

THE ECONOMIC IMPACT OF RISING THE RETIREMENT AGE: LESSONS FROM THE SEPTEMBER 1993 LAW*

Pedro Martins**

Álvaro Novo***

Pedro Portugal***

1. INTRODUCTION

In most developed countries, pension systems have been facing longer life spans and lower birth rates. Almost without exception, political decision makers have responded by increasing the legal retirement age (the age when workers are entitled to collect their full retirement pension); by increasing the contributions to the pension funds; and by restricting the access to early retirement.¹

Despite its importance, there is very little research about the economic impact of these policies. This essay investigates the legal change that occurred in the Portuguese labour market when, from 1994 onwards, the retirement age of women (initially set at 62) was gradually adjusted in order to converge to the retirement age of men (65 throughout). We exploit the richness of the individual records of “*Inquérito ao Emprego*” and “*Quadros de Pessoa*”, which also allow one to follow individuals and their firms.

In this context, it is particularly insightful to draw on matching estimators, which can establish an accurate comparison between the “treatment” groups (those individuals or firms subject to an intervention) and the “control” group. We follow these procedures, in order to estimate the impact of the increased retirement age of women upon their wages and hours worked. We also obtain estimates of the impact of the reform upon the personnel policies (hiring and firing) of firms. Finally, we also evaluate the impact of the higher retirement age upon the performance of the firms.

2. THE NEW LEGAL RETIREMENT AGE

In the early 1990s, the Portuguese social security system was facing the sustainability problems common in pay-as-you-go systems. These problems were related to the ageing of the population, as life expectancy increased and as birth falls fell. In 1993, those aged 65 or more corresponded to 21.6% of the active population (Banco de Portugal, 1994). In this context, the Portuguese government decided to raise the legal age of retirement (LAR) of women from 62 to 65 years, making that equal between the two sexes (“*Decreto-Lei*” 329/93). This law raised the retirement age by six months in every civil year, up to 1999, when the two retirement ages converged (see Table 1). For instance, while a woman born in 31st December 1939 would be entitled to retire in 31st December 1993 (on the day of her 62nd birthday), a woman born one day later (1st January 1932) would only be eligible to collect her full pension six months later, at 1st July 1994, when she would be 62 years and 6 months old. However, given the

* The analyses, opinions and findings of this article represent the views of the authors, they are not necessarily those of the Banco de Portugal. The authors are grateful to Mário Centeno for his comments and suggestions.

** Queen Mary, University of London.

*** Economics and Research Department, Banco de Portugal.

(1) Fourteen OECD countries have recently increased their retirement age and are gradually approaching that new level.

Table 1

TREATMENT GROUPS: BEFORE AND AFTER THE NEW LEGAL AGE OF RETIREMENT (LAR)							
Year:	Treatment groups						
	1992	1994	1995	1996	1997	1998	1999
LAR:	62	62.5	63	63.5	64	64.5	65
	[57.5, 58)						[64.5, 65)
	[58, 58.5)					[64, 64.5)	
	[58.5, 59)				[63.5, 64)		
	[59, 59.5)			[63, 63.5)			
	[59.5, 60)		[62.5, 63)				
	[60, 60.5)	[62, 62.5)					

gradual implementation of the increase in the LAR, women aged 6 months later (1st July 1933), would only reach their retirement age on the 1st July 1995 (when aged 63). The following analyses exploit these gradual adjustments in the law as a source of identification of the impact of the increased LAR upon the labour market.

There are two additional aspect of the pension system in Portugal that need to be taken into account. The first one is that LAR refers to the age at which the worker is entitled to his/her retirement pension. At that point in time, the labour contract between that worker and his/her employer expires, although the worker can establish a new contract with the same (or a different) employer. The retirement earnings and the work earnings (in a possible new contract) are independent. The second additional aspect is that the social security system included some exceptions to the standard LAR, as in early retirement. The most conspicuous cases concerned the long-term unemployed, workers in firms undergoing major restructuring processes, and workers in particularly demanding jobs (e.g. air-traffic controllers), thus limiting the impact of the new law.

3. IDENTIFICATION AND ESTIMATION

Given the non-experimental nature of the law reform, a quality of our evaluation depends crucially on the quality of the groups used for the construction of counterfactuals. The following analysis selects carefully control groups (counterfactuals) and follows simultaneously two methodologies that have been suggested in the literature on non-experimental identification: difference-in-differences (Meyer, 1995) and matching (Rubin, 1977, and Rosenbaum and Rubin, 1983). These two methodologies are combined in a single one: difference-in-differences matching (DDM).

3.1. Statistical methodology

Let Y_{it}^D be the potential value for individual i in period t if in state D , in which $D=1$ if exposed to treatment and $D=0$ if not. Assume that treatment occurs in period t . The fundamental identification problem is that it is impossible to observe in period t the value of individual i in both states. It is therefore impossible to estimate the individual effect of treatment, $Y_{it}^1 - Y_{it}^0$. It is possible, however, to estimate the average impact of treatment on the treated if there is a suitable control group.

The idea behind the difference-in-differences estimator is that once can use information from a group of individuals that were not subject to the treatment in order to identify the time change in Y that is not due to treatment, i.e. that would result from the simple time difference. The identification hypothesis in this method can be described as

$$E[Y_{it}^0 - Y_{it}^1 | D=1] = E[Y_{it}^0 - Y_{it}^1 | D=0],$$

in which t' is the period before the implementation of the program. This hypothesis states that, throughout time, the evolution in Y for the treated individuals ($D=1$), in the case in which they were not treated, would have been the same than the evolution observed for the individuals that were not exposed to the treatment ($D=0$). The difference-in-differences estimator provides estimates of the mean effect of treatment on the treated and can be obtained from the sample moments in

$$DdD = \{ E[Y_{it} | D=1] - E[Y_{it} | D=0] \} - \{ E[Y_{it'} | D=1] - E[Y_{it'} | D=0] \}.$$

The limitations of this estimator are related to how comparable the two groups are in terms of their observed characteristics. However, it is possible to combine this method with the matching method suggested by Rubin (1977), in order to make sure that the groups share common characteristics by removing from the sample those individuals whose characteristics are not common to the two groups (only those in the common support are kept). The method that combines these two methodologies, suggested by Heckman *et al.* (1997, 1998) is known as difference-in-differences matching. The feasibility of this identification strategy depends on the availability of a rich set of individual characteristics in the data, which in our case are present at both the individual- and the firm-level.

In the case of panel data, the DDM estimator takes the following form:

$$MDD = E[(Y_t^1 - Y_t^0) - \hat{E}(Y_t^0 - Y_{t'}^0 | P)],$$

in which $\hat{E}(\cdot | P)$ represents the expected value of the time change in Y for the individuals in the control group that are statistically close to those in the treatment group (i.e. whose probability of participation in the treatment, P , is also high). This probability is computed by a probit model (dichotomous dependent variables), and conditioning the dependent variable (undertaking or not the treatment) upon individual observed characteristics. It is possible to show that, if the selection to treatment is independent on the potential result of the treatment, conditionally on the observed characteristics, then conditioning on P is equivalent, but computationally simpler. In practical terms, the first step of the procedure involves estimating for each individual in the sample the difference in the behaviour over time (from t' to t), so to compute the first differences (separately for the treatment and the control groups). In the second step, one estimates P , following which one matches units in the treatment group to units in the control group, in order to compute the (second) difference amongst comparable units. The mean of these differences represents the mean impact of treatment upon the treated group.

4. DATA

The analysis of the impact of the new law exploits two data sets that allow one to follow workers and firms over time. The impact upon the inactivity transition rates is based on the “*Inquérito ao Emprego*” data set, from *INE* (Statistics Portugal). In this data set it is possible to follow workers over the 6 quarters in which they are surveyed. The impact upon the wages and hours worked of the targeted women and upon the flows of workers and firm performance are based on the “*Quadros de Pessoa*” data set (Ministry of Employment). The latter data set includes matched information about the worker (e.g. age, schooling, tenure) and the firm (e.g. equity, sales, industry, firm size) over time, so that one can study

the impact of the new law at both the individual- and the firm-level, controlling for the characteristics of each unit under analysis.

5. THE NEW LAW AND LABOUR MARKET PARTICIPATION

The extension of the length of activity for women is the most immediate impact of the new law. In fact, the data confirm that the probability that a woman in the treatment group is employed increased by 31.1% with respect to the group not affected by the law. Symmetrically, the probability that such woman is inactive fell by 27.9%.² These values are further corroborated by the developments in the employment rates of women aged 62-65, from 23.2% in 1992 to 30.4% in 2000.³ From these figures, one can conclude that the new law had a significant impact upon the activity level of the women targeted by the law. Absent any other effects, the new law contributed towards the sustainability of social security. However, it is possible that those firms that held these targeted women in their workforces adjusted their personnel policies, namely in terms of their pay, hirings and separations.

6. THE IMPACT UPON WAGES AND HOURS WORKED

6.1. Treatment and control groups

Strictly speaking, the new law affected all women aged 62 or less. However, the law affected some women more immediately than others, in such a way that one can expect that those women played a larger role in the consequences of the reform. Amongst such women, one can include those that would have retired in $t+1$ if the LAR had been unchanged at t . For instance, women aged [60; 60.5) by the end of 1992 would have retired in 1994 under the old LAR (62); however, as LAR increased to 62.5 by then, these women had to postpone their retirement to 1995. All women aged [55; 60.5) at the end of 1992 had to postpone their retirement during the period 1994-1999 (1999 being the year in which LARs were equalised for the two genders). Those women will correspond to our treatment group. Table 1 describes as different age groups in 1992 were affected differently by the law throughout the period in which the LAR converged completely.

The control group is composed of men at the same age as the women included in the treatment group. The LAR for men was already 65 when the new law became effective in the 1994-1999 period. Therefore, the two groups are comparable in this age dimension. On the other hand, it is clear that the choice of this control group can raise questions in terms of the comparability across genders. Such a problem is, however, mitigated to the extent that one is willing to accept the hypothesis of time invariance of the difference-in-differences estimator. In other words, if the differences across the two genders are constant throughout the period under analysis, then focusing on men as the counterfactual for women is no longer an issue. In fact, the data support this hypothesis: between 1991 and 1993, the difference in the logarithm of wages between men and women was 0.39, 0.39 and 0.38, respectively, and 0.098, 0.093 and 0.100 for the logarithm of hours worked. We therefore conclude that our control group is a legitimate one.

(2) The results reported are obtained from the coefficient identifying the impact of the treatment in a multivariate logit model. See Pedro Martins, Álvaro Novo and Pedro Portugal, 2007, "Increasing the legal retirement age: The impact on workers' wages, hours, worker flows and firm performance", *mimeo*, Banco de Portugal.

(3) As a benchmark, the activity rate of women aged 15-64 is around 61.5% in the same period. For women aged 55-64, the activity rate in 1992 was 35.5% (66.2% for men), increasing to 41.8% (64.4%) in 2000. The difference-in-differences method controls for the difference in the time trends that are observed prior to the years in which the new law was introduced.

In order to implement the difference-in-differences method, one also needs to define the “before” and “after” periods. One should bear in mind that, beginning in 1994, women’s LAR increased six months for each civil year, up to 65 in 1999. Therefore, there are two obvious choices for the “before” period: 1992 and 1993. Our preference for 1992 is based on the fact that the new law was promulgated in 1993, so that some individuals and firms may have reacted before 1994, for instance through early retirement procedures; moreover, the government policy was not known in 1992, thus making it less likely that individuals’ behaviour was affected on that year.

6.2. The mean individual impact of the increase in LAR

Given the non-experimental nature of the event under analysis, we begin by focusing our attention on the quality of the matching procedure, which is crucial for the validity of the reported effects. As argued previously, this is particularly important in a non-experimental environment, which in the present case is even more important given the gender differences. In order to assess the quality of the matching, Table 2 reports the mean value for a set of variables (characteristics) used in order to estimate the participation probability (i.e. to be targeted by the new law). One can expect that, before the matching, there are (statistical) differences between the treatment and control groups. In fact, the first lines in Table 2 confirm that there are differences between the two groups.⁴ One should notice that, after the control group units are restricted to those share (under a statistical metric) the same probability of participation upon treatment, the differences disappear and one cannot any longer find any differences in the mean characteristics between the two groups. This procedure ensures that the comparability of the two groups was achieved, so that one can attribute any differences in behaviour to the impact of the treatment. There may still be differences in unobserved variables, but these will also be eliminated (controlled for), in this case by the difference-in-differences method, assuming those are constant over the period under analysis.

Table 3 presents a new set of DDM estimates concerning the effects of the new law upon the group of women targeted in terms of their pay and hours worked. The global assessment of the results suggests that the impact of the increase in the LAR is very weak in terms of these labour market variables. In statistical terms, none of the estimated impacts concerning wages and hours worked is significant, while in economic terms the impact is also very minor.

Choosing men as a counterfactual comparison group may raise some criticism. In order to assess the sensitivity of the estimates to the definition of the control group, we consider two alternative control groups. The most obvious choice for the control group would be a comparison of women with other women. However, this is problematic, as any woman aged less than 62 in 1993 was affected by the new LAR. In this context, a natural choice would be to consider as a control group those women aged above 62 in 1993 and that are still employed although they were already entitled to retire with a full pension under the earlier law. Yet another choice would be to take into account as the control group those women that were not forced to postpone their retirement from 1994 to 1999, although this may raise endogeneity issues. The last two columns in Table 3 present our results. The conclusions do not seem to depend on the specific definition of the control group. Neither pay nor hours worked are significantly affected by the postponement of the retirement age.

(4) These are the lines corresponding to the title “Unmatched” in the “Sample” columns.

Table 2

AVERAGE CHARACTERISTIC VALUES FOR SAMPLE BEFORE AND AFTER MATCHING PROCEDURE				
Variable	Sample	Group		p-value ^(a)
		Treatment	Control	
Experience	Unmatched	52.44	52.19	0.000
	Matched	52.43	52.34	0.163
Experience ²	Unmatched	2759.2	2734.4	0.000
	Matched	2758.4	2749.3	0.162
Tenure	Unmatched	15.58	17.80	0.000
	Matched	15.60	15.51	0.704
Tenure ²	Unmatched	381.6	484.9	0.000
	Matched	382.4	382.5	0.994
Sales (in logs)	Unmatched	7.02	7.78	0.000
	Matched	7.03	7.09	0.227
Education	High school	Unmatched	0.03	0.093
		Matched	0.03	0.739
	College degree	Unmatched	0.03	0.001
		Matched	0.03	0.402
Year	1994	Unmatched	0.17	0.191
		Matched	0.17	0.471
	1995	Unmatched	0.20	0.603
		Matched	0.20	0.765
	1996	Unmatched	0.14	0.006
		Matched	0.14	0.746
	1997	Unmatched	0.18	0.071
		Matched	0.18	0.824
	1998	Unmatched	0.15	0.246
		Matched	0.15	0.818
	1999	Unmatched	0.17	0.107
		Matched	0.17	0.909

Source: *Quadros de Pessoal*, with authors' computations.

Notes: Unmatched: The reported values use the total sample of individuals before the matching procedure. Matched: The reported values for treatment and control groups use only individuals matched in the probability of participating in treatment (using the kernel matching algorithm). This table omits the value for the dummy variable for sector of activity and district. (a) A p-value greater than 0.05 indicates that the difference of average value between the two groups is not statistically significant at the 5 percent level.

Table 3

LABOR MARKET IMPACT: WAGES AND HOURS WORKED			
Variable	Men	Women +62 years	Women [50, 55]
Log(Wages)	0.008	-0.023	-0.003
	(0.011)	(0.022)	(0.011)
	10 204	4 953	6 788
Log(Hours worked)	0.009	0.006	-0.010
	(0.009)	(0.016)	(0.009)
	9 823	5 041	6 850

Source: *Quadros de Pessoal*, with authors' computations.

Notes: The reported values for each variable are: point estimate, standard deviation (in parentheses) and number of observations. The estimation method is difference-in-differences matching for longitudinal data. The treatment group consists of women that in 1992 belonged to the age group [57.5, 60.5]. The control groups consist of, respectively, men in the same age group, women aged 62 or more years old and women aged 50 to 55 years old in 1992.

7. THE IMPACT UPON PERSONNEL POLICIES AND FIRM PERFORMANCE

Although the new law had an immediate impact upon the increase in female labour supply, as driven by the extension of the range of ages in which women are active, it is important to assess any effects in terms of firms' labour demand and production decisions. One area of particular interest concerns firms' personnel policies, namely hirings and separations/firings. In fact, the increase in the LAR, in a context of severe restrictions to the adjustment of employment levels – as is abundantly known for the Portuguese case – may significantly block the flow of workers.

7.1. Treatment and control groups

In order to examine separations, hirings and the net flow of workers at the firm level, we considered as the treatment group the set of firms that employed at least one female worker aged [55; 60.5) in 1992. Once the control group firms are identified – those that did not employ any woman in that age group –, the effect of the new law was assessed over a period of five years. In other words, the estimates of the effects of the new law in terms of hirings (or other flows) are based upon the sum of all hirings (or other flows) observed between 1994 and 1999.

7.2. Mean impact of the increase in the LAR upon firms

As before, the comparison between treatment and control group firms is established by drawing on a matching method based on the probability that each firm belongs to the treatment group. These estimates are then used in order to establish the matching between treatment group firms and those control groups firms that most closely resemble the former. This method seeks to restrict the comparison between the two groups to a set of firms that are effectively comparable in terms of a desirably long set of observed characteristics. The variables considered for the matching include the firm size, five qualitative variables for different firm size categories, the percentage of female workers, the mean total salary, the mean hours of work, the percentage of workers that are men aged 60 or above, the percentage of the equity of the firm held by foreign investors and by domestic private investors, 57 qualitative variables for different industries and 29 qualitative variables referring to the geographical units in which the firms are located.⁵

From the estimation of the mean effect of the increase of the LAR, it is found that the accumulated effect over the five years (1994-1999) lead to a retention of about 1.6 workers by the treatment-group firms (which employed, on average, 1.2 women affected by the law). This fall in the level of separations induced a decrease of 2.6 hirings (see Table 4). From the joint effect of the change in these flows, there is a net loss of one employee when one compares the evolution of the level of employment in the treatment group firms with the equivalent evolution in the control group firms. The reduction in hirings is particularly strong amongst younger workers, particularly females. From the decomposition of the hiring flows it is found that about two thirds of the decrease in hirings affected women aged 25 or less (see Martins, Novo and Portugal, 2007).

There is also evidence of a slight but significant decrease in the level of sales in firms that employed workers affected by the increase in the LAR, once a comparison with the equivalent change in the level

(5) Some variables are included as polynomials of the second or third degree. The data set includes only firms with less than 100 employees, since it proved particularly difficult to find large firms that did not employ a single woman worker affected by the new law.

Table 4

AVERAGE TREATMENT EFFECT ON TREATED FIRMS: WORKERS FLOW, SALES AND SALES PER WORKER					
Flow	Period	Impact	Standard deviation	Number of observations	
				Treatment group	Control group
Hiring (number of workers)	1995-99	-2.59	0.49	4048	40197
Separations (number of workers)	1995-99	-1.64	0.40	4048	40197
Net job creation (number of workers)	1995-99	-0.95	0.25	4048	40197
Sales (thousands of euros)	1995-99	-12.63	2.74	5310	57330
Sales per worker (thousands of euros)	1995-99	-0.04	0.17	5359	57104

Source: *Quadros de Pessoal, 1991 to 1999 with authors' computations.*

Note: Point estimates based on difference-in-differences matching with kernel matching.

of sales is established (Table 4). This result may be interpreted as a consequence of a scale effect in production related to the increase in labour costs unleashed by the forced retention of older female workers. However, this result may not have consequences in terms of the average firm productivity, when calculated as the ratio between sales and firm size, since firm size has also fallen. All in all, these results seem to be consistent with the evidence that the firms affected did not allow that the increase in the LRA affected their decision regarding their wage bills.

8. CONCLUSIONS

This essay studied the microeconomic impact of the increase in the legal retirement age of women that took place in September 1993, using statistical techniques that allow one to establish an appropriate comparison between the treatment group and a group of individuals that exhibit similar characteristics, except that they were not affected by the new law (the control group).

While evidence that the new law was binding is obtained, based on the result of a significant increase in the employment rate of women aged between 62 and 65, there is no evidence that wages or hours worked changed.

However, the increase in the legal retirement age decreased significantly the hiring of new workers, particularly young women. Over an horizon of five years, for each female worker affected by the increase in the retirement age, up to two new workers were not hired as a consequence. There was also evidence that the level of sales fell, while average sales per worker has not changed, since net job creation was also negative.

The evidence about a net loss of employment amongst the firms affected may imply, absent any general equilibrium effects that counteract upon the decrease in hirings, that the positive effect upon the sustainability of social security may have been attenuated by a decrease in the amount of social security revenues and by an increase in the amount of unemployment benefits for younger workers.

In the institutional setting of the Portuguese labour market, stimulating employment amongst the older workers ("active ageing") may be, in part, counteracted by the rigidity of the wage determination mechanisms, especially through the widespread usage of tenure-related pay increases, and by employment protection, which in the Portuguese case corresponds to the high firing costs.

REFERENCES

- Banco de Portugal (1994), *Annual Report, Banco de Portugal*.
- Heckman, J., Ichimura, H., Smith, J. & Todd, P. (1998), "Characterizing selection bias using experimental data", *Econometrica* 66(5), 1017–1098.
- Heckman, J., Ichimura, H. & Todd, P. (1997), "Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme", *Review of Economic Studies* 64(4), 605–654.
- Martins, Pedro, Álvaro A. Novo e Pedro Portugal (2007) "Increasing the legal retirement age: The impact on workers' wages, hours, worker flows and firm performance", *mimeo*, Banco de Portugal.
- Meyer, B. D. (1995), "Natural and quasi-experiments in economics", *Journal of Business & Economic Statistics* 13(2), 151–162.
- Rosenbaum, P. & Rubin, D. (1983), "The central role of the propensity score in observational studies for causal effects", *Biometrika* 70(1), 41–55.
- Rubin, D. (1977), "Assignment to a treatment group on the basis of a covariate", *Journal of Educational Statistics* 2(1), 1–26.