Research Institute for the Evaluation of Public Polices

http://irvapp.fbk.eu

FBK-IRVAPP Working Paper series

Research Institute for the Evaluation of Public Policies

Firm-Level Effects of Reductions in Working Hours

Marta C. Lopes

Alessandro Tondini

July 2022

FBK-IRVAPP Working Paper No. 2022-05

Firm-Level Effects of Reductions in Working Hours

Marta C. Lopes
Universidad Carlos III de Madrid and IZA
macardos@eco.uc3m.es

Alessandro Tondini
FBK-IRVAPP
atondini@irvapp.it

FBK-IRVAPP Working Paper No. 2022-05

July 2022

Research Institute for the Evaluation of Public Policies
Bruno Kessler Foundation
Vicolo dalla Piccola 2, 38122 Trento (Italy)

Phone: (+39) 0461.314209
Fax: (+39) 0461.314240

E-mail: info@irvapp.it
Website: http://irvapp.fbk.eu

The purpose of the IRVAPP Working Papers series is to promote the circulation of working papers prepared within the Institute or presented in IRVAPP seminars by outside researchers with the aim of stimulating comments and suggestions. Updated reviews of the papers are available in the Reprint Series, if published, or directly at the IRVAPP.

The views expressed in the articles are those of the authors and do not involve the responsibility of the Institute.

Firm-Level Effects of Reductions in Working Hours *

Abstract

How do legislative reductions in hours impact firms? In this paper, we use matched employer-employee data to evaluate a policy reform in Portugal that unexpectedly reduced the usual weekly working hours from 44 to 40 hours. Using a difference-in-differences approach that exploits initial heterogeneity across collective agreements covering workers, we show that the reform led to a significant drop in working hours in treated firms, while salaries did not adjust, resulting in higher wages per hour. We observe only a small and insignificant negative effect on employment, as treated firms are able to maintain or even increase sales despite the fall in labor input (total hours worked within a firm). We show that this partly reflects higher prices rather than higher (or constant) volumes, whereby firms are able to shift the higher labor costs onto consumers.

JEL codes: J22, J23, J31
Keywords: working time, hours, wages, labor demand, labor cost

* *Corresponding author*: atondini@irvapp.it. The authors are thankful to all comments and suggestions received from Thomas Crossley, Luc Behaghel, Michele Belot, Andrea Ichino, Arianna Tassinari, John P de New, Jan Kabátek and all other participants at the internal Microeconometrics Working Group of the European University Institute, the Online Workshop Labour Market and Institutions - Multidisciplinary Dialogue, the Melbourne Institute Seminar Series, Carlos III economic department seminar, and the 4th EUI Alumni conference. The usual disclaimer applies.

1 Introduction

Working hours are a key economic variable; if workers and hours are not perfect substitutes, working hours (should) enter directly into the production function (Feldstein (1967)). This implies a direct link between working hours and output, employment and productivity. Consequently, working hours *legislation*, i.e. regulatory limits to the number of hours one can/should work, is likely to have important impacts on economic outcomes. This type of legislation is pervasive across countries, particularly so in Europe, and often debated in the public discussion. Yet, empirical evidence of its effects is still very limited, and in particular little is known of the impact on firms of limits to weekly working hours.

The paper aims to fill this gap by looking at Portugal, a country that offers a unique case study in response to these questions. Using comprehensive matched employer-employee data, we estimate the effects, at the firm level, of a national reduction in weekly working hours from 44 to 40 hours, made in 1996. Importantly, this reform was implemented without compensating measures and impacted most but not all workers and firms.[1] We use initial heterogeneity in workers' exposure to different collective agreements - which can and do set hours lower than the national standard - to define a treatment and control group. This allows us to compare firms more and less affected by the reform. We find a large reduction in working hours in more affected firms, with no adjustment in terms of monthly salary, which leads to a proportional increase in hourly wages. Contrary to the predictions of a standard labor demand model, this does not result in significant employment losses in treated firms. The negative employment effect is lower than -1% up to four years after the reform. Against the prediction of negative scale effects, firms are able to maintain and even significantly increase sales, even if labor input, i.e. the total number of hours worked within a firm, falls significantly. This means that sales per hour, i.e. the output that firms are able to produce with one unit of labor input, increases as much as hours decrease. This finding could reflect two mechanisms at play: in line with diminishing returns, workers become more productive per hour as hours fall. Alternatively, this could also reflect price increases, whereby firms, in a non-perfectly competitive market, are able to shift the

[1] One other key aspect of this reform - which has not been acknowledged by the previous literature on it - is that this reduction occurs while the standard daily hour is set at 8, which implies a discontinuous effect on the number of usual working days from 6 to 5.

cost onto consumers. We use disaggregated sectoral price data from national accounts to show that more treated sectors experience higher growth in prices relative to control sectors. When deflating sales, the effect on output becomes negative and significant, and the effect on sales per hour decreases by one fourth, but remains strongly positive. While our data does not capture the full potential extent of price effects, this shows that part of the effect certainly goes through prices.

Partly because of its salience in the policy debate, the employment effects of a reduction in working hours have received the most attention in the literature. We summarize these findings in a concise way in Table OA1 in the Online Appendix, by reform and level of analysis (workers, firms, sectors, regions). Overall, while there are well-identified but conflicting estimates of the impact of a reduction in working hours on *incumbent* workers, evidence of the effect on labor demand is scarce. Crépon and Kramarz (2002) find that, in France, a reduction in standard time from 40 to 39 hours made in 1982 increased the probability of *incumbent* workers being fired. This does not, by definition, imply that the total employment effect at the firm level is negative, as it does take hirings into account. Using the same approach, Gonzaga et al. (2003) look at the effect of a reduction from 48 to 44 hours in Brazil in 1988, and estimate no effect on job loss, while Raposo and van Ours (2010) find ambiguous effects on the separation rate when Portugal reduced standard working hours from 44 to 40.[2] Between 2001 and 2005 in Chile, Sánchez (2013) finds no impact on employment transitions for a reform that allowed for a 4-year adjustment period. Estevão and Sá (2008) look at aggregate employment in large versus small firms in France, which, while impacted by the reduction to 35 hours at different times, have no visible difference in their employment dynamics.

Empirical evidence at the firm level, which should allow us to capture the overall effect on labor demand, is more scarce and less convincing. Kawaguchi et al. (2017) look at the reduction in standard hours in Japan in the 1990s, but find no significant fist stage overall (i.e. average hours were not significantly impacted). For a subset of firms with a significant drop in hours, they estimate a negative but insignificant employment effect. Crépon et al. (2004) analyze the employment and productivity effects of the French reduction in standard time to 35 hours. They find that firms affected earlier by the change in hours had a relative increase in employment. However, evaluation of the French reform is

[2]Separation rate decreases for workers directly impacted by the reform, but increases for workers indirectly impacted.

made difficult by the simultaneous implementation of important cuts in social security contributions (SSC) meant to ease the transition to the shorter working week. The authors argue that the full relative increase in employment can be explained by the lower labor cost. In this paper, we argue that the case of Portugal, which also experienced a large and sudden reduction in standard time but without compensating measures, is clearer.[3]

Lastly, some studies have focused on the sector and/or regional-level, in an attempt to capture aggregate equilibrium effects not limited to labor demand (firms) or incumbent workers. Hunt (1999) shows that in Germany, in the late 80s and early 90s, sectors that adopted agreements regulating working time experienced a *relative* decrease in employment. Both Skuterud (2007) and Chemin and Wasmer (2009) use regional legislative specificities to capture the effect of a reduction in working hours. Skuterud (2007) shows that, while Quebec (Canada) reduced standard hours from 44 to 40 hours, there has been no positive effect on employment, despite an adjustment in monthly wages (as opposed to most European reforms). Chemin and Wasmer (2009) show that Alsace-Moselle (France), which for historical reasons experienced a relatively smaller reduction in working hours than the rest of France, had similar employment dynamics to other regions after the reform. The only study indicating a positive correlation between employment and a reduction in standard time is Raposo and Van Ours (2010), who show that local labor markets (region $\times$ sector) that were more impacted by the 1996 reform in Portugal subsequently experienced higher employment growth.

Overall, while empirical evidence does not support the work-sharing argument, the literature remains ambiguous on the extent to which reductions in standard hours have a detrimental effect on employment, due to the lack of clear evidence on what happens to firms. In our view, this debate is deeply linked to the discussion about the "elusive" employment effects of minimum wages (Cengiz et al. (2019), Manning (2020)). In a way, a reduction in working hours at a constant monthly salary can be thought of as theoretically akin to a minimum wage increase (with some important differences), as are the conventional theoretical predictions that it should negatively impact employment. However, this prediction is always based on the assumption that labor markets are perfectly competitive, which is not necessarily the case. Firms may adjust on other margins than employment, such as, for example,

[3]Our paper complements preliminary work by Varejao (2005), who looked at how standard hours reduction in Portugal impacted on firms' employment dynamics.

prices (Clemens (2021)). Our results show that firms are able to adjust to the higher labor cost per hour by equally increasing sales per hour, partly through higher prices, which is a result that contradicts the predictions of a standard model and could explain why the employment effect is small and not statistically significant.

Furthermore, there is no evidence, to our knowledge, about how a nation-wide reduction in standard hours impacts output at the firm level and workers' (and hours) productivity. Do we observe a negative scale effect on output? Do workers become more productive (per hour worked) as working hours decrease? In the literature, evidence on the relationship between reductions in standard hours and productivity is virtually absent. A few studies have examined how working hours and productivity are related in very specific occupations: Brachet et al. (2012) look at the performance of paramedics with respect to shift-length, Pencavel (2014) examines this in the context of munition workers in the UK in the 1950s, while Collewet and Sauermann (2017) focus on the case of call center workers in the Netherlands. All these studies find evidence of a linear relationship between hours and output up to a certain point, followed by diminishing marginal returns above said level. Other papers have looked at how part-time employment affects firms' productivity, finding that firms with a higher share of part-time workers tend to be more productive in specific sectors (e.g. pharmacies, Künn-Nelen et al. (2013)), or when part-time employees work more than a certain level of hours (Garnero et al. (2014)). In this paper, we are able to provide evidence that reductions in working hours strongly increase output per hour at the firm level. We argue, however, that the step from output per hour to productivity is not straightforward, as the role of price effects makes it difficult to quantify the real increase in productivity.

The paper proceeds as follows. Section 2 provides an overview of the working time legislation and describes its chronological evolution in Portugal. Section 3 discusses how the theoretical predictions around a reduction in hours differ. Section 4 presents the data that is used in the analysis. Section 5 describes the empirical strategy based on a difference-in-differences approach, while Section 6 presents the results. Section 7 concludes.

2 Working Time Legislation in Portugal

2.1 General Aspects of Working Time Legislation

Working time regulations include all the legislation that limits the number of working hours a worker can work, and regulate the organization of the working week. There are many aspects to working time legislation, including, for example, regulations on night-shifts, weekend work, paid leave, national holidays etc. Arguably, the most relevant are standard hours, overtime and the overtime rate. Standard hours refer to the length of the *usual* working week, i.e. how many hours a worker usually works, and are usually averaged over a certain reference period. In other words, standard hours set the daily and/or weekly limit at which overtime hours begin. Overtime hours are the hours worked on top of standard hours, and are usually also limited (for example, maximum weekly hours are capped at 48 hours in the European Union). Moreover, overtime hours are paid at a higher rate, referred to as the *overtime rate*, which sets the wage increase a worker should earn on each extra hour.

2.2 Working Hours Legislation in Portugal over Time

Working time in Portugal is regulated by both the national legislation and collective agreements that can vary by sector and district, covering approximately 80% of workers. Generally, national legislation sets the upper bound, while collective agreements can specify either lower levels or exceptions. The national legislation concerning working time in Portugal dates back to 1971, when the *Decreto-lei 409/71* set standard working time at 8 hours daily, and 48 hours weekly. It also allowed for 2 daily hours of overtime, and a maximum of 240 per year. The first hour of overtime was paid a 25% premium, and the second a 50% premium. We summarize the many changes in the national legislation that have occurred since then in Table 1.

Yearly overtime was first reduced from 240 to 160 hours in 1983, and the premium increased to 50% for the first hour, and 75% for the second. It is interesting to note that the text of this law explicitly mentioned "work-sharing as the rationale" for this reform.[4] The first significant reform in

[4]"A necessidade de distribuir o trabalho existente pelo maior número possível de trabalhadores impõe que a prestação de trabalho fora do horário normal só seja permitida nos casos em que se mostre necessário" (Decreto-lei 421/83, Portugal, 1983).

Table 1: Working Time Legislation in Portugal, 1971–2003

Year	Standard Hours		Overtime		Overtime Rate
	Daily	Weekly	Daily	Yearly	
1971	8	48	2	240	25%, 50%
1983	8	48	2	**160**	**50%, 75%**
1991	8	**44**	2	**200**	50%, 75%
1996	8	**40**	2	200	50%, 75%
2003	8	40	2	**150-200**	50%, 75%

This information was collected by the authors using the national legislation. *Year* refers to when the law was published in the official gazette, not the date of effective implementation. *Standard Hours* indicates the maximum usual hours specified in the national legislation, both at the daily and weekly level. *Overtime* refers to the maximum number of hours that can be worked on top of standard hours, by paying the *Overtime premium*: the first number refers to the first hour of overtime, and the second to the second hour. All reforms are in bold, see text for some remaining aspects of working time legislation not covered in the Table.

standard time took place in 1991, when standard working hours were reduced from 48 hours weekly to 44, while maximum daily standard hours remained at 8 daily. This was partly compensated by an increase in yearly overtime hours, which were raised from 160 hours to a maximum of 200 per year. *De facto* this reform only decreased hours by around 1 hour, workers were already at 45 hours per week (and not 48), because of the regulation imposing days off and breaks. The reform was passed by national legislation in 1991, but anticipated by a tripartite agreement between the government, unions and the employers' association, which, in 1990, had already implemented this change.

A much larger reform, the focus of this paper, took place in 1996, and was implemented in 1997 to 1998, when standard hours where suddenly reduced from 44 hours to 40. The reform allowed for an adjustment period of 1 year; the limit was first lowered to 42 hours in 1997 and to 40 in 1998. This time the reform did not allow for any increase in overtime, but only for a longer reference period over which to average standard hours. Importantly, contrary to the 1991 legislative change, this reform was the result of the left-wing party winning the elections in 1996 and unilaterally approving a decrease in hours to hasten convergence with other European countries. The law had no specific provision for

what should happen to salaries and wages as hours fell, but, as we will see later, monthly salaries did not adjust and this resulted in significant wage increases. Lastly, it is key to reiterate that, contrary to the French reform, no compensating measures were put in place (such as cuts in social security contributions.) This allows for, we argue, a much cleaner interpretation of the effects of reductions in standard hours at the firm-level.

A key aspect of the reform - which has not been underlined by the previous literature (Raposo and Van Ours (2010), Raposo and van Ours (2010), Lepinteur (2019)) - is that the reduction from 44 hours to 40 occurred while the daily limit was fixed at 8 hours. This implies that for those workers at 8 hours per day, the reform also had a discontinuous effect on the number of working days, which moved from 6 to 5. This was not necessarily the case if workers were working less than 8 hours a day before the reform. However, it is likely that for the majority of workers this implied one half-day less per week. It is important to acknowledge this when interpreting the results: while difficult to include in a model, there is no reason to expect the same effect when a reduction in hours reduces the number of working days than when it does not. Given that much of the recent political debate around working time legislation revolves around the reduction in the number of days of work (rather than hours), we argue that this aspect makes the evidence presented in this paper particularly salient.

Figure 1: Share of Full-Time Workers Working More than 40 Hours, 1994–2000, QP and LFS

(a) QP (Matched Employer-Employee Data)

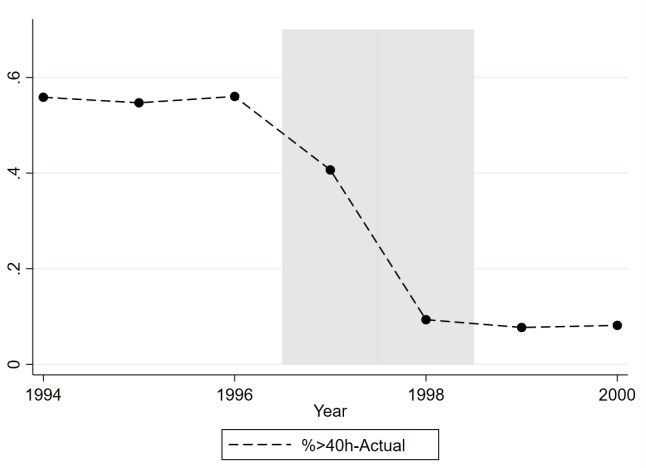

(b) Labour Force Survey

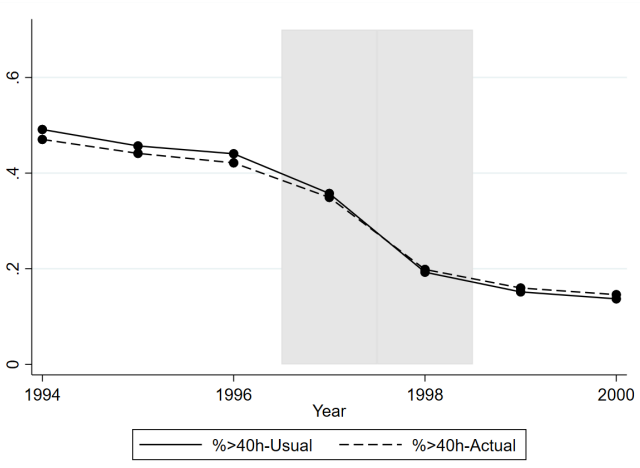

Note: The shaded area indicates the period of the reform. Hours in the QP refer to contracted actual hours in the month of October. The LFS measures both weekly hours usually worked throughout the year, and actual hours worked in the past week.
Source: Authors' calculations on QP and LFS data.

Other relatively less important changes in working time legislation took place in 1998, when the European Working Time Directive of 1993 was ratified, thereby setting a limit to maximum weekly working hours (i.e. standard time plus overtime) at 48 hours. *De facto*, this did not introduce a binding threshold. Lastly, the 2003 Labour Code set yearly overtime at 150 hours maximum for firms above 50 employees, 175 for those below. However, this limit could still be raised to 200 hours by collective agreements. In any case, Castro and Varejão (2007) documented that overtime is very rarely employed by Portuguese firms, probably because of the high overtime premium. Our data shows that, in the period of interest, less than 3% of firms made use of paid overtime, and this share did not change even as standard hours were reduced.

3 Conceptual Framework

The theoretical underpinnings of the effects of reductions in standard hours are well-documented and understood. In its simplest version, without fixed costs of employment and exogenous wages, a standard labor demand model would predict that a reduction in hours would unambiguously increase employment.[5] Profit-maximizing firms would substitute hours for workers, such that a reduction in hours would certainly result in a higher number of workers. This argument is what is known as "work-sharing", and its simplicity and intuitiveness can probably explain its appeal in the public and political debate. However, it is enough to add a fixed cost component to employment and firms' endogenous overtime response to significantly complicate the predictions. This was shown by Calmfors and Hoel (1988), a seminal paper in the literature. The authors argue that, if overtime and the overtime premium are added in the model alongside a fixed cost for each worker, then the predictions of the effect on labor demand become much more ambiguous. The effect on employment will depend on hours before the reform, and whether firms are cost-minimizers, hence facing a fixed demand, or profit-maximizers, which can adjust output accordingly. The cost-minimization case, where output is fixed, is more likely to give rise to work-sharing. However, even in this scenario, the effect on employment

[5] As Estevão and Sá (2008) state: "In a partial equilibrium model of labour demand where average hours of work and employment are perfect substitutes and the only relevant labour cost is the hourly wage, a reduction in the standard workweek reduces average hours and raises employment." Such a model can be found in textbooks such as Hart and Sharot (1978) or Hamermesh (1996).

is ambiguous once overtime is considered. This is because lowering standard hours reduces the cost of an overtime hour relative to the cost per worker, with the perverse effect that firms may actually substitute overtime for workers. This implies that a reduction in standard hours could result in being counterproductive and lead to workers working more actual hours, as firms increase overtime and employ their workers more intensively. However, given the high overtime premium and the scarcity of overtime in Portugal, we argue that this channel is unlikely to be important in our setting. On the contrary, the high overtime premium makes work-sharing more likely, as firms would be more prone to hiring new workers. The other main take-away from Calmfors and Hoel (1988) is that, if firms are profit-maximizing and output is not fixed, then reducing working hours produces a negative scale effect on output, which in turn makes the employment effect even more ambiguous, and likely to be negative. This is because the cost per worker, which includes a fixed cost, increases even if hourly wages remain fixed.

Up until now we have considered models where the hourly wage is exogenous, and does not change as a result of the change in hours, but this does not need to be the case. Trejo (1991) proposes a "fixed-job" model as opposed to a "fixed-wage" model. In this setting, firms are able to fully adjust hourly wages and overtime such that the monthly salary of workers does not change. In this model, a legislative reform in hours would have no real effects, not even on the level of hours worked. But this is not what we observe in the case of the European reforms, where, on the contrary, hours were usually lowered at constant earnings, such that the hourly wage increased significantly. As we will see later, this is also clearly the case of the Portuguese reform in 1996 and of most reforms to standard hours, with the exception of the Canadian case (Skuterud (2007)), where monthly salaries could and did adjust to compensate for the lower hours. Formally, the scenario in which nominal monthly wages do not adjust is treated by Crépon and Kramarz (2002). When nominal monthly salaries do not adjust, and hourly wages go up, this further exacerbates the negative scale effect on employment shown by Calmfors and Hoel (1988). Moreover, as Crépon and Kramarz (2002) point out, it also has an impact on worker flows, as firms have the incentive to fire workers hired with the old standard, and hire new workers at lower wages with the new standard. This last point reiterates the importance of studying employment at the firm level: an increase in the separation rate for incumbent workers, which is what

some studies have shown, does not necessarily mean an overall negative effect on labor demand if the hiring rate also increases.

As already pointed out by Boeri and Van Ours (2013) and Raposo and Van Ours (2010), monopsonistic power in the labor market could provide a justification for legislative reductions in standard hours, as wages are less than the marginal product of labor and working hours are longer than what would be optimal for the worker.[6] Moreover, the predictions of standard models could again change drastically if we move away from perfect labor markets and perfectly competitive product markets. This relates to the debate about the "elusive employment effect of the minimum wage" (Manning (2020)).[7] In other words, if we move away from perfectly competitive labor markets, and take into account the potentially positive effects on labor supply, the effects of an increase in the hourly wage due to lower working hours become even less obvious. To this logic, we add that the models cited above assumed perfectly competitive goods and services markets where firms are price-takers, which made the negative scale effect on output unambiguously negative. If, instead, firms enjoy some market power and are able to adjust their prices to compensate for the higher labor costs, then the predicted effects again become ambiguous.

To sum up, the theoretical predictions of the effects of reductions in standard hours are less than clear-cut and necessarily depend on many assumptions. The work-sharing argument, while theoretically possible, holds only under very strong conditions. A negative employment effect, on the other hand, seems likely in competitive markets but more dubious otherwise. In short, this is an empirical question, which we try to address in the paper. This is even more the case for predicted effects on productivity. The effect on workers' productivity will crucially depend on assumptions about the shape of the production function, and on whether and where marginal returns to hours are increasing or diminishing.

[6] "Another reason for a mandatory reduction of working hours arises when employers have monopsony power. Manning (2003) argues that in a monopsony not only the wage rate is less than the value of marginal product but the firm can also induce the worker to work more than would be optimal for the worker given the monopsony wage. In the same way as a minimum wage can be welfare improving in case of a monopsony, working hours reduction can be welfare improving." (Raposo and Van Ours (2010))

[7] "The strong a priori belief held by many that a rise in the minimum wage must cost jobs ultimately derives from the assumption that the low-wage labor market is close to perfectly competitive." (Manning (2020))

4 Data and Descriptive Statistics

The bulk of the analysis is carried out on *Quadros de Pessoal* (QP - "Lists of Personnel"), administrative, matched employer-employee data collected every year by the Ministry of Employment. This data covers the universe of workers and firms with at least one worker. The data at our disposal covers the period from 1985 to 2016. The QP is collected in a specific month every year; until 1994, this snapshot took place in the month of March, since then, it has changed to October. The full information available in the QP is specified in the Online Appendix: it includes hours, wages, firm and establishment code, sales at the firm level, and several characteristics of the worker and the firm. Importantly, it also collects information on which collective agreement covers a given worker. The data at our disposal has two gaps, in 1990 and 2001. For the purpose of our analysis, we focus on the period 1994 to 2000 in our working sample, which allows us: i) not to have any gaps; ii) always have the data collected in the same period (i.e. month of October); iii) have a balanced window around the time of the reform ($\pm$ 3 years).

Importantly, the QP does not have a measure of *usual* hours until 2003. As discussed before, usual hours means the hours habitually worked in a week, which can be averaged over a certain period, called the reference period. Instead, the QP measures *contracted* actual hours, i.e. the hours actually worked during the month of October. While there is certainly a strong correlation between the two, they are not identical. Moreover, actual hours can (and do) fluctuate much more than usual hours.[8] This difference will be important in our definition of firms' treatment status later in the paper. The data also contains information about overtime hours, meaning that we know the number of hours paid at an overtime premium rate, and whether firms increase overtime after the reduction in standard hours.

For descriptive purposes, we also make use of the *Labour Force Survey*, in the harmonized version available through Eurostat. Contrary to the QP, this data has information on usual hours worked over the entire year (rather than a single month), and a measure of *self-reported* actual hours worked. This allows us to show that the reform reflected in "real" hours, and not only in those reported in

[8]For example, in the year 2000, when both measure are available in the data, we observe that - among full-time workers, in the month of October - 48% of workers have actual hours equal to usual hours, 37% have higher usual than actual hours, 14% have higher actual than usual hours.

administrative datasets (Figure 1). These relevant variables are also outlined in the Online Appendix.

4.1 Descriptive Statistics

Figure OA1 shows the average weekly working hours over time, both in the administrative data and in the labour force surveys. Overall, we can see that hours are always slightly higher in the LFS, which is consistent with the fact that the survey measures actually worked hours rather than contracted hours as in the administrative data. Alternatively, the small difference in level could be attributed to the period of measurement, given that QP only collects information in October. This may also explain the higher variability observed in the administrative data. The trends are, however, comparable, and the effect of the reform, in 1997-98, is very clear. In Figure OA2, we show the distribution of actual hours before and after the reform, both in the QP and in the LFS.

5 Empirical Strategy

5.1 Definition of Treatment Status

For each worker in our data, we have information about the code of the collective agreement that covers him/her. With the worker-level data, we construct information about the most frequent level of hours within each collective agreement code, i.e. the mode, for each year prior to the reform (1994–1996).[9] This allows us to observe which collective agreements are more exposed to the national reform in working hours. The reason we do not want to use directly the share of workers working more than 40 hours to define treatment is out of concern of possible reversion to the mean. Because we observe actual hours in the month of October and not usual contract hours, if we were to select as treated those firms with more workers with high hours before the reform, we would suffer the risk of selecting groups as treated with unusually high hours in that given month. This could lead to mean-reversion and violation of the common trend assumption. Our approach, instead, allows us to determine treatment by a characteristic that is virtually fixed, i.e. the collective agreement covering each worker, and that is not affected by the business cycle.

[9] We do not use years prior to 1994 because of the break in the series explained in Section 4.

As we do not have the provision legally specified in each collective agreement, because we do not see the text of the agreement, but only the realized value in actual hours, we define a collective agreement as treated if, for the period 1994 to 1996, it has a mode above 40 hours. Instead, we define a collective agreement as control if, for the period 1994 to 1996, the mode is at or below 40 hours. We calculate the mode separately for each year because we want to exclude those collective agreements that change hours over the period. Moreover, this allows us to limit the extent of measurement error in the definition of treatment status, if the mode that is calculated is sensible to yearly variation. For 1996, the year of the reform, we end up with 178 treated collective agreements, 179 control collective agreements, 132 missing values. The missing values represent those collective agreements that either: i) did not exist in 1994 (i.e. no worker had that code in 1994); ii) changed the mode hours between 1994–1996 from more than 40 to below, or the other way around.[10]

The process above defines what collective agreement codes are considered as affected or not. Each worker is coded in the sample with one of these codes. Formally, we define a worker as treated in 1996 if:

$$w_i = \begin{cases} 1 \; if \; Mo(CA_i) > 40 \\ 0 \; if \; Mo(CA_i) \leq 40 \end{cases}$$

where Mo indicates the most frequent value within that collective agreement code. In other words, workers are considered as treated if they are covered by a collective agreement where the most frequent value in the population is above 40; they are considered control if the most frequent value is at or below 40. In 1996, we thus have 23% of workers with a value of 0, 44% with a value of 1, and 37% are undefined.[11]

The last step is to define treatment at the firm-level. In this case, we then define a measure of treatment exposure at the firm level, j, as (a function of) the share of workers whose collective agreement is exposed, taking into account that workers in the same firms can be covered by different

[10] 55 did not exist in 1994, 41 went from a mode of above 40 hours to below from 1994 to 1996, and 36 went from a mode of below 40 hours to above.

[11] Those coded as missing are the ones whose collective agreement did not exist in 1994, or changed the mode of hours between 1994 and 1996.

collective agreements:

$$E_j = \frac{\sum_{i=1}^{W_j} w_{ij}}{W_j}$$

where W_j is the total number of workers in firm j. We plot this measure of exposure in Figure 2. As the graph shows clearly, an overwhelming majority of firms have either all or none of their workers exposed. In other words, there is not a lot of variation within firms in terms of the hours specified by the different collective agreements (or not a lot of variation in the collective agreements themselves within firms).

Figure 2: Firms by Share of Workers with a Collective Agreement Affected by Reform, 1996

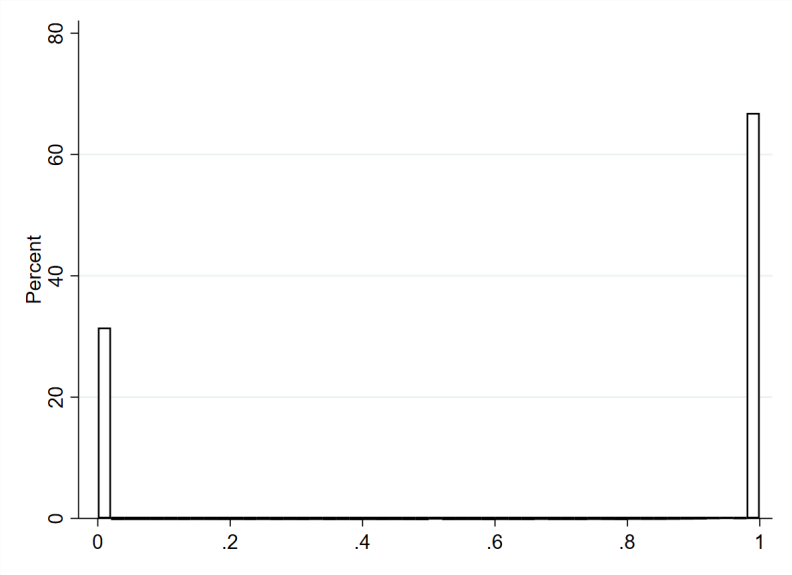

Note: This graph plots the distribution of firms by share of workers whose collective agreement is impacted by the reform. 31% of firms have no workers impacted, 66% have all workers impacted, only 3% in-between.
Source: Authors' Calculation on QP

Given the distribution observed in Figure 2, we decided to discretize the treatment variable, and define as treated those firms with more than 50% of their workers exposed, and as control those with less than 50% of their workers exposed, in the following way:

$$T_j = \begin{cases} 1 \; if \; E_j \geq 0.5 \\ 0 \; if \; E_j < 0.5 \end{cases}$$

the value of 0.5 here is arbitrary but irrelevant for the results of the estimation. Setting the threshold

at 0.25 or 0.75 would result in virtually identical treatment and control groups, as this is essentially equivalent to setting as treated those with $E_j = 0$ and as control those with $E_j = 1$. Indeed, only 3% of firms are in the in-between situation, with only some workers impacted, but we keep them in the estimation as these are likely to be larger firms. T_j is undefined for around 28% of firms in 1996. For the reasons mentioned above, workers in these firms are classified as neither treated nor control. We exclude these from the estimation.

5.1.1 Empirical Specification

After defining treatment status in the way described above, we run the following regression:

$$Y_{jt} = \gamma_j + \delta_t + T_j \sum_{t=1}^{T} \beta_t \mathbb{1}\{year = t\} + \varepsilon_{jt} \tag{1}$$

where Y_{jt} are the outcomes of interest at the firm-level; γ_j and δ_t are firm and year fixed effects respectively. We are interested in the coefficients β_t, which captures the dynamic effects of the reduction in standard hours when comparing firms whose workers are covered by a collective agreement affected by the reform versus those whose workers are not affected. We cluster standard errors at the firm level, in order to account for panel nature of the data.

The identification assumption necessary for a correct identification of the coefficient β_t is that more and less affected firms would have followed the same trend in the absence of the reform. This is the "classic" diff-in-diff assumption, but, as some recent papers have highlighted (Kahn-Lang and Lang (2019)), its validity always deserves a deeper discussion. As the definition of the treatment status itself highlights, affected and non- affected firms have to be initially different in at least some dimensions (here, hours) as the difference itself is what determines that some firms are more treated and some less. In our setting, this is the result of different workers/firms being covered by different collective agreements. Therefore, it is to be expected that treated and control firms would differ across other dimensions too, such as employment levels and productivity, which are likely to impacted directly by collective agreements and to be correlated with the same firm and worker-specific characteristics that determine hours. For example, firms with different hours cannot have, by definition, the same

output and, at the same time, the same workers' and hours' productivity. The relevant question for identification then becomes to what extent we could have expected the *evolution* of affected and non-affected firms to have been parallel in the absence of the intervention, *considering that these firms were starting from different levels*. As Kahn-Lang and Lang (2019) rightly emphasize, this also implies, indirectly, a structural assumption: when the starting level is different, a "common trend" assumption cannot hold in both absolute or relative value at the same time, i.e. treated and control firms cannot evolve in the same way in log *and* level at the same time.[12] In the case of employment, for example, it is important to determine whether we expect the same employment change in percentage or in number of workers. Other than in the case of hours, where the absence of a trend allows us to estimate the effect both on units or logs, we put all our outcomes in logarithmic form, as analysis of pre-trend shows that relative evolution is more likely to hold for firms starting from different levels.[13]

In a recent paper De Chaisemartin and d'Haultfoeuille (2020) warn against the use of linear regressions with both group and time fixed effects. They show that these estimations provide a weighted average of ATTs (average treatment effect on the treated), where some of these ATTs can receive negative weights when effects are heterogeneous. These can have important consequences for the estimate, even flipping the sign of the effect. As a robustness check, we follow the procedure they suggest to estimate the weights and find that in our setting no negative weights arise. Consequently, when running their proposed estimator (in Stata *did_multiplegt*), which is robust to heterogeneous effects, results are qualitatively identical to those presented below. We present these results among the robustness checks.

5.1.2 Sample Selection

Our sample from the QP is made up of the universe of for-profit private enterprises and workers. We exclude agriculture, the self-employed, associations and non-standard employers. We also exclude all those sectors (i.e. education, health, social work) in which there is a strong public component. In our estimations, as by definition we can only define treatment for existing firms, we restrict the sample to

[12] The only case in which this is not true is if there is no trend, i.e. the evolution is flat.

[13] Taking the log also has the advantage of helping us to deal with outliers due to large firms, which are particularly troublesome for wages and sales, without having to significantly trim the sample.

Table 2: Descriptives of Treated and Control Firms in 1996

	N	Avg. Workers	Med. Workers	Salary	Wage	Weekly Hours	$H_i > 40$
$T_j = 1$	57,051	11.5	4	373	2.07	41.99	66%
$T_j = 0$	25,085	14.0	4	479	2.78	40.58	43%
	Manufacturing	Retail	Lisbon	Porto	Med. Sales	Overtime	Years Present
$T_j = 1$	28%	42%	17%	15%	107 590	1.1%	5.6
$T_j = 0$	11%	65%	50%	26%	163 715	0.9%	5.5

The sample includes all for-profit firms for which treatment status is defined, excluding agriculture, self-employment and non-standard employers. *Treated* is defined as described in the text, around 18% of firms present in 1996 are excluded because treatment is not defined. $H_i > 40$ indicates the share of workers with actual hours above 40. *Years Present* indicates the average number of years the firm appears between 1994 and 2000. Wages and sales are in current euros.

firms that exist in the year of the reform (1996). We focus exclusively on firms whose workers have more than 35 hours on average over the period, so we exclude firms where average hours are below the full-time level. We exclude firms that report zero sales over the whole period (1994–2000); when zero sales arise in a given year, to avoid dropping after the log transformation, we impute the deflated median over the period. Because of the highly volatile nature of the variable, we trim those firms in the bottom 1% and the top 99% of the distribution of sales growth over the period 1994–2000. We do the same with employment growth.

We are left with a sample of 431,712 observations, of which 82,136 unique firms, 25,085 in the control group, and 57,051 in the treatment group. The average characteristics of treated and control firms are in Table 2. Control firms are on average larger, both in terms of workers and sales, pay higher salaries and are less likely to be in manufacturing and more likely to be in the retail sector. They are also more likely to be located in the large Portuguese cities: Porto and Lisbon. Treated and control firms do not seem to have differential attrition (which indicates a firm closes) over the period: over the seven years from 1994 to 2000, these are likely to be present roughly 5.5 years on average. Average hours are around 1.5 hours higher in treated firms than in control firms. The share of workers whose hours are higher than 40 hours is 66% in treated firms, and around 43% in control firms, which in part indicates the fluctuation of actual hours relative to usual hours. Very few firms make use of overtime in Portugal: in our working sample, only around 1% of firms have any worker

working extra hours.

6 Results

A summary of the results is available in Table 3, taking as a reference point the year 2000, three years after the full implementation of the reform. We discuss and show here below the impact on each outcome in a dynamic way, separating between the first stage on hours and wages, the impact on employment and labor input, and the impact on sales and sales per hour.

6.1 Impact of the Reform on Hours and Wages

We present the effect of the reform on working hours in Figure 3. The effect is clearly visible: at the time the reform was implemented, average working hours within treated firms drop by around 1.5 hours in the two years of the implementation. The full drop is not immediate in 1997 because the reform allowed for an adjustment period of one year, as the limit is first lowered to 42 hours in 1997 and to 40 hours in 1998. When transforming the outcome in log, we observe that the drop in hours is around 4% in treated firms relative to control firms. We do not observe any significant increase in the use of overtime for treated firms relative to controls firms, nor any significant increase over the period in absolute terms.

The size of the first stage is less than what might be expected if all treated firms had workers at 44 hours, and all controls firms at 40 hours. In that case, we would have expected a drop of 4 hours (instead of 1.5) and of around 9% (instead of 4%). This is due to the fact that we measure actual hours and not usual hours, but, also to the fact that we measure treatment imperfectly: we define a treated firm as one whose workers have a collective agreement that, on average, is exposed, and vice versa for control firms. This can lead to treated firms not being fully affected and, symmetrically, to control firms being partly affected if the collective agreement being affected *on average* does not fully capture the situation of that worker. This clearly emerges if, instead of the coefficients, we plot the level of average hours for the treated and control group over time: there is also a visible drop in hours for firms in the control group (Figure OA3). Our estimate should therefore be understood as a

relative effect rather than an absolute one.

The text of the reform did not mention explicitly that salaries could not be adjusted when working hours decreased. However, a certain stickiness of nominal monthly salaries is to be expected, as adjusting the already contracted nominal wage might be very difficult. Moreover, it should be taken into account that: i) Portugal has a national minimum wage scheme; ii) each collective agreement can also specify several wage scales for the different occupations and sectors. Inflation over the period 1996 to 2000 was relatively low and stable, always at or slightly below 3% a year, which could have made it more difficult for real salaries to adjust. What we observe is that the reform translated into a significant increase in hourly wages for workers, as monthly salaries did not adjust. This is clearly visible in Figure 4: the difference in monthly salary stays constant between treated and control firms over the period, such that the drop in hours shows up fully in hourly wages, which also increase by roughly 4%.

Figure 3: First Stage on Average Hours by Firm

(a) Hours

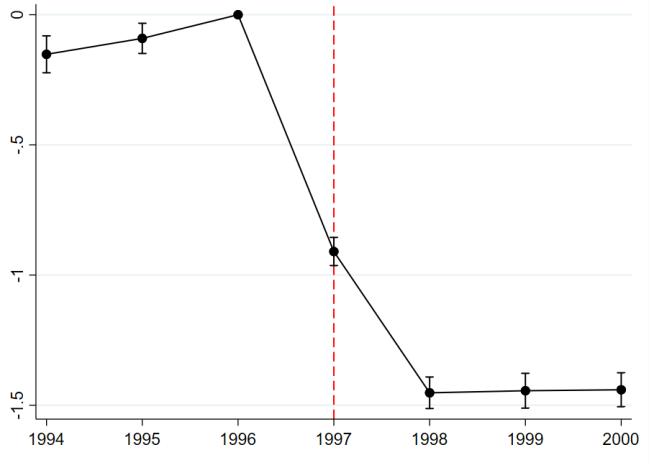

(b) Log of Hours

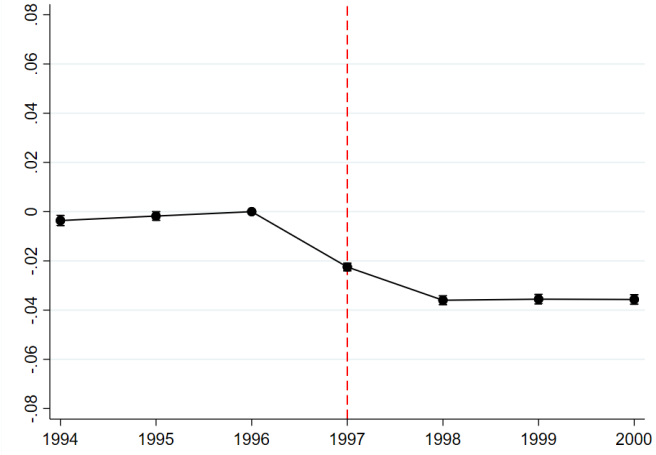

Note: these graphs plot the coefficients of Equation 1 on average hours within the firm both with hours measured in level, panel (a), and in log, panel(b). The red line indicates the year of implementation of the reform. The reference year for the coefficient is 1996, the vertical bars indicate 95% confidence intervals.
Source: Authors' estimations on QP

Figure 4: Wages

(a) Monthly Wages

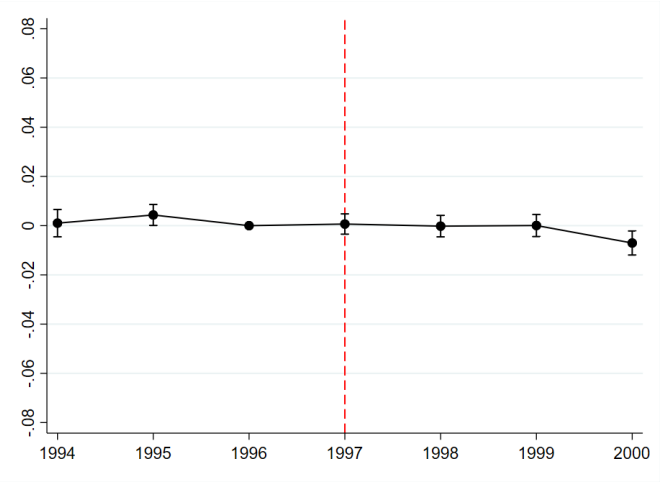

(b) Hourly Wages

Note: these graphs plot the coefficients of Equation 1 on average monthly wages within the firm (panel a), and hourly wages (panel b). Both outcomes are in logs. The reference year for the coefficient is 1996, the vertical bars indicate 95% confidence intervals.
Source: Authors' estimations on QP

6.2 Impact on the Number of Workers and Labor Input

For the effect on labor demand, we present the coefficients of equation 1 in Figure 5. As in the case of hours and wages, pre-reform evolution between treatment and control firms seems to be roughly similar. With the employment variable in logs, this implies that treatment and control firms were experiencing the same growth trend in employment over the period. We observe no significant effect on employment up to the four year mark after the reform. This shows that there is no "work-sharing" logic at play: firms do not substitute hours for workers to compensate for the decrease in weekly hours. Yet, we also observe no significant negative effect on employment, although the coefficients are negative and reach up to -1% three to four years after the reform. We only have information about the type of contracts of workers after 2000, so we cannot check directly whether this delayed effect partly reflects the expiration of temporary contracts that were not renewed.

In the same figure, we also present the effect on labor input, which is simply the total number of hours worked within a firm (equivalent to the average number of hours per worker times the number of workers.) Given that hours per worker fall, and the number of workers stays roughly constant (if not in slight decline), it is mechanical to observe that the labor input falls significantly, by roughly

the same amount of hours per worker (around 4%). This reiterates the point that treated firms do not substitute hours for workers, but instead let labor input fall, in proportion by the same amount that the drop in average hours.

Figure 5: Number of Workers and Labor Input

(a) Number of Workers

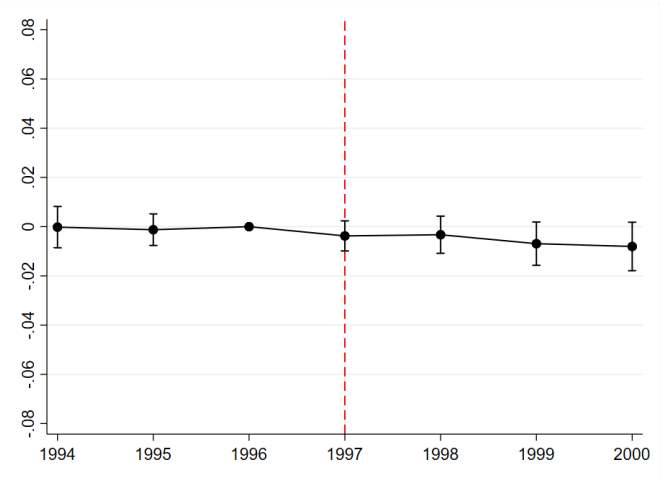

(b) Firm Total Hours

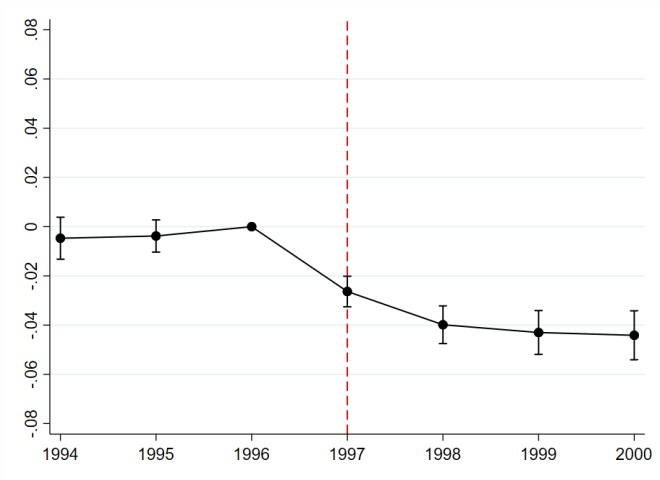

Note: these graphs plot the coefficients of Equation 1 on total employment within a firm (panel a), and total hours worked within a firm (panel b). Both outcomes are in logs. The reference year for the coefficient is 1996, the vertical bars indicate 95% confidence intervals.
Source: Authors' estimations on QP

6.3 Impact on Total Sales and Sales per Hour

The impact on output is presented in Figure 6. Here, output is proxied by the value of sales. It should be underlined that, with this measurement: i) potential changes in the price of intermediate inputs are included in the outcome; and ii) we measure sales value and not sales volume. As we discuss later, this implies that any estimate on sales potentially captures effects through prices and not only through quantity.

The effect that we observe is of a moderate increase in sales in treated firms relative to control firms. This is surprising, considering the large decrease in labor input, for which one would expect, if anything, sales to decrease rather than increase. Moreover, the theoretical prediction was of an unambiguous decrease in output (Calmfors and Hoel (1988), Crépon et al. (2004)), if firms are profit maximizing and output is not fixed by demand. Instead, we observe a moderate increase in total

sales, which disappears around 2 years after the reform to go back to the same level as control firms. The increase followed by relative stability in sales results in a large increase in sales per hour, which derives mechanically from the simultaneous drop in labor input. Sales per worker, which we do not show here, also increase, although not significantly.

Figure 6: Effect on Sales - Coefficients

(a) Total Sales

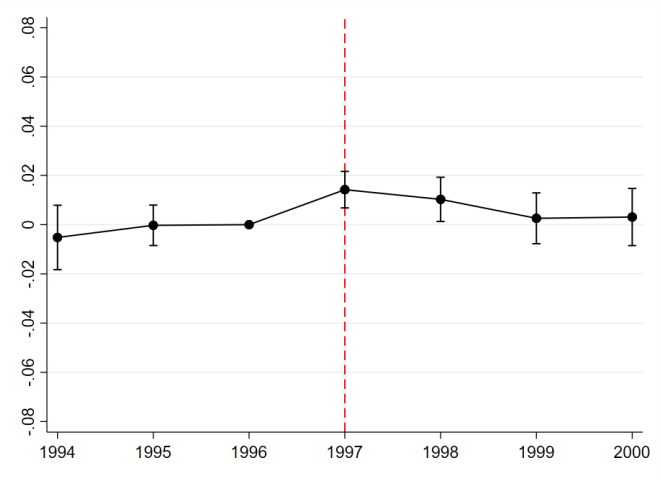

(b) Sales per Hour

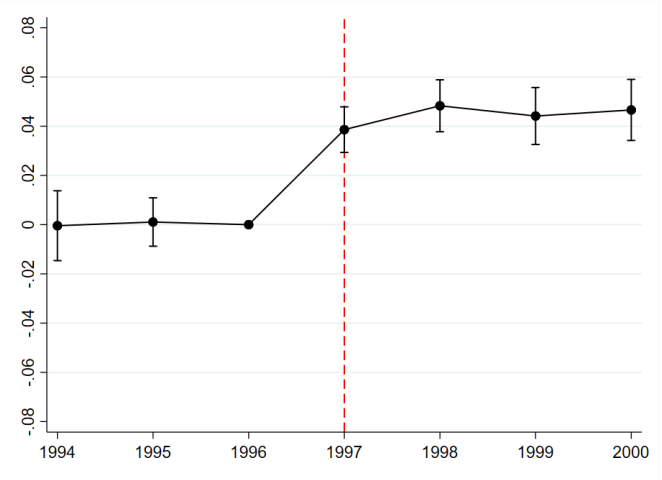

Note: these graphs plot the coefficients of equation 1 on total sales within a firm, panel (a), and sales per hour, panel (b). Both outcomes are in logs. The reference year for the coefficient is 1996, the vertical bars indicate 95% confidence intervals.
Source: Authors' estimations on QP

These results on sales are counter-intuitive, and are not consistent with the predictions of standard economic theory. The question that we are unable to test directly with this data is whether this is a volume or price effect, i.e. whether firms actually produce more (and more per hour) as a result of the reform, or whether they are able to charge higher prices and hence increase and then maintain the same growth of sales as control firms. We argue that, to some extent, some price effects must be at play here. Surely, reducing working hours is likely to have beneficial effects on productivity, and part of the increase that we observe in Figure 6, in terms of sales per hour, is likely to be due to that. However, it is unlikely that a decrease in hours per worker would actually result in increases in the volume of output at the firm level, this would imply that pre-reform firms were in an equilibrium with such high hours that returns to an extra hour of work were not diminishing but actually zero.

6.3.1 Price Effects

In order to test this mechanism, we make use of price data at the sectoral level that is available through national accounts. We merge into our data output deflators at the NACE Rev. 1.1 2-digit level available from the EU KLEMS project (http://www.euklems.net/). This gives us disaggregated industry price indexes from 1995 for 49 different industries.[14] Note that we do not use value-added price indexes, but the ones for output, as this more closely resembles the concept of sales measured in our data, which include intermediate inputs. In this way, we are able to test whether part of the effect on sales, and sales per hour, is due to changes in prices in more treated sectors.

Figure 7: Evolution of Average Sectoral Price Indexes by Firm Treatment Status

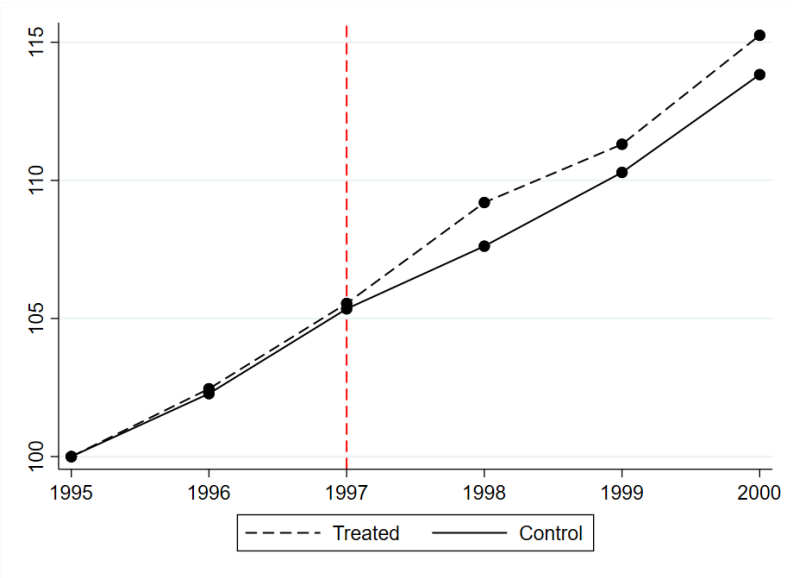

Note: This graph plots the evolution of average sectoral price indexes for treated and control firms between 1995 and 2000. Price indexes are at the 2-digit NACE Rev. 1.1 sector level. The index is equal to 100 in 1995.
Source: Authors' Calculation on KLEMS data

We plot the evolution of prices for firms in treated and control sectors in Figure 7. While we see clearly that these sectors experience a very similar evolution of prices from 1995 to 1997, the year after the reform kicks in we observe a relative increase in prices in the sectors of treated firms. The full extent of the increase in prices as a result of the reform might be larger if there are also within-sector price effects, which we would not capture in this way. Nonetheless, it clearly shows that part of the higher labor cost resulting from the higher wages is shifted onto prices.

[14] We exclude the year 1994 because it was covered by a different industrial classification.

Figure 8: Effect on Sales with and without Price Effects - Coefficients

(a) Total Sales

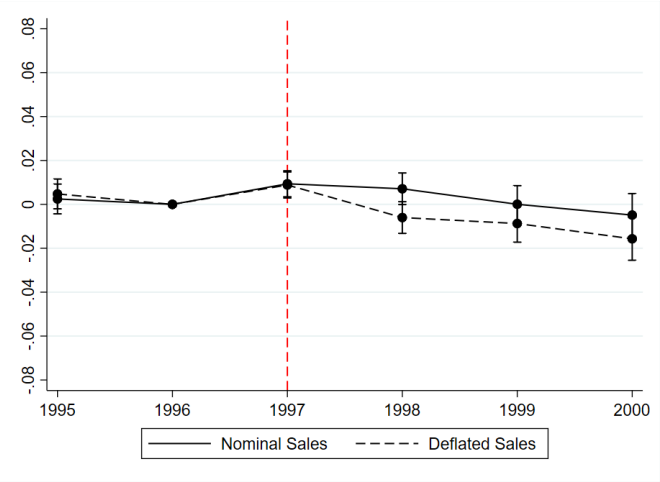

(b) Sales per Hour

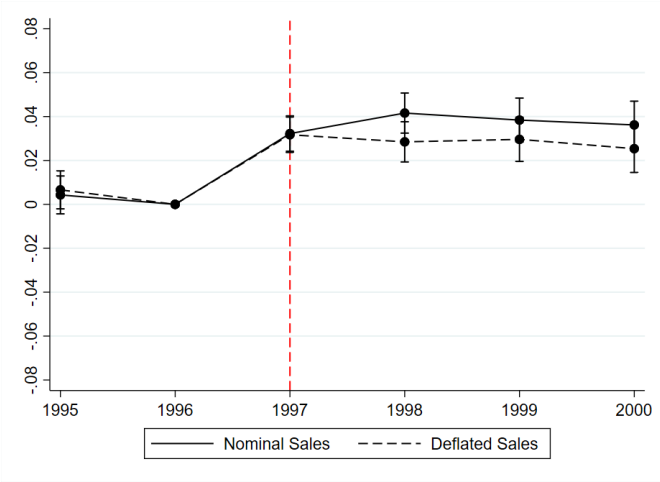

Note: these graphs plot the coefficients of equation 1 on total sales within a firm, panel (a), and sales per hour, panel (b). Both outcomes are in logs. The solid line indicates nominal sales, not adjusted for prices, while the dashed line indicates deflated sales, adjusted for output deflators at the 2-digit NACE Rev. 1.1 sector level. The reference year for the coefficient is 1996, the vertical bars indicate 95% confidence intervals.
Source: Authors' estimations on QP

Using these sectoral output deflators, we can check how the effect on sales and sales per hour changes when price effects are taken into account. We plot these results in Figure 8 for the period 1995–2000, showing the coefficient both for nominal sales and sales deflated by the sectoral price index. We can see clearly that, taking into account price effects, the effect on output becomes negative and significant starting from 1998, and reaches up to -2% in the year 2000. In the same way, the price effect eats away around one fourth of the positive effect on sales per hour, which, however, remains strongly positive and significant.

These results show that part of the effect we observe on output certainly goes through prices. It is important to reiterate that we cannot test the full extent of the price effects, because we only observe relative price changes across sectors. Sectors fixed effects explain around 50% of the variation in treatment status, so it is possible that the total price effects would be larger, and would push the effect on deflated sales further down, and the effect on sales per hour towards zero.

This result is consistent with the recent trend in the literature putting the emphasis on imperfect

labor markets. The literature on minimum wages has long pointed out that there is a certain pass-through - albeit small - of higher minimum wages on prices (Aaronson (2001), Lemos (2008), Renkin et al. (2020)). If firms are not price takers, then we can expect output prices to be an important margin of adjustment (Clemens (2021)). Overall, the evidence that we observe is not consistent with perfect competition: the prediction of a standard model would have been that an increase in the hourly wage would have led to a negative scale effect on both output and the number of workers, instead we observe no significant adjustment in terms of employment and that output stays constant or even increases, in part because more treated sectors experience relative price increases.

6.4 Robustness Checks

The robustness checks to the estimation are presented in Table 3, along with a summary of the results at $t + 4$, i.e. the year 2000. As treated firms are much more likely to be found in the manufacturing sector and less in retail, we estimate equation 1 with manufacturing×year and retail×year fixed effects. The results are very similar when we take out the variation across these two large sectors from the estimation. However, the coefficient on total sales becomes negative, consistently with the prediction of a negative scale effect, and the one on sales per hour decreases. This can be seen as confirming the story on price adjustments, which would be partly captured away by the sector×year fixed effects. Up to now we have presented unweighted regression, where each firm receives a weight of 1 irrespective of its size. As a robustness check, we run the estimation by weighting each firm by its log of employment in 1996.[15] Again, this does not qualitatively change the results. Lastly, as suggested by De Chaisemartin and d'Haultfoeuille (2020), we check whether our estimation gives negative weights to some ATTs (average treatment on the treated) through the Stata command *twowayfeweights* and find no negative weights. Consistently, when we run the estimator proposed by De Chaisemartin and d'Haultfoeuille (2020), the estimates are virtually identical to those of estimating Equation 1 through OLS. The coefficients generated by this estimation, and related standard errors, are presented in the Online Appendix in Figures OA4 to OA7.

[15] We use the log of employment as weight because otherwise large firms introduce a lot of noise into the estimation.

Table 3: Summary of Effects at $t+4$

Outcome in Logs:	Avg. Hours	Salary	Wages	Workers	L. Input	Sales	H. Sales/Hour
$\mathbb{1}\{E_j >= 0.5\}$	-0.036	-0.005	0.030	-0.007	-0.043	-0.001	0.042
	(0.001)	(0.002)	(0.002)	(0.005)	(0.005)	(0.006)	(0.006)
$\mathbb{1}\{E_j >= 0.5\}$	-0.036	-0.003	0.034	-0.005	-0.042	-0.014	0.028
(+ M×t, R×t)	(0.001)	(0.002)	(0.003)	(0.005)	(0.005)	(0.006)	(0.006)
Weighted by 1996 Ln(Employment)							
$\mathbb{1}\{E_j >= 0.5\}$	-0.036	-0.005	0.031	-0.011	-0.047	-0.003	0.044
	(0.001)	(0.002)	(0.002)	(0.005)	(0.005)	(0.006)	(0.006)
De Chaisemartin and d'Haultfoeuille (2020) Estimator							
$\mathbb{1}\{E_j >= 0.5\}$	-.036	-0.007	0.029	-0.008	-0.045	-0.001	0.044
	(0.001)	(0.003)	(0.003)	(0.006)	(0.006)	(0.005)	(0.006)
Observations	431,712	431,712	431,712	431,712	431,712	431,712	431,712
Number of Firms	82,136	82,136	82,136	82,136	82,136	82,136	82,136

Note: This table presents the coefficients of equation 1 on outcomes at the firm level, in logs. Each row indicates a different specification, and each column a different outcome: the first row presents the coefficients of equation 1 with E_j defined as in the text; the second row presents the same specification with the addition of manufacturing×year and retail×year fixed effects. The third row weighs the estimation by log of employment in 1996. The last row presents the coefficient following the estimator suggested by De Chaisemartin and d'Haultfoeuille (2020) and related standard errors. All outcomes are in logs and evaluated at $t+4$ from the reform. Standard errors, clustered at the firm level, in parentheses.

7 Conclusion

Working hours are a key economic variable, directly linked to first-order economic outcomes, such as employment and productivity. Working time regulations are omnipresent across countries, in particular in Europe. Yet, we still know very little about their effects. This paper studies how reductions in standard time, through changes in national legislation, affect economic outcomes at the firm level. We exploit a national reform in Portugal that lowered hours from 44 to 40 hours without putting in place compensating measures for firms. Initial heterogeneity in the mix of workers who are covered by different collective agreements gives us different intensities to firms' exposure to this reform that can be used to identify the effects of a larger adjustment in terms of working hours. Overall, we find that the decrease in hours was not matched by a decrease in salaries, such that hourly wages increased significantly in treated firms. We observe no work-sharing at play: labor input drops significantly in treated firms, but they are still able to maintain output levels as sales per hour increase significantly. Part of this effect goes through prices, as we show that more treated sectors experience higher price growth. These sectoral price effects are enough to make the effect on output volumes negative and significant, but eat away only one fourth of the productivity effect.

Our results indicate, similarly to the literature on minimum wages, that the theoretical predictions of conventional models do not necessarily need to hold in practice. The lack of a drop in employment - and the fact that firms are able to adjust output and prices instead - suggests a move away from perfect competition in the explanation of these results. Much remains to be investigated about the effect of reductions in standard hours, as there is a surprising gap in the literature, given the importance of this topic in theory and in the public debate. While we provide clear evidence on employment, in this paper we are only able to hint at the effect on productivity, because it is mingled in with the effect on prices and sales. However, testing for the potentially beneficial effects on productivity of lower hours, without damaging employment, has key economic and policy implications.

References

Aaronson, D. (2001). Price Pass-Through and the Minimum Wage. *Review of Economics and statistics 83*(1), 158–169.

Boeri, T. and J. Van Ours (2013). *The Economics of Imperfect Labor Markets*. Princeton University Press.

Brachet, T., G. David, and A. M. Drechsler (2012). The Effect of Shift Structure on Performance. *American Economic Journal: Applied Economics 4*(2), 219–46.

Calmfors, L. and M. Hoel (1988). Work Sharing and Overtime. *The Scandinavian Journal of Economics*, 45–62.

Castro, A. and J. M. Varejão (2007). Operating Hours, Working Times and Employment in Portugal. In *Operating Hours and Working Times*, pp. 147–167. Springer.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The Effect of Minimum Wages on Low-Wage Jobs. *The Quarterly Journal of Economics 134*(3), 1405–1454.

Chemin, M. and E. Wasmer (2009). Using Alsace-Moselle Local Laws to Build a Difference-in-Differences Estimation Strategy of the Employment Effects of the 35-hour Workweek Regulation in France. *Journal of Labor Economics 27*(4), 487–524.

Clemens, J. (2021). How Do Firms Respond to Minimum Wage Increases? Understanding the Relevance of Non-employment Margins. *Journal of Economic Perspectives 35*(1), 51–72.

Collewet, M. and J. Sauermann (2017). Working Hours and Productivity. *Labour Economics 47*, 96–106.

Crépon, B. and F. Kramarz (2002). Employed 40 Hours or Not Employed 39: Lessons from the 1982 Mandatory Reduction of the Workweek. *Journal of Political Economy 110*(6), 1355–1389.

Crépon, B., M. Leclair, and S. Roux (2004). RTT, Productivité et Emploi: Nouvelles Estimations sur Données d'Entreprises. *Economie et Statistique 376*(1), 55–89.

De Chaisemartin, C. and X. d'Haultfoeuille (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review 110*(9), 2964–96.

Estevão, M. and F. Sá (2008). The 35-Hour Workweek in France: Straightjacket or Welfare Improvement? *Economic Policy 23*(55), 418–463.

Feldstein, M. S. (1967). Specification of the Labour Input in the Aggregate Production Function. *The Review of Economic Studies 34*(4), 375–386.

Garnero, A., S. Kampelmann, and F. Rycx (2014). Part-time Work, Wages, and Productivity: Evidence from Belgian Matched Panel Data. *ILR Review 67*(3), 926–954.

Gonzaga, G. M., N. A. Menezes Filho, and J. M. Camargo (2003). Os Efeitos da Redução da Jornada de Trabalho de 48 Para 44 Horas Semanais em 1988. *Revista Brasileira de Economia 57*(2), 369–400.

Hamermesh, D. S. (1996). *Labor Demand*. princeton University press.

Hart, R. A. and T. Sharot (1978). The Short-run Demand for Workers and Hours: a Recursive Model. *The Review of Economic Studies 45*(2), 299–309.

Hunt, J. (1999). Has Work-Sharing Worked in Germany? *The Quarterly Journal of Economics 114*(1), 117–148.

Kahn-Lang, A. and K. Lang (2019). The Promise and Pitfalls of Differences-in-Differences: Reflections on 16 and Pregnant and Other Applications. *Journal of Business & Economic Statistics*, 1–14.

Kawaguchi, D., H. Naito, and I. Yokoyama (2017). Assessing the Effects of Reducing Standard Hours: Regression Discontinuity Evidence from Japan. *Journal of the Japanese and International Economies 43*, 59–76.

Künn-Nelen, A., A. De Grip, and D. Fouarge (2013). Is Part-time Employment Beneficial for Firm Productivity? *ILR Review 66*(5), 1172–1191.

Lemos, S. (2008). A Survey of the Effects of the Minimum Wage on Prices. *Journal of Economic Surveys 22*(1), 187–212.

Lepinteur, A. (2019). The Shorter Workweek and Worker Wellbeing: Evidence from Portugal and France. *Labour Economics 58*, 204–220.

Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press.

Manning, A. (2020). The Elusive Employment Effect of the Minimum Wage. *Journal of Economic Perspectives*.

Pencavel, J. (2014). The Productivity of Working Hours. *The Economic Journal 125*(589), 2052–2076.

Raposo, P. S. and J. C. Van Ours (2010). How a Reduction of Standard Working Hours Affects Employment Dynamics. *De Economist 158*(2), 193–207.

Raposo, P. S. and J. C. van Ours (2010). How Working Time Reduction Affects Jobs and Wages. *Economics Letters 106*(1), 61–63.

Renkin, T., C. Montialoux, and M. Siegenthaler (2020). The Pass-Through of Minimum Wages into US Retail Prices: Evidence from Supermarket Scanner Data. *Review of Economics and Statistics*, 1–99.

Sánchez, R. (2013). Do Reductions of Standard Hours Affect Employment Transitions? Evidence from Chile. *Labour Economics 20*, 24–37.

Skuterud, M. (2007). Identifying the Potential of Work-Sharing as a Job-Creation Strategy. *Journal of Labor Economics 25*(2), 265–287.

Trejo, S. J. (1991). The Effects of Overtime Pay Regulation on Worker Compensation. *The American Economic Review*, 719–740.

Varejao, J. (2005). Work-Sharing Revisited-Lessons from a Natural Experiment. Technical report.

For Online Appendix

OA.1 Extra Tables and Figures

Table OA1: Standard Time Reduction and Employment: Overview of the Literature

Paper	Country/Year	Reform	Level of Analysis	Sign on Emp.
Crépon and Kramarz (2002)	France - 1982	40h to 39	Worker	Higher firing (negative*)
Gonzaga et al. (2003)	Brazil - 1988	48h to 44h	Worker	Null
Raposo and van Ours (2010)	Portugal - 1996	44h to 40h	Worker	Ambiguous
Sánchez (2013)	Chile - '01-'05	48h to 45h	Worker	Null
Estevão and Sá (2008)	France - 1998	40h to 35h	Worker	Null
Varejao (2005)	Portugal - 1996	44h to 40h	Firm	Null**
Kawaguchi et al. (2017)	Japan - 1997	44h to 40h	Firm	Negative***
Crépon et al. (2004)	France - 1998	39h to 35h	Firm	Ambiguous
Hunt (1999)	Germany - '84–'95	Various	Sector	Negative
Skuterud (2007)	Canada - '97–'00	44h to 40h	Sector/Region	Null
Raposo and van Ours (2010)	Portugal - 1996	44h to 40h	Sector × Region	Positive
Chemin and Wasmer (2009)	France - 1998	39h to 35	Region	Null

*This does not, by definition, imply that the total employment effect is negative, as it does not account for potential changes in hiring.** Varejao (2005) finds a null effect on employment when defining treatment and control firm in a binary way for the period '96-'99, he estimates a negative coefficient when including treatment as a continuous variable. Kawaguchi et al. (2017) do not find a significant first stage on hours overall: for a subsample of firms with a significant first stage, they find a negative but insignificant effect on new hires..

Figure OA1: Average Usual and Actual Weekly Working Hours of Full-Time Workers, 1994–2000, QP and LFS

(a) QP

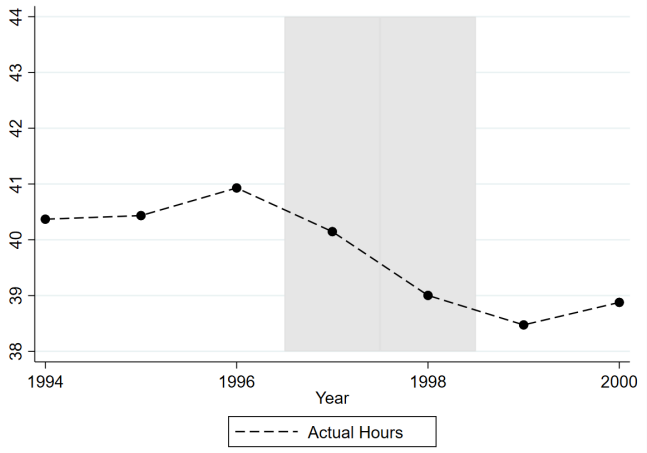

(b) LFS

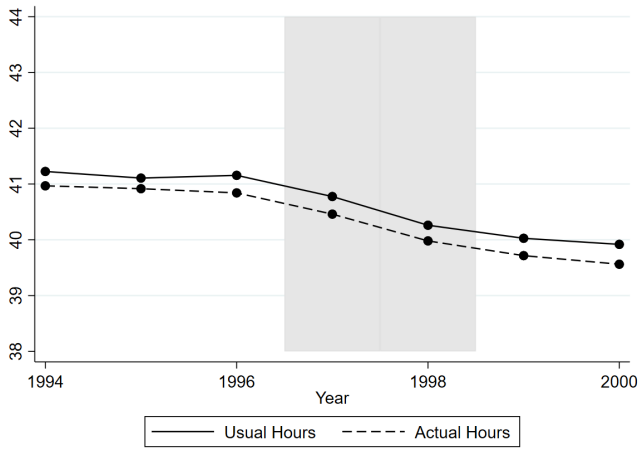

Note: The shaded area indicates the period of the reform. Hours in the QP refer to contracted actual hours in the month of October. The LFS measures both weekly hours usually worked throughout the year, and actual hours worked in the past week.
Source: Authors' calculations on QP data and LFS

Figure OA2: Hours Distribution of Full-Time Workers, Before and After Reforms, QP and LFS

(a) QP

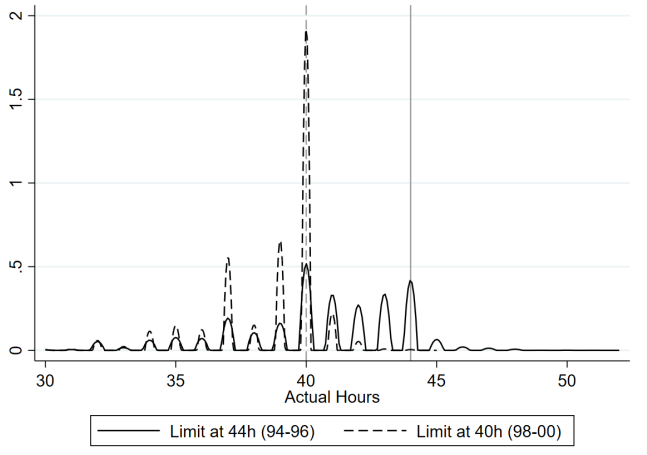

(b) LFS

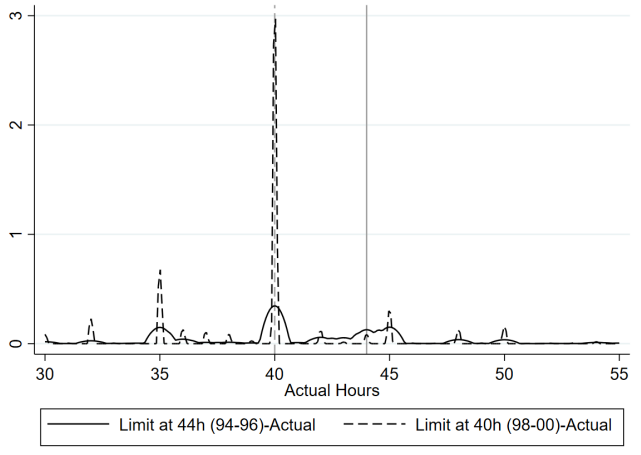

Note: These graphs show the distribution of working hours in the years before the reform, and after the reform both in the QP and in the LFS. Hours in the QP refer to contracted actual hours in the month of October. The LFS measures both weekly hours usually worked throughout the year, and actual hours worked in the past week.
Source: Authors' calculations on QP and LFS data

Figure OA3: First Stage on Average Hours by Firm

(a) Hours

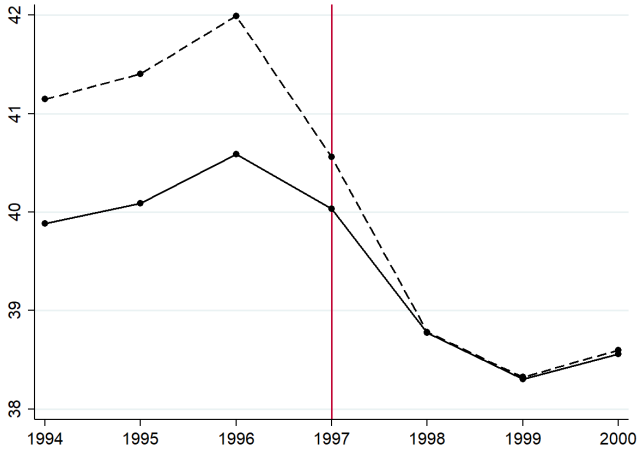

(b) Log of Hours

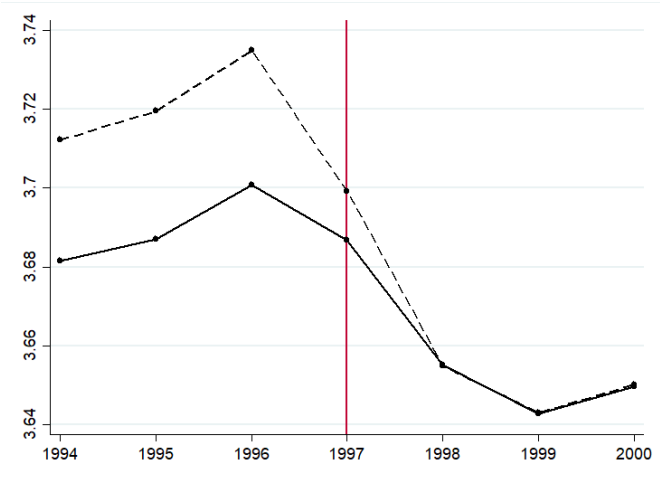

Note: these graphs plot the evolution of average hours within the firm both with hours measured in level, panel (a), and in log, panel(b). The red line indicates the year of implementation of the reform.
Source: Authors' calculation on QP

OA.2 Graphs with De Chaisemartin and d'Haultfoeuille (2020) Estimator

Figure OA4: First Stage on Average Hours by Firm

(a) Hours

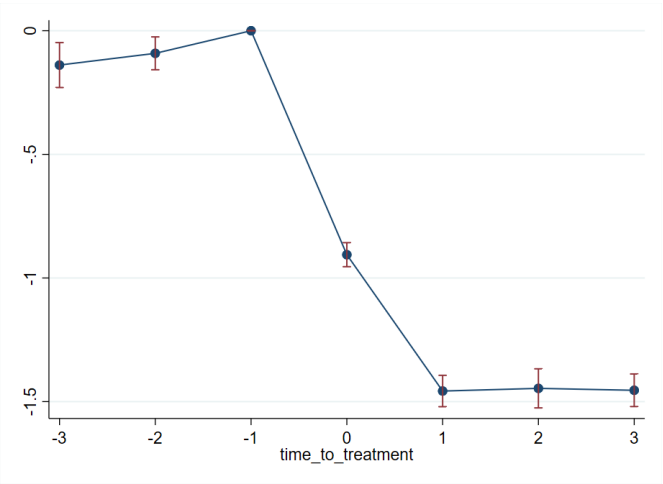

(b) Log of Hours

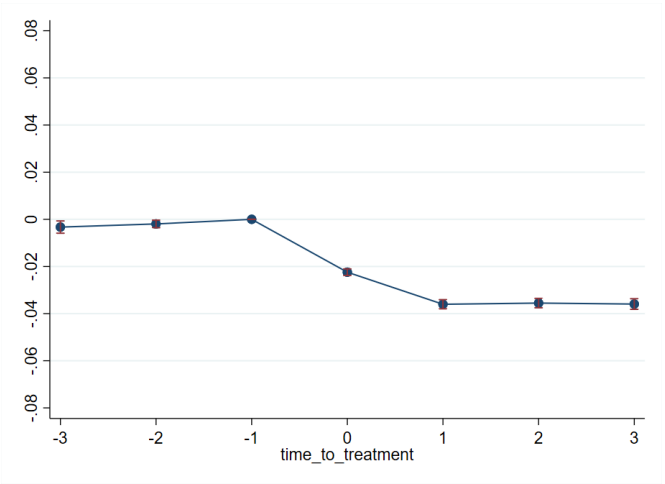

Note: these graphs plot the coefficients of the De Chaisemartin and d'Haultfoeuille (2020) estimator on average hours within the firm, both with hours measured in level, panel (a), and in log, panel(b). The red line indicates the year of implementation of the reform. The reference year for the coefficient is 1996, the vertical bars indicate 95% confidence intervals.
Source: Authors' estimations on QP

Figure OA5: Wages

(a) Monthly Wages

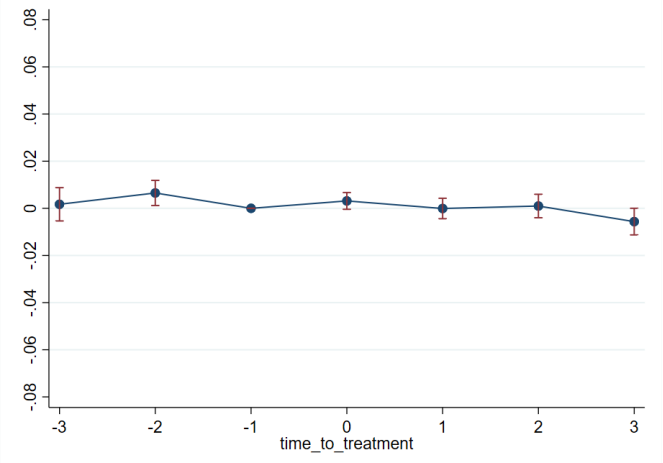

(b) Hourly Wages

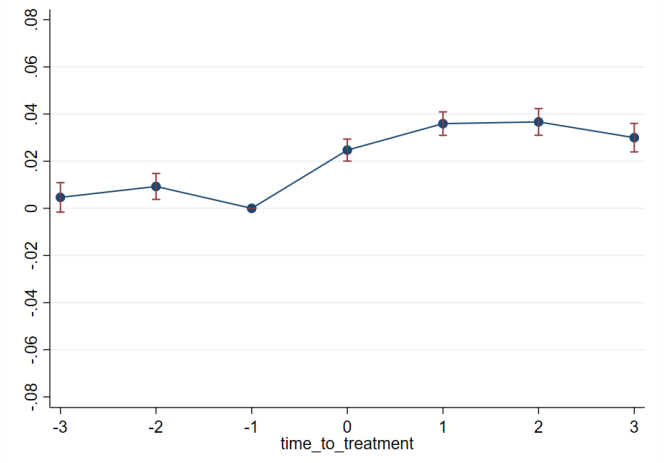

Note: these graphs plot the coefficients of the De Chaisemartin and d'Haultfoeuille (2020) estimator on average monthly wages within the firm (panel a), and hourly wages (panel b). Both outcomes are in logs. The reference year for the coefficient is 1996, the vertical bars indicate 95% confidence intervals.
Source: Authors' estimations on QP

Figure OA6: Number of Workers and Labor Input

(a) Number of Workers

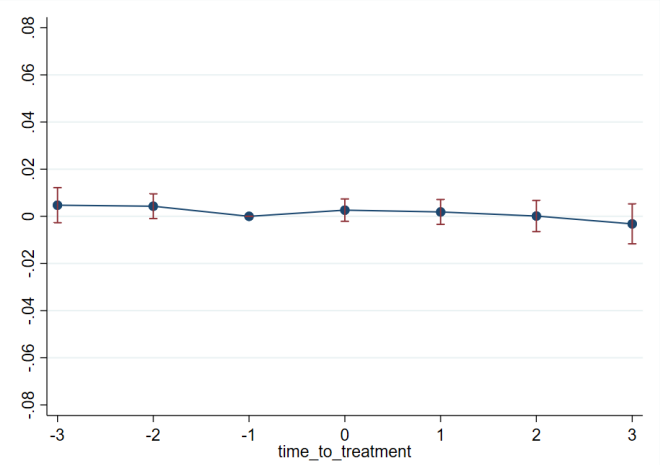

(b) Total Number of Hours

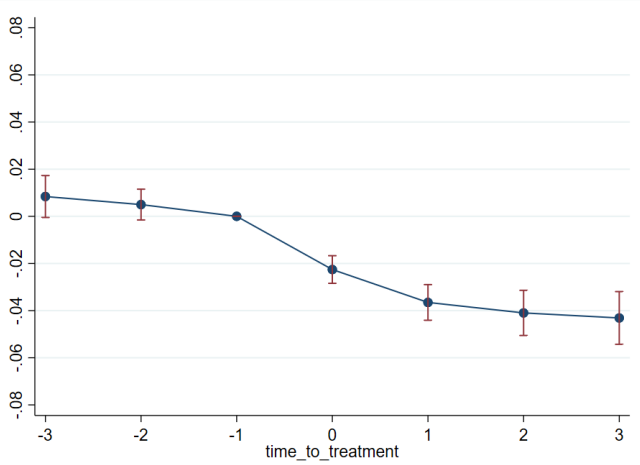

Note: these graphs plot the coefficients of the De Chaisemartin and d'Haultfoeuille (2020) estimator on total employment within a firm (panel a), and total hours worked within a firm (panel b). Both outcomes are in logs. The reference year for the coefficient is 1994, the vertical bars indicate 95% confidence intervals.
Source: Authors' estimations on QP

Figure OA7: Effect on Sales - Coefficients

(a) Total Sales

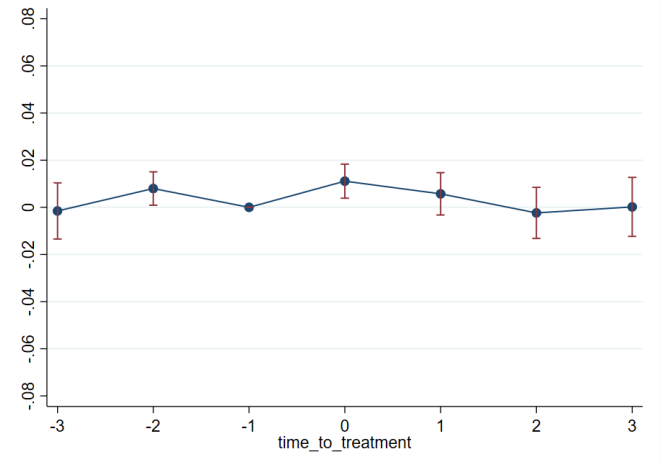

(b) Sales per Hour

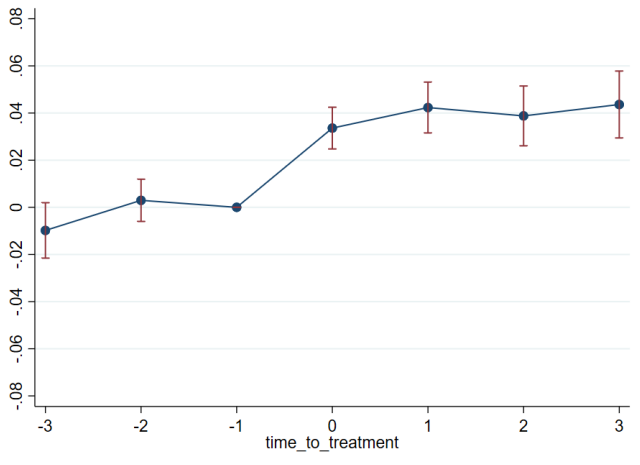

Note: these graphs plot the coefficients of the De Chaisemartin and d'Haultfoeuille (2020) estimator on total sales within a firm, panel (a), and sales per hour, panel (b). Both outcomes are in logs. The reference year for the coefficient is 1996, the vertical bars indicate 95% confidence intervals.
Source: Authors' estimations on QP

OA.3 Data Appendix

This appendix lists how the variables are measured in the administrative data (QP), and also in the Labour Force Survey.

OA.3.1 Quadros de Pessoal (QP - "Lists of Personnel")

Years available: 1985 to 2016

Firm-Level variables

firmbirth: year of firm creation; *legal*: firm legal status; *capital*: firm social capital, in euros; *capitalpriv*: firm share of private domestic capital; *capitalpub*: firm share of public capital; *capitalfor*: firm share of foreign capital; *nut1firm to nut3firm*: firm region at the NUT 1, 2, 3 level; *distfirm*: district location of the firm's headquarter; *municipfirm*: municipality location of firm's headquarter; *caef1 to 6*: economic activity of firm from 1 to 6 digits of disaggregation; *sales*: sales value from October t-1 to October t in euros; *nest*: number of establishments; *workersfirm*: number of workers employed by firm

Establishment-Level

headquarter: dummy equal 1 if establishment is headquarter; *nut1estab to nut3estab*: firm establishment at the NUT 1, 2, 3 level; *distestab*: district location of the establishment; *municipfirm*: municipality location of the establishment; *caest1 to 5*: economic activity of establishment; *workersest*: number of workers of establishment

Worker-Level

nationality: nationality of the worker; *gender*: gender of the work; *workerbirth*: year and month of birth; *age*: age in years of the worker; *hiring date*: year and month of hiring; *tenure*: tenure in years of the worker; *promdate*: year and month of last promotion; *colective*: collective agreement covering the worker; *employment*: worker employment type (employee, employer, self-employed, family-worker); *contract*: worker's contract type (fixed-term contract, permanent contract); *schedule*: worker's schedule (part-time, full-time); *educ1 to 3*: worker's education from 1 to 3 digits; *prof 1 to 6*: worker's profession 1 to 6 digits of disaggregation; *wage*: worker's monthly wage (base plus extra time, plus

bonus); *hours month*: worker's hours in March (before 1993), or October (after 1993); *hours extra*: worker's overtime hours in March (before 1993), or October (after 1993);

OA.3.2 Labour Force Survey

Years available: 1985 to 2016

Worker-Level

ILOSTAT: ILO employment status; *stapro*: worker's professional status (employee, self-employed, employer); *ftpt*: worker's schedule (part-time, full-time); *hwusual*: working hours usually worked in a given week.