4
Quantitative Analysis

This section assesses quantitatively the equilibrium effects of public employment on welfare. The algorithm used to solve numerically for the stationary recursive equilibrium is standard. We use value function iterations to solve the household problem and a variant of the algorithm suggested by Imrohoroglu et al. (1999), augmented with an extra loop to pin down, for each θ, the value of $\underline{z}(\theta)$ that implies $\lambda(\theta)$ type-θ public employees.

4.1.
Calibration

We calibrate the model to match some characteristics of the Brazilian economy. Whenever we calibrate a parameter to target a specific aggregate variable, we consider its annual average for the period between 2000 and 2009. We also use the 2005 *Pesquisa Nacional por Amostra de Domicílios* (PNAD) – an annual cross-sectional household data survey – to calibrate the parameters associated with the distributions of workers and wages. In Appendix C we describe the sample of workers we use to tabulate these distributions.

4.1.1.
Demography

We assume agents are born (i.e. enter the labor market) with 25 years old. They may live up to the age of 80, when death is certain. Each period corresponds to a five years interval, so that $T = 12$. The agents retire at the age of 65, that is $Tr = 9$. We calculate the age-dependent probability of survival π_t from mortality data provided by the *Instituto Brasileiro de Geografia e Estatística* (IBGE) – the government department responsible for collecting data and processing official statistics.[1]

4.1.2.
Productivity and Public Wage Setting

In order to specify the productivity process, one must proxy the level of human capital θ with an observable variable. In particular, we proxy θ by the degree of education an individual acquired before entering the labor market. We consider three levels of θ: (1) at most 10 years of schooling, which includes basic education and incomplete secondary education; (2) between 11 and 14 years of schooling, which includes secondary education and incomplete college education; and (3) those with at least 15 years of schooling, which includes college education. The distribution of θ is obtained from the PNAD. In particular, we calculate the share of workers in each education group: $\mu_{\theta_1} = 0.59$ (basic or no education), $\mu_{\theta_2} = 0.31$ (secondary education), and $\mu_{\theta_3} = 0.10$ (college education).

Recall that the productivity profile in the private sector is given by:

$$q_y(t,\theta,z) = \exp\left\{\gamma_1^y t + \gamma_2^y t^2 + \gamma_3^y(\theta) + \gamma_4^y(z)\right\}.$$

Under the assumption that markets behave competitively, by using data at the individual level on wages, experience in the labor market and schooling, obtained from the PNAD, we estimate γ_1^y = 0.124, γ_2^y = -0.009, $\gamma_3^y(\theta_1)$ = 0, $\gamma_3^y(\theta_2)$ = 0.53 and $\gamma_3^y(\theta_3)$ = 1.47. The estimation procedure is described in Appendix C.

However, even in a competitive equilibrium, the government may not reward productivity. Hence, an analogous estimation procedure for public workers does not represent their productivity profile. Instead, we interpret it as the wage-setting rule in the public sector, given by:

$$\hat{q}_g(t,\theta,z) = \exp\left\{\hat{\gamma}_1^g t + \hat{\gamma}_2^g t^2 + \hat{\gamma}_3^g(\theta) + \hat{\gamma}_4^g(z)\right\}$$

Using data from the PNAD, we estimate $\hat{\gamma}_1^g$ = 0.048, $\hat{\gamma}_2^g$ = -0.005, $\hat{\gamma}_3^g(\theta_1)$ = 0, $\hat{\gamma}_3^g(\theta_2)$ = 0.54, and $\hat{\gamma}_3^g(\theta_3)$ = 1.24. It remains to specify $\hat{\gamma}_4^g(\cdot)$, to which we turn later.

In the absence of a good strategy to estimate the productivity profile in the public sector, we assume that productivity profiles are the same in both sectors

[1] In particular, $\pi_t \in \{1, 0.991, 0.990, 0.987, 0.982, 0.975, 0.964, 0.948, 0.927, 0.895, 0.844, 0.775\}$.

but the government does not remunerate productivity. That is, $q_g(t,\theta,z)$ $= q_y(t,\theta,z)$. We acknowledge this is an extreme assumption. Hence, we check sensitivity by reporting results when the productivity profile varies across sectors and the government remunerates productivity. That is, $q_g(t,\theta,z) = \hat{q}_g(t,\theta,z)$. In practice, reality should be in between these extremes scenarios.

4.1.3.
Idiosyncratic Risk

The Markov process $\Pi(z',z)$ follows from an approximation of an AR(1) process in logs:

$z' = \rho z + \varepsilon$, where $\varepsilon \sim N(0,\sigma^2)$

In Brazil, due to the lack of a household panel data survey, such as the Panel Study of Income Dynamics in the U.S., we cannot estimate ρ properly. As an alternative strategy, we set ρ = 0.82 based on evidence for the U.S. economy,[2] and then we use the distribution of residual wages in the private sector to estimate σ^2 = 0.17. The estimation procedure is described in Appendix C. Importantly, since we use residual wages, the idiosyncratic risk does not absorb some permanent components of the actual productivity profile that are not properly modeled in this paper.

We use Rouwenhorst's (1995) algorithm with 17 states to approximate this AR(1) process using a Markov chain. We assume that the initial distribution of the idiosyncratic risk is the invariant distribution associated with this Markov chain. The Rouwenhorst (1995) method has a property that is useful to define $\hat{\gamma}_4^g(\cdot)$, i.e. the function that maps the risk z into public wages. In particular, the transition matrix associated with the Markov chain does not depend on the variance of the AR(1) process. Hence, by reducing σ, the values of the states get more compressed, but the transition probabilities remain the same.

Many empirical studies have found that wages in the public sector are more compressed and less dispersed than their counterparts in the private sector. Hence, these empirical regularities can be captured by associating $\hat{\gamma}_4^g(z_i)$, $i = 1,\ldots,$

[2] The literature estimates this process to be very persistent. Flodén and Lindé (2001), for example, estimate ρ = 0.91, whereas French (2005) estimates ρ = 0.98 using annual data. Since a period in the model encompasses five years, we set $\rho = 0.96^5$.

n, with the i-th state generated by the Rouwenhorst (1995)'s algorithm applied to an AR(1) process with the same persistence ρ = 0.82 but a smaller standard deviation than σ, say $\hat{\sigma}$.[3] As the states get more compressed, whoever draws a low (high) z in the private sector might be paid more (less) in the public sector. Hence, the possibility to enter the public sector is a source of insurance in this economy. As described in Appendix C, we use the distribution of residual wages in the public sector to estimate $\hat{\sigma}$. In particular, we find that $\hat{\sigma}^2$ = 0.12, which corresponds to 71 percent of its counterpart in the private sector, σ^2.

4.1.4.
Preferences and Private Production

We set the coefficient of relative risk aversion γ at 2.5, which is within the range used in the literature. In addition, we set β to match the annual ratio of capital to output of 3, which is obtained from national accounts provided by the IBGE. The capital share α in Brazil is around 0.4 (e.g. Paes and Bugarin (2006)). The productivity of public goods ξ is set to 0.1. In the absence of a consensus on the magnitude of this coefficient, with estimates ranging from zero (e.g. Holtz-Eakin (1994)) to 0.2 (e.g. Lynde and Richmond (1993)), we perform sensitivity analysis on ξ. Finally, δ_y is set to match the annual ratio of investment to capital of 0.05, obtained from national accounts provided by the IBGE.

4.1.5.
Public Sector

The production function in the public sector is calibrated as follows. We set δ_g to match the annual ratio of public investment to public capital of 0.04, which is obtained from the national accounts provided by the IBGE. Since public goods are not tradable in the market, their value is proxied by the IBGE through information on production costs. In particular, the ratio of public goods to output is 0.14. We normalize A_g to match this figure. In the absence of information on η, which is the parameter in the public production technology, we set it equal to its counterpart in the private sector α, which is 0.4. We later perform sensitivity analysis on η.

[3] Notice that in an extreme scenario in which $\hat{\sigma} = 0$, $\hat{\gamma}_4^g(\cdot)$ becomes constant.

We follow Pereira and Ferreira (2010) to calibrate some tax instruments. In particular, by using data on tax revenues and macroeconomic variables, we calculate the average consumption, labor income and capital tax rates, which are $\tau_c = 0.23$, $\tau_h = 0.21$ and $\tau_a = 0.14$, respectively. We follow the tax code to set the tax rate on bequests τ_{beq} at 0.04 and the contribution to the pension system τ_{ss} at 0.11, whereas the flat benefit b is set to match the pension deficits as a percentage of output, obtained from the *Ministério da Previdência e Assistência Social* – the government branch responsible for managing the pension system.

The ratios of public investment I_g to output, lump-sum transfers Υ to output and debt D to output are set to 2.2, 8.4 and 47 percent, respectively. These figures are provided by the IBGE and the National Treasury. Note that public consumption C_g is left free to balance the government budget.

Finally, we consider parameters related to public employment and wage policies. The public wage rate is set to match the ratio of the public wage bill to the private wage bill, i.e. $w_g H_g / w_y H_y = 0.3$, provided by the IBGE. Recall that $\lambda(\theta_1) + \lambda(\theta_2) + \lambda(\theta_3)$ is the share of public workers, which is 13.5 percent according to the PNAD. Hence, it remains to calibrate $\lambda(\theta_1)$ and $\lambda(\theta_2)$ to match the share of public workers with basic or no education (i.e. 27 percent) and secondary education (i.e. 45 percent). These figures are also obtained from the PNAD.

4.1.6. Summary

Table 2 summarizes the values assigned to internally calibrated parameters.

Parameters	Target variable	Data	Model
$\beta = 0.85$	annual K_y / Y	3	2.96
$\delta_y = 0.23$	annual I_y / K_y	0.05	0.05
$A_g = 0.74$	G/Y	0.14	0.13
$\delta_g = 0.18$	annual I_g / K_g	0.04	0.04
$w_g = 0.44$	$w_g H_g / w_y H_y$	0.30	0.31
$b = 0.29$	pension deficits / Y	0.014	0.014

Table 2: Internally calibrated parameters.

The model is also able to generate some statistics, other than targeted variables, that represent the Brazilian economy during the 2000s. The Gini coefficient for earnings, for instance, is 0.48 in the calibrated model, which is close to 0.53, calculated with data from the PNAD. The ratio of the average wage paid in the public sector to the average wage paid in the private sector is 1.99 in the calibrated model, which is close to 1.77, also calculated with data from the PNAD. The calibrated model also generates a high Gini coefficient for wealth, 0.72, and a high annual interest rate, 7.6 percent, which characterize the Brazilian economy during the 2000s.

4.2. Results

This section reports the results. First, we discuss whether the model is able to replicate some dimensions of the distribution of public workers across age and education groups. Second, we study the welfare implications of different public employment policies. Finally, we perform some sensitivity analysis.

4.2.1. Public Employment

The main objective of this paper is to study the welfare effects of public employment accounting for its role in improving the insurance degree in the economy. Hence, it is desirable that the model replicates some features in the data associated with public employment. Due to data availability, we only consider the distribution across age and education groups.[4]

In Figure 1, we compare the distribution of public workers across age groups in the model against the data, tabulated from the PNAD.

[4] As each period in the model encompasses five years, we fit the data to age intervals. For the age of 25, for example, we group agents who have between 21 and 25 years of age; for the age of 30, we group the agents who have between 26 and 30 years of age; and so on.

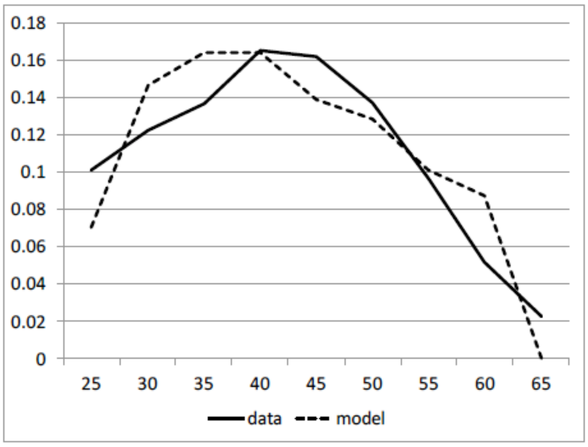

Figure 1: Distribution of public workers across age groups.

The share of workers increases up to a certain age, 35 in the model but 40 in the data, remains nearly flat for one period, and then, declines. Hence, the model can replicate the general pattern of the distribution. However, the model predicts higher shares of both young and old workers in the public sector.

Figure 2 plots de distribution of public workers across age groups for each level of human capital. Notice that the higher share of young workers in the model is partially due to those with college education, whereas the higher share of old workers is due to those with secondary education.

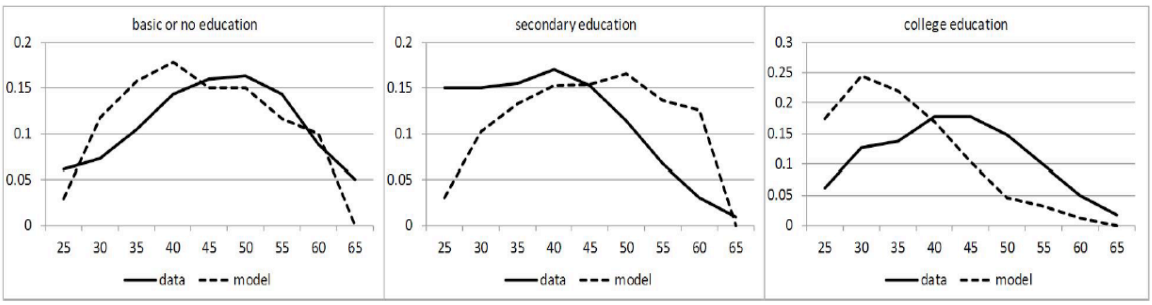

Figure 2: Distribution of public workers across age and education groups.

We conjecture that these discrepancies are due to three facts: (1) the model does not allow retirement at earlier ages, so that the share of old workers is higher than in the data;[5] (2) in Brazil during the 2000s, agents had a more generous pension benefit if they stayed longer in the public sector, so that the share of

[5] See Ferreira and dos Santos (forthcoming) for a model with endogenous choice of retirement.

middle-aged workers is higher in the data;[6] and (3) some public jobs require more previous training and experience (e.g. judges) than others, which might explain a smaller share of young workers with college education in the data. Since the aim of this paper is to study to role of the public sector as insurance provider through a less uncertain wage schedule, we abstract from these features. We conclude that, albeit imperfectly, the model can replicate some dimensions of the distribution of public employment across age groups.

Finally, Figure 3 plots the distribution of private and public workers across levels of idiosyncratic risk z_i, i = 1, …, 17, for each level of education. In particular, the thresholds $\underline{z}(\theta)$ to enter the public sector associated with basic or no education, secondary education and college education are 12, 11 and 8, respectively.

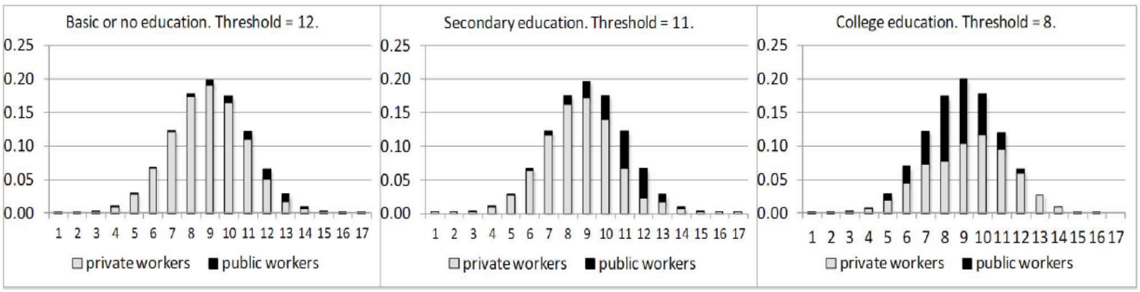

Figure 3: Distribution of public workers across productivity shocks.

Notice that public employment is effective in increasing the welfare of intermediate types. Indeed, most of the agents with very high shocks prefer to work in the private sector, whereas those with low shocks cannot enter the public sector. Moreover, a sizable number of individuals benefits from the fact that they cannot be fired.

4.2.2.
Welfare Effects

This section shows our main results. In particular, we report the welfare implications of different sizes of public employment. Once the government

[6] See Gloom et al. (2009) for a description of the convoluted Brazilian pension system during the 2000s. This paper develops a macroeconomic model to study a reform that induces civil servants to retire later.

changes the size of public employment, it affects the public wage bill and, thus, has to adjust its fiscal policy in order to balance its budget. We consider three types of policy adjustments: (1) consumption taxes τ_c; (2) capital taxes τ_a; and (3) lump-sum transfers Υ.[7]

Results considering a lump-sum tax adjustment should be read with caution. As we argue above, lump-sum transfers capture the role of large welfare programs, which require public workers to operate them. Hence, in practice, exchanging public employment for lump-sum transfers might not be feasible. In contrast, a simple change in the capital or consumption tax rate can be designed without an effective change in public employment.

Recall that in our benchmark calibration, public employment is set to 13.5 percent of the workforce. Figure 4 plots the welfare gains (*y-axis*) against the size of public employment ranging from 2 to 16 percent of the workforce (*x-axis*). If the government tries to hire more than 16 percent of the workforce, it would not be able to fill all open vacancies that require college education.

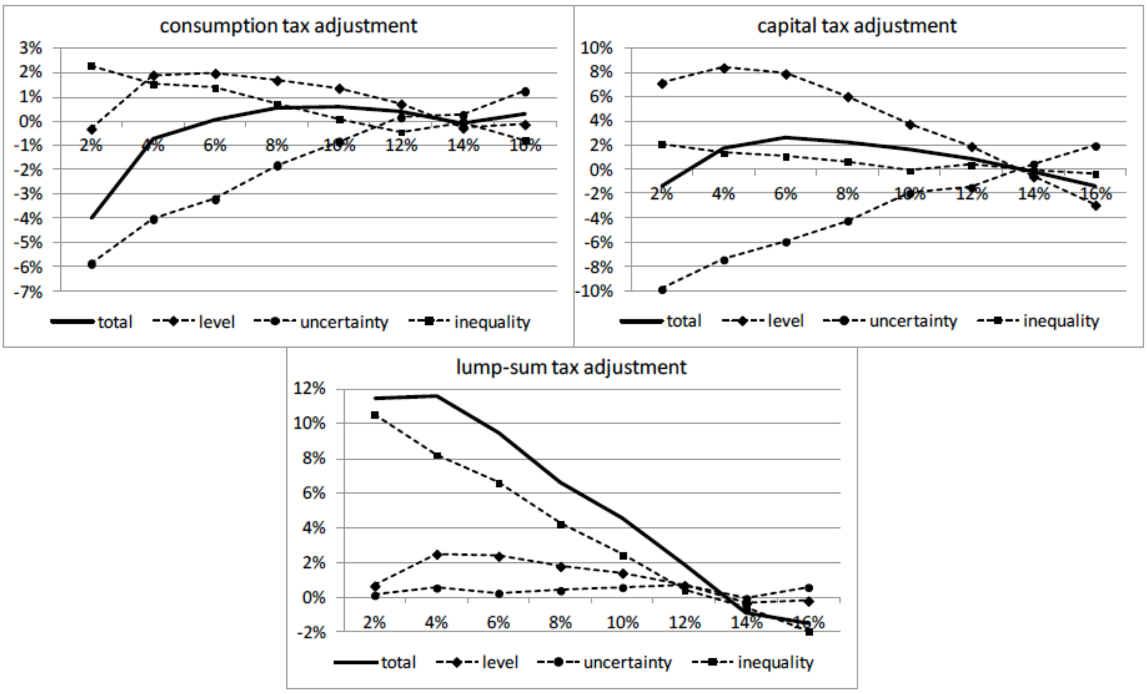

Figure 4: Welfare implications (*x-axis*: share of public workers; *y-axis*: welfare effects).

First, consider the experiment with consumption tax adjustment (top-left

[7] We do not consider income taxes τ_h because labor is supplied inelastically. Hence, adjustments in τ_h are not distortive.

plot). Social welfare is maximized at any share of public employment ranging from 8 to 12 percent, which is associated with a total welfare effect of nearly 0.5 percent. However, welfare losses due to uncertainty increase at a fast pace as public employment drops. In particular, if the government reduced public employment to 8 percent, welfare losses due to uncertainty would be 1.8 percent. These losses are counteracted by welfare gains of 1.7 and 0.7 percent due to level and inequality, respectively, effects.

It is worth noting that the uncertainty effect increases monotonically with the level of public employment. Hence, a larger government is associated with a higher degree of insurance in the economy. In contrast, the inequality effect decreases monotonically with public employment. Intuitively, as Figure 3 highlights, a sizeable government benefits mostly individuals with intermediate levels of productivity. Hence, consumption inequality tends to increase with the size of public employment.

Second, results considering a capital tax adjustment (top-right plot) are qualitatively similar. Quantitatively, the optimal public employment is 6 percent of the workforce, which represents a total welfare gain of 2.6 percent. Reducing the public workforce to the optimal level, however, would bring a welfare loss of 6% associated with the smaller degree of insurance. It is also worth noting that, if compared to the previous case, the welfare gains due to the level effect and the losses due to the uncertainty effect are amplified for all levels of public employment. If the government reduced the public employment from 13.5 percent to 8 percent of the workforce, for example, these figures would be 6 percent and 4.2 percent, respectively.

Finally, welfare gains can be considerably high if the government is allowed to exchange public employment for lump-sum transfers (bottom plot). In this case, the optimal level of public employment ranges from 2 to 4 percent of the workforce, which represents welfare gains of nearly 11.5 percent. These gains are due mainly to the inequality effect. Intuitively, a large public sector benefits mostly individuals with intermediate levels of productivity. Once the government's size becomes smaller, the extra resources obtained from the reduction in the public wage bill are distributed lump sum, which particularly improves the welfare of those agents at the bottom of the consumption distribution. Hence, consumption is distributed from intermediate to low types, increasing social welfare.

Due to the insurance properties of the lump-sum transfers, the uncertainty effect remains fairly constant along all levels of public employment. The decreased insurance due to the smaller government is compensated by an increased insurance from the higher level of lump-sum transfers, so that the overall degree of uncertainty in the economy remains roughly the same.

4.2.3.
Welfare Effects by Education Groups

In this section, we decompose the welfare effects by education groups. In particular, for each education group, we calculate the total welfare and uncertainty effects. The level effect reported in Figure 4, for instance, does not vary across education groups.

Figure 5 plots total welfare and uncertainty effects (*y-axis*) conditional on different levels of θ against the size of public employment ranging from 2 to 16 percent of the workforce (*x-axis*).

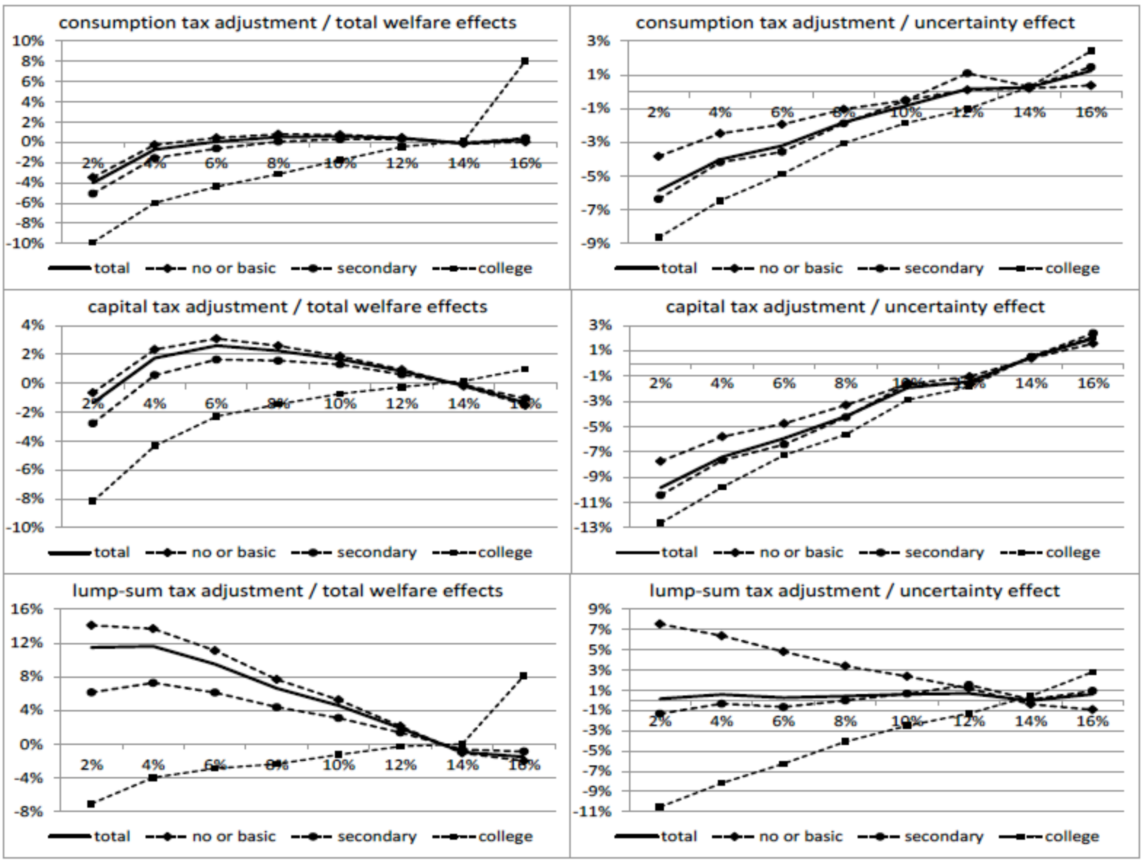

Figure 5: Welfare implications by education groups (*x-axis*: share of public workers; *y-*

axis: welfare effects).

Notice that, independent of the tax instrument used to balance the budget, total welfare gains for college graduates increase monotonically with the size of the government (left-plots). Moreover, individuals with college education are the group most benefited by an increase in public employment. This result is due to the fact that (i) the government hires proportionately more college graduates[8] and (ii) the public sector provides a more effective insurance scheme, through public wages, to this group (right-plots). Indeed, as Figure 3 highlights, the productivity threshold for a college graduate to enter the public sector is the 8th highest possible realization of the idiosyncratic risk z, whereas this threshold for an individual with no or basic (secondary) education is the 12th (11th) highest possible realization.[9]

4.3.
Sensitivity Analysis

In this section, we check whether our results are sensitive to: (i) a different productivity profile in the public sector; (ii) different values of ξ, which capture the productivity of public goods; (iii) different values of η, which capture the productivity of public capital. In all cases, we recalibrate the model to match the targets in Table 2.

4.3.1 Different Productivity Profiles

In the absence of a good strategy to estimate the productivity profile in the public sector, our benchmark results consider an extreme case in which productivity profiles are the same in both sectors but the government does not remunerate productivity. That is, $q_g(t,\theta,z) = q_y(t,\theta,z)$ for all t, θ, z. In this section, we assume that the productivity profile varies across sectors and government remunerates productivity. That is, $q_g(t,\theta,z) = \hat{q}_g(t,\theta,z)$ for all t, θ, z. In practice, reality should be in between these extremes scenarios.

In a stationary equilibrium, $q_g(t,\theta,z)$ only affects the aggregate level of

---

[8] 37 percent of workers with college education work in the public sector, versus 20 percent and 6 percent of workers with secondary and basic education, respectively.

[9] Recall that we consider n = 17 possible realizations of the idiosyncratic risk.

efficient labor units employed at the public sector H_g. Indeed, since we normalize the public total factor productivity A_g to match the ratio of public goods to product G/Y we observe in the data, any reduction in H_g is absorbed by an increase in A_g.[10] Hence, except for A_g and H_g, the stationary equilibrium has the same properties in both this and the benchmark cases. Figure 6 shows the welfare implications under this alternative productivity profile. It plots the welfare gains (*y-axis*) against the size of public employment ranging from 2 to 16 percent of the workforce (*x-axis*).

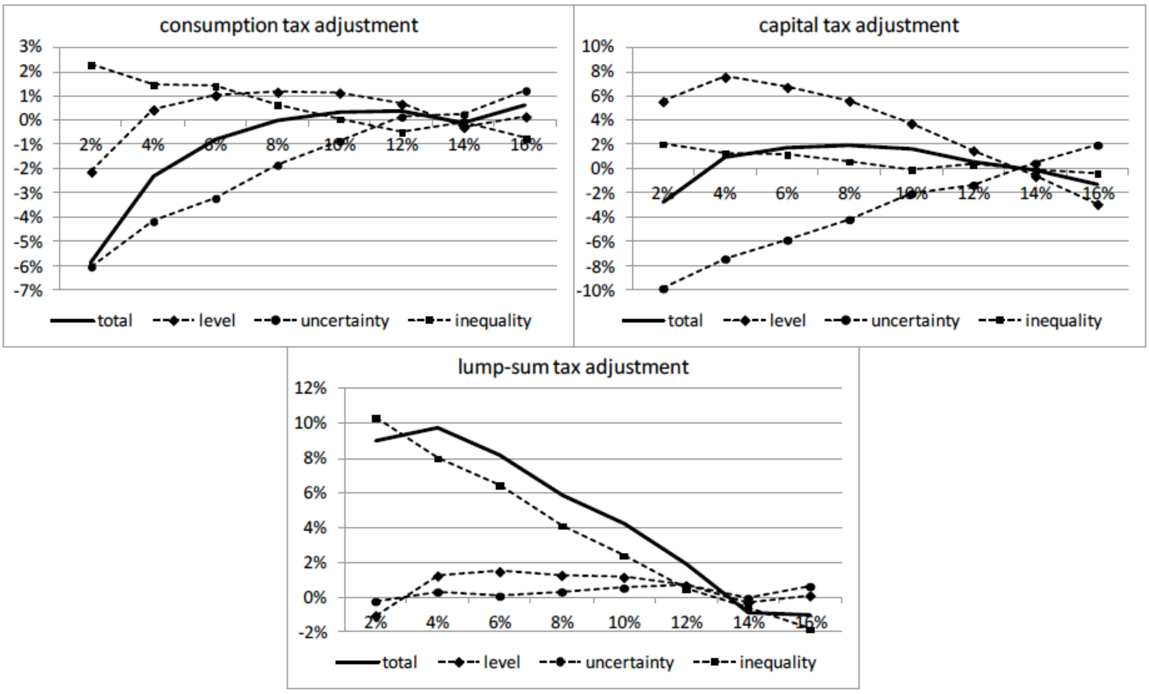

Figure 6: Welfare implications (*x-axis*: share of public workers; *y-axis*: welfare effects). Alternative productivity profile.

A comparison between Figures 4 and 6 shows that welfare implications due to both uncertainty and inequality effects are almost the same as in the benchmark case. Since agents are remunerated according to $w_g\hat{q}_g(t,\theta,z)$ when working in the public sector, the productive profile $q_g(t,\theta,z)$ does not enter directly in their optimization problems. However, they depend indirectly on $q_g(t,\theta,z)$, as it may affect prices r and w and thresholds $\underline{z}(\theta)$ through A_g and H_g. Hence, we conclude that general equilibrium and selection effects are not strong enough to modify the welfare implications due to uncertainty or inequality.

---

[10] In particular A_g increases from 0.74 to 0.90

In contrast, welfare gains due to the level effect are smaller in this case. Intuitively, once the government reduces public employment, workers leave a relatively more productive public sector than in the previous case, which mitigates welfare gains. Nonetheless, the difference of level effects is not that large. For example, if the government reduced public employment to 8 percent of the workforce, the level effect would be nearly 0.5 percent higher in the benchmark case independent of the tax instrument used to balance its budget.

4.3.2 Productivity / Production of Public Goods (ξ and η)

In this section, we analyze the role of ξ, which governs the productivity of public goods, and η, which governs the productivity of public capital.

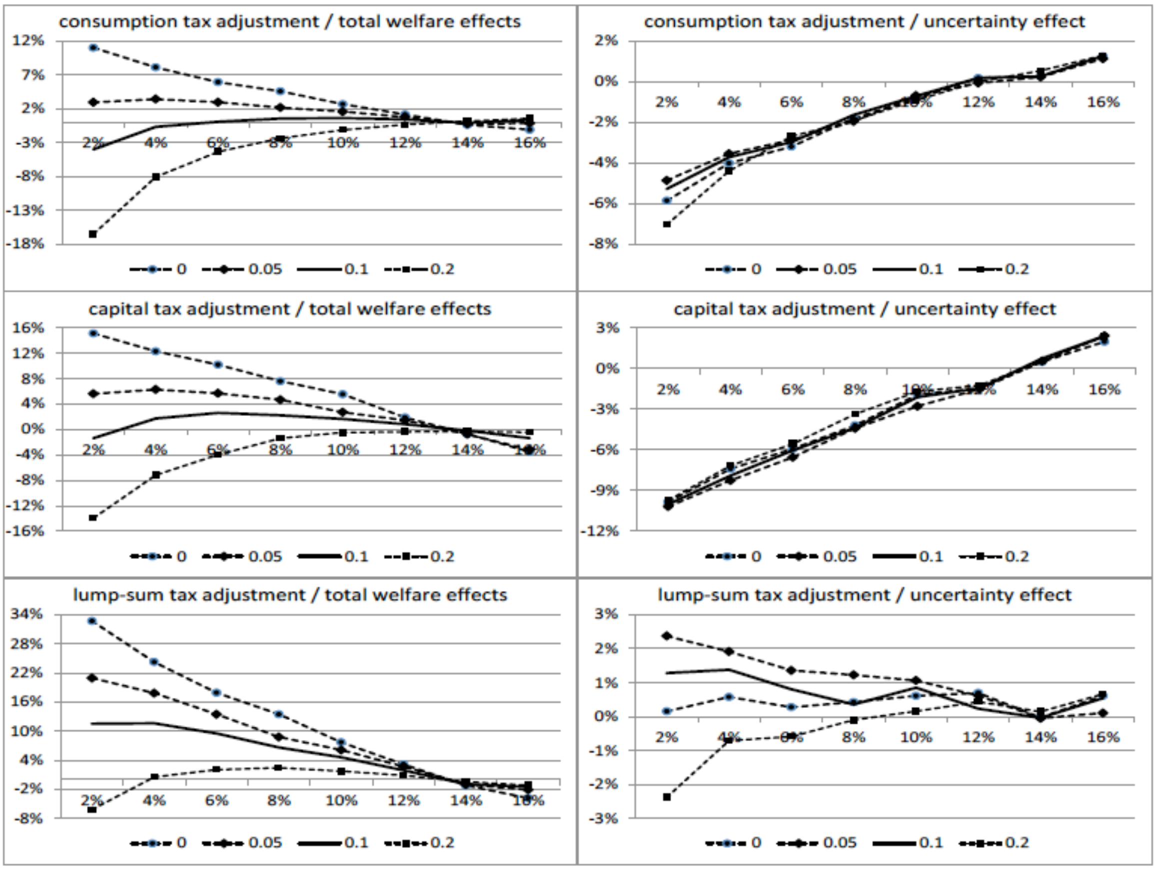

Figure 7: Welfare implications (*x-axis*: share of public workers; *y-axis*: welfare effects). Alternative values for ξ.

Figure 7 plots total welfare and uncertainty effects (*y-axis*) for different values of ξ against the size of public employment ranging from 2 to 16 percent of the workforce (*x-axis*). In particular, we consider $\xi = 0$, $\xi = 0.05$ and $\xi = 0.2$.

Recall that the benchmark value of ξ is set at 0.1. Figure 8 plots total welfare and uncertainty effects (*y-axis*) for different values of η against the size of public employment ranging from 2 to 16 percent of the workforce (*x-axis*). In particular, we consider η = 0.3 and η = 0.5. Recall that the benchmark value of η is set at 0.4.

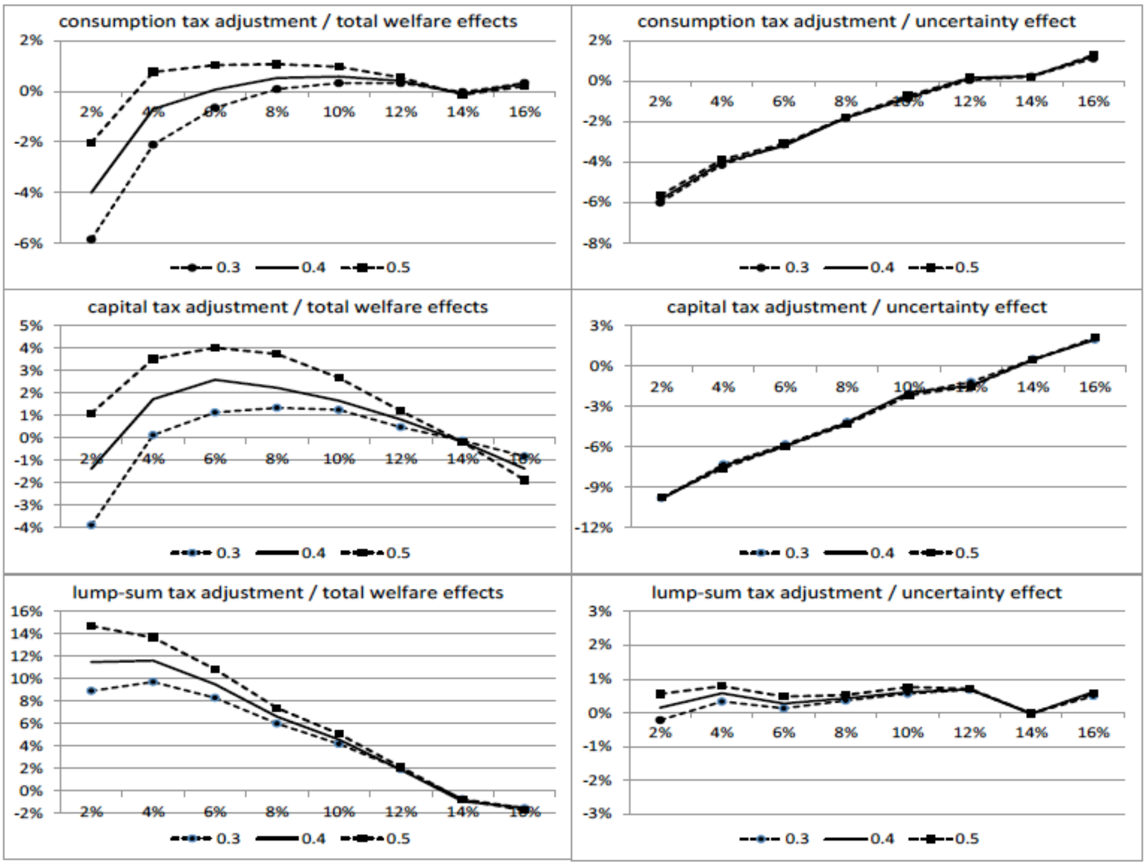

Figure 8: Welfare implications (*x-axis*: share of public workers; *y-axis*: welfare effects). Alternative values for η.

As in the previous section, the level effect explains most of the changes in welfare gains due to different values of ξ or η (not shown in Figures 7 and 8). A similar intuition applies. Moreover, as Figures 7 and 8 highlight, except for the lump-sum tax adjustment case, the uncertainty effect due to a smaller size of the government does not change much as we vary ξ or η.

We conclude from these sensitivity analyses that, although a misspecification of the technology associated with the public sector might bias social welfare evaluations, its part due to the uncertainty effect is fairly robust to misspecification.