Investment	4.6	3.6
Hours	0.93	0.38

Notes: Statistics for Consumption and investment under the column -Data- are computed using NIPA data.

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).

[33]Statistics are computed after having HP-filtered the data with parameter 6.28.

Table 3.8: Correlation with output

	Data	Model
Output	1	1
Consumption	0.90	0.86
Investment	0.94	0.86
Hours	0.85	0.83

3.9 Counterfactual Experiment

In this model, long run trends in hours by gender, age and education come partly from the dynamics of the wage gap by gender, age and education, and partly through the exogenous long run trends in the amount of women housework, the fertility rate dynamics and the cost of acquiring education. Counterfactual experiments consist of removing these latter exogenous trends, while maintaining all the shocks as in the original simulation. This simulation generates some counterfactual time-series in which the amount of labour reallocation across groups is greatly curtailed. [34] It is then possible to see how total output volatility evolves over time in comparison with the original simulation. This very exercise is carried out in the next subsection 3.9.1. Subsequently, in subsection 3.9.2, these long-run trends are removed one by one in an attempt to asses the relative importance of the labour reallocation respectively by age, education and gender.

3.9.1 Removing all trends

In this first counterfactual experiment, fertility, educational costs and women housework are kept at their steady state levels. Figures 3.16, 3.17, and 3.18 show hours shares by

[34] The part of these trends that comes from changes in wage gaps could be removed by changing the productivity processes. This would however make the comparison with the original simulation less transparent. The fact that some of the long run trends in labour shares remain in the counterfactual experiments make the quantitative results conservative: being able to remove the remaining trends would foster the results found in this paper.

gender, age and education in the counterfactual and original simulations (in the original simulation, trends in the amount of women housework, the fertility rate dynamics and the cost of acquiring education, are obviously included). [35] As it can be seen from comparing with the original simulation, most of the trend in these shares has been removed. Table 3.9 contains the standard deviation of output during the sub-samples of interest: the period before the great moderation 1967-1983, and the period of the great moderation: 1984-2000. For completeness, the table also reports the initial period of high turbulence, 1967-1977, and the last period of the sample: 2001-2009.

Table 3.9: Standard deviation of output over time

	$67-77$	$67-83$	$84-00$	$01-09$
Original simulation	1.48	1.39	0.93	0.53
Counterfactual simulation	1.41	1.33	0.99	0.54

Notes: numbers represent the standard deviation of the percent distance to the HP filter trend with HP parameter 6.28.

As one can see from the table, volatility would have been about 5% lower during the turbulent initial years of the sample and reduced much less in the counterfactual than in the original simulation. The lower volatility in the earlier years depend upon the fact that labour is distributed in favor of more stable groups in the counterfactual. [36] The lack of trends in these labour shares explains the smaller drop in output volatility over time. The fact that volatility in the original and counterfactual simulations converges as one moves to the late part of the sample depends upon the fact that the shares in the counterfactual are very similar to the ones of the latest part of the sample.

[35] Consistently with the original simulation, initial conditions are computed by solving the model at a steady state where the shocks have the values observed at the beginning of the sample.

[36] This fact highlights how the labour composition not only explains the volatility slowdown over time, but is also responsible for part of the very high volatility of the early seventies.

To get a visual sense of how this mechanism is affecting volatility over time, figure 3.19 shows the trend of actual and counterfactual cyclical volatility measured as a 3 years roll over standard deviation of output.[37] As can be seen, counterfactual volatility would have been lower in the early seventies, and higher in the 90s. To get a syntectic statistic that quantifies the amount of the great moderation explained, we proceed as in section 3.8.1, comparing the standard deviations in the first sub-sample (1968-1983) and in the second one (1984-2000). Between the two sub-samples, aggregate output volatility decreased by 40 log points. Had the shares remained stable as in the counterfactual, we would have observed a reduction in volatility of 29 log points. Therefore, changes by gender account for (40-29)/40 or approximatively 26% of the moderation in output.

Last, figure 3.20 shows how labour reallocation played an important role for output levels: in the counterfactual, not only the economy is more stable, but output levels are much higher in the early part of the sample: this depends on the fact that, in the counterfactual experiment, the labour distribution is roughly constant at the end of sample levels, with a higher level of education and prime age workers.[38]

3.9.2 Removing trends one by one

Table 3.10 summarizes the various statistics as in table 3.9 in the previous sub-section. Each line represents a counterfactual experiment where only one of the long run trends are removed and substituted with their steady state values.

Table 3.11 column 1, reports the various contributions to the moderation and to the high volatility in the early seventies. Age trends seem to be the most important, although the other two are non negligible. Columns 2 to 4 report the evolution of output volatility;

[37] The point estimate, say, in 1980, is the standard deviation of the relative deviation of output from HP-trend between 1979 and 1981.

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

Table 3.10: Standard deviation of output over time

	$67-77$	$67-83$	$84-00$	$01-09$
Original simulation	1.48	1.39	0.93	0.53
Age	1.46	1.36	0.99	0.52
Education	1.48	1.39	1.01	0.54
Gender	1.44	1.35	0.95	0.54

it is interesting to notice how removing trends by age would have implied lower output volatility in the period 01-09. This is because in the counterfactual, the share of old does not increase and the share of prime age does not decrease (see figure 3.17).

These results help to get a sense of the relative importance of each component, however they have to be handled with caution. The fact that summing up the contribution of each group does not obtain the result in the previous subsection (reported in the first row of the table), shows how the contribution of each group cannot be disentangled from the others. This is because changes in the labour composition by one group may have implications for the other groups. An interesting case is the one where educational costs are removed, figure 3.21 shows the implications for hours shares by the 3 groups: the gender gap at the beginning of the sample would have been even wider. This is due to the presence of the housework duty for women, which induces men to acquire more education and to work more hours.

Table 3.11: Standard Deviations

	Great Moderation	Counterf.-Actual St.Dv. 69-76	Counterf.-Actual St.Dv. 84-00	Counterf.-Actual St.Dv. 01-09
All	26.3	-4.7	6.5	2.1
Gender	10.7	-2.9	1.6	-0.3
Age	19.6	-1.5	6.2	-2.4
Education	18.6	-0.2	8.0	3.3

Notes: numbers are expressed in percent 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-Actual St.Dv. measures the percentage difference between output and counterfactual output volatility

3.10 Extensions

This study has focused on the business cycle implications of demographic changes by gender, age and education. It would be interesting to investigate wether other partitions of the labour force may show the two characteristics - compositional changes and stable volatility differences. Changes by sector for instance, may have aggregate volatility implications. It is well known that the service sector has expanded over time: figure 3.8 plots per capita hours worked by sector: durables, non durables, services and public sector. The share of employment over population in the service sector increased relative to the other sectors, moving from an average of 21 % in the first sub-sample, to 29 % in the second one. As reported in table 3.12, third column, the cyclical volatility of service hours is lower than for the other groups. Figure 3.9 plots the relative volatility between hours in the service sector and manufacturing; the relative volatility remains fairly stable over the entire sample and significantly below one.

The last column of the table below shows how much of the moderation occurred within groups; this is the part of the moderation that cannot be accounted for simply by the reallocation of workers across these four groups, but it may be explained by the interactions with the other considered grouping variables.

Table 3.12: Hours Volatility Sector

	workers'share 68-84	workers'share 84-07	st. dev	$\frac{st.dev(84-07)}{st.dev(68-84)}$
Manufacturing	12.14	10.24	3.12	76.52
Services	20.61	29.07	1.10	65.33
Public Sect	2.03	2.38	2.79	48.75

Notes: numbers are expressed in percent terms. Workers' share is the ratio of workers over total population. St.deviation measures the standard deviation of the percent change of HP filtered hours with HP parameter 6.28

Figure 3.28 show realized and counterfactual employment by sector. Figures 3.23, 3.25, 3.27 and 3.29 show observed and counterfactual aggregate hours by sector.

Following Jaimovich and Siu (2009), comparing the standard deviations of the growth rates in the first sub-sample (1968-1983) and in the second one (1984-2007), one gets a measure of the moderation in aggregate hours: changes by sector accounts for 21% of the moderation in aggregate hours.

From this simple exercise one may hazard that sectorial changes are the most important contributors to the great moderation, followed by changes in education, age and gender composition. Some words of caution are however necessary to state that this numbers do not quantify the relative importance of each demographic change considered. Suppose, as it is the case, that women spend more hours working in the service sector than in other sectors. An increase in women participation rate would then induce an increase in the share of hours worked in the service sector. Changes by sector would then be a consequence of changes by gender and the exercise above would understate the importance of gender composition for aggregate volatility. This said, the result above seems to be sufficient to conjecture that compositional changes by sector may also be important on top of the ones already considered. [39]

3.11 Conclusion

This paper documents that while the composition of the labour force by gender, age and education has changed substantially over the last 40 years, the relative volatility within these groups has remained remarkably stable. These facts lead to the conjecture of this study: that changes in the composition of the labour force have a causal impact on the evolution of aggregate volatility over time.

The endogenous interaction between these demographic changes and the possible presence of omitted variables and reverse causality between these groups and aggregate volatility make accounting exercises and regression analysis not reliable methods and motivate a structural approach. To this aim, a general equilibrium model of the business cycle with overlapping generations, educational and marriage choices is developed. The model is able to replicate at the same time the observed demographic and aggregate volatility changes and therefore, is a useful framework to quantify the role of demographic changes on aggregate volatility. This is done by running counterfactual experiments aimed at removing the demographic changes, while maintaining the shocks that generate the business cycle.

If the labour force composition remained stable over time at the levels of the most recent years, volatility would have been 5 % lower than what was observed in the early 1970s and 6.5% more volatile in the late 80s and 90s. Therefore, the exercise shows how a part of the

[39] Stock and Watson (2005) found that sector reallocation accounts for 8% of the moderation in the U.S. They considered total production by sector and shifted average sector shares to what they were on average in the first sub-sample, and then applied to these averages the growth rates observed in the second sub-sample. Simply maintaining the observed growth rates implies that the service sector, by the end of the sample, gains much of the share that was artificially removed, thereby understating the average share redistribution. Furthermore, this exercise does not take into account the endogenous price changes that follow such redistributions.

high volatility of the early seventies and of the slowdown in the 1980s is accounted for by the considered changes in the composition of labour.

The methodology adopted in this paper could be applied to other partitions of the labour force, for instance, hours in the service sector are less volatile than in other sectors, and the sharp and steady increase in this sector may have implications for aggregate volatility.

The presence of a causal relationship between the composition of the labour force and the business cycle is relevant for policies that affect gender inequality, educational attainments and social security (which affects the labour distribution by age). More generally, this paper suggests that micro an macro policies can be related, and while it is often useful to study them separately to maintain tractability, this my come at a cost. Furthermore, this causal relationship may help improving the predictability of aggregate volatility, this is especially true given the fact that trends in the composition of labour appear to be quite predictable.[40] For instance, with the ageing of the baby boom generation, the share of prime age workers is bound to decline over time. Other things equal, this will increase aggregate volatility.

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

To summarize, this paper introduces some facts on the relative volatility of hours by demographic groups, develops a methodology that shows how these facts can help relating the labour force composition to aggregate volatility, introduces a computational technique able to solve large scale models with accuracy over a wide transition path.

[40] At present, it becomes policy relevant to predict future aggregate volatility: how much should we increase capital requirements for banks depends on the risk of future crisis.

3.12 Appendix

3.12.1 Computational Algorithm

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

$$E_t[f(x_{t+1}, y_{t+1}, x_t, y_t)] = 0 \tag{3.35}$$

where E_t is the expectation operator given information at time t, x_t is a vector of state variables sorted so that all shocks enter at last, y_t is a vector containing all other variables of the model. [41] Solutions to equation 3.35 which satisfy transversality conditions are

$$y_t = g(x_t, \sigma) \tag{3.39}$$

and

$$x_{t+1} = h(x_t, \sigma) + \sigma\eta_{t+1}. \tag{3.40}$$

where, following Schmitt-Grohe´ and Uribe (2004), σ is a parameter that scales the variance of $\eta_t = [0, u_t]$, where u_t is a vector containing the shocks'innovations. An approximated solution can be found by Taylor expanding equation 3.35 around the deterministic steady state where $x_t = x_{t+1} = \bar{x}$, $y_t = y_{t+1} = \bar{y}$ such that

$$f(x_{t+1}, y_{t+1}, x_t, y_t) = 0. \tag{3.41}$$

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

$$MaxE_0 \sum_{t=0}^{\infty} \beta^t log(c_t)$$

subject to feasibility

$$k_{t+1} + c_t = k_t(1-\delta) + e^{A_t} k_t^{\theta} \tag{3.36}$$

and to the productivity process

$$A_t = \rho A_{t-1} + \sigma u_t, u_t \sim d(0, var_u). \tag{3.37}$$

The equilibrium conditions are the last two equations 3.36-3.37 and

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}} (1 + A_{t+1} \alpha k_{t+1}^{\alpha-1} - \delta) \right). \tag{3.38}$$

With $x_t = [k_t, A_t]$ and $y_t = c_t$, the three equilibrium conditions 3.36-3.38 are easily casted into equation 3.35.

This can be done, for instance, by applying the algorithm of Klein and Gomme (2011). Taylor expansions of equation 3.41 are done around the steady state. This is because the steady state is a point which is typically easy to find and where equation 3.41 holds and equations 3.39-3.40 hold with $\sigma = 0$. [42] If there is a point $\widehat{x}$ in the state space for x where one knows the values $\widehat{x_1} = g(\widehat{x}, 0)$, $\widehat{y} = h(\widehat{x}, 0)$ and $\widehat{y_1} = h(\widehat{x_1}, 0)$ so that

$$f\left(\widehat{x_1}, \widehat{y_1}, \widehat{x}, \widehat{y}\right) = 0,$$

one could take the Taylor expansion there and have a solution well approximated around that point. [43]

The algorithm that I am going to introduce seeks to find points outside the steady state and on the transition path from a given initial condition x_0 and for a given sequence of shocks u_t, with the aforementioned characteristics (that satisfy equation 3.41 as well as equations 3.39-3.40 with $\sigma = 0$). This is done by backward induction from the steady state.

Call

$$F(x, \sigma, h, g) \equiv E_x\left(f[h(x, \sigma) + \sigma\eta, g(h(x, \sigma) + \sigma\eta, \sigma), x, g(x, \sigma)]\right), \quad (3.42)$$

Where E_x is the expectation over η given the state variables x. Pick $\hat{\varepsilon} > 0$ small. The following algorithm aims at providing policies h_x, g_x with precision

$$|F(x, 0, h_x, g_x)| < \hat{\varepsilon}.^{44}$$

for any x on the equilibrium path.

1. Generate a sequence of shocks for T periods, which starts with the given sequence of shocks u_t and that converges to the steady state of the shocks' processes.

[42] A point x_1,y_1,x,y that satisfies equation 3.41 but not 3.39-3.40, is on a path that violates transversality conditions.

[43] The solution is not perfect at $\widehat{x}$ because a solution of equation 3.41 doesn't exactly solve 3.35 because of the Jensen's inequality. This is independent of whether $\widehat{x}$ is the steady state or not.

[44] Note that $\sigma = 0$, that is because the expectation operator in equation 3.42 is replaced by the assumption of zero innovations.

2. Taylor expand $f(x_{ss}, g(x_{ss}), x_{ss}, g(x_{ss}))$ where $x_{ss}, g(x_{ss}, 0)$ is the deterministic steady state of the model, and obtain the the policy functions $h_{ss}(x, \sigma)$, $g_{ss}(x, \sigma)$. If those are stable, then go to step 3 (for stability see for instance Blanchard and Kahn (1980)). [45]

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

4. Simulate from x_0 with these policy functions and with $\sigma = 1$ generating time series $\{x_t\}_0^T$, $\{y_t\}_0^T$. If the model does not converge to the steady state, increase T and go back to step 1.

5. Put $t = T$.

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

7. if $|F(\hat{x}, 0, \widetilde{h}, \widetilde{g})| > 0$, find $\widehat{x_1}$ such that $f(\widehat{x_1}, \widehat{y_1}, \widehat{x}, \widehat{y}) = 0$ where $\widehat{y}_1 = \widetilde{g}(\widehat{x}_1, 0)$.

8. Derive the functions $h_{t-1,\alpha}(\cdot)$, $g_{t-1,\alpha}(\cdot)$ Taylor expanding $f(x_1, y_1, x, y)$ around the latter point $(\widehat{x_1}, \widehat{y_1}, \widehat{x}, \widehat{y})$.

9. Put $\widetilde{h}(\cdot) = h_{t-1,\alpha}(\cdot)$ and $\widetilde{g}(\cdot) = g_{t-1,\alpha}(\cdot)$. if $\alpha < 1$ increase it to a number smaller or equal to 1 and such that $|F(\hat{x}, 0, \widetilde{h}, \widetilde{g})| < \hat{\varepsilon}$ where $\hat{x}$ has been updated accordingly: $\hat{x} = \alpha x(t-1) + (1-\alpha)x(t)$. Go back to step 10. If $\alpha = 1$, and $t > 2$ store $h_{x(t-1)}(\cdot) = h_{t-1,\alpha}(\cdot)$, put $t = t-1$ and go back to step 6. If $\alpha = 1$, and $t = 2$, $h_{x(t-1)}(\cdot) = h_{t-1,\alpha}(\cdot)$ and go to the next step.

10. Simulate the model from x_0 and with the given sequence of shocks u_t, using at each point $t = 0, ..., T$ the policies $g_{x(j)}(\cdot, 1)$, $h_{x(j)}(\cdot, 1)$ with the smallest $|x_t - x_j|$.

[45]This algorithm is described for stable models, it might be possible to extend it for models that are locally unstable in some regions of the state space. In fact, it could be extended to models that do not have a steady state provided that a point $(\widehat{x_1}, \widehat{y_1}, \widehat{x}, \widehat{y})$ such that $f(\widehat{x_1}, \widehat{y_1}, \widehat{x}, \widehat{y}) = 0$ is known.

11. Iterate from point 5 until the time series $\{x_t, y_t\}_0^T$ do not coincide with the ones of the previous iteration. [46]

Variations of this algorithm can be conceived; for instance, to increase speed one could avoid going backward through all the points on the equilibrium path, but make larger jumps from the steady state until x_0. Furthermore, the researcher is free to choose the degree of the Taylor expansion at each point, but there is not much gain from orders higher than 1; this is because the Taylor expansion is always very close to the point of interest. I hence used first order Taylor expansions.

To evaluate the accuracy of this algorithm, I test it on the model in the note at the beginning of the section and with full depreciation, for which the analytical solution to the equilibrium conditions, equations 3.36-3.38, is known. I then compare the true equilibrium path $\{x_t^*, y_t^*\}_0^T$ with the one generated through this algorithm, $\{x_t^{**}, y_t^{**}\}_0^T$ and with the one generated by a second order expansion around the steady state $\{x_t^{***}, y_t^{***}\}_0^T$. For an initial condition quite far from the steady state, $x_0 = [.2k_{ss}, -.5]$, with variance of the shock equal to 0.007 [47] the maximum error

$$Max_t[Max\left(|x_t^* - x_t^{***}|, |y_t^* - y_t^{***}|\right)] \tag{3.43}$$

using second order approximation around the steady state is 0.0081. Using the proposed algorithm, the maximum error

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

[46] Although I don't have a proof that the algorithm converges to a time series $\{x_t, y_t\}_0^T$, this has been the case for any model I tried. It typically takes less than 5 iterations. The iteration procedure over the equilibrium path is reminiscent of the Parametrized Expectation Approach (see Den Haan and Marcet (1990) and Marcet and Lorenzoni (1999)): both algorithm break the curse of dimensionality by only approximating the global policy function over the equilibrium path rather than over the entire state space. In practice however, the Parametrized Expectation Approach may show some convergence problems that make its implementation hard, especially for high-dimensional applications. See on this respect the improvements made by Judd et al. (2009) (typically, this approach it is also less accurate because it interpolates across the points).

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

is 6.4568e-004, which is 12.5 times smaller than taking the expansion only around the steady state. Although this result seems impressive, in some sense this exercise understates the improvement, this is because both methods make the same steady state error due to the assumption of certainty equivalence and this underscores the improvement over the transition path, far from the steady state. To isolate the error over the transition from the one due to the Jensen's inequality, I run the experiment for the deterministic case, for which there is no error at the steady state and hence the only error would be due to a bad approximation when far from the steady state. Using the proposed method, the maximum error is 3.2116e-007. That is 2.5121e+004 times smaller than the one made by second order perturbation around the steady state. I conclude that this algorithm makes a terrific improvement respect to perturbation around the steady state.

In this example the code takes a few seconds to run on a laptop. Solving the main model of the paper in section 3.5 takes more time, I therefore avoid going backward between the points on the equilibrium path, but make larger jumps (Taylor expand every point and maintain α equal to one) from the steady state until x_0. With these larger jumps, the model can be solved in a laptop in 37 minutes and the approximated solution is 2.15 times more accurate than local perturbation around the steady state in the sense of equation 3.43.

3.13 Figures

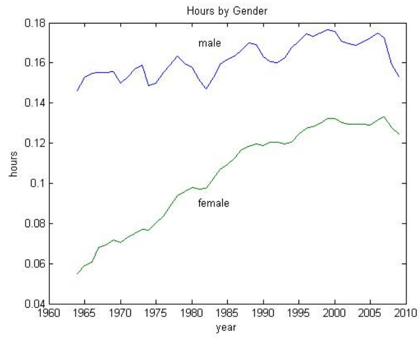

Figure 3.1: Hours by gender

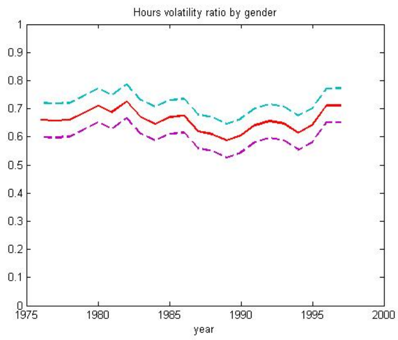

Figure 3.2: Hours volatility ratio by gender.

Notes: In each period t, the figure plots the ratio of the standard deviation of hours over a period of 18 years centered at year t. Confidence intervals are calculated assuming that the time series follows an AR 1 process.

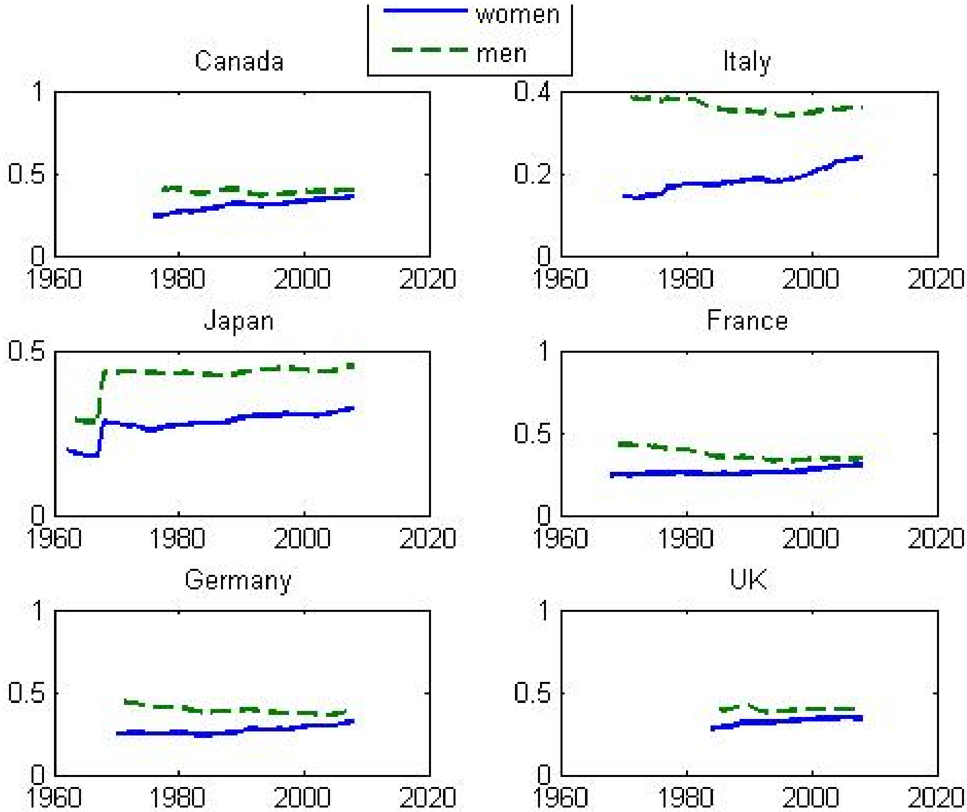

Figure 3.3: Employment over population by gender for remaining G7

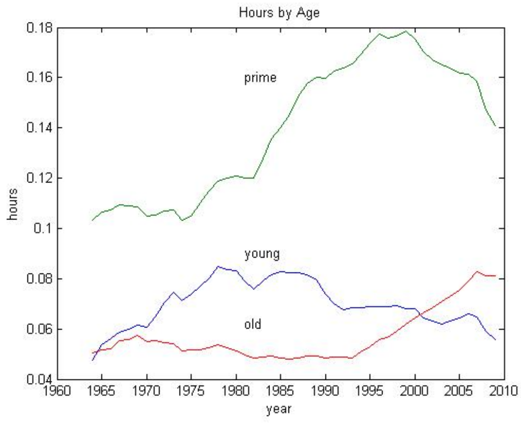

Figure 3.4: Hours by age.

Notes: young workers range from 15-29 years old. Prime age ranges from 30-50. Old are the ones above 50.

Figure 3.5: Hours volatility ratio between prime age and young.

Notes: In each period t, the figure plots the ratio of the standard deviation of hours over a period of 18 years centered at year t.

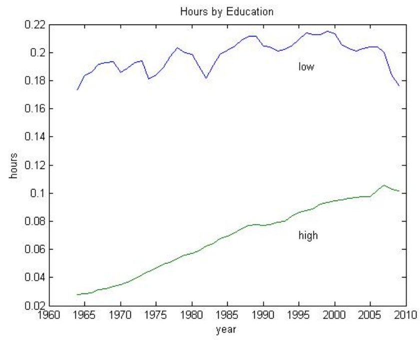

Figure 3.6: Hours by education.

Notes: By high education it is meant at least four years of college.

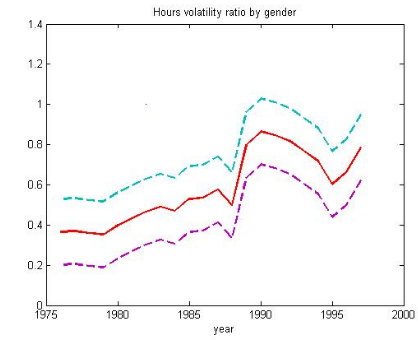

Figure 3.7: Hours volatility ratio by education groups.

Notes: a. In each period t, the figure plots the ratio of the standard deviation of hours over a period of 18 years centered at year t.

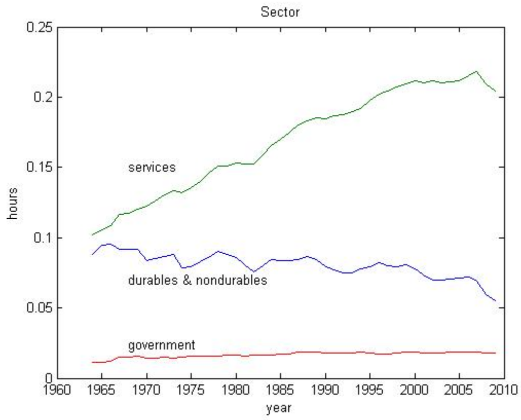

Figure 3.8: Hours by sector

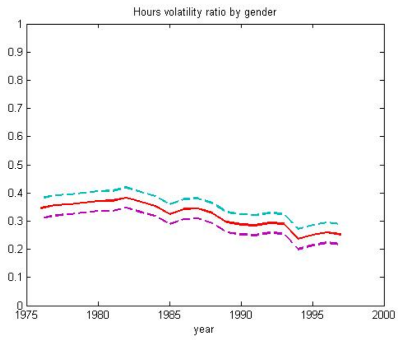

Figure 3.9: Hours volatility ratio between services and manufacturing.

Notes: In each period t, the figure plots the ratio of the standard deviation of hours over a period of 18 years centered at year t.

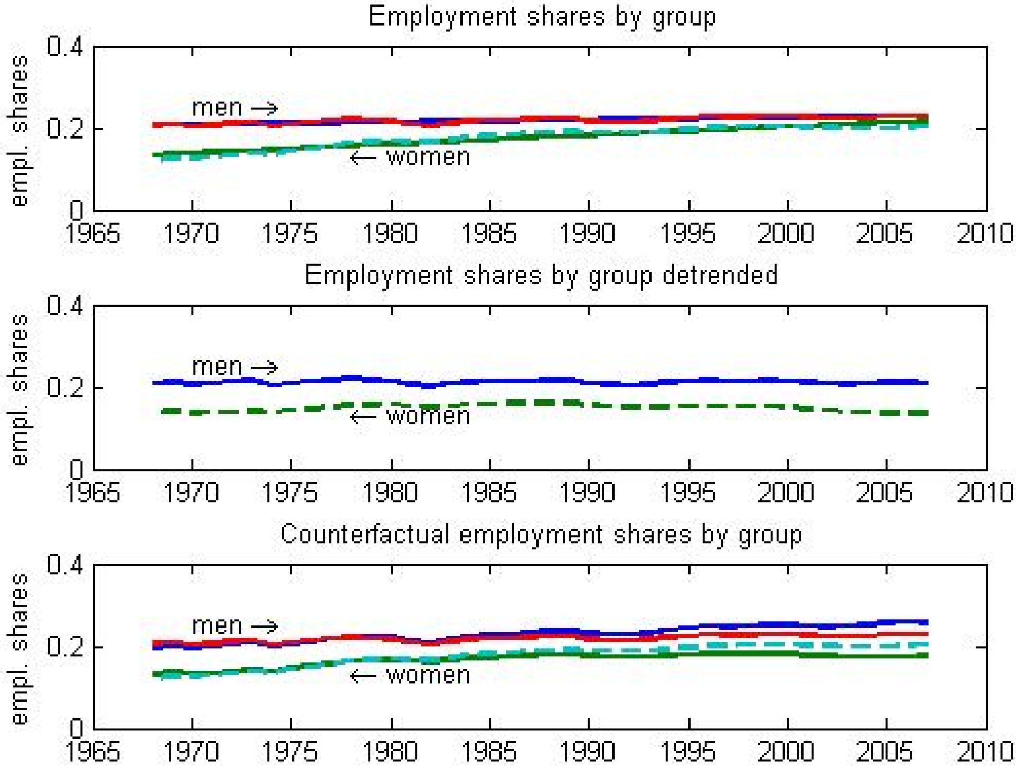

Figure 3.10: Construction of counterfactual employment shares by gender

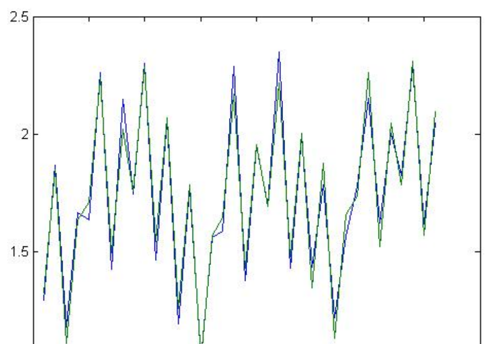

Figure 3.11: Group factors versus labour specific shocks Z in year 2009.

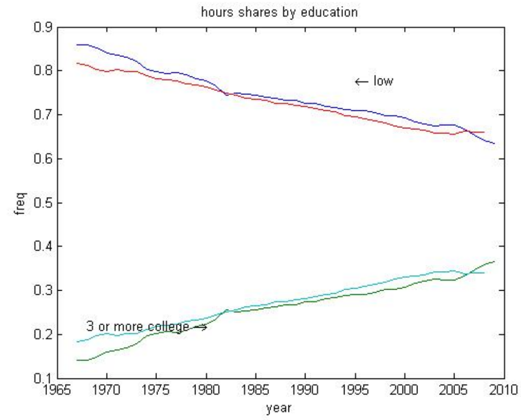

Figure 3.12: Data vs model labour shares by education

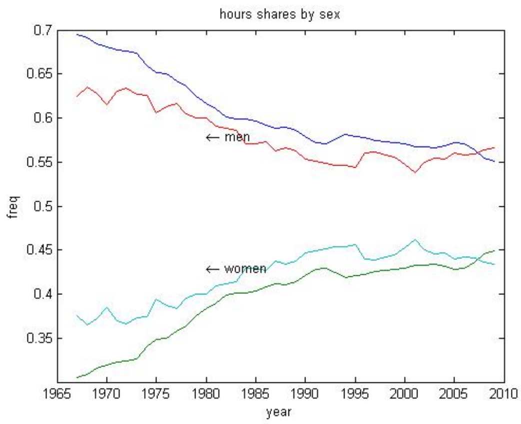

Figure 3.13: Data vs model labour shares for by gender

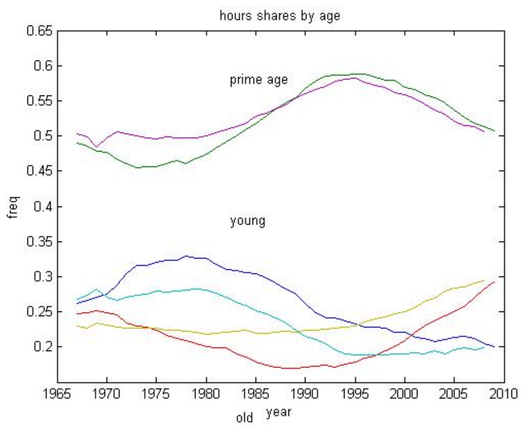

Figure 3.14: Data vs model labour shares for by age groups

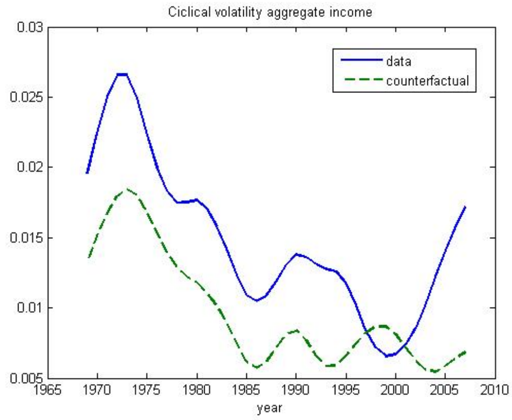

Figure 3.15: Output volatility over time, data versus simulation

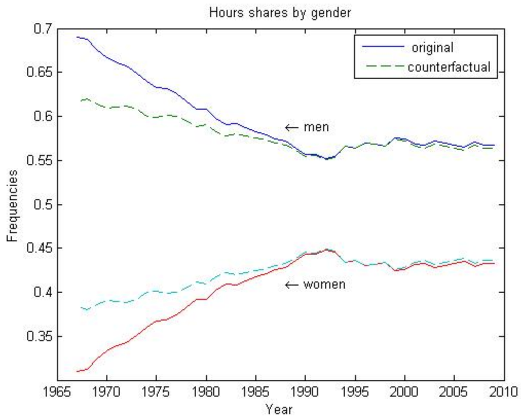

Figure 3.16: Hours shares by gender, original versus counterfactual simulation

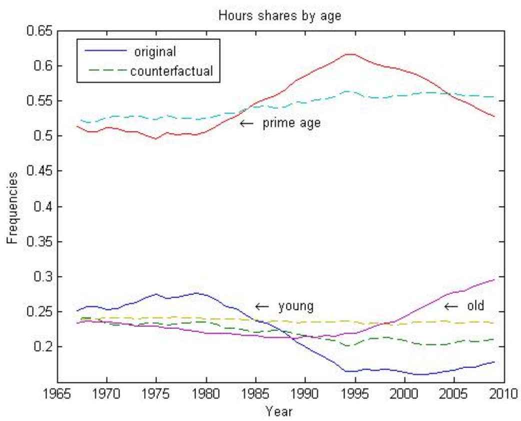

Figure 3.17: Hours shares by age, original versus counterfactual simulation

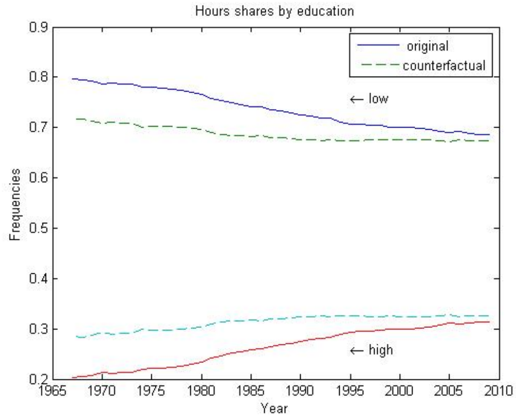

Figure 3.18: Hours shares by education, original versus counterfactual simulation

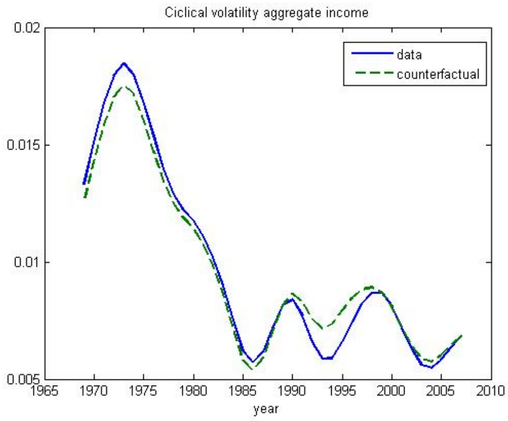

Figure 3.19: Output volatility over time, original versus counterfactual simulation

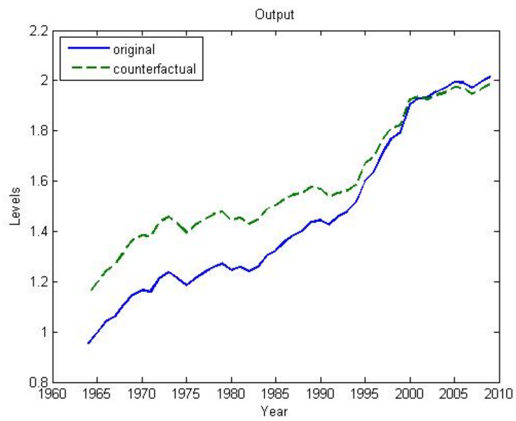

Figure 3.20: Output, original versus counterfactual simulation

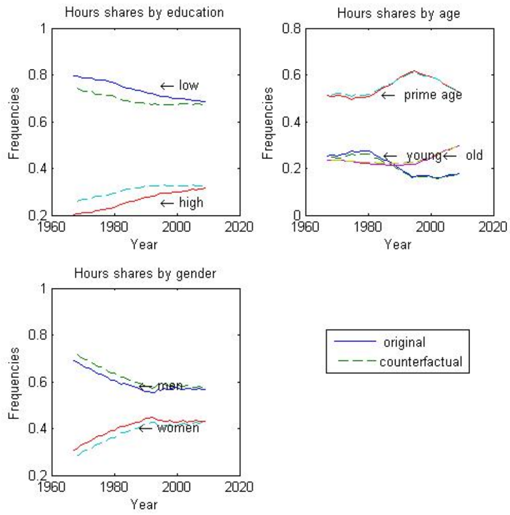

Figure 3.21: Hours shares, original versus counterfactual simulation without educational costs but maintaining female homework and population trends.

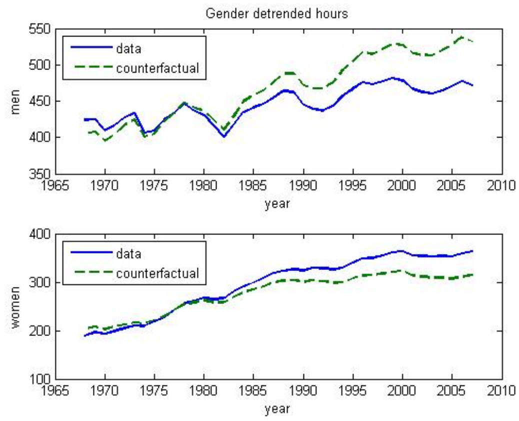

Figure 3.22: Actual vs counterfactual hour without trend in gender

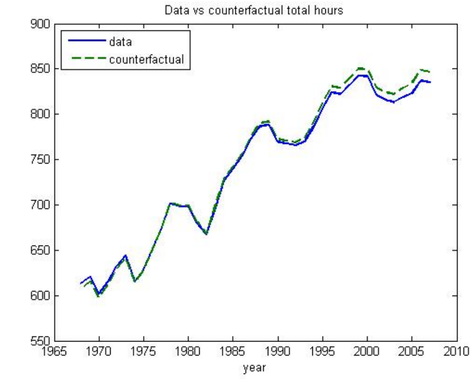

Figure 3.23: Actual vs. counterfactual aggregate hours without trend in gender

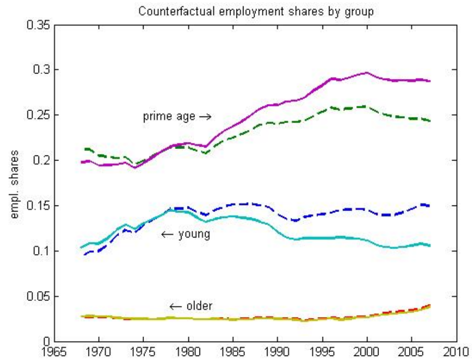

Figure 3.24: Actual vs counterfactual employment without trend in age.

Notes: Counterfactuals are in dashed line.

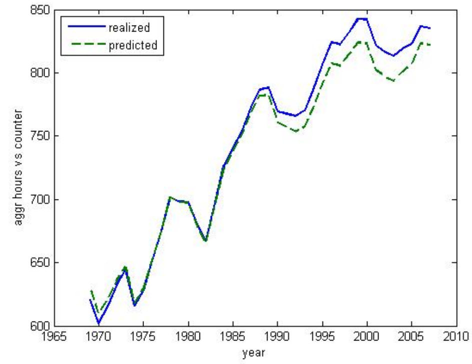

Figure 3.25: Actual vs.counterfactual aggregate hours without trend in age

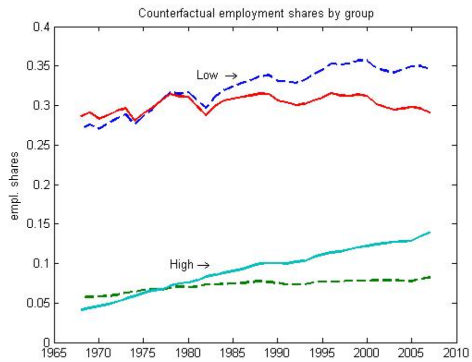

Figure 3.26: Actual vs counterfactual employment without trend in education.

Notes: Counterfactuals are in dashed line.

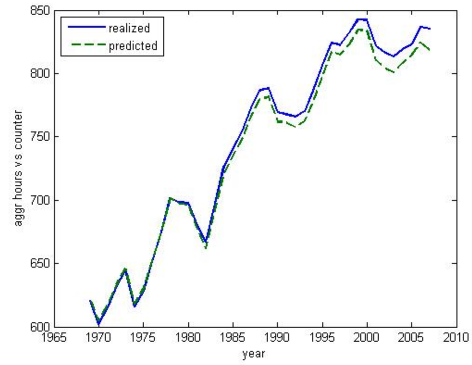

Figure 3.27: Actual vs.counterfactual aggr. hours without trend in education

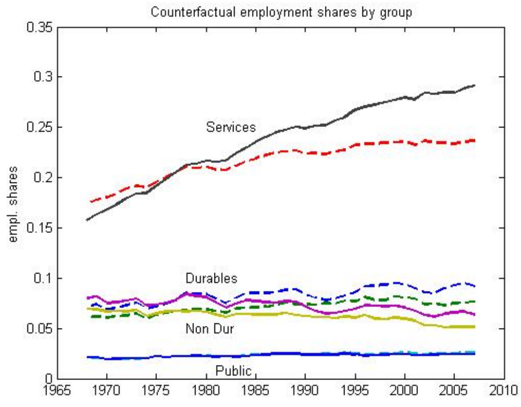

Figure 3.28: Actual vs counterfactual employment without trend in sector

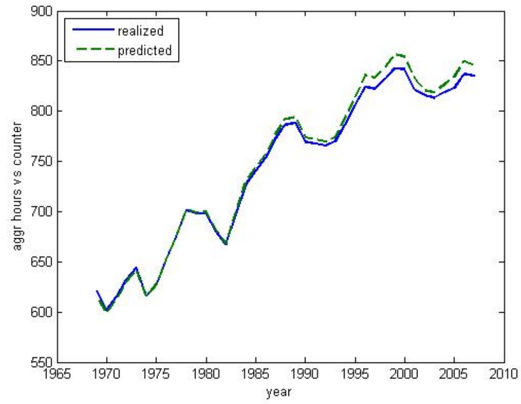

Figure 3.29: Actual vs.counterfactual aggr. hours without trend in sector

Chapter 4

Optimal Fiscal Policy in the Neoclassical Growth Model Revisited

Abstract

This paper studies optimal Ramsey taxation in a version of the neoclassical growth model in which investment becomes productive within the period, thereby making the supply of capital elastic in the short run. Because taxing capital is distortionary in the short run, the government's ability/desire to raise revenues through capital income taxation in the initial period or when the economy is hit with a bad shock is greatly curtailed. Our timing assumption also leads to a tractable Ramsey problem without state-contingent debt, which gives rise to debt-financed budget deficits during recessions.

4.1 Introduction

This paper studies optimal fiscal policy in a version of the neoclassical growth model in which capital is elastically supplied even in the short run. This is accomplished by letting investment in capital become productive within the period.

It is well understood that the conventional timing in the neoclassical growth model, in which the size of the capital stock today is the result of past investment decisions, implies that capital is inelastically supplied in the short run. It should be equally clear that a by-product of this conventional timing assumption is at the heart of many well-established results within the optimal taxation literature. A prominent example is the well-known prescription to tax initial holdings of assets at confiscatory rates, a result that Chamley (1986) and much of the subsequent literature tries to circumvent by imposing bounds on tax rates: without these exogenous bounds, a first-best allocation obtains, an obviously uninteresting problem. Tax rates over the business cycle are similarly dictated by the conventional timing of the neoclassical growth model. Every period, the government promises not to distort the return to investment while at the same time announcing that recessions will be financed with unusually high taxes on capital income, and *vice versa* during booms. This strategy is clearly optimal as the government can avoid distorting investment decisions *ex ante* while at the same time having the ability to exploit the fact that since the stock of capital is fixed *ex post*, taxing/subsidizing its return represents a non-distortionary way to absorb shocks.

This paper shows that changing the timing of events in the neoclassical growth model in such a way as to make the supply of capital elastic in the short run drastically alters the prescription that emanates from standard Ramsey problems. Our assumption that investment in capital becomes productive within the period gives individuals an alternative to supplying capital, namely consuming, which is not present under the conventional timing.

Knowing that this alternative exists limits the ability and desire of the government to use capital income taxes to finance government expenditures, either in the initial period or over the business cycle.

One of our main results, as eluded to above, is that the solution to the Ramsey problem features a unique non-trivial level of distortions. While the level of distortions depends on individuals' initial holdings of debt and capital, it does not rely on the presence of bounds exogenously imposed on the Ramsey problem. As such, the trivial result that the solution to the Ramsey problem without imposing exogenous bounds is time-consistent does not hold in our environment.[1]

Next we offer a complete characterization of the behavior of tax rates in a stochastic environment in which the government has access to state-contingent debt. Under a class of utility function in which consumption and leisure are separable, we show that neither the labor nor the capital income tax varies over time, and that the tax on capital is zero in all but the initial period. Under Cobb-Douglas utility, both tax rates become pro-cyclical, that is, they are low during recessions. In either case, the government uses state-contingent government debt as a shock absorber, much like the *ex post* capital income tax is used for that purpose in Chari et al. (1994). As a result, debt and the primary deficit move in opposite directions, a result which Marcet and Scott (2009) showed was pervasive in models in which the government has access to state contingent government debt. This leads us to study a Ramsey problem under incomplete markets.

The Ramsey problem without state-contingent debt is quite tractable in our framework. Technically, this tractability emanates from the fact that our first order conditions can be expressed in terms of prices as a function of quantities. This allows us to write down a

[1] The conventional solution entails taxing the initial return on capital at confiscatory rates, and to finance all future government expenditures through the return on that capital. This solution turns out to be highly distortionary under our environment. The contrast in results across the two environments is reminiscent of the Lucas (1980) vs Svensson (1985) timing issue in cash-in-advance models, as shown in Nicolini (1998).

version of the Ramsey problem, known as the primal, in which the the government chooses quantities subject to a sequence of *implementability* constraints which can be studied using techniques developed in Marcet and Marimon (1995). The upshot of this problem is that in this environment the government runs debt-financed primary deficits during recessions.

Our work complements that of Farhi (2005), who uses the conventional timing but imposes that the government sets capital income tax rates one period ahead in order to mitigate the free lunch associated with volatile *ex post* capital income tax rates. In our model, the ability of agents to build and use capital within the period, makes capital elastic enough to obtain less volatile capital taxes compared to Farhi (2005). In fact, unlike Farhi (2005), the prescription of our model is to finance recessions with deficit, in line with the empirical findings of Marcet and Scott (2009). In Farhi (2005), although taxes on capital are set in advance, they are not distortive enough to prevent the government to run surpluses during recessions. To the same end, Scott (2007) and Marcet and Scott (2009) rule out capital income taxes altogether and show that the implications of their model without state contingent debt is more consistent with the data than models with state contingent debt. In particular, with incomplete markets government debt and the labor tax rate inherit a unit root component which, as emphasized by Aiyagari et al. (2002) in a model without capital, lends some support to Barro (1990)'s conjecture. Qualitatively, our simulations confirm that these results hold even when the government sets capital tax rates optimally.

Before moving to the description of our economic environment, our central assumption that investment becomes productive within the period deserves some comments. First, we show in the appendix that this assumption can be viewed as the opposite from the equally extreme conventional assumption that today's investment only becomes productive in the next period. Second, we view this assumption more as a way to introduce some elasticity to the supply of capital rather than a way of improving the realism of the neoclassical growth model. There are countless issues for which the conventional timing assumption is either desirable or, at least, innocuous.[2] Optimal taxation is just not one of those issues.

The rest of the paper is organized as follows. The next section presents our general economic environment, which consists of the neoclassical growth model with an alternative assumption. In Sections 4.3 and 4.4 we set up and analyze a deterministic and a stochastic Ramsey problem, respectively. Section 4.5 is devoted to the analysis of a Ramsey problem without state contingent debt. A brief conclusion is offered in Section 4.6.

4.2 General Economic Environment

The economic environment we consider is similar to that of Chari et al. (1994): a one-sector stochastic neoclassical growth model. As emphasized in the introduction, the main distinguishing feature of our model is that current investment in capital becomes productive immediately. In this section, we introduce the general economic environment. We later study special cases of this environment, starting with a deterministic version, followed by stochastic versions with and without state-contingent government debt.

Time is discrete and lasts forever. Each period the economy experiences one of finitely many events $s_t \in S$. We denote histories of events by $s^t = (s_0, s_1, \ldots, s_t)$, where s_0 is taken as given. As of date 0, the probability that a particular history s^t will be realized is denoted $\pi(s^t)$.

Production The production technology is represented by a neoclassical production function with constant returns to scale in capital (k) and labor (l)

$$y(s^t) = f\big(k(s^t), l(s^t), s_t\big) = A(s^t)k(s^t)^{\alpha}l(s^t)^{1-\alpha}, \tag{4.1}$$

[2]In fact, the first-best allocations under both timing assumptions are essentially indistinguishable.

where $A(s_t)$ represents the state of technology in period t, $y(s^t)$ denotes the aggregate (or per capita) level of output, and $k(s^t)$ and $l(s^t)$ denote capital and labor used in production. Output can be used either for private consumption ($c(s^t)$), public consumption ($g(s^t)$), or as investment ($i(s^t)$). Feasibility thus requires that

$$c(s^t) + g(s^t) + i(s^t) = f\big(k(s^t), l(s^t), s_t\big). \tag{4.2}$$

What distinguishes this paper from others in the literature is our law of motion for capital, defined via

$$i(s^t) = k(s^t) + \delta k(s^t) - k(s^{t-1}). \tag{4.3}$$

Noticed that since investment or capital becomes productive immediately, it is used in production and depreciates within the period. In this way, the supply of capital is elastic even in the short run.

The usual properties of the neoclassical model hols: the capital to labor ratio is independent of scale, firms make zero profits in equilibrium, and factors are paid their marginal products:

$$\hat{r}(s^t) = f_k\big(k(s^t), l(s^t), s_t\big) - \delta = f_k(s^t) - \delta; \tag{4.4}$$

$$\hat{w}(s^t) = f_l\big(k(s^t), l(s^t), s_t\big) \qquad = f_l(s^t). \tag{4.5}$$

Households The economy is populated by a large number of identical individuals who live for an infinite number of periods and are endowed with one unit of time every period. Individuals' preferences are ordered according to the following utility function

$$\sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi(s^t) U\big(c(s^t), l(s^t)\big), \tag{4.6}$$

where $c(s^t)$ and $l(s^t)$ represent consumption and hours worked at history s^t. We assume that the felicity function is increasing in consumption and leisure $(1 - l^{s^t})$, strictly concave, twice continuously differentiable, and satisfies the Inada conditions for both consumption and leisure.

Each period individuals face the budget constraint

$$c(s^t) + k(s^t) + \sum_{s_{t+1}} q(s_{t+1}|s^t) b(s_{t+1}|s^t) = w(s^t) l(s^t) + r(s^t) k(s^t) + k(s^{t-1}) + b(s_t|s^{t-1}) \quad (4.7)$$

where $w(s^t) = [1 - \tau^w(s^t)]\hat{w}(s^t)$ and $r(s^t) = [1 - \tau^k(s^t)]\hat{r}(s^t)$, and where 'hats' denote pre-tax prices. The fiscal policy instruments τ^w and τ^k, as well as government debt $b(s_{t+1}|s^t)$ will be discussed in detail below. Notice that capital and government debt are treated rather symmetrically in budget constraint (4.7)—except of course for the fact that the size of the capital stock and its return cannot depend on the state tomorrow. In other words, the price of one unit of capital tomorrow is $1 - r(s^t)$, much like the price of a bond which pays one unit of consumption good tomorrow in state s_{t+1} costs $q(s_{t+1}|s^t)$. As we will see later, the symmetry is even clearer without uncertainty or in the absence of state-contingent government debt.

Letting $p(s^t)$ denote the Lagrange multiplier on the budget constraint at history s^t, the first order necessary (and sufficient) conditions for a solution to the consumer's problem are given by (4.7) and

$$\beta^t \pi(s^t) U_c(s^t) = p(s^t), \quad (4.8)$$

$$\beta^t \pi(s^t) U_l(s^t) = -w(s^t) p(s^t), \quad (4.9)$$

at all dates t and histories s^t for consumption and labor,

$$-p(s^t)\left(1 - r(s^t)\right) + \sum_{s_{t+1}} p(s^{t+1}) = 0, \quad (4.10)$$

at all dates t and histories s^t for capital,

$$-p(s^t) q(s_{t+1}|s^t) + p(s^{t+1}) = 0, \quad (4.11)$$

at all dates t, histories s^t, and all states s_{t+1} tomorrow for bond holdings, as well as the transversality conditions

$$\lim p(s^t)k(s^t) = 0, \tag{4.12}$$

$$\lim \textstyle\sum_{s_{t+1}} p(s^{t+1})b(s_{t+1}|s^t) = 0. \tag{4.13}$$

Under complete markets, it is well knows that these first order conditions and the budget constraint can be conveniently combined into a single present value constraint, as stated next:

Proposition 1 *Under complete markets, an allocation solves the consumer's problem if and only if it satisfies equations* (4.7)–(4.13)*, or, equivalently, if and only if it satisfies the implementability constrain* [3]

$$\sum_{t,s^t} \beta^t \pi(s^t) \left[U_c(s^t)c(s^t) + U_l(s^t)l(s^t) \right] = A_0, \tag{4.14}$$

where $A_0 = U_c(s_0)[k_{-1}+b_{-1}]$*, and* k_{-1} *and* b_{-1} *are initial amounts of capital and government debt held by individuals.*

Proof. The proof is standard. [See for example Chari et al. (1994).] ■

The Government The government's role in this economy is to finance an exogenous stream of government expenditures, $g(s^t)$. The fiscal policy instruments available to the government consist of a proportional labor income tax $\tau^w(s^t)$; a proportional capital income tax $\tau^k(s^t)$; and issuance of new government debt $b(s_{t+1}|s^t)$.[4] At date t, the government's budget constraint is as follows:

$$g(s^t) + b(s_t|s^{t-1}) = \sum_{s_{t+1}} q(s_{t+1}|s^t)b(s_{t+1}|s^t) + \tau^w(s^t)\hat{w}(s^t)l(s^t) + \tau^k(s^t)\hat{r}(s^t)k(s^t). \tag{4.15}$$

[3] To obtain the implementability constraint, multiply the budget constraint (4.7) by $p(s^t)$, add them up, and use the first order conditions (4.8)–(4.11) to replace prices.

[4] Although we use the term capital income tax throughout the paper, the tax can just as well be thought of as applying to interest on government debt simply by defining these prices appropriately. We chose the current formulation as it is simpler.

The government thus has to finance government expenditures as well as debt issued yesterday that promised to pay in the event that s_t would occur today. In addition to taxing capital and labor income, the government can raise revenues by issuing new (state contingent) debt.

4.3 Deterministic Ramsey Problem

Before analyzing the general stochastic model introduced in the previous section, it will prove instructive to study a deterministic version of the model first. The intuition from this simpler model will in some sense carry over to the more complicated stochastic environment.

Accordingly, we set up a standard Ramsey problem for a deterministic version of the model. As is well known, there is an equivalence between choosing fiscal policy instruments directly and choosing allocations among an appropriately restricted set of allocations.[5] The government's problem consists of maximizing the utility of the representative individual (4.6) subject to the implementability constraint (4.14) and feasibility (4.2).[6] If we denote λ the Lagrange multiplier on the implementability constraint, we can then define the pseudo-welfare function W by

$$W\left(c_t, l_t, \lambda\right) = U\left(c_t, l_t\right) + \lambda \left(U_{c_t} c_t + U_{l_t} l_t\right)$$

The Lagrangian associated with the Ramsey problem, given k_{-1} and b_{-1}, is thus given by:

$$L(k_{-1}, b_{-1}) = \min_{\lambda} \max_{\{c_t, l_t, k_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t W\left(c_t, l_t, \lambda\right) - \lambda U_{c_0}(k_{-1} + b_{-1})$$

subject to the feasibility constraint

$$c_t + g_t + k_t = f(k_t, l_t) - \delta k_t + k_{t-1}.$$

[5]See Chari and Kehoe (1999) or Erosa and Gervais (2001).

[6]It is well known that if an allocation satisfies the implementability constraint and the feasibility constraint, it must also satisfy the government budget constraint (4.15)—see Chari and Kehoe (1999) or Erosa and Gervais (2001) for example. Accordingly, we omit the proof.

It should be clear that one can replace the feasibility constraint into the objective function, and that the labor supply can be assumed to satisfy an optimality condition. Accordingly, slightly abusing notation, the Ramsey problem can be rewritten as

$$L(k_{-1}, b_{-1}) = \min_{\lambda} \max_{\{k_t\}_{t=0}^{\infty}} \left\{ W\big(k_{-1}, k_0, \lambda\big) - \lambda U_{c0}(k_{-1}, b_{-1}) + \sum_{t=1}^{\infty} \beta^t W\big(k_{t-1}, k_t, \lambda\big) \right\}$$

Notice that the last term inside the maximand is a standard recursive problem: if we define $V(k, \lambda)$ via

$$V\big(k, \lambda\big) = \max_{k'} \Big\{ W\big(k, k', \lambda\big) + \beta V(k', \lambda) \Big\},$$

then the problem becomes

$$\begin{aligned} L(k_{-1}, b_{-1}) &= \min_{\lambda} \max_{k_0} \big\{ W\big(k_{-1}, k_0, \lambda\big) - \lambda U_{c0} a_{-1} + \beta V\big(k_0, \lambda\big) \big\} \\ &= \min_{\lambda} \widehat{V}(k_{-1}, b_{-1}, \lambda), \end{aligned}$$

where $\widehat{V}$ is the value of the maximand evaluated at the optimum for any given value of λ.

Figure 4.1 shows the shape of the value function $\widehat{V}$ as a function of λ, given some values for initial assets.[7] What this Figure shows is that without any restrictions on the fiscal policy instruments or otherwise, the optimal degree of distortions, as represented by λ, is non-zero. Indeed, labor income is taxed at a rate of 19% in the long run. Capital income is not taxed in the long run: this can be shown formally as we will see in the next section.

The fact that it is optimal to distort this economy is in sharp contrast to results obtained under the more conventional timing whereby investment made during the period only becomes productive the next period. The reason is well known: under conventional timing, taxing initial assets represents a lump-sum way to raise revenues for the government, as these assets were accumulated in the past. Accordingly, the optimal fiscal policy entails taxing these initial assets at 'confiscatory' rates, or just enough that the government can

[7] Initial capital is set to 1.5 and debt to 0. See section 4.5.3 for the calibration of the parameters.

finance government the present value of its spending. In terms of Figure 4.1, the value function $\widehat{V}$ would be a strictly increasing function, with its minimum at exactly zero, meaning that a first-best outcome would be attained.

The intuition for our result comes directly from our timing assumption. Since investment becomes productive immediately, and its return is realize during the period, taxing capital at date zero becomes distortionary: individuals do not have to supply their capital accumulated from the past. They can, and will, consume large amounts if the government tries to tax their capital away. Realizing that fact, the government does not confiscate initial assets. Nevertheless, in the numerical example underlying Figure 4.1, the tax rate on capital is very high, around 696%.[8] As a result, consumption at date 0 is around 50% higher than in period 1, which is itself slightly below its steady state level. The tax rate on labor at date 0, however, is negative 20%: this makes leisure relatively expensive in that period, thereby increasing the labor supply.

The general message of this analysis is that the government's ability to use capital income taxes in a lump-sum fashion disappears once the supply of capital is elastic. This simple yet powerful message will also be at the heart of our findings in a stochastic economy, to which we now turn our attention.

4.4 Stochastic Ramsey Problem

To study optimal policy in this environment, we proceed as in the previous section and set up a standard Ramsey problem. With λ still denoting the Lagrange multiplier on the implementability constraint, the pseudo-welfare function W now reads

$$W\big(c(s^t), l(s^t), \lambda\big) = U\big(c(s^t), 1 - l(s^t)\big) + \lambda \left[U_c(s^t)c(s^t) + U_l(s^t)l(s^t)\right].$$

[8] While capital income taxes are very high in the initial period, they are clearly not sufficiently high to eliminate all future distortions, as discussed above.

The Ramsey problem is thus as follows:

$$L(k_{-1}, b_{-1}) = \min_{\lambda} \max_{\{c(s^t), l(s^t), k(s^t)\}_{t,s^t}} \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi(s^t) W\big(c(s^t), l(s^t), \lambda\big) - \lambda U_c(s_0)[k_{-1} + b_{-1}]$$

subject to the feasibility constraint (4.2), keeping in mind that the government budget constraint must holds and so does not constraint the solution to this problem.

The government typically has more instruments than it needs, in the sense that many tax codes can decentralize any given allocation (e.g. see Zhu (1992) or Chari et al. (1994)). Such is not the case in our environment: our tax code is unique, in the sense that any given allocation can only be decentralized by a single tax system. Technically, this comes from the fact that the tax rate on capital is pined down by the marginal product of capital (4.4) as well as the optimality conditions (4.8) and (4.10): for any implementable allocation, there exists a single value of the capital tax which makes these two equations hold. Intuitively, the indeterminacy under conventional timing comes from the fact that an allocation can, for example, be implemented with a tax rate on capital that varies with the state tomorrow and risk-free debt, or with a flat tax on capital income tomorrow and state-contingent debt. Here, the capital income tax applies to the return to investment made during the period, so it is uniquely determined even with state-contingent debt. It follows that ruling out state-contingent debt is not innocuous in our environment, as we will see in the next section.

The optimality conditions of this Ramsey problem are quite simple, and can be analyzed analytically. Let $\beta^t \phi(s^t)$ represent the Lagrange multipliers on the feasibility constraint (4.2). The first order conditions with respect to consumption, labor, and capital, are, respectively,

$$\pi(s^t)W_c(s^t) = \phi(s^t); \tag{4.16}$$

$$\pi(s^t)W_l(s^t) = -f_l(s^t)\phi(s^t). \tag{4.17}$$

$$\phi(s^t)\big[1-(f_k(s^t)-\delta)\big] = \beta\sum_{s_{t+1}}\phi(s^{t+1}). \tag{4.18}$$

4.4.1 Optimal Fiscal Policy

The rest of this section is devoted to characterize optimal fiscal policy. Our characterization, which requires making assumptions about the form of the utility function, involves in turn the labor income tax and the capital income tax.

Our first two Propositions show that while the labor tax does not depend on the state of the economy if the per-period utility is separable between consumption and *labor* and both part exhibit constant elasticity of substitution (CES), it becomes pro-cyclical when individual care about leisure, even if the utility function is CES in *leisure*.

Proposition 2 *Assume that the felicity function is separable, $U(c,l) = u(c) + v(l)$, with $u(c)$ and $v(l)$ both exhibiting constant elasticity of substitution. Then the tax rate on labor income is invariant to the productivity shock.*

Proof. Combining the first order conditions with respect to consumption (4.16) and labor (4.17) from the Ramsey problem and using (4.5), we get

$$-\frac{W_l(s^t)}{W_c(s^t)} = \hat{w}(s^t). \tag{4.19}$$

The derivatives W_c and W_l are given by

$$W_c(s^t) = (1+\lambda)U_c(s^t) + \lambda U_c(s^t)H_c(s^t),$$

$$W_l(s^t) = (1+\lambda)U_l(s^t) + \lambda U_l(s^t)H_l(s^t),$$

where

$$\begin{aligned} H_c(s^t) &= \frac{U_{c,c}(s^t)c(s^t) + U_{c,l}(s^t)l(s^t)}{U_c(s^t)}, \\ H_l(s^t) &= \frac{U_{l,c}(s^t)c(s^t) + U_{l,l}(s^t)l(s^t)}{U_l(s^t)}. \end{aligned}$$

Now pick two histories as of date t, s^t and $\tilde{s}^t$. From (4.19), it must be that

$$\frac{W_l(s^t)}{W_c(s^t)\hat{w}(s^t)} = \frac{W_l(\tilde{s}^t)}{W_c(\tilde{s}^t)\hat{w}(\tilde{s}^t)},$$

or, equivalently,

$$\frac{\left[1 + \lambda + \lambda H_l(s^t)\right]U_l(s^t)}{\left[1 + \lambda + \lambda H_c(s^t)\right]U_c(s^t)\hat{w}(s^t)} = \frac{\left[1 + \lambda + \lambda H_l(\tilde{s}^t)\right]U_l(\tilde{s}^t)}{\left[1 + \lambda + \lambda H_c(\tilde{s}^t)\right]U_c(\tilde{s}^t)\hat{w}(\tilde{s}^t)}.$$

Since the felicity function is separable, the functions H_c and H_l become

$$\begin{aligned} H_c(s^t) &= \frac{U_{c,c}(s^t)c(s^t)}{U_c(s^t)}, \\ H_l(s^t) &= \frac{U_{l,l}(s^t)l(s^t)}{U_l(s^t)}. \end{aligned}$$

And since the sub-utilities for consumption and labor are both from the constant elasticity of substitution class of utility, that means both H_c and H_l are constant. Accordingly, the last expression reduces to

$$\frac{U_l(s^t)U_c(\tilde{s}^t)}{U_c(s^t)U_l(\tilde{s}^t)} = \frac{\hat{w}(s^t)}{\hat{w}(\tilde{s}^t)}.$$

But the first order conditions for consumption and labor from the household's problem (equations (4.8) and (4.9)) under histories s^t and $\tilde{s}^t$ imply

$$\frac{U_l(s^t)U_c(\tilde{s}^t)}{U_c(s^t)U_l(\tilde{s}^t)} = \frac{w(s^t)}{w(\tilde{s}^t)} = \frac{(1 - \tau^w(s^t))\hat{w}(s^t)}{(1 - \tau^w(\tilde{s}^t))\hat{w}(\tilde{s}^t)}.$$

For the last two equations to hold it must be the case that $\tau^w(s^t) = \tau^w(\tilde{s}^t)$. ■

The intuition for this result is that because the elasticity of the labor supply does not vary with the shock, there is no reason for the government to tax labor at rates that vary with the shock.[9] Note that the previous result does not apply when individuals care about leisure, as opposed to disliking labor. The following proposition shows that indeed labor income taxes will in general not be constant when individuals care about leisure.

Proposition 3 *Assume that $\lambda > 0$ and that the felicity function is given by $U(c,l) = u(c)v(l)$, with $u(c) = (1-\sigma)^{-1}c^{1-\sigma}$ and $v(l) = (1-l)^{\nu(1-\sigma)} = (1-l)^{\eta}$, with $\sigma > 1$ and $\nu > 0$, and $\ln(c) + \eta \ln(1-l)$ for $\sigma = 1$. Assume that there exist two states s^t and $\tilde{s^t}$ such that $l(s^t) > l(\tilde{s}^t)$. Then $\tau^w(s^t) > \tau^w(\tilde{s}^t)$ if and only if*

$$\lambda < \frac{-1}{(1-\sigma)(1+\nu)}. \tag{4.20}$$

Proof. From equations (4.8)–(4.9) and (4.19), the tax rate on labor income is given by

$$\tau^w(s^t) = \frac{\lambda\big(H_l(s^t) - H_c(s^t)\big)}{1 + \lambda + \lambda H_l(s^t)}. \tag{4.21}$$

Under the stated utility function, H_c and H_l are such that

$$H_l(s^t) - H_c(s^t) = \frac{-1}{1 - l(s^t)},$$
$$H_l(s^t) = -\sigma + \frac{1 - \eta l(s^t)}{1 - l(s^t)}.$$

Using these expression in equation (4.21) we have

$$\tau^w(s^t) = \frac{\lambda}{1 - \lambda(\sigma - 2) - l(s^t)\big(1 + \lambda(1-\sigma)(1+\nu)\big)}.$$

It follows that the tax rate is higher under state s^t than $\tilde{s}^t$ if the term in from of labor is positive, which is the condition given above. ■

Note that we need to assume that the economy is distorted ($\lambda > 0$), otherwise all taxes are zero. This Proposition establishes that whenever condition (4.20) is satisfied, if labor is pro-cyclical, so will the tax rate on labor income. Note that under logarithmic utility,

[9] Evidently, the same argument can be made using s^{t-1} and s^t as the two histories, which means that the tax rate on labor is not only state-independent, but also constant over time.

i.e. when $\sigma = 1$, the condition is always satisfied. It becomes less likely to be satisfied as individuals become more risk averse, i.e. as σ increases. As such, this Proposition is useful to interpret the finding in Chari et al. (1994) that the correlation between the shock and labor taxes changes sign as they change the risk aversion parameter. Finally, note that what is key is whether the utility function exhibits constant elasticity of substitution in labor or in leisure. When it is CES in leisure, the labor supply elasticity varies with the *level* of the labor supply, becoming more inelastic as the labor supply increases. This is in contrast to our previous proposition, where the labor supply elasticity was invariant to the level of the labor supply.

Our next results pertain to the tax on capital. We first show that capital income should not be taxed if the utility function is separable and exhibits constant elasticity of substitution in consumption. We then argue that under non-separable preferences, the tax rate on interest income is likely to be pro-cyclical.

Proposition 4 *Assume that the felicity function is separable, $U(c,l) = u(c) + v(l)$, and that $u(c)$ exhibits constant elasticity of substitution. Then the capital income tax rate is zero at all dates and histories (other than the first period).*

Proof. Recall that the first order conditions (4.8) and (4.10) from the households' problem imply that

$$(1 - r(s^t)) = \sum_{s_{t+1}} \frac{\beta\pi(s^{t+1})U_c(s^{t+1})}{\pi(s^t)U_c(s^t)}. \tag{4.22}$$

Similarly, combining first order conditions (4.16) and (4.18) from the Ramsey problem we have

$$\left[1 - (f_k^p(s^t) - \delta)\right] = (1 - \hat{r}(s^t)) = \sum_{s_{t+1}} \frac{\beta\pi(s^{t+1})W_c(s^{t+1})}{\pi(s^t)W_c(s^t)}. \tag{4.23}$$

But with separable utility and constant elasticity of substitution,

$$W_c(s^t) = (1 + \lambda + \lambda H^c(s^t))U_c(s^t) = (1 + \lambda - \lambda\sigma)U_c(s^t),$$

where σ is the inverse of the intertemporal elasticity of substitution. Hence we can replace W_c with U_c in equation (4.23). But then the only way for both equation (4.23) and equation (4.22) to hold is if we have $\tau^k(s^t) = 0$. ∎

This Proposition is in sharp contrast to the results in Chari et al. (1994), where the *ex post* tax rate on capital income is extremely volatile.[10] The intuition is that in their set up, the return on investment made today is taxed tomorrow. Since the investment decision has already been made when the tax authority sets the tax rate on capital income, this instrument is extremely useful to absorb shocks to the budget of the government. For example, if the economy experiences a bad shock today, then the government will tax capital income at a high rate to absorb the loss in revenue. The more persistent the shock is, the higher the tax rate. In fact, under standard parameter specifications, the increase in capital income taxes is so large that the government runs a primary surplus in the period of a negative shock, thereby absorbing the future path of low government revenues with very little change to the tax rate on labor income. Of course, the tax authority always promises individuals that *on average* capital income will not be taxed. This is what Chari et al. (1994) refer to as the *ex ante* tax rate on capital income, which, under the assumptions of proposition 4, is zero.

In our setup, the return on capital is known at the time individuals make their investment decision, thereby eliminating the distinction between *ex ante* and *ex post* taxes on capital. In particular, the tax authority no longer has the ability to absorb shocks in an essentially non-distortionary fashion using the capital income tax as a shock absorber.

Under more general preferences, the tax rate on capital will not in general be equal to zero. For instance, if $U(c,l) = u(c)v(l)$, with $u(c) = (1-\sigma)^{-1}c^{1-\sigma}$ and $v(l) = (1-l)^{\nu(1-\sigma)} = (1-l)^{\eta}$, with $\sigma > 1$ and $\nu > 0$, then capital income will tend to be subsidized in bad times

[10] As pointed out at the beginning of this section, however, one should keep in mind that this statement implicitly picks one of many potential tax codes.

and taxed in good times. To see this, note that the function $H_c(s^t)$ under this utility function is given by

$$H_c(s^t) = -\sigma - \eta \frac{l(s^t)}{1 - l(s^t)},$$

which, since $\eta < 0$, is increasing in l. Now from equations (4.22) and (4.23), we have

$$\frac{1 - r(s^t)}{1 - \hat{r}(s^t)} = \frac{\sum_{s_{t+1}} \pi(s^{t+1}|s^t)\big(1 + \lambda + \lambda H^c(s^t)\big)U_c(s^{t+1})}{\sum_{s_{t+1}} \pi(s^{t+1}|s^t)\big(1 + \lambda + \lambda H^c(s^{t+1})\big)U_c(s^{t+1})}. \tag{4.24}$$

When this ratio is smaller than 1, capital income is subsidized, and capital income is taxed if the ratio is greater than 1. In particular, capital income is subsidized when $H_c(s^t)$ is relatively low, i.e. when the labor supply is relatively low. Much like the labor income tax, the capital income tax is thus likely to be pro-cyclical as long as labor is pro-cyclical.

The results of this section tell us that depending on the form of the utility function, labor and capital income taxes can either be acyclical or pro-cyclical. However, these results are silent as to the behavior of government debt over the business cycle, even if taxes are pro-cyclical. This is because with state contingent government debt, it may be optimal for the government to commit to a policy that involves repaying a lower amount of debt during recessions—a partial default of debt in the words of Chari and Kehoe (1999). This can easily be established by deriving a present value budget constraint for the government. By substituting forward $b(s^{t+1}|s^t)$ into the government budget constraint (4.15), letting $ps(s^t) = \tau^w(s^t)\hat{w}(s^t)l(s^t) + \tau^k(s^t)\hat{r}(s^t)k(s^t) - g(s^t)$ denote the primary surplus, one obtains the following representation for debt:

$$b(s^t) = ps(s^t) + \sum_{\tau=t}^{\infty} \sum_{s_{\tau+1}} q(s^{\tau+1}|s^t) ps(s^{\tau+1}|s^t).$$

The equation above clearly states that any shock that reduces the present value of primary surpluses induces a reduction in debt.

To conclude, our model implies that while the primary deficit can be counter-cyclical (i.e. tax revenues are low in bad times and high in good times), the presence of state-contingent government debt can make government debt pro-cyclical and thus negatively correlated with the primary deficit, a phenomenon which we typically do not observe (see Marcet and Scott (2009)). Accordingly, we now turn our attention to a situation in which the government only has access to risk-free debt.

4.5 Ruling out State-Contingent Debt

Ruling out state-contingent debt in the standard neoclassical growth model has proven difficult (e.g. see Chari and Kehoe (1999)). In our framework, however, this task is quite tractable. To see this, consider the consumer's budget constraint without state contingent debt:

$$c(s^t) + k(s^t) + q(s^t)b(s^t) = w(s^t)l(s^t) + r(s^t)k(s^t) + k(s^{t-1}) + b(s^{t-1}). \tag{4.25}$$

It should be clear that the first order conditions for consumption, labor, and capital, equations (4.8)–(4.10), remain valid under this budget constraint. These equations imply that

$$\begin{aligned} w(s^t) &= -\frac{U_l(s^t)}{U_c(s^t)}; \\ 1 - r(s^t) &= \beta \sum_{s_{t+1}} \pi(s^{t+1}|s^t) \frac{U_c(s^{t+1})}{U_c(s^t)}, \end{aligned}$$

which can be replace in the budget constraint to obtain

$$c(s^t) + (k(s^t) + b(s^t))\beta \sum_{s_{t+1}} \pi(s^{t+1}|s^t) \frac{U_c(s^{t+1})}{U_c(s^t)} = -\frac{U_l(s^t)}{U_c(s^t)} l(s^t) + k(s^{t-1}) + b(s^{t-1}). \tag{4.26}$$

Of course, without state-contingent debt these budget constraints can no longer be expressed as a single present-value budget constraint. Ruling out state-contingent debt amounts to imposing a sequence of budget or implementability constraints of the form above. The difficulty in the neoclassical growth model under conventional timing is that the interest rate cannot merely be substituted out because it appears within an expectation sign in the Euler equation.

Given the form of the implementability constraint (4.26), we can use the methodology developed in Marcet and Marimon (1995) to obtain the following Ramsey problem in Lagrangian form:

$$
\begin{aligned}
L(k_{-1}, b_{-1}) = \min_{\{\lambda(s^t)\}_{t,s^t}} \max_{\{c(s^t),l(s^t),k(s^t),b(s^t)\}_{t,s^t}} \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi(s^t) \Big\{ U\big(c(s^t), l(s^t)\big) \\
- \lambda(s^t) \left(c(s^t) + \frac{U_l(s^t)}{U_c(s^t)} l(s^t) - k(s^{t-1}) - b(s^{t-1}) \right) U_c(s^t) \\
- \lambda(s^{t-1}) \big(k(s^{t-1}) + b(s^{t-1}) \big) U_c(s^t) \Big\}
\end{aligned} \tag{4.27}
$$

subject to feasibility (4.2) at all dates and histories, given k_{-1} and b_{-1}, with $\lambda_{-1} = 0$.

4.5.1 Analysis

We first establish that the evolution of the multiplier λ, which reflects the distortionary nature of taxation over time, contains a permanent component—a result first discussed in Aiyagari et al. (2002) in a model without capital, and more recently by Scott (2007) in a model with capital in which capital income taxation is ruled out. To establish this result, notice that the first-order condition for government debt states that

$$
\sum_{s^{t+1}|s^t} \beta^{t+1} \pi(s^{t+1}) \left(\lambda(s^t) U_c(s^{t+1}) - \lambda(s^{t+1}) U_c(s^{t+1}) \right) = 0. \tag{4.28}
$$

Since $\lambda(s^t)$ is known at history s^t, it can be taken out of the expectation, establishing that

$$
\lambda(s^t) = \frac{\sum_{s_{t+1}} \pi(s^{t+1}|s^t) U_c(s^{t+1}) \lambda(s^{t+1})}{\sum_{s_{t+1}} \pi(s^{t+1}|s^t) U_c(s^{t+1})}, \tag{4.29}
$$

so that the multiplier λ follows a risk-adjusted Martingale. An interesting special case, which we study in more details below, is one where the felicity function is quasi-linear, i.e. $U(c, l) = c + v(l)$. In this case, the marginal utility of consumption is constant at unity, and so the stochastic process for the multiplier λ becomes a martingale. Indeed, Farhi (2005) shows that if the government faces natural debt limits and the stochastic process governing the state s_t converges to a unique (non-degenerate) stationary distribution, then

λ_t converges to zero, which implies that the Ramsey allocation converges to a first-best allocation (i.e. all taxes are zero in the long run). This result holds in our economy as well.

In general not much can be said analytically about the behavior of optimal taxes in this environment. In particular, nothing can be said about the labor income tax, at least as far as we can tell. For the capital income tax, we will now establish one special case where it is always zero. If we let $\beta^t \pi(s^t)\phi(s^t)$ be the multipliers on the feasibility constraint, the first order condition with respect to capital reads

$$+ \sum_{s^{t+1}|s^t} \beta^{t+1}\pi(s^{t+1}) \left(\lambda(s^{t+1}) - \lambda(s^t)\right) U_c(s^{t+1}) \\ + \beta^t \pi(s^t)\phi(s^t)\Big(1 - \big(f_k(s^t) - \delta\big)\Big) + \sum_{s^{t+1}|s^t} \beta^{t+1}\pi(s^{t+1})\phi(s^{t+1}) = 0,$$

which, given (4.28), implies that

$$1 - \big(f_k(s^t) - \delta\big) = 1 - \hat{r}(s^t) = \frac{\sum_{s^{t+1}|s^t} \beta\pi(s^{t+1})\phi(s^{t+1})}{\pi(s^t)\phi(s^t)}. \tag{4.30}$$

As usual, recalling equation (4.22)—which holds here as well—interest income should not be taxed if the shadow value of resources is equal to marginal utility at all dates and states, i.e. if $\phi(s^t) = U_c(s^t)$.[11] This will in general not be the case, even under a per-period utility function separable between consumption and leisure. In this case, the value of the multiplier ϕ, from the first order condition for consumption, is given by

$$\phi(s^t) = U_c(s^t) + \lambda(s^t). \tag{4.31}$$

This equation highlights the role of λ as a measure of distortions. Clearly, whenever λ is not zero, taxes on capital may not be zero. There is, however, one special case under which we can establish that capital income should not be taxed, as we state in the following proposition.

[11] Note that this is a necessary condition, but there can be cases in which the equation does not hold and yet the tax on capital is zero.

Proposition 5 *If the per-period utility function is quasi-linear in consumption, i.e.* $U(c,l) = c + v(l)$*, then the tax rate on capital income is zero.*

Proof. First note that under this utility function, because the marginal utility of consumption is fixed at unity, (4.22) implies that $1 - r(s^t) = \beta$. From (4.31), the value of the multiplier on the feasibility constraint is given by $\phi(s^t) = 1 + \lambda(s^t)$. Furthermore, (4.29) implies that $\lambda(s^t) = \sum_{s_{t+1}} \pi(s^{t+1}|s^t)\lambda(s^{t+1})$. Using these facts in equation (4.30) imply that $1 - \hat{r}(s^t) = \beta$. ■

4.5.2 Numerical Example

We now solve the model numerically to characterize more the behaviour of our economy. The computation of the present model presents a challenge that comes from the fact that an ergodic set that bounds the behaviour of the variables of the model may depend on initial conditions for the shocks, bonds and capital, if it exists at all. This is perhaps why very good initial guesses are needed for convergence. The model is therfore solved using homotopy starting from a deterministic model without capital. From there we move to the stochastic and incompete market model without capital of Aiyagari et al. (2002) with quasi-linear preferences on consumtpion. Subsequently we introduce decreasing marginal utility of consumption, capital with full depreciation, and then introduce partial depreciation. The model is solved by policy function iteration. Policy functions are approximated by cubic splines in the proximity of the grid points. We consider 15 grid points for k, b and λ, and three for the shock A. This leads to a total of 10125 points on the state space.

4.5.3 Calibration

Preference are given by $u(c,l) = log(c) + \nu log(1-l)$. $\nu = 1.5$. The production function is Cobb-Douglas as in equation (4.1) and $\alpha = 0.33$. Since capital can be used within the period we take a period to be a year. The discount factor β is set equal to 0.958, which corresponds to a steady state interest rate of 4-5 percent. Capital depreciation δ is 0.07. Productivity is modeled as a three state first order Marcov chain which approximates an AR 1 process with persistence .5, which corresponds to about 2 years of recessions and booms, and standard deviation of the innovations equal to 0.014, which is the annualized standard deviation of the one used by Fuentes-Albero et al. (2009). Government spending g is equal to 0.1067, which implies an average spending-output ratio of 16.7 %.

4.5.4 Simulation

Perhaps the two most interesting aspects that simulations can clarify are the responses of taxes and bonds to shocks, and the long run properties of λ. We have shown that λ converges almost surely to zero when preferences are quasi-linear in consumption. It is interesting to see whether this is also the case with more general preferences.

To analyze fiscal policy impulse responses, we resemble a two period recession by letting the model converge to a steady state and then hit the economy with a negative shock for 2 periods, after which the productivity shock goes back to its mean (second grid point). The first lession that can be learned from this exercise is that impulse responses depend on the point on the state space we are at: in particular the response of government debt depends on the initial level of government debt itself. We therefore show what happens when starting with a negative and with a positive value of debt. Figure 4.2 shows impulse responses when the model started with bonds set at zero, which imply a steady state with negative debt. As it can be noticed, when a negative shock hits, the tax rate on labour reduces by very little; the fact that labour tax rates are very stable is a well established result with these models, with both complete and incomplete markets: see for instance Chari et al. (1994) and Farhi (2005). Taxes on capital increase, however this increase is very mild: it is in fact not sufficient to generate an increase in primary surplus. Government debt increases.

This is in sharp contrast to what happens in Figure 4.3, where debt over output is about 60% in the steady state before the negative shock. As the figure shows, debt responds very differently: it goes down. This happens notwithstanding the fact that primary surplus is also reduced. No matter how perverse this result may appear, there is a clear explanation for it: the following formula describes the reaction of bonds to a shock:

$$\frac{\partial b'}{\partial A} = -\frac{1}{1-r}\frac{\partial ps}{\partial A} + \frac{(b-ps)}{(1-r)^2}\frac{\partial r}{\partial A}. \tag{4.32}$$

b' means newly issued bonds, ps primary surplus, and all the other notation is conventional with the rest of the paper. The formula shows that the reaction of bonds to a productivity shock is the sum of two components. The first one is the effect on primary deficit. Impulse responses suggest that a negative shock induces a reduction in primary surplus, which means $\frac{\partial ps}{\partial A} < 0$. Hence the first term in the right hand side of equation (4.32) is positive. The second one however depends on the size of $b - ps$. $\frac{\partial r}{\partial A}$ is positive. Hence after a negative shock, the second component contributes to decrease debt if $b - ps < 0$. Otherwise this term pushes bonds down during recession. When bonds are sufficiently positive, this second component becomes dominant, implying that bonds go down in recession. Intuitively, the reduction in the interest rate reduces the cost of debt. The larger is debt, the more important this effect is. For sufficiently high debt, it becomes predominant over the increase in deficit that comes with a recession.[12] It should be noticed that this result is not just an artifact of our timing, it is true also with the usual timing and it is in fact a general accounting rule reminiscent of the Tobin effect, that when interests go down, the higher debt, the lower debt repayments.[13]

To conclude the analysis of debt we notice that even with high initial indebtedness, debt increases above the initial level after two periods. Accordingly, our measure of distortions,

[12] We have not been able yet to identify a threshold level of debt such that the two components compensate each other. But just to give a sense of what this threshold might be, we found that with debt at 22 % of output, debt goes up after a negative shock. So the threshold is larger than that.

[13] It is also worth emphasizing that the measure typically used for indebtedness -debt over GDP- goes up after a negative shock in our model, even in the experiment with higher debt, this can be seen in Figure 4.4.

λ remains persistently higher than before the recession. The trade-off between present well being and future costs from higher distortions is balanced by a policy that finances recessions by deficit and lead to an ever worse situation there after.

In our second experiment, the economy is again hit by a negative shock lasting two periods, but this time is followed by a long sequence of positive shocks. Letting this second recession end with a boom may highlight some non linearities in the response of the capital tax rate. Furthermore, the protracted period of high shocks (about 8000 periods) serves to see whether λ reaches the zero lower bound, where no distortions have to be imposed because spending can be completely financed by interests from the negative debt issued by the government.

As shown in Figures 4.5 and 4.6, impulse responses to the shock in the first two periods after the negative shock are obviously identical to the first experiment. When the positive shock hits, taxes on capital are reduced and λ and bonds start decreasing. They however reach a lower bound which is well above the planner solution with no distortions ($\lambda = 0$) as shown in Figures 4.9 and 4.10. This exercise therefore suggests that it may be very hard, if not impossible, to cancel distortions in this model when utility of consumption is concave.

We conclude this section by showing what happens at the beginning of times. The results confirm what found for the deterministic case in section 4.3. Despite the fact that capital is elastic, the lack of commitment by the initial Ramsey planner induces a very high tax on capital, up to almost 1000 % of the return on capital. This is smaller when starting with $b = 0$. Despite the very high tax rate on capital, the economy does not jump to the Pareto optimum as it would happen with usual timing and no exogenous tax constraints: this can be seen by the fact that λ jumps to a positive number.

4.6 Conclusion

This paper studies optimal fiscal policy in a neoclassical model where capital is elastically supplied in the short run. This is accomplished by assuming that investment becomes productive immediately.

Elastic supply of capital in the short run drastically alters the prescription that emanates from standard Ramsey problems. The solution to the Ramsey problem features a unique non-trivial level of distortions which depends on individuals' initial holdings of debt and capital and it does not rely on the presence of bounds exogenously imposed on the Ramsey problem.

The paper offers an analytical characterization of the behavior of tax rates for a variety of utility functions in a stochastic environment in which the government has access to state-contingent debt. Independently of whether taxes are pro-cyclical or not, the government uses state-contingent government debt as a shock absorber, much like the *ex post* capital income tax is used for that purpose in Chari et al. (1994). As a result, debt and the primary deficit move in opposite directions. Marcet and Scott (2009) showed that this is counterfactual and pervasive in models in which the government has access to state contingent government debt.

This leads us to study a Ramsey problem under incomplete markets. The upshot of this problem is that in this environment the government runs debt-financed primary deficits during recessions. The guiding principle underlying this result, which is key to bring the prescription from the model closer to what is typically observed in the data, is the classic trade-off between the gains of short-run counter-cyclical policy and the burden associated with government debt in the long run. Trade-off that is not present with the usual timing, where the fact that taxing capital in reaction to shocks is non distortive, implies no short run benefits from expansionary fiscal policy.

4.7 Figures

Figure 4.1: Value function, deterministic case

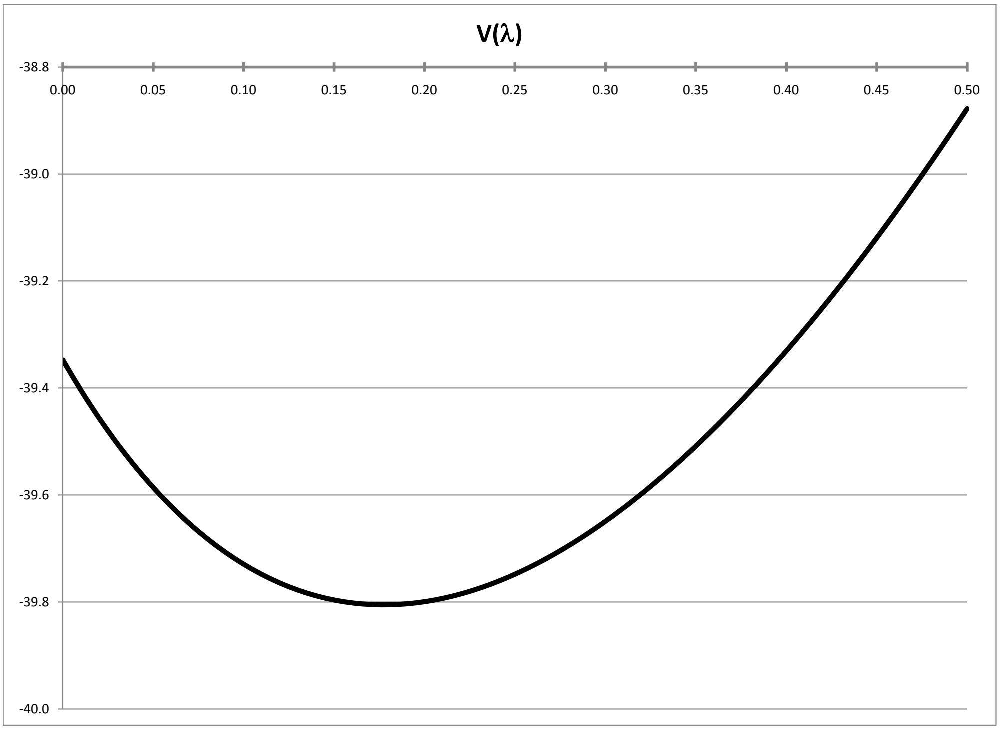

Note: initial capital is set to 1.5 (slightly below its steady state value of 1.8) and initial debt is set to 0.

Figure 4.2: First experiment with initial $b = 0$

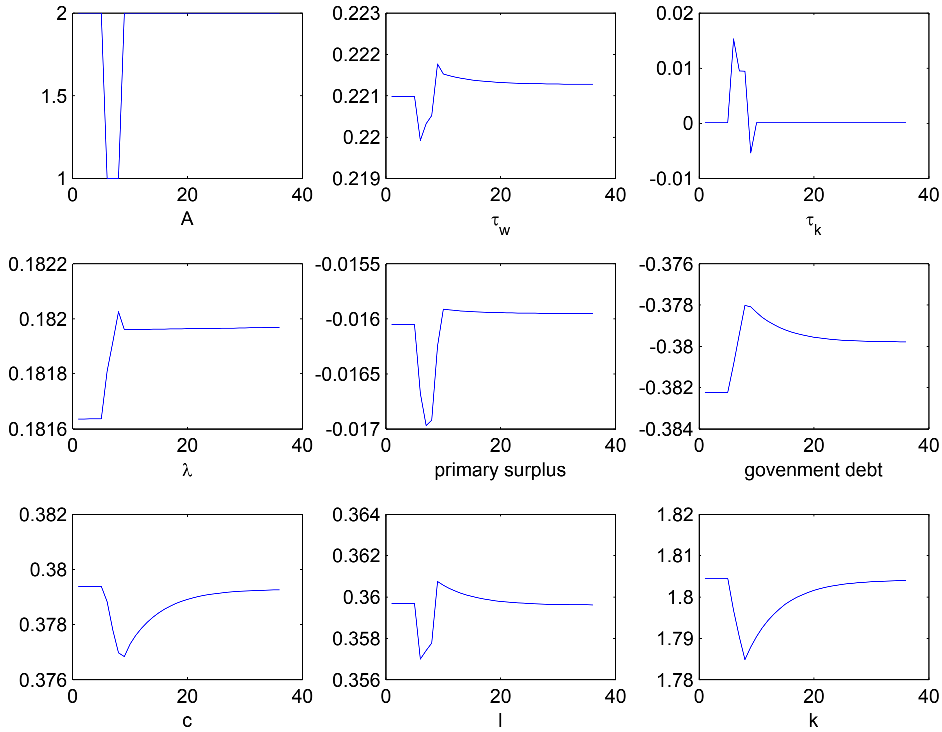

Impulse responses to a 2 period negative shock.

Note: initial $b = 0$

Figure 4.3: First experiment with initial $b = 1.3$

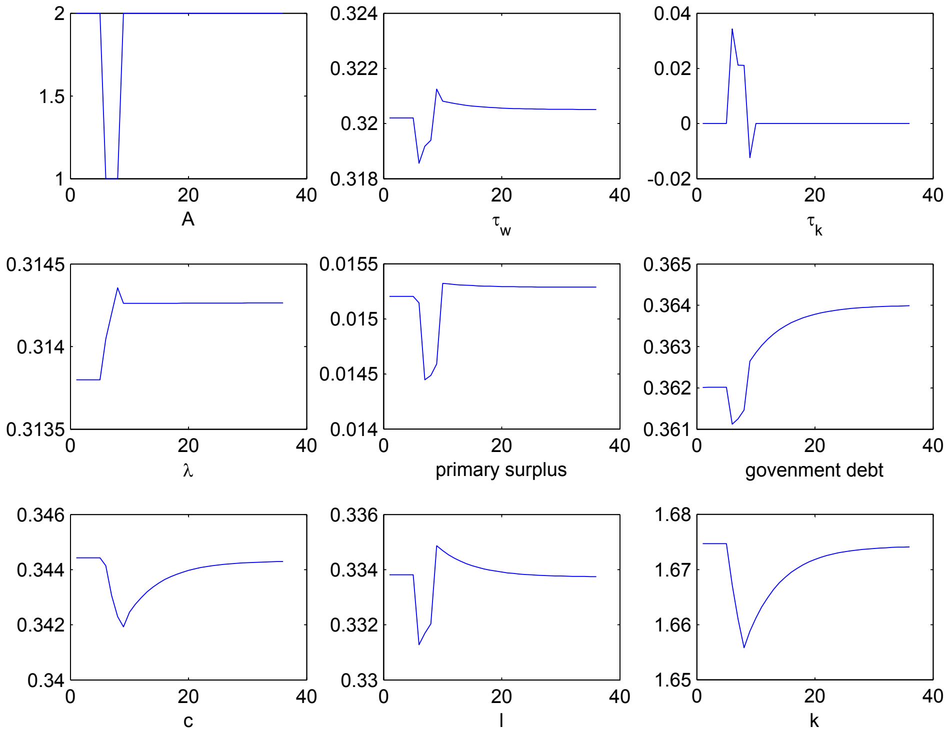

Impulse responses to a 2 period negative shock.

Note: initial $b = 1.3$, which implies a steady state of debt over output of around 60%.

Figure 4.4: First experiment with initial $b = 1.3$

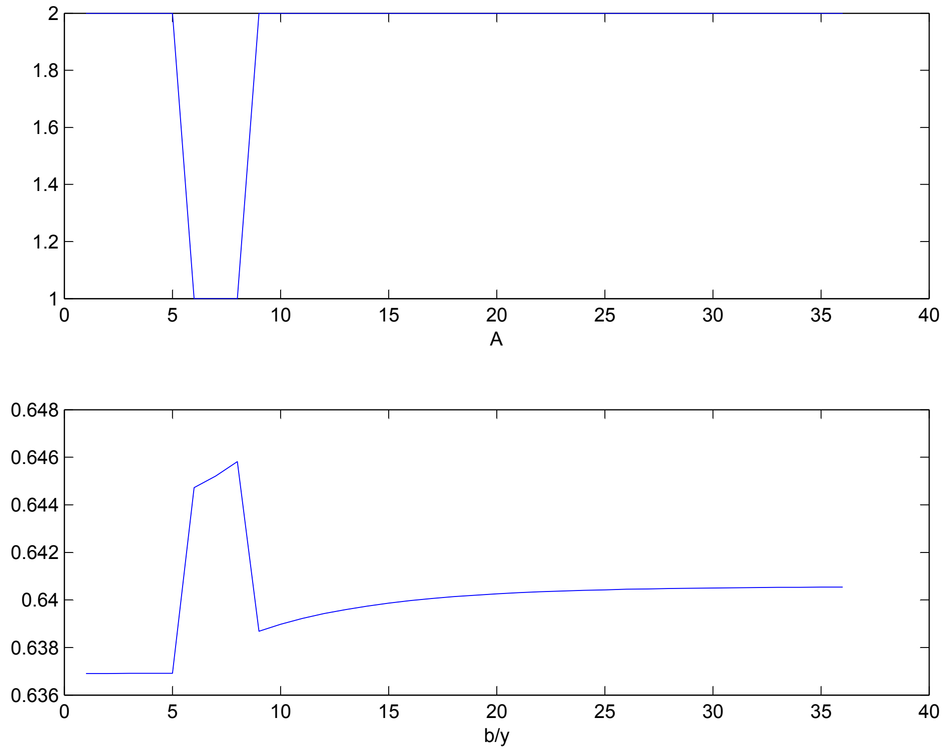

Impulse response of debt over output to a 2 period negative shock.

Note: initial $b = 1.3$.

Figure 4.5: Second experiment with initial $b = 0$

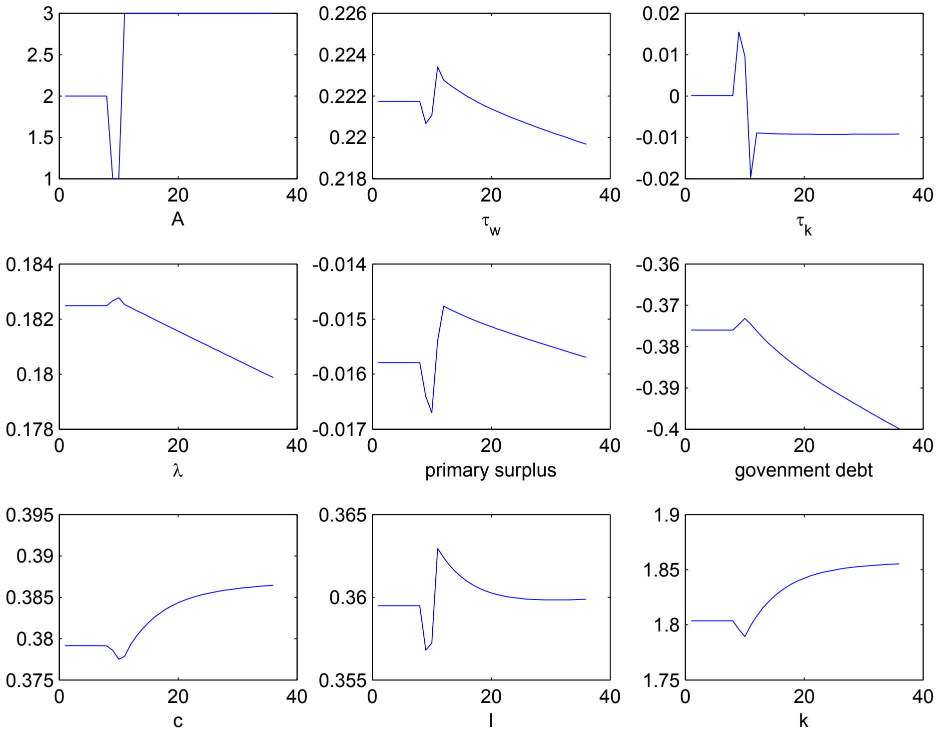

Impulse responses to a 2 period negative shock followed by a protracted period of positive shocks. Note: initial $b = 0$, which implies a steady state with negative debt over output.

Figure 4.6: Second experiment with initial $b = 1.3$

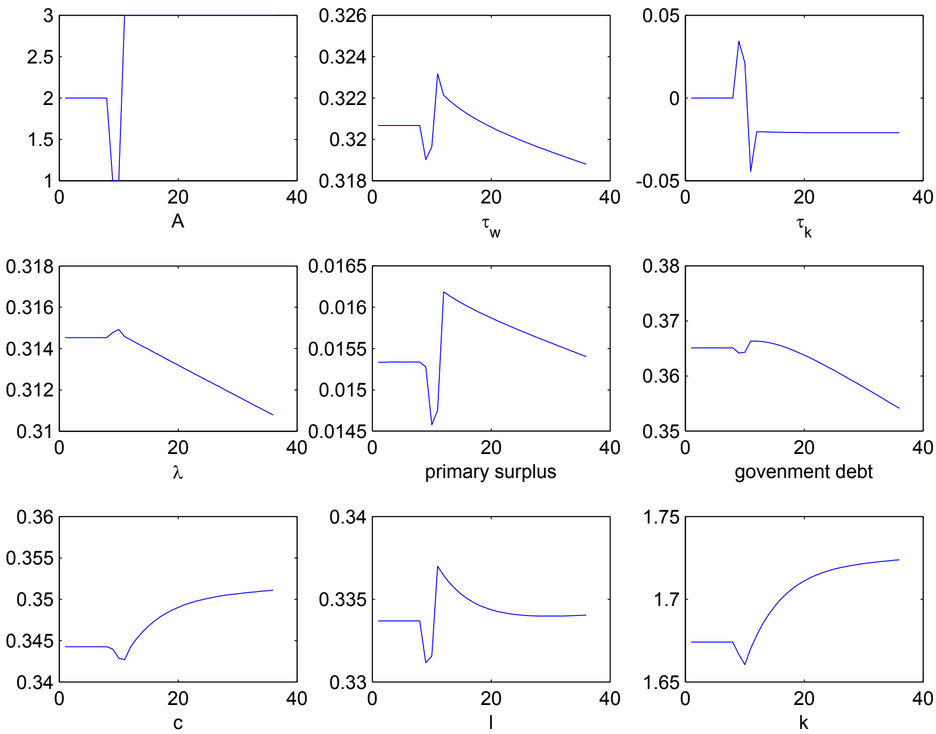

Impulse responses to a 2 period negative shock followed by a protracted period of positive shocks. Note: initial $b = 1.3$, which implies a steady state of debt over output of around 60%.

Figure 4.7: Initial periods with $b = 0$

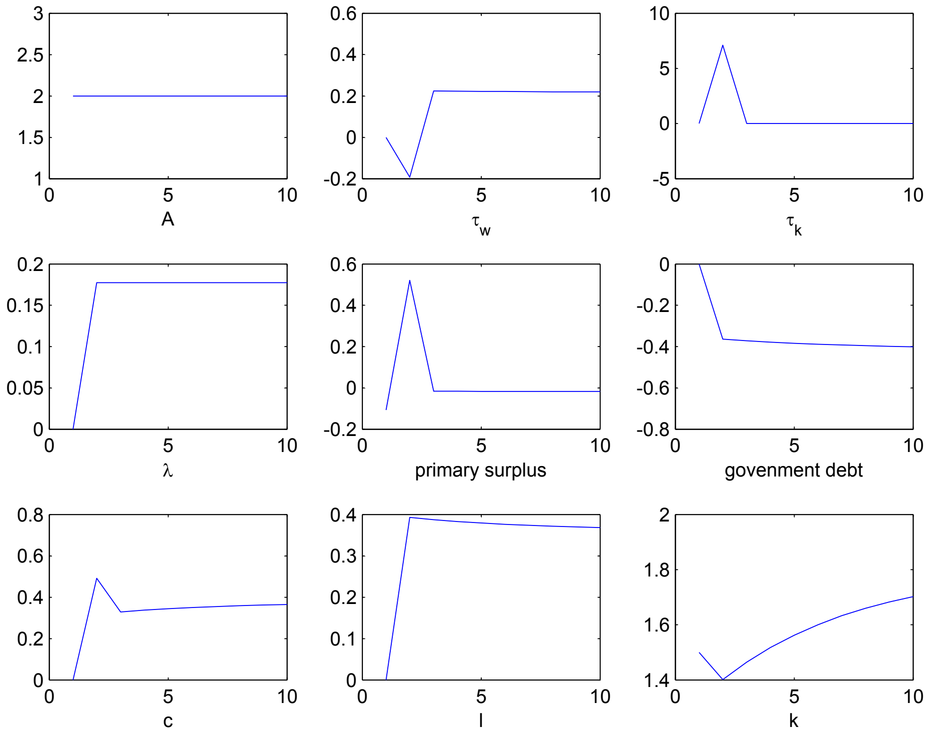

Initial times.

Note: initial $b = 0$

Figure 4.8: Initial periods with $b = 1.3$

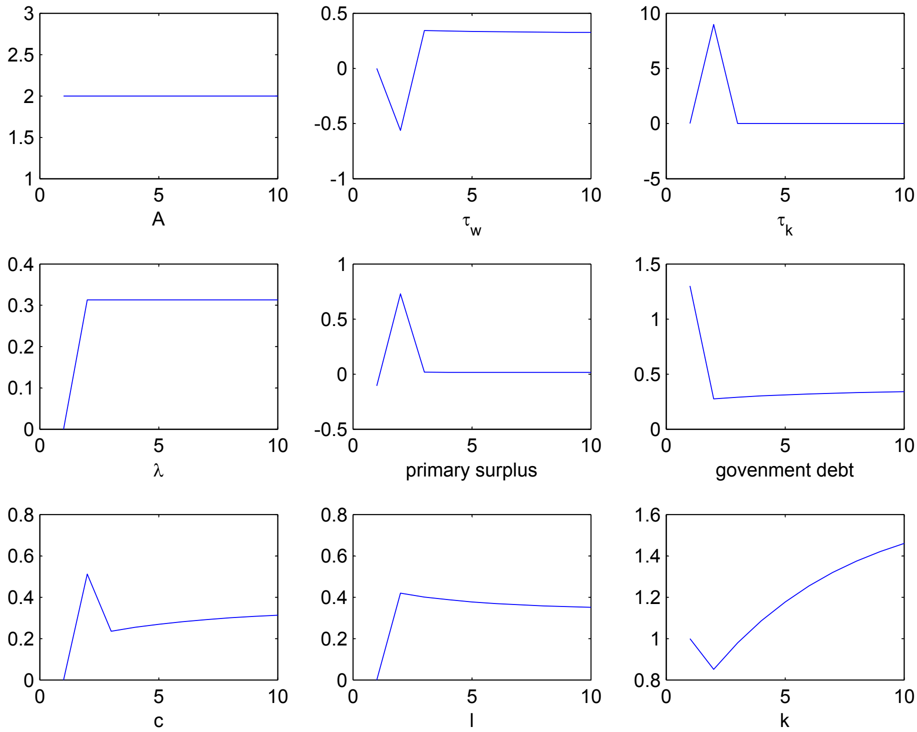

Initial times.

Note: initial $b = 1.3$

Figure 4.9: Long run with initial $b = 0$

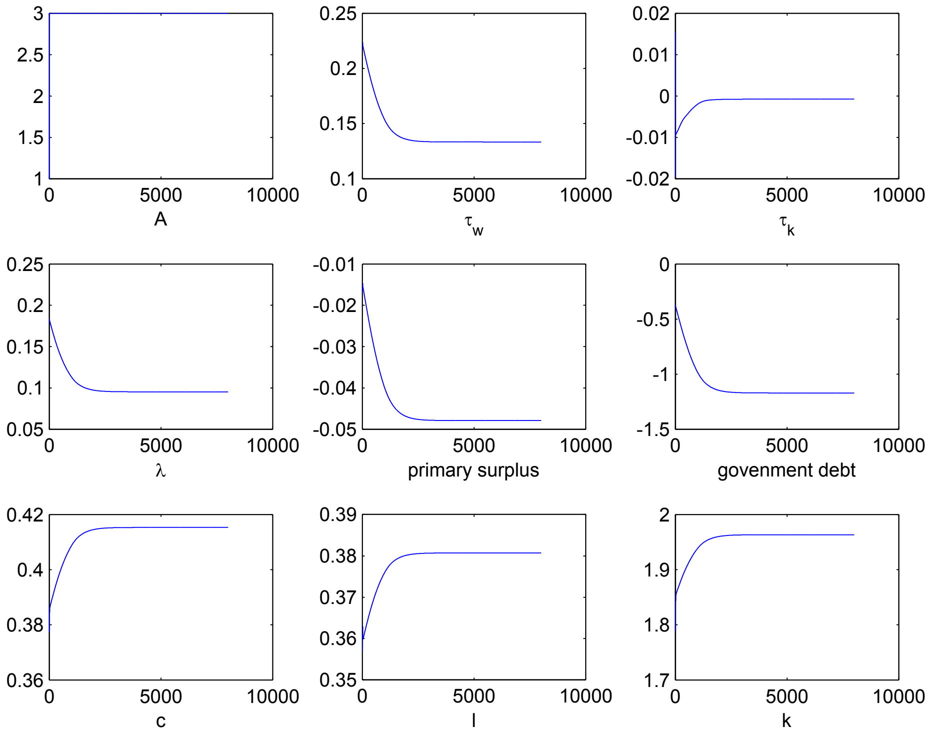

Long run with protracted positive shocks.

Note: initial $b = 0$

Figure 4.10: Long run with initial $b = 1.3$

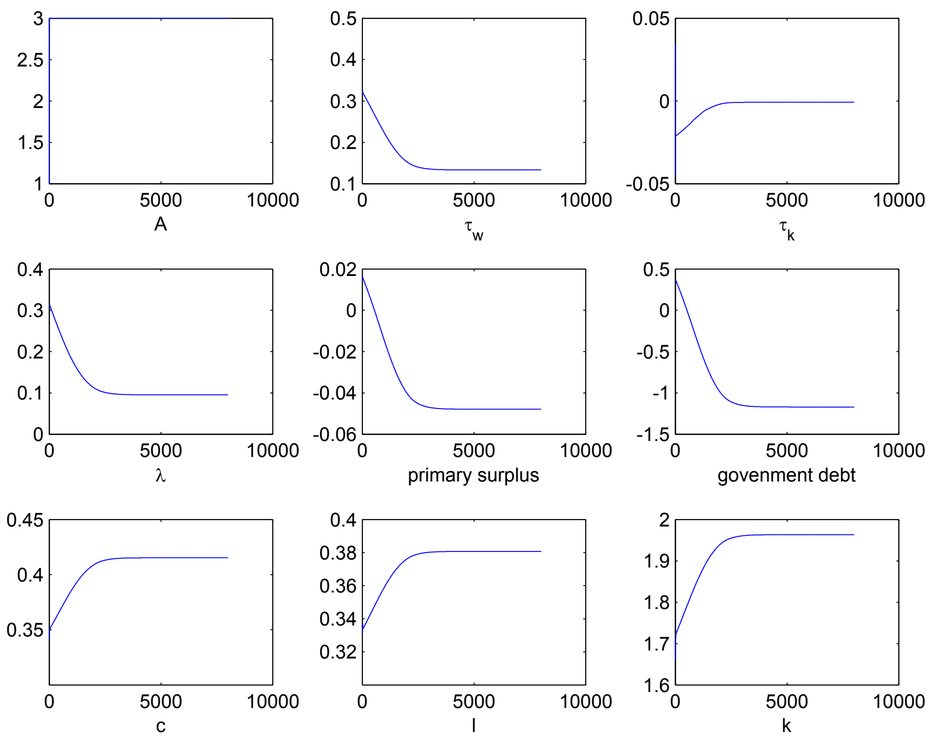

Long run with protracted positive shocks.

Note: initial $b = 1.3$

Chapter 5

Conclusion

The goal of my research is to contribute to the understanding of the causes of the business cycle, its welfare consequences and its policy implications.

In light of the pressing agenda on fiscal policy prompted by the current crisis, my work on fiscal policy could be of immediate interest for policy makers interested in designing an appropriate fiscal policy in the aftermath of the crisis. The change in timing introduces the key trade-off at the center of the current policy debate between short-run benefits form deficit and subsequent costs. Perhaps the absence of this trade-off prevented this literature from being of any use in the current policy debate. The paper shows how the results found by the previous literature that allowed for capital taxation, which where counterfactual and against common wisdom, only depend on the timing typically assumed in the neoclassical growth model; these results are not inherent to the neoclassical growth model, nor to the optimality criterium adopted - Ramsey taxation. By showing this, the paper might help at making this literature of more direct applicability to practical issues.

An interesting question that comes from this research is the quantitative amount of short run capital elasticity present in actual economies. This quantitative investigation would be crucial to predict the optimal responses of capital taxation over the business cycle. One virtue of this study is that, being grounded on the neoclassical growth framework, it is able to simultaneously study the short run and the long run. This makes it possible to address the related issue of what the average government debt should be in the long run.

My positive study, which relates the labour force composition to the business cycle, may contribute to enhancing our understanding of how shocks that hit groups within the labour force, such us particular sectors, and other groups identified by age, education, occupation or gender, propagate to the entire economy. This could shed light on how the business cycle actually comes about. This knowledge would help in identifying the groups that have a stronger impact on the business cycle and its costs for each group in the population. This may permit the design of targeted policies, more effective and less expensive than aggregate policies designed for the entire population.

A key assumption made in the paper was that groups specific labour elasticities are unaffected by changes in the labour force composition. This assumption leads to a benchmark that disentangles the immediate role that changes in the composition of labour has on changes in aggregate volatility, abstracting from the endogenous reaction of labour supply. This permitted to merge the structural neoclassical framework, with a thorough investigation of micro data in relation to the business cycle.

Perhaps the most obvious extension of this work would be to solve a general equilibrium model and see how the endogenous changes in labour supply amplify or mitigate the results. One difficulty with this approach is that it involves solving a rich overlapping generations model with agents differentiated by gender education and occupation. This leads to a model with a large number of state variables, that may only be solved with perturbation methods, which are accurate around the steady state of the economy, but may be inaccurate along the transition, which is the focus of this paper. To this aim, I am currently working on an algorithm that allows to use perturbation methods around points along the equilibrium path which can be far from the steady state. This may allow for a higher degree of accuracy along the transition path.

The work on the measurement of technological shocks shows how simply allowing for curvature in the transformation frontier between consumption and investment goods, improves the performance of the neoclassical growth model on some key targets such as the correlation between consumption and investments. As it is shown in the paper, it has some important implication for the size of the business cycle that can be accounted for by technology shocks.

This paper may have immediate applications relevant for policy making: if the research community reacts positively to the proposed mechanism, then it could easily be incorporated into a fully-fledged Dynamic Stochastic General Equilibrium model that has a direct use for policy analysis, such as the ones used by central banks to support their monetary policy decision making.

Bibliography

Aiyagari, A., A. Marcet, T. Sargent, and J. Seppala (2002). Optimal taxation without state contingent debt. *Journal of Political Economy 110*, 1220–1254.

Arias, A., G. D. Hansen, and L. E. Ohanian (2007). Why have business cycle fluctuations become less volatile? *Economic Theory*.

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

Attanasio, O. and G. Weber (1995). Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey. *Journal of Political Economy*.

Barro, R. (1990, October). Government spending in a simple model of endogenous growth. *Journal of Political Economy*.

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

Bernanke, B. S. (2004). Output fluctuations in the united states: What has changed since the early 1980s. *Remarks by Governor Ben S. Bernanke At the meetings of the Eastern Economic Association, Washington, DC February 20.*

Blanchard, O. and C. Kahn (1980). The solution of linear difference models under rational expectations. *Econometrica 48(5).*

Blanchard, O. and J. Simon (2001). The long and large decline in u.s. output volatility. *Brookings Papers on Economic Activity*.

Boldrin, M., L. Christiano, and J. Fisher (2001). Habit persistence, assets returns and the business cycle. *American Economic Review 91.*

Carvalho, V. and X. Gabaix (2010). The great diversification and its undoing. *Unpublished manuscript*.

Castro, R. and D. Coen-Pirani (2008). Why have aggregate skilled hours become so cyclical since the mid-1980s? *International Economic Review*.

Chamley, C. (1986). Optimal taxation of capital income in general equilibrium with infinite lives. *Econometrica 54*(3), 607–622.

Chari, V. V., L. J. Christiano, and P. J. Kehoe (1994). Optimal fiscal policy in a business cycle model. *Journal of Political Economy 102*(4), 617–652.

Chari, V. V. and P. J. Kehoe (1999). Optimal fiscal and monetary policy. In J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, Volume 1C of *Handbooks in Economics*, Chapter 26. North-Holland: Elsevier Science, North-Holland.

Choi, C. S.-K. and V. Rios-Rull (2009). Understanding the dynamics of labor share: The role of noncompetitive factor prices. *Working Paper*.

Cummins, J. and G. Violante (2002, April). Investment-specific technical change in the u.s. (1947-2000): Measurement and macroeconomic consequences. *Review of Economic Dynamics 5(2).*

Davis, S. and J. A. Kahn (2008). Interpreting the great moderation: Changes in the volatility of economic activity at the macro and micro levels. *NBER Working Papers 14048, National Bureau of Economic Research, Inc..*

Den Haan, W. and A. Marcet (1990). Solving the stochastic growth model by parameterized expectations. *Journal of Business and Economic Statistics 8.*

Diamond, P. A. (2011). When a nobel prize isnt enough. *The New York Times 5 May.*

Erosa, A. and M. Gervais (2001). Optimal taxation in infinitely-lived agent and overlapping generations models: A review. *Federal Reserve Bank of Richmond Economic Quarterly 87*(2), 23–44.

Farhi, E. (2005). Capital taxation and ownership when markets are incomplete. *MIT mimeo*.

Fisher, J. (2006, June). The dynamic effects of neutral and investment-specific technology shocks. *Journal of Political Economy 114(3).*

Freeman, R. B. and L. F.Katz (1994). Rising wage inequality: The united states vs. other advanced countries. In R. B. Freeman (Ed.), *Working under different rules.*

Fuentes-Albero, C., M. Kryshko, J. V. Ríos Rull, R. Santaeulàlia-Llopis, and F. Schorfheide (2009). Methods versus substance: Measuring the effects of technology shocks on hours. *NBER Working Paper 15375.*

Gordon, R. (1990). The measurement of durable goods prices. *University of Chicago Press*.

Greenwood, J., Z. Hercowitz, and G.Huffman (1988, June). Investment, capacity utilization and the business cycle. *American Economic Review 78.*

Greenwood, J., Z. Hercowitz, and P. Krusell (1997). Long-run implications of investment-specific technological change. *American Economic Review 87(3).*

Greenwood, J., Z. Hercowitz, and P. Krusell (2000). The role of invest-ment speci.c technological change in the business cycle. *European Economic Review*.

Greenwood, J., A. Seshadri, and M. Yorukoglu (2005). Engines of liberation. *Review of Economic Studies 72(1)*.

Guerrieri, L., D. Henderson, and J. Kim (2009). Interpreting investment-specific technology shocks. *mimeo, Federal Reserve Board*.

Hansen, G. D. (1985). Indivisible labor and the business cycle. *Journal of Monetary Economics*.

Heathcote, J., K. Storesletten, and G.L.Violante (2010). The macroeconomic implications of rising wage inequality in the united states. *Journal of Political Economy, forthcoming*.

Huffman, G. and M. Wynne (1999). The role of intratemporal adjustment costs in a multisector economy. *Journal of Monetary Economics 43(2)*.

Jaimovich, N., S. Pruitt, and H. E. Siu (2009). The demand for youth: Implications for the hours volatility puzzle. *NBER Working Paper*.

Jaimovich, N. and S. Rebelo (2009). Can news about the future drive the business cycle? *American Economic Review 99(4)*.

Jaimovich, N. and H. E. Siu (2009). The young, the old, and the restless: Demographics and business cycle volatility. *American Economic Review*.

Johnson, M. and M. Keane (2007). A dynamic equilibrium model of the u.s. wage structure, 19681996. *Mimeo, Arizona State University*.

Judd, K., L. Maliar, and S. Maliar (2009, August). Numerically stable stochastic simulation approaches for solving dynamic economic models. *NBER Working Paper 15296*.

Justiniano, A., G. Primiceri, and A. Tambalotti (2008). Investment shocks and business cycles. *Federal Reserve Bank of New York Staff Reports 322*.

Justiniano, A., G. Primiceri, and A. Tambalotti (2009). Investment shocks and the relative price of investment. *CEPR 7597*.

Katz, L. F. and D. H. Autor (1999). Changes in the wage structure and earnings inequality. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Handbooks in Economics, pp. 1463–1555. North-Holland: Elsevier Science.

Kim, C.-J. and C. Nelson (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, 608 616*.

Klein, P. and P. Gomme (2011). Second-order approximation of dynamic models without the use of tensors. *Journal of Economic Dynamics and Control*.

Krusell, P., L. E. Ohanian, J. V. Rìos-Rull, and G. L. Violante (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*.

Lucas, J. R. E. (1980). Equilibrium in a pure currecy economy. *In Karaken and Wallace (Eds.), Models of Monetary Economies*.

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

Marcet, A. and R. Marimon (1995). Recursive contracts. Unpublished manuscript.

Marcet, A. and A. Scott (2009, March). Debt and deficit fluctuations and the structure of bond markets. *Journal of Economic Theory 144*(2), 473–501.

Marimon, R. and F. Zilibotti (1998). Actual versus virtual employment in europe is spain different? *European Economic Review 42(1).*

McConnell, M. and G. Perez-Quiros (2000). Output fluctuations in the united states: What has changed since the early 1980s. *American Economic Review 90, No. 5, 1464-1476.*

Nelson, R. and C. Plosser (1982). Trends and random walks in macroeconomic time series. *Journal of Monetary Economics*.

Nicolini, J. P. (1998). More on the time consistency of monetary policy. *Journal of Monetary Economics*.

Owyang, M., J. Piger, and H. J. Wall (2008). A state-level analysis of the great moderation. *Regional Science and Urban Economics*.

Ravn, M. and H. Uhlig (2002). On adjusting the hp-filter for the frequency of observations. *Review of Economics and Statistics*.

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

Schmitt-Grohe, S. and M. Uribe (2008). What's news in business cycles. *National Bureau of Economic Research NBER Working Papers 14215.*

Scott, A. (2007, April). Optimal taxation and OECD labor taxes. *Journal of Monetary Economics 54*(3), 925–944.

Smets, F. and R. Wouters (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review*.

Stock, H. and M. W. Watson (2002). Has the business cycle changed and why? *NBER Macroeconomics Annual*.

Stock, J. H. and M. W. Watson (2005). Understanding changes in international business cycle dynamics. *Journal of the European Economic Association 3*(5), 968–1006.

Svensson, L. E. O. (1985). Money and asset prices in a cash-in-advance economy. *Journal of Political Economy 93*(5).

Zhu, X. (1992). Optimal fiscal policy in a stochastic growth model. *Journal of Economic Theory 58*(2), 250–289.