Labour Market Effects of Declining Teen Employment †

Afonso C. Leme
Aix-Marseille School of Economics
Job Market Paper

Filipe B. Caires
European University Institute

This version: November 11, 2025

Click Here for Updated Version

Abstract

Teen employment has fallen sharply across advanced economies, reducing firms' access to a traditionally flexible and low-cost source of labour, yet little is known about how they adapt to this shift. This paper studies how firms and labour markets adjust to an exogenous contraction in teen labour supply, exploiting Portugal's 2009 compulsory schooling reform, which raised the minimum school-leaving age to 18. Using matched employer–employee data covering all private-sector firms between 2002 and 2016, we construct sector–municipality exposure measures capturing both the intensity and persistence of pre-reform reliance on teen workers and estimate event-study models comparing more and less exposed local labour markets. The reform led to a sharp and lasting reduction in teen employment and on-the-job training. Firms compensated primarily by hiring slightly older workers, without upgrading skills or wages, suggesting adjustment along cost rather than productivity margins. Ongoing work explores implications for productivity, capital expenditures, and technology adoption.

JEL Classification: J21; J22; J24; I28

†We are grateful for helpful discussions with Thomas F. Crossley, Habiba Djebbari, Arthur Guillouzouic, Andrea Ichino, Morgan Raux, Anna Vitali, and Roberta Ziparo. We thank conference and seminar participants at Aix-Marseille School of Economics, the European University Institute, and Lisbon Micro Group. Filipe B. Caires gratefully acknowledges funding from *Fundação para a Ciência e a Tecnologia* (SFRH/BD/150666/2020) and the Portuguese DG for European Affairs (Bolsa Mário Soares). The project leading to this publication has received funding from the French government under the "France 2030" investment plan managed by the French National Research Agency (reference :ANR-17-EURE-0020) and from Excellence Initiative of Aix-Marseille University - A*MIDEX. We thank access to remote-access data files by Statistics Portugal (*Instituto Nacional de Estatística*) and the Nova School of Business and Economics Social Sciences Datalab for access to QP data in the context of a safecenter.
Leme: afonso.camara-leme@univ-amu.fr.
Caires: filipe.b.caires@eui.eu

1 Introduction

Over the past few decades, teenage employment has declined sharply across advanced economies. In the United States, the labour force participation rate (LFP) of individuals aged 16–19 fell from 53.7% in 1990 to 35.3% in 2019 — a decline of more than one third. A similar pattern emerged in the United Kingdom, where the rate dropped from 58.4% to 34.1% over the same period. Southern Europe stands out as an extreme case: in Italy, Spain, Portugal, and Greece, teen LFP fell by more than half (Figure 1.) This secular decline has been linked to a combination of supply- and demand-side forces, including the expansion of educational opportunities and rising returns to schooling, tighter regulations on the employment of minors, increases in minimum wages, and greater competition from low-skilled immigrant labour (Smith, 2012; Neumark and Shupe, 2019).

While several studies have examined how early work experience affects human capital formation and subsequent labour market outcomes (Kroupova et al., 2024; Le Barbanchon et al., 2023), much less is known about how firms and local labour markets adjust to the disappearance of this segment of the workforce. Teen workers traditionally provided a flexible and low-cost source of labour, often occupying entry-level positions or participating in firm-based training programs. A sustained contraction in teen labour supply thus represents a potentially important shock to firms' labour input composition, training investments, and production decisions. Understanding these adjustment margins is crucial for evaluating the broader consequences of educational expansion and teen labour regulations on firm outcomes and labour market equilibrium.

This paper studies how firms and labour markets adjust to an exogenous decline in the availability of teen workers, exploiting a major compulsory schooling reform in Portugal as a natural experiment. The reform, enacted in 2009, raised the minimum school-leaving age from 15 to 18, effectively banning full-time employment of 16–17 year-olds from 2012 on. The policy induced a sharp and permanent reduction in teen labour supply, providing an ideal setting to study firm-level and local labour market adjustments to the removal of an entire age group from the workforce.

To study the consequences of this drop, we use matched employer–employee data covering the universe of Portuguese private-sector firms between 2002 and 2016 with detailed information on firms' and workers' characteristics. We construct measures of pre-reform exposure to teen labour at the sector–municipality level, capturing both the intensity and persistence of reliance on teen workers, and estimate difference-in-differences

event-study models comparing more and less exposed local labour markets before and after the reform.

The reform led to a pronounced and persistent decline in teen employment in more exposed sector–municipality cells, consistent with a large and exogenous contraction in the supply of teen labour. Firms in high-exposure cells reduced the share of teen workers (aged 16-20) by roughly 4 percent of total employment relative to low-exposure areas. Given that most teens were employed as trainees, these firms also experienced a substantial decline in training intensity.

To compensate for the loss of teen workers, firms primarily substituted toward slightly older cohorts, particularly workers aged 25–35, without increasing educational attainment or wages. This substitution pattern indicates that teen and low-skilled young adult workers are close substitutes in production, and that firms responded to the reform by relying on the next cheapest available labour source rather than upgrading their workforce.

Overall, these results suggest that the decline in teen employment did not induce firms to reorganize production or shift toward higher-skill labour, but instead triggered a reallocation of employment within the lower end of the age and skill distribution. Work in progress examines longer-term consequences for firm dynamics, including productivity, capital expenditures, technology adoption, and profits.

We contribute to three main strands of the literature. First, we contribute to the literature on the decline of teen employment across advanced economies. Existing studies have primarily focused on identifying the causes of this secular trend (Smith, 2012; Neumark and Shupe, 2019). Related work on the effects of early work experience has examined how changes in teen employment affect young people themselves — including their educational attainment, skill formation, and life-cycle earnings (Kroupova et al., 2024; Le Barbanchon et al., 2023). In contrast, much less is known about how firms and labour markets adjust to the disappearance of this segment of the workforce. Our paper fills this gap by examining how firms respond to the sharp and permanent contraction in teen labour supply generated by a compulsory schooling reform, showing that this adjustment primarily occurred through reduced on-the-job training and changes in firms' organization and production decisions.

Second, we contribute to the literature on substitution patterns between low-skilled labour and other production inputs. Prior work shows that children and adults can be

substitutes rather than complements in production (Doran, 2013), and that declines in low-skilled labour supply can induce capital–labour substitution or spur automation (Andersson et al., 2022; San, 2023; Danzer et al., 2024). Our results complement this evidence by showing that when low-skilled labour becomes scarcer, firms primarily substitute toward slightly older and similarly low-skilled workers rather than higher-skill labour. This suggests that adjustment occurred narrowly within the lower end of the skill and age distribution. Ongoing work explores longer-term implications for capital expenditures, and technology adoption.

Finally, our study adds to the large literature on the effects of educational expansion and compulsory schooling reforms. Most of this work focuses on the pecuniary and non-pecuniary returns to schooling for individuals whose educational attainment increased exogenously (Angrist and Krueger, 1991; Oreopoulos, 2007). We instead highlight a novel margin: the impact of such reforms on firms and labour markets through the contraction of youth labour supply. By documenting how a compulsory schooling reform can generate sizeable labour supply shocks with consequences for firms, we bridge the literatures on education policy, child labour regulation, and firm adjustment.

This paper is organized as follows: Section 2 describes the Portuguese compulsory schooling reform and data used. Section 3 describes the main facts characterizing teen employment in Portugal, and Section 4 details the empirical strategy. Section 5 presents the main results, with robustness of these results assessed in Section 6. Finally, Section 7 concludes and discusses the implications of the findings.

2 Institutional Setting and Data

2.1 Compulsory schooling reform

In 2009, Portugal implemented a major education reform that raised the minimum school-leaving age from 15 to 18 and made completion of 12 years of schooling compulsory. The reform was motivated by persistently high rates of early school leaving throughout the 2000s, which placed Portugal among the worst performers in the European Union. In 2008, the early school leaving rate stood at 35%, more than twice the EU average of 14.4%. By 2024, this figure had fallen to 6.6%, below the EU average of 9.4%. Câmara Leme (2024) shows that the reform's impact was particularly strong among disadvantaged students — those from lower socioeconomic backgrounds or with weaker prior academic

performance — indicating that the policy played an important role in extending schooling for the most at-risk groups and reducing educational inequality.

The reform represented a substantial tightening of compulsory schooling requirements. Before its introduction, students could legally leave education after completing the 9th grade — the final year of lower-secondary schooling — or upon turning 15. Under the new regime, compulsory education extends until completion of the 12th grade — the final year of upper-secondary schooling — or until the student turns 18[1]. The law applied to all cohorts entering grade 7 in the 2009/10 academic year (typically aged 12–13), such that the first fully exposed cohort reached upper-secondary education in 2012/13. Upper-secondary schooling in Portugal comprises grades 10 to 12 and includes both academic and vocational tracks. The academic year begins in September and ends in late August of the following calendar year, with the main school holidays taking place between July and September. Compliance with the new law was enforced through potential sanctions for parents or legal guardians, including fines, loss of welfare benefits, or intervention by child protection services in cases of persistent non-attendance.

Implications for teen employment. As full-time employment became incompatible with full-time compulsory schooling during the academic year, the government amended the labour code in 2012 to align with this new reality. The revised code allowed minors under 18 to work only if they had completed compulsory schooling (for example, students still under the previous regime or those who graduated at age 17 under the new law) or if they remained enrolled in upper-secondary education[2]. The minimum legal working age for standard employment remained 16 years, unchanged by the reform. Below that age, minors may be employed only to *light* work[3].

In practice, this legislative change led to a gradual disappearance of 16–17-year-olds from the labour market starting in 2009. The timing depended on students' age and progression through school: those entering grade 7 at the standard age (around 12) were

[1]Students in Portugal can complete upper-secondary school at age 17 if they started school in the year they turned six, progressed without grade repetition, and have birthdays after the end of the school year in July.

[2]The 2009 revision of the Labour Code had already made this restriction implicit by conditioning the employment of 16–17-year-olds on their attendance in education or training leading to completion of compulsory schooling (Lei n.º 7/2009, Artigo 69). The 2012 amendment clarified this provision to remove any remaining ambiguity.

[3]Light work refers to simple tasks that, by their nature or the physical and mental effort required, do not endanger the minor's health, safety, school attendance, or physical, moral, and intellectual development. Before the reform, minors under 16 could be hired for light work only if they had completed compulsory schooling as defined by the previous regime; after the reform, this exception was extended to include those enrolled in upper-secondary education.

first restricted from working in 2013, while students with prior grade retentions — aged 13 or 14 upon entering grade 7 — were affected one or two years earlier, in 2011 or 2012. Any decline in teen employment observed as early as 2009 or 2010 likely reflects firms' anticipatory adjustments to the new legislation or behavioural responses from teens still under the previous regime. Consequently, the supply of teen labour contracted progressively across cohorts, culminating in a complete ban on full-time work by minors under 18 from 2013 onward.

The legal restrictions governing the employment of 16–17-year-olds relative to adults remained largely unchanged following the reform. Minors are prohibited from working overtime or performing night work (between 10 p.m. and 7 a.m.), and they cannot be employed in tasks considered hazardous or detrimental to their health or safety. In addition, they must be granted two consecutive days of weekly rest, whereas adults are only entitled to one, allowing greater flexibility in the organization of the work week for the latter group. Apart from these protections, the standard working time limits are identical for both groups: eight hours per day and forty hours per week.

2.2 Data

This paper leverages *Quadros de Pessoal* (QP), a longitudinal matched employer-employee dataset collected annually by the Ministry for Employment (*Ministério do Trabalho, Solidariedade e Segurança Social*). The survey is mandatory for all private-sector establishments with at least one wage earner and the data reflects the reference month of October each year. QP provides detailed information on firms, establishments, and workers, uniquely identified and traceable over time.

The dataset includes rich firm- and establishment-level information such as location , sector, ownership structure, legal status, and total employment. Worker-level information includes age, gender, educational attainment, qualifications, hiring date, last promotion, contract type (full- vs part-time), occupation, and hours worked, distinguishing between usual and extra hours. It also provides granular wage data, including base pay, regular supplements, and extraordinary payments (e.g., overtime compensation or annual bonuses). Moreover, QP records the collective agreement governing each worker and their professional category, allowing for precise job title definitions.

Beyond its richness, QP offers key additional advantages. Its mandatory nature ensures universal coverage, while employer-reported data mitigates self-reporting bias and

attrition, common in wage earner data. Nevertheless, the dataset has limitations. It captures only a reference month per year, preventing within-year reallocation analyses. Workers who are unemployed or out of the labour force during October are unobserved, and the dataset does not distinguish between exits to unemployment, inactivity, or public-sector employment (excluded from QP).

2.2.1 Sample of firms

We compile and clean the raw datasets to ensure consistency of the longitudinal data, following the procedure of Card and Cardoso (2022).[4]

We restrict the analysis to the 2002–2016 period to examine the impact of the 2009 reform. Publicly participated firms — identified as those with a non-zero share of public capital — are excluded, representing only 0.15% of the QP total sample of firms. We also exclude firms located in the autonomous island regions, which account for 3.5% of all firms of all firms, due to their distinct labour market conditions.

Finally, we exclude outliers in terms of employment and sales growth. We drop the observation if the firm's growth from the previous year in employment or sales is at the top 1% or the bottom 1%. Excluding these outliers is necessary to account for the strong dynamics in some firms that would otherwise significantly increase the noise in the estimations.

3 Teen Employment in Portugal

We begin by presenting several facts that characterize teen employment in Portugal before the compulsory schooling reform, extending the analysis to the post-reform period when relevant. We focus on two groups: (i) minors aged 16–17, and (ii) all teens up to 20 years old.

Sectoral distribution. In 2002, the first pre-reform year in our analysis, teen employment was heavily concentrated in five broad sectors: manufacturing, construction, retail, accommodation and food services, and support service activities[5]. Together, these sectors

[4] Data quality checks and selection criteria are detailed in Appendix A of Card and Cardoso (2022).

[5] According to the NACE 1-letter classification these sectors represent: C - Manufacturing, F - Construction, G - Wholesale and retail trade, I - Accommodation and food service activities, and N - Administrative and support service activities.

accounted for about 95% of minor employment, roughly 91% of total teen employment, 75% of total employment, and around 72% of the total number of firms in the QP data (Table 1). Given this concentration, our analysis throughout the paper focuses on these five sectors.

Skill distribution. Turning to skill composition, the QP classifies workers into five qualification categories based on job-title definitions: trainees, low-skilled, qualified, high-skilled, and top management. In 2002, around half of all minors were trainees, 30% were low-skilled, and 20% were in qualified positions (left-side panel of Figure 2). As expected, minors were virtually absent from high-skilled or management roles. Looking from the opposite perspective — the share of minors within each qualification group — about 3% of all trainees were minors, compared to less than 0.5% among low-skilled or qualified workers (right-side panel of Figure 2). This pattern, consistent with minors' limited experience and education, motivates our focus on lower-skilled segments of the labour market — specifically, trainee and low-skilled positions.

Contract type. Regarding contract types, minor employment was overwhelmingly full-time in the pre-reform years. As seen in Figure 3, the share of part-time contracts never exceeded 10%, indicating that the vast majority of employed minors worked full-time rather than combining work and schooling in the month of October, when QP is measured.

Teen employment trends. Finally, aggregate employment trends show a steady decline in the share of minors and teens in total employment between 2002 and 2016 (Appendix Figure A1). While the downward trend was already visible before the reform, it accelerated markedly from 2011 onward — coinciding with the gradual implementation of the new compulsory schooling requirements.

4 Empirical Strategy

4.1 Local labour market exposure

To analyse the impact of the reform, we look at local labour market units, and create measures of pre-reform exposure to teen labour that capture both the intensity and persistence of exposure, which we use to attribute treatment and control status. Treatment

status is given to units that had persistent high exposure to teen labour, while control status is given to units that only had sporadic high exposure. We exclude units that never employed teens during the pre-reform period, as these represent structurally different local labour markets with no direct exposure to the policy shock.

4.1.1 Local labour market

Our preferred unit of analysis is conducted at the local labour market (LLM) level, defined by three-digit sector–municipality pairs. This definition captures both the geographical and sectoral dimensions of teen employment in Portugal, which exhibited substantial heterogeneity across regions and sectors throughout the 20th century (Goulart and Bedi, 2017).

Based on the evidence presented in the previous section, we focus our analysis on lower-skilled LLMs — those composed of trainee and low-skilled positions — in the five main sectors employing teens: manufacturing, construction, retail, accommodation and food services, and support service activities. To avoid noise from extremely small units, we restrict the sample to LLMs with a minimum pre-reform employment size of at least 10 workers. This ensures that each unit includes at least one *small* firm or several *microenterprises* with a comparable aggregate employment level[6].

Focusing on LLMs rather than individual firms offers several advantages for identifying the effects of the decline in teen employment generated by the compulsory schooling reform. First of all, this aggregation is appropriate given our focus on outcomes such as employment and wages, which are determined in equilibrium at the market level rather than within individual firms. Although LLMs encompass multiple firms with heterogeneous exposure to the shock, the impact of a significant labour supply shock extends beyond directly treated firms — particularly in highly exposed markets — altering both the pool of available workers and the outside options of those currently employed. Restricting attention to variation in firm exposure within LLMs would therefore risk biasing our estimates, since even firms with limited direct exposure may be indirectly affected through market-wide adjustments. In ongoing work, we test the robustness of our results using firms as the unit of exposure and find broadly consistent patterns.

[6]Firms in Portugal are classified according to the European Commission Recommendation 2003/361/EC: micro firms employ fewer than 10 workers and have annual turnover or total assets below €2 million; small firms employ 11–50 workers with turnover or assets below €10 million; and medium-sized firms employ 51–250 workers with turnover below €50 million or total assets below €43 million. For our sample selection we only take into account total employment size.

Second, this approach mitigates potential bias from under-reporting of teen employment in administrative records coming from the QP. As long as reported and actual exposure are positively correlated within each LLM, and informal teen labour supply also responds to the reform, the analysis should still capture meaningful variation in firm outcomes, even if the true intensity of the teen labour supply shock is underestimated due to under-reporting.

Finally, the LLM level provides a more stable measure of exposure to teen labour by aggregating over firms. Portuguese firms are on average very small, and many employ minors only intermittently. As a result, firm-level exposure is highly volatile over time: a small firm may appear "exposed" in one year and "unexposed" in the next, even though it operates in a labour market consistently dependent on teen workers. Aggregation to the LLM level smooths this volatility and captures both the intensity and the persistence of reliance on teen labour within these units.

4.1.2 Pre-reform exposure: intensity and persistence

To measure pre-reform exposure to teen labour, we focus on employment of minors aged 16–17, as this age group was directly affected by the reform. Treatment status is assigned to LLM units that exhibited persistently high exposure to minors before the reform, while control status is attributed to those with only sporadic high exposure.

Specifically, our baseline treatment definition classifies an LLM as treated if it appears above the 90th percentile of the within-year distribution of minor-employment intensity in more than one pre-reform year. Conversely, an LLM is considered a control unit if it reaches the 90th percentile in only one year. Appendix Figure A2 illustrates the distribution of exposure in 2002 — the first pre-reform year — showing the large mass of LLMs with zero exposure in a given year, and the high-exposure tail above the 90th percentile.

This approach identifies groups that are comparable in their underlying reliance on teen labour, but differ in the persistence of that reliance. While both treated and control LLMs share characteristics that make them dependent on teen workers to some extent, those with persistently high exposure are expected to face a sharper contraction in teen labour supply after the reform and, consequently, stronger responses along firm adjustment margins.

To assess robustness, we consider alternative definitions of treatment and control status. Specifically, we examine two additional classifications: (i) LLMs appearing at least

twice in the top quartile of exposure versus those in the top quartile only once; and (ii) LLMs appearing at least three times above the 90th percentile versus those at most twice. Results reported in the robustness section show that the estimated effects become stronger as the treatment definition is made more stringent, consistent with our interpretation.

Finally, an alternative approach would be to divide LLMs into those that ever employed teens and those that never did. However, units that never hired teens during the pre-reform period differ markedly in pre-trends from those that did. For this reason, we restrict our analysis to comparing LLMs with varying degrees of teen labour dependence, rather than including structurally unrelated markets.

4.1.3 Treated and Control Units: Trends and Comparison

Equipped with a treatment definition, we next compare treated and control LLMs. Among the 4,084 local labour market units, 693 are classified as treated and 698 as control, corresponding to roughly 17Figure 4 plots trends over the analysis period in several key variables: the share of minors (aged 16–17) and of all teens (aged ≤ 20) in employment, the number of minors employed, total employment (in both absolute and logarithmic terms), and the number of firms. These variables, along with all others used in the subsequent analysis, are described in Table 2. The first vertical dashed line marks the introduction of the 2009 compulsory schooling reform, while the second marks 2013, when the reform's phased implementation culminated in a complete ban on full-time work by minors under 18.

Both treated and control LLMs were already on a gradual downward trend in teen employment prior to the reform. However, the decline in the share of teen workers accelerates markedly in treated LLMs after 2009, consistent with the contraction in teen labour supply induced by the reform. Treated LLMs are also larger on average, both in total employment and in the number of firms. While these variables display an upward trend until around 2007, followed by a decline beginning in 2010 — likely reflecting the Great Recession — these aggregate movements appear similar across treated and control LLMs.

To estimate the effects of the reform more precisely, we implement this analysis in a difference-in-differences event-study model, which we detail below.

4.2 Identification

We estimate the effects of the reform using an event-study difference-in-differences framework. The specification includes a rich set of fixed effects, allowing potential outcomes for each local labour market unit (sector–municipality pair) to evolve flexibly over time through additive sector and municipality components, following Bergeaud et al. (2025). The identifying variation comes from differences in pre-reform exposure to teen labour across municipalities within the same sector, relative to their municipality-specific trends. Based on the treatment and control definitions described above, we estimate the following regression:

$$Y_{smt} = \alpha_{sm} + \gamma_{st} + \lambda_{mt} + \sum_{\substack{k=2002 \\ k \neq 2008}}^{2016} \delta_k \cdot \mathbb{1}\{t = k\} \cdot D_{sm} + \epsilon_{smt} \tag{1}$$

where Y_{smt} is the outcome of sector s, in municipality m, in year t, such as the share of workers in a given age group relative to total employment. α_{sm} are sector–municipality fixed effects, controlling for time-invariant differences across local labour market units; γ_{st} are sector-year fixed effects, capturing sector-specific time variations flexibly; while λ_{mt} are municipality-year fixed effects, capturing municipality-specific time variations. D_{sm} is the treatment indicator equal to 1 for treated units and 0 for control units. The coefficients δ_k estimate the year-specific effects of the reform relative to the base year 2008, the last pre-reform period. We cluster standard errors at the sector–municipality level, the dimension at which exposure varies, taking into account serial correlation of shocks within units over time.

For expositional clarity, we first describe the identifying assumption in a model that excludes municipality–year fixed effects (λ_{mt}): within each sector, treated and control municipalities would have followed similar trends absent the reform. Including λ_{mt} relaxes this requirement, implying instead that, in the absence of the reform, treated municipalities within a sector would have evolved similarly *relative to their municipality-specific trend* compared to control ones.

Although the parallel-trends assumption cannot be directly tested, finding δ_t not to be significantly different from zero for pre-2009 years provides evidence consistent with the absence of differential pre-trends between treated and control units.

5 Results

We begin by quantifying the magnitude of the teen labour supply shock induced by the reform. After establishing that the share of all workers aged up to 20 fell by roughly 4 percentage points (a 26% decrease), we examine how firms adjusted their workforce composition in response. Finally, we assess how these adjustments translated into changes in key firm and labour market outcomes, such as productivity and wages.

5.1 Labour supply shock

Appendix Figure A3 presents event-study estimates from equation 1 for changes in the share of workers by age group: minors (16–17), 18-, 19-, and 20-year-olds. By 2013, the share of workers in each of these age groups had declined by about 1 percentage point in treated LLMs relative to control ones. In contrast, as shown in Appendix Figure A4, we find no significant effects for older age groups (21–24, analysed individually). While the drop in minor employment is immediate after the reform — likely reflecting firms' anticipatory adjustments or behavioural responses from cohorts still under the previous regime — the decline for 18–20-year-olds unfolds more gradually.

Based on these results, Figure 5 presents the main estimates of the labour supply shock. Panel A reports effects for minors (aged 16–17), while Panel C pools all affected age groups (≤ 20), showing that the total teen employment share declined progressively until 2013 — when the reform's phased implementation culminated in a full ban on full-time work by minors — amounting to a total reduction of about 4 p.p. (a 26% decrease from the pre-reform level of 15Using an alternative measure of exposure — the share of total hours worked by workers in these age groups — Panels B and D reveal declines of nearly identical magnitude. This consistency across measures indicates that our estimated labour supply shock is robust and that the reduction in full-time employment was not offset by increases in part-time work by the same age groups. Finally, the absence of differential pre-trends across treated and control LLMs for all these outcomes supports the validity of our identification strategy.

There are several reasons why employment among 18–20-year-olds also declined, even though these age groups were not directly targeted by the reform. First, some students compelled to remain in school until age 18 because of the reform may have continued beyond the minimum requirement, either because of realized returns to school-

ing or improved educational opportunities. Indeed, Câmara Leme (2024) finds that the probability of remaining in school until at least age 18 or 19 increased by roughly 2.5 percentage points for the first affected cohort, with potentially larger effects for later cohorts as schools became better equipped to accommodate them and public spending on education increased. Second, older teens not directly constrained by the reform may have voluntarily reduced their labour force participation to improve their their educational attainment in response to a new equilibrium where higher levels of education became the norm. Finally, part of the decline may occur mechanically through the loss of *retained* workers: the stock of 18–20-year-olds employed in any given year reflects not only new hires in that age group but also those who began working as minors and remained employed. As the latter pool disappeared following the reform, the share of 18–20-year-old workers naturally fell. To shed light on these possible mechanisms, work in progress looks at hiring and separation margins of these age groups.

5.2 Firms' Adjustment and Workforce Composition

5.2.1 Total employment and substitution pattern

Having established that the reform led to a contraction in the supply of teen labour, we now turn to how firms adjusted their overall employment. In particular, we consider three possible adjustment patterns: less than one-for-one, one-for-one, and more than one-for-one replacement of teen workers.

A first possibility is that total employment declines (a less-than-one-for-one replacement): firms may struggle to compensate for the loss of teen labour due to skill mismatches, hiring frictions, or because the reduction in workers for the average firm is too small to justify replacement. Alternatively, they may respond by substituting low-skilled labour with capital. A second possibility is a one-for-one substitution, whereby firms replace each lost teen worker with a relatively older employee, leaving total employment unchanged. Finally, total *reported* employment could even increase (a more-than-one-for-one replacement). This could occur if the reform effectively removed both reported and unreported teen workers from the labour market, prompting firms to compensate for these losses by hiring additional, older workers whose employment is formally reported.

Results in 6 show no evidence of a change in the logarithm of total employment following the reform, consistent with a one-for-one substitution of teen workers by other

employees. However, adjustment patterns may vary across local labour markets. Ongoing work explores heterogeneity by sector and other LLM characteristics, as well as the potential labour–capital substitution channel.

5.2.2 Workforce Composition

Age composition. Having shown that firms replaced teen workers with other employees, we now examine how the age composition of the workforce adjusted. Specifically, we analyse changes in employment shares across two broad age groups: (i) younger adults aged 25–34, and (ii) older adults aged 35 and above. Results in Panel A of Figure 7 indicate that the share of younger adults increased significantly by around 2.5 p.p. following the reform, while the share of older adults remained essentially unchanged (Panel B, Figure 7).

Demographic composition. We next examine how the demographic composition of the workforce adjusted, focusing on the gender balance and the share of foreign-born workers. As shown in Panels C and D of Figure 7, we find no significant changes in either of these dimensions.

Educational composition. We then turn to the effects of the reform on the educational composition of the workforce. Namely, we analyse changes in the share of workers with (i) at most lower-secondary education, (ii) at most an upper-secondary diploma, and (iii) a university degree. Panel E of Figure 7 shows estimates for the share of workers with at most an upper-secondary diploma, while Appendix Figure A5 shows estimates for the other two shares. We find no evidence of significant changes in any of these shares following the reform. While the reform likely increased the supply of workers with upper-secondary education, the first affected cohort would have completed that level no earlier than 2016. Hence, our observation window may be too short to capture these effects, and extending the analysis beyond 2016 could reveal longer-term adjustments.

Skill composition. Next, we examine how the skill composition evolved after the reform. Because our analysis focuses on lower-skilled LLMs — composed of trainee and low-skilled positions — any change in the share of one qualification category mechanically mirrors the other. Panel F of Figure 7 shows that the share of trainees declined by nearly 4 p.p. As discussed in Section 3, since most teens were employed as trainees, this

reduction implies a substantial drop in training intensity among firms in treated LLMs.

Mechanism behind compositional change. For all the significant changes in workforce composition documented above, it is important to distinguish to what extent they arise mechanically from the disappearance of teens or reflect compositional differences among newly hired workers who replaced them. To shed light on this mechanism, ongoing work examines the characteristics of new hires following the reform.

5.3 Firm Performance and Wage Effects

Having shown that firms replaced teen workers with other employees — and that this substitution altered workforce composition, notably through an increase in the share of young adults (aged 25–34) and a decline in the share of trainees — we now turn to examining how these adjustments affected firm survival and wages. Ongoing work explores implications for productivity, profits, capital expenditures, and technology adoption.

Firm survival. Figure 8 presents the estimated effects on the number of firms (in log levels) and on the share of firms exiting the market in each unit and year. Firm survival could be affected for two main reasons. First, the contraction in teen labour supply may have led firms that relied heavily on this workforce to face hiring constraints and eventually exit the market. Second, the Great Recession may introduce bias if, conditional on the fixed-effects structure of equation (1), exposure to teen labour is correlated with the likelihood of being more severely affected by the crisis. We find no significant effects on either outcome, suggesting that firm survival was largely unaffected by the reform.

Wages. Figure 9 presents the estimated effects of the reform on the logarithm of average wages. In a standard competitive framework, a negative labour supply shock would be expected to raise wages. However, we find no evidence of such an increase. Although there is a positive and significant effect in 2016 — the final year of our analysis period — its magnitude is similar (in absolute terms) to a significant effect for the pre-reform year 2002. Overall, the pattern of results is not consistent with a progressive rise in wages accompanying the decline in teen labour supply shown in Figure 5.

One plausible explanation is the presence of wage rigidities in these low-skilled labour markets, where a substantial share of workers earn the minimum wage. Such rigidity could have prevented wages from adjusting upward despite the reduction in labour

supply.

6 Robustness

6.1 Robustness to Treatment Definition

Our baseline treatment definition considers an LLM as treated if it appears above the 90th percentile of the within-year distribution of minor-employment intensity in more than one pre-reform year, and as a control unit if it reaches the 90th percentile in only one year. Because our objective is to compare LLMs that exhibited persistently high exposure to minors before the reform (the treated group) with those that only had sporadic exposure, we assess how sensitive our results are to alternative thresholds of exposure persistence.

Figures A6 and A7 respectively present estimates for the main outcomes under two alternative classifications: (i) LLMs appearing at least twice in the top quartile of exposure versus those in the top quartile only once; and (ii) LLMs appearing at least three times above the 90th percentile versus those at most twice. Estimated effects strengthen as the treatment definition becomes more stringent, consistent with our interpretation that higher pre-reform exposure implies a sharper contraction in teen labour supply after the reform and, consequently, stronger firm-level adjustment responses.

7 Conclusion

This paper set out to understand how firms and labour markets adjust to a sharp and permanent contraction in the supply of teenage workers. Teen employment has declined steadily across advanced economies, yet the literature has focused on its causes and on the consequences of early work experience for the individuals directly affected. Much less is known about how firms and markets reorganize when an entire segment of the low-skilled workforce disappears. We address this question by exploiting Portugal's 2009 compulsory schooling reform, which raised the minimum school-leaving age from 15 to 18 and effectively banned full-time work by minors from 2013 onward. The reform provides a natural experiment that generated a contraction in the availability of teen labour, without confounding changes in demand.

Using matched employer–employee data covering the universe of private-sector firms

between 2002 and 2016, we measure pre-reform exposure to teen labour across sector–municipality cells and implement an event-study difference-in-differences design. Our approach compares local labour markets that relied more or less intensively and persistently on teen workers before the reform, controlling flexibly for sectoral and municipal trends. This framework allows us to trace firm and labour market responses to the removal of an entire age group from the workforce.

We find that the reform induced a sharp and lasting decline in teen employment. The share of workers aged 16–20 fell by around four percentage points — a 26 percent reduction relative to the pre-reform average — while on-the-job training intensity declined parallelly. Firms compensated for this loss primarily by hiring slightly older, low-skilled young adults aged 25–34, suggesting that teen and young adult workers are close substitutes in production. Despite this compositional shift, we find no evidence of changes in average wages or educational attainment. Nor do we observe effects on firm survival, implying that firms were able to adjust through substitution within the lower end of the skill and age distribution.

Our findings highlight the importance of considering firm behaviour when evaluating the broader consequences of educational expansion and child labour regulation. While such reforms can benefit young individuals through higher schooling and human capital accumulation, they can also reshape the composition of local labour markets and firms' training investments. Ongoing work examine the longer-term implications of this adjustment process, including effects on firm productivity, capital investment, and the evolution of training practices once the reform became fully embedded in the workforce. By shedding light on how firms adapt to structural shifts in the availability of low-skilled labour, this research contributes to a broader understanding of the interaction between education policy, labour supply, and firm dynamics.

References

Andersson, D., M. Karadja, and E. Prawitz (2022). Mass migration and technological change. *Journal of the European Economic Association 20*(5), 1859–1896.

Angrist, J. D. and A. B. Krueger (1991). Does compulsory school attendance affect schooling and earnings? *The Quarterly Journal of Economics 106*(4), 979–1014.

Bergeaud, A., A. Guillouzouic, E. Henry, and C. Malgouyres (2025). From public labs to private firms: Magnitude and channels of local r&d spillovers. *The Quarterly Journal of Economics 140*(4), 3233–3282.

Câmara Leme, A. (2024). Too young to quit school? increasing the compulsory schooling leaving age and students' educational paths. *AMSE Working paper series WP 2024 - Nr 36*.

Card, D. and A. R. Cardoso (2022). Wage flexibility under sectoral bargaining. *Journal of the European Economic Association 20*(5), 2013–2061.

Danzer, A. M., C. Feuerbaum, and F. Gaessler (2024). Labor supply and automation innovation: Evidence from an allocation policy. *Journal of Public Economics 235*, 105136.

Doran, K. B. (2013). How does child labor affect the demand for adult labor?: Evidence from rural mexico. *Journal of Human Resources 48*(3), 702–735.

Goulart, P. and A. S. Bedi (2017). The evolution of child labor in portugal, 1850–2001. *Social Science History 41*(2), 227–254.

Kroupova, K., T. Havranek, and Z. Irsova (2024). Student employment and education: A meta-analysis. *Economics of Education Review 100*, 102539.

Le Barbanchon, T., D. Ubfal, and F. Araya (2023). The effects of working while in school: Evidence from employment lotteries. *American Economic Journal: Applied Economics 15*(1), 383–410.

Neumark, D. and C. Shupe (2019). Declining teen employment: minimum wages, returns to schooling, and immigration. *Labour Economics 59*, 49–68.

Oreopoulos, P. (2007). Do dropouts drop out too soon? wealth, health and happiness from compulsory schooling. *Journal of public Economics 91*(11-12), 2213–2229.

San, S. (2023). Labor supply and directed technical change: Evidence from the termination of the bracero program in 1964. *American Economic Journal: Applied Economics 15*(1), 136–163.

Smith, C. L. (2012). The impact of low-skilled immigration on the youth labor market. *Journal of Labor Economics 30*(1), 55–89.

Figures

Figure 1

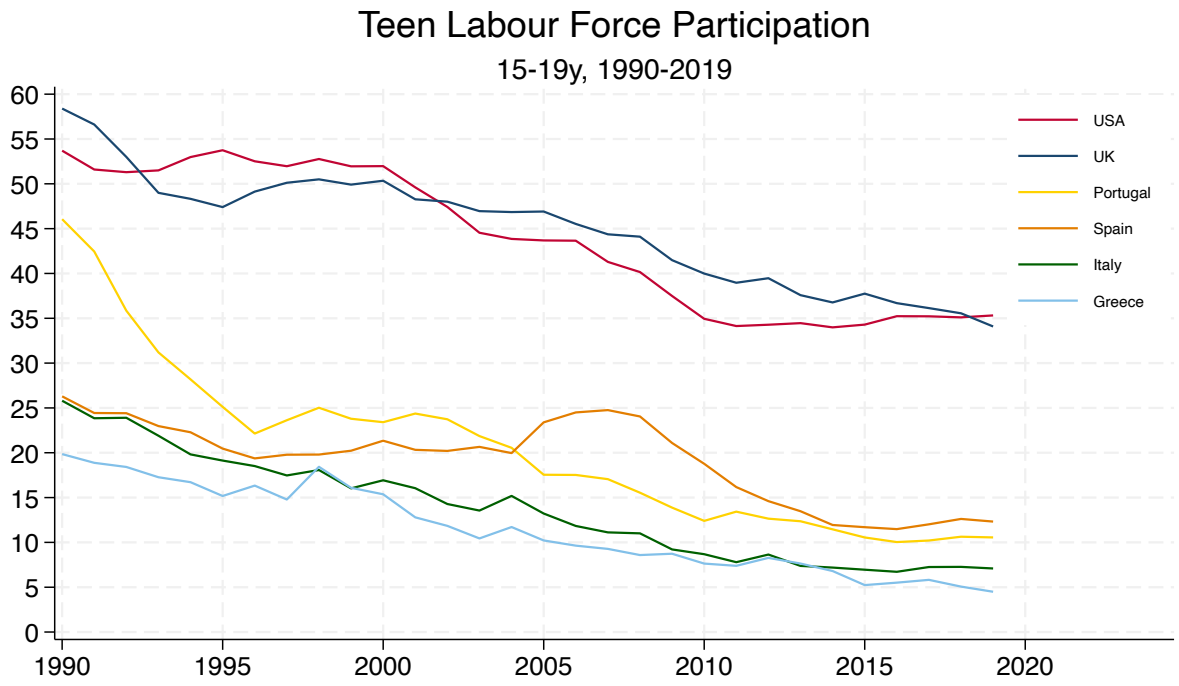

Teen labour force participation rates (%) for individuals aged 15–19 across selected countries, 1990–2019. Annual rates are calculated as averages of quarterly or monthly data, depending on each country's reporting frequency. Source: International Labour Organization (ILO).

Figure 2

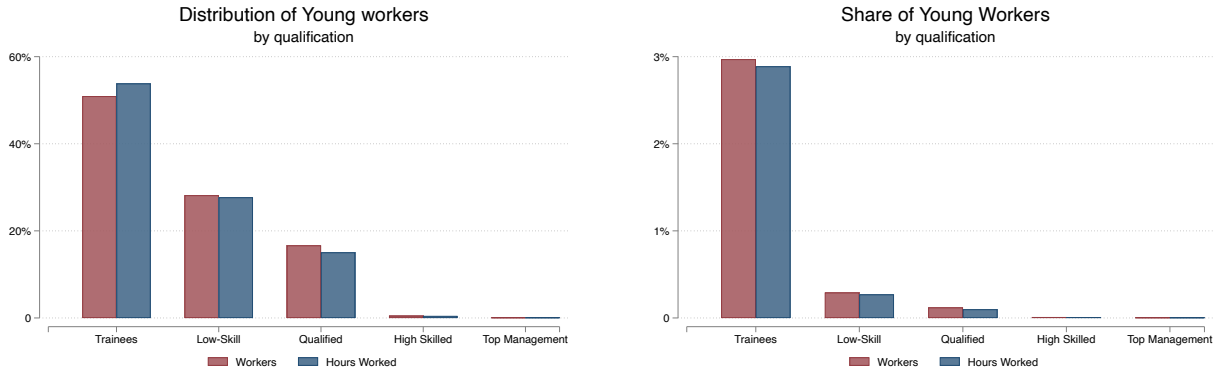

The left panel shows the distribution of minors workers (age 16-17) across qualification categories, while the right panel displays their share within each qualification group in 2002. The *Quadros de Pessoal* data classifies workers into five qualification categories based on job-title definitions: trainees, low-skilled, qualified, high-skilled, and top management.

Figure 3: Share of minors in full- and part-time contracts

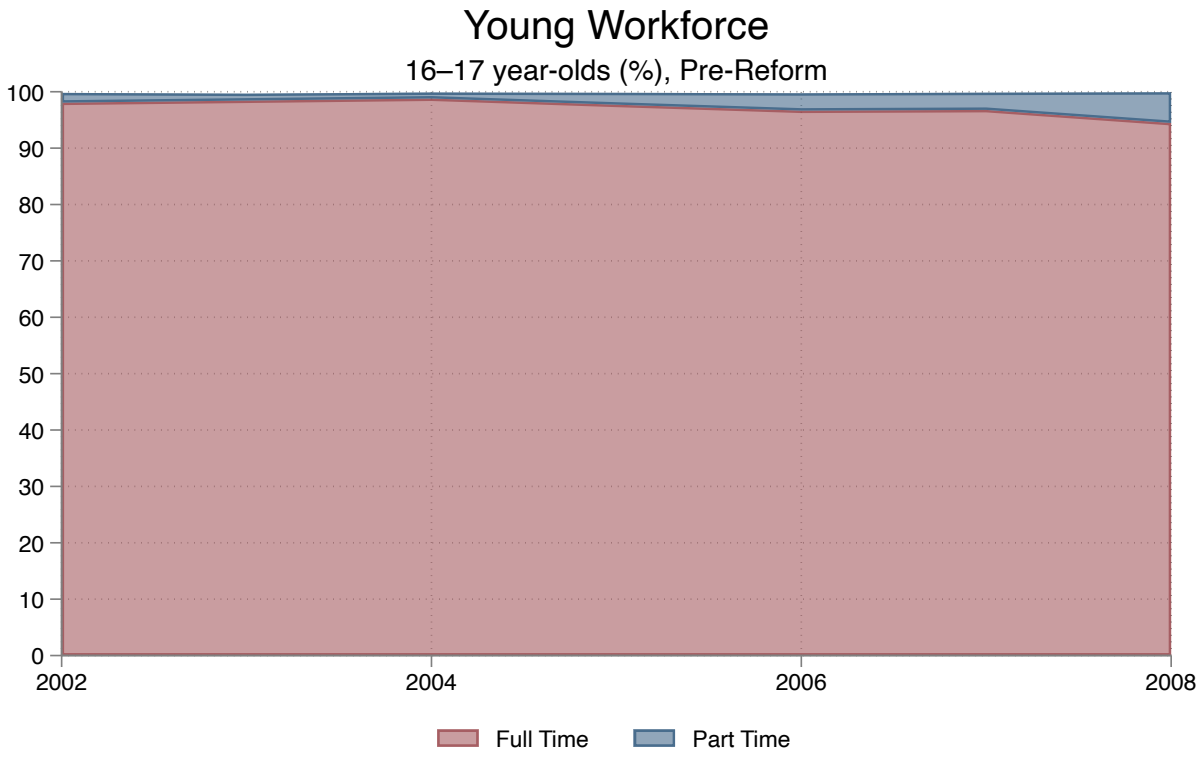

Share of minor workers (age 16-17) employed full-time and part-time in the pre-reform period.

Figure 4: Trends in key variables for treated and control local labour markets

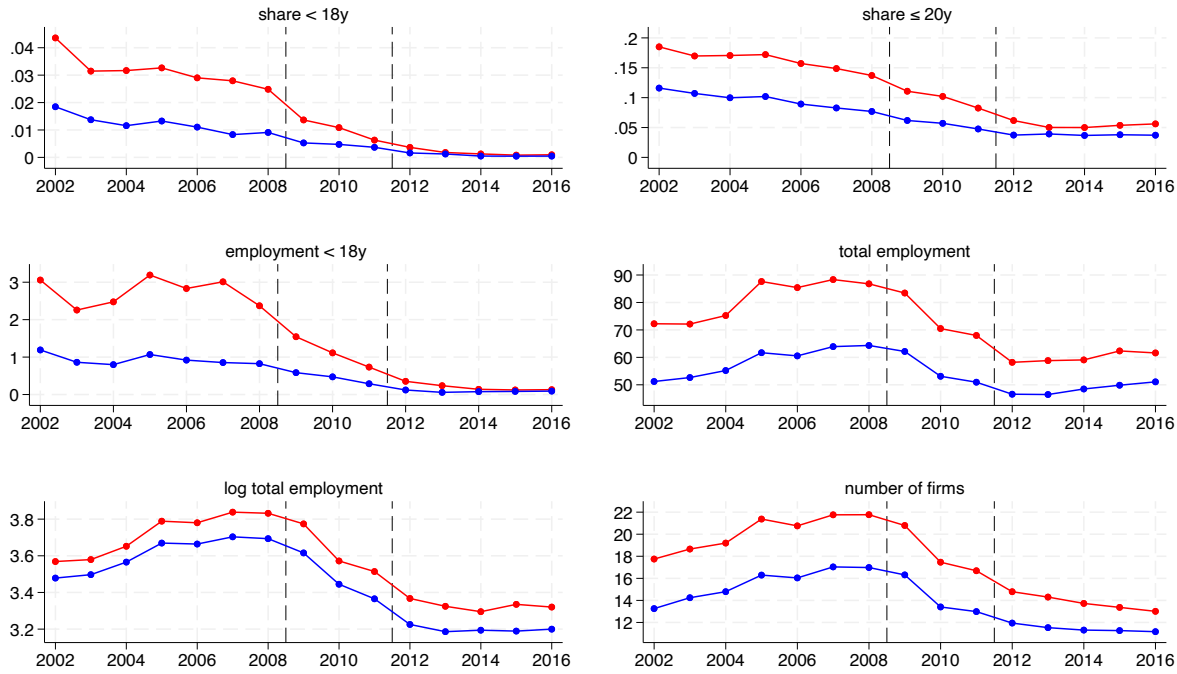

Trends over the analysis period (2002–2016) for key variables, shown in red for treated units and blue for control units. Variables are defined in Appendix Table 2.

Figure 5: Teen labour supply shock

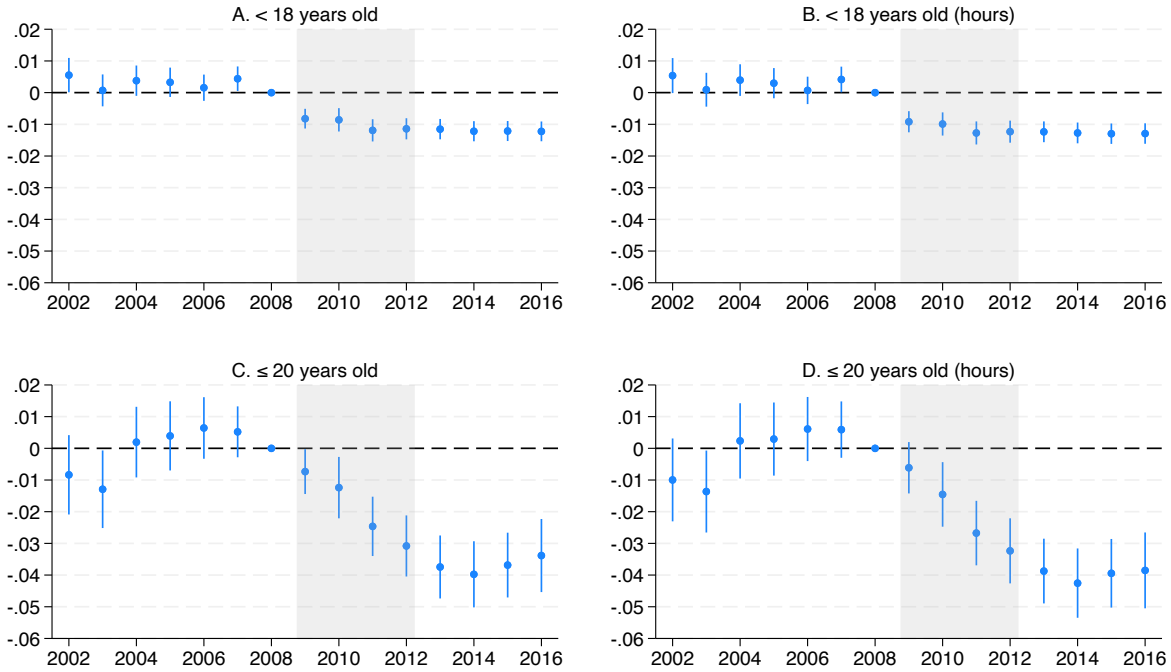

Panels A–D plot event-study difference-in-differences estimates from equation 1 for the share and total hours of teen employment in treated relative to control local labour markets. Panels A and B show effects on workers aged below 18, measured respectively in shares and total hours worked. Panels C and D present analogous estimates for all workers aged 20 or below. Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around the point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.

Figure 6: Total employment effects

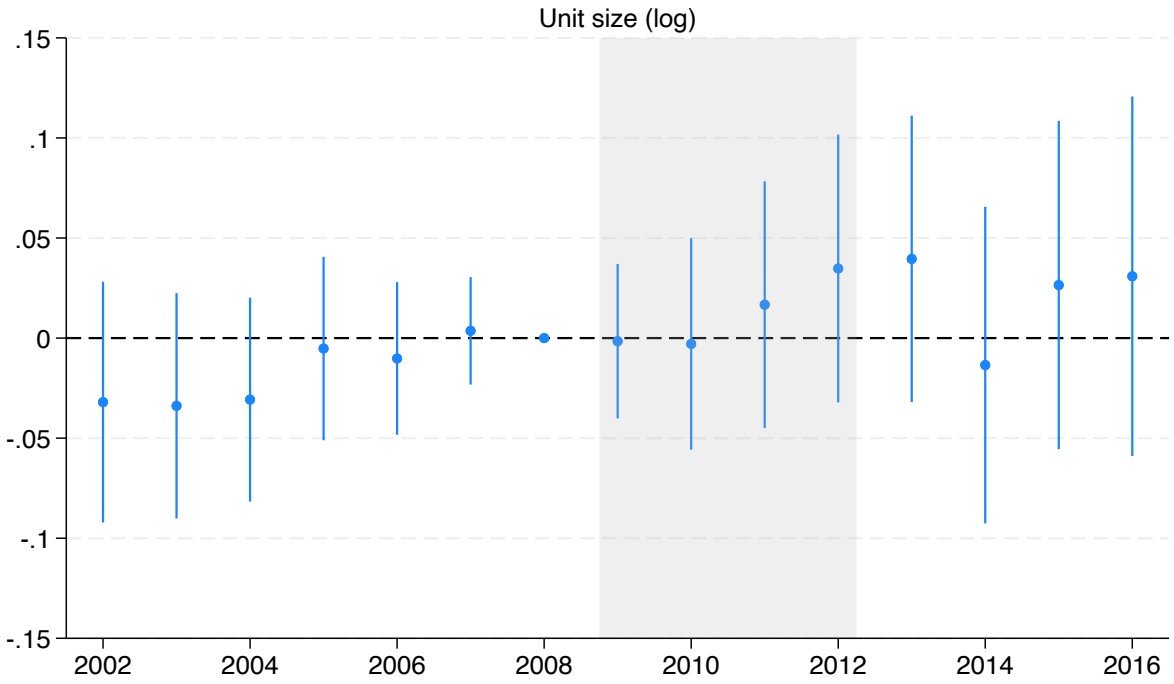

Event-study difference-in-differences estimates from equation 1 for the logarithm of total employment in treated relative to control local labour markets. Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around the point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.

Figure 7: Workforce Composition effects

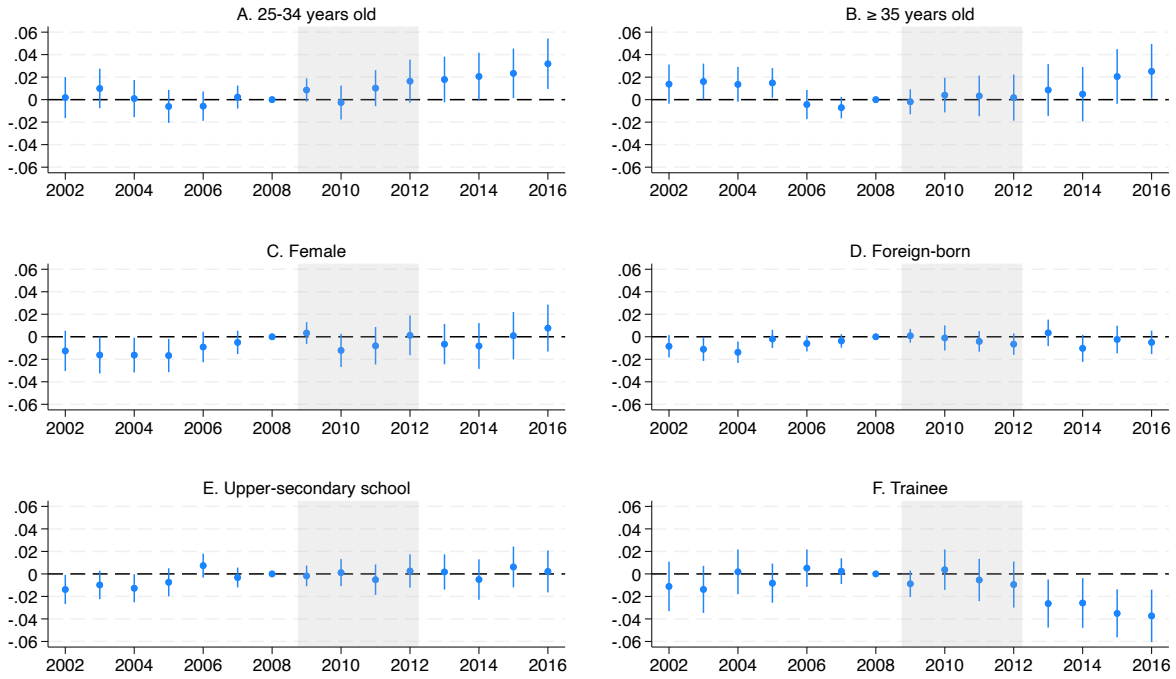

Panels A–F plot event-study difference-in-differences estimates from equation 1 for the share of workers in selected demographic and qualification groups in treated relative to control local labour markets. Panels A and B show effects by age group (25–34 and ≥35 years old), Panels C and D by demographic characteristics (female and foreign-born), and Panels E and F by education and qualification (upper-secondary school and trainee). All variables are defined in Appendix Table 2. Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around the point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.

Figure 8: Firm survival effects

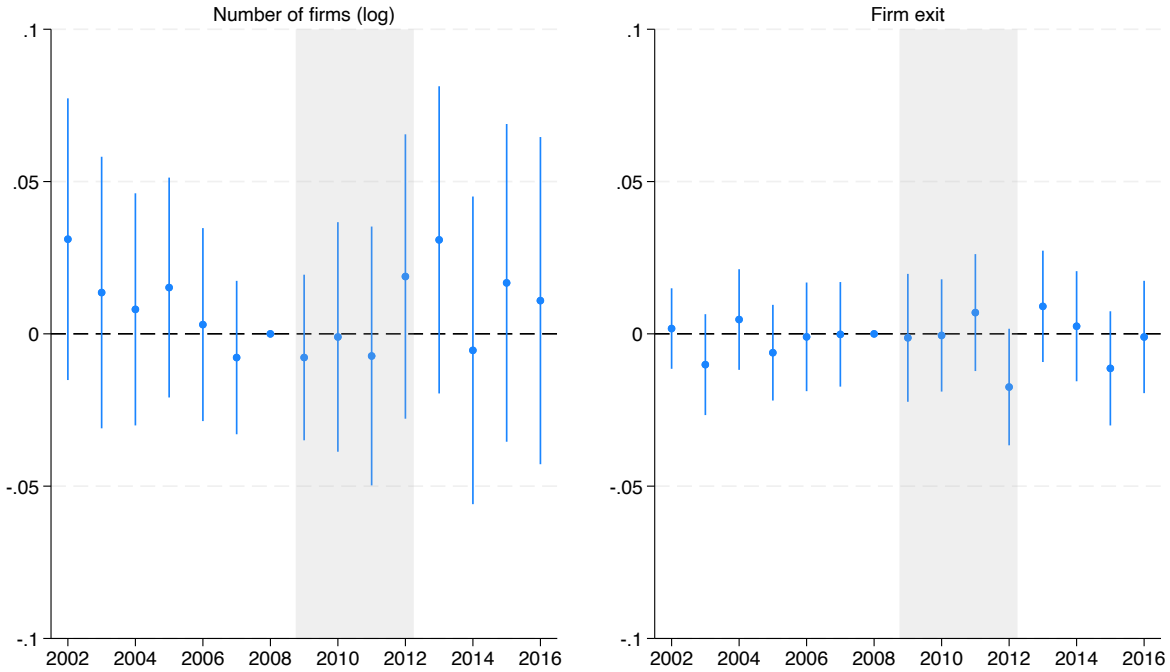

Panels A and B plot event-study difference-in-differences estimates from equation 1 for firm survival outcomes in treated relative to control local labour markets. Panel A shows effects on the logarithm of the number of firms, while Panel B reports effects on the share of firms exiting the market from year $t-1$ to year t. Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around the point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.

Figure 9: Wage effects

Event-study difference-in-differences estimates from equation 1 for the logarithm of average wages in treated relative to control local labour markets. Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around the point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.

Tables

Table 1: Sectoral Distribution of Employment and Firms in 2002

1-Letter Sector	Share minors	Share ≤ 20 years old	Share Total Employment	Share Firms	Cumulative share Minors	Cumulative share ≤ 20 years old	Cumulative share Total Employment	Cumulative share Firms
C - Manufacturing	0.445	0.344	0.303	0.169	0.445	0.344	0.303	0.169
F - Construction	0.205	0.148	0.090	0.153	0.649	0.492	0.393	0.322
G - Wholesale & Retail	0.157	0.258	0.162	0.257	0.806	0.750	0.555	0.578
I - Accommodation & food service activities	0.093	0.076	0.079	0.119	0.899	0.826	0.634	0.697
N - Administrative and support service activities	0.047	0.085	0.117	0.021	0.947	0.912	0.750	0.718
A - Agriculture, forestry and fishing	0.013	0.012	0.033	0.060	0.960	0.923	0.783	0.779
S - Other services activities	0.009	0.015	0.031	0.050	0.969	0.939	0.814	0.829
M - Professional, scientific and technical activities	0.008	0.016	0.028	0.042	0.977	0.955	0.841	0.871
Q - Human health and social work activities	0.005	0.016	0.084	0.052	0.982	0.970	0.926	0.923
H - Transporting and storage	0.005	0.007	0.013	0.017	0.987	0.977	0.939	0.940

Shares of minors (aged 16–17), all teens (aged ≤ 20), total employment, and firms across the top 10 one-letter NACE sectors ranked by minor employment in 2002. Cumulative shares are calculated sequentially by sector, ordered by the share of minor employment.

Table 2: Description of Variables Used in the Analysis

Variable	Description	Type
Share of age group y	Share of workers aged y to total workers in the local labour market (e.g., share of minors aged 16–17).	Share [0,1]
Share of hours worked by age group y	Share of hours worked by workers aged y to total hours worked in the local labour market (e.g., share of hours worked by minors)	Share [0,1]
Total employment	Total number of workers in the local labour market	Discrete
Share basic education	Share of workers with at most lower-secondary education out of total workers in the local labour market	Share [0,1]
Share upper-secondary school	Share of workers with at most an upper-secondary school diploma out of total workers in the local labour market.	Share [0,1]
Share university	Share of workers with a university degree out of total workers in the local labour market.	Share [0,1]
Share trainee	Share of workers classified as trainees out of total workers in the local labour market.	Share [0,1]
Share female	Share of female workers out of total workers in the local labour market.	Share [0,1]
Share foreign-born	Share of foreign-born workers out of total workers in the local labour market.	Share [0,1]
Number of firms	Number of firms in the local labour market	Discrete
Firm exit	Share of firms exiting the local labour market from year t-1 to year t	Share [0,1]
Wages	Average wages in the local labour market	Continuous

Definitions and types of all variables used in the empirical analysis. All variables are computed at the local labour market (sector–municipality) level unless otherwise noted.

Appendix

Additional figures

Figure A1: Trends in share of workers aged ≤20 and <18

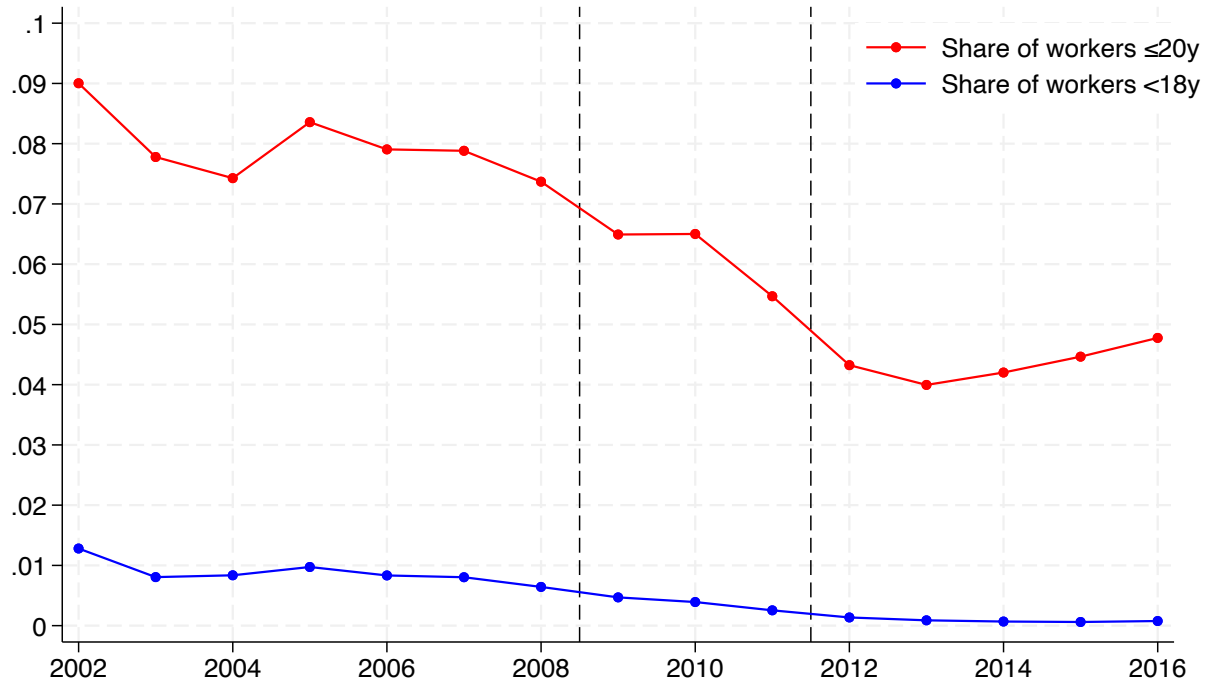

Share of workers aged ≤20 and <18 to total workers in the five main 1-letter sectors employing minor workers: manufacturing, construction, retail, accommodation and food services, and support service activities. The first vertical dashed line marks the introduction of the 2009 compulsory schooling reform, while the second marks 2013, when the reform's phased implementation resulted in a complete ban on full-time work by minors under 18.

Figure A2: Distribution of the share of workers aged 16-17 in 2002

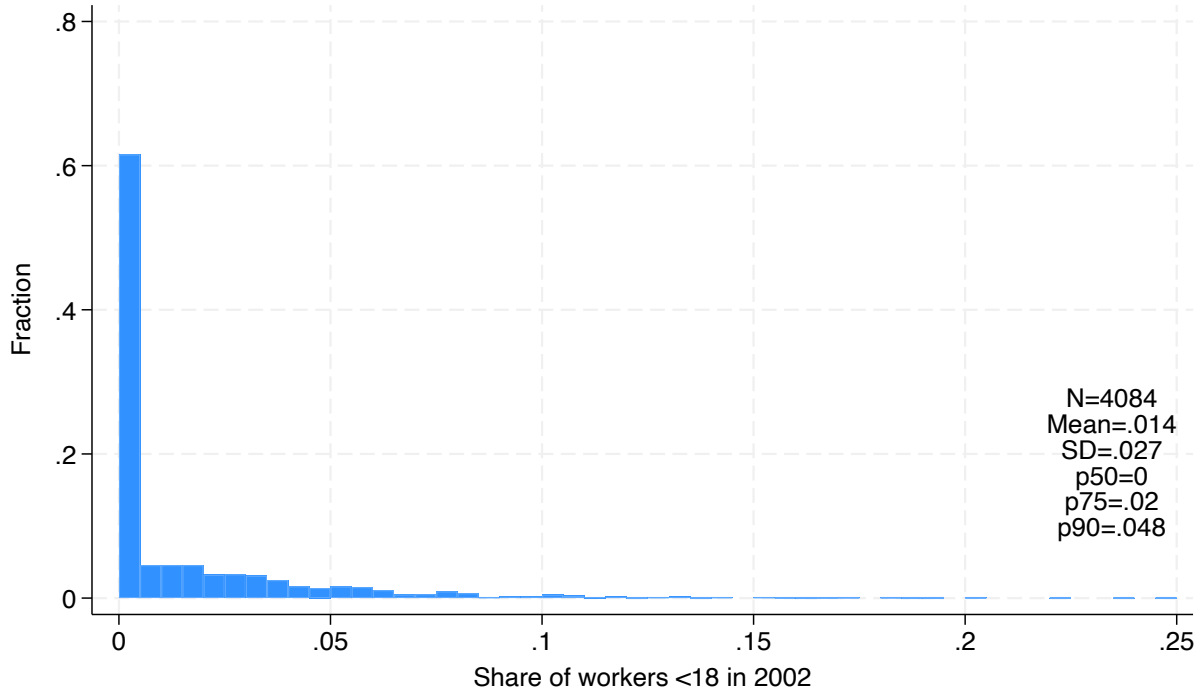

Distribution of the share of workers aged 16-17 to total workers across local labour markets in 2002. The mass at zero corresponds to 60%.

Figure A3: Teen labour supply shock: individual age groups

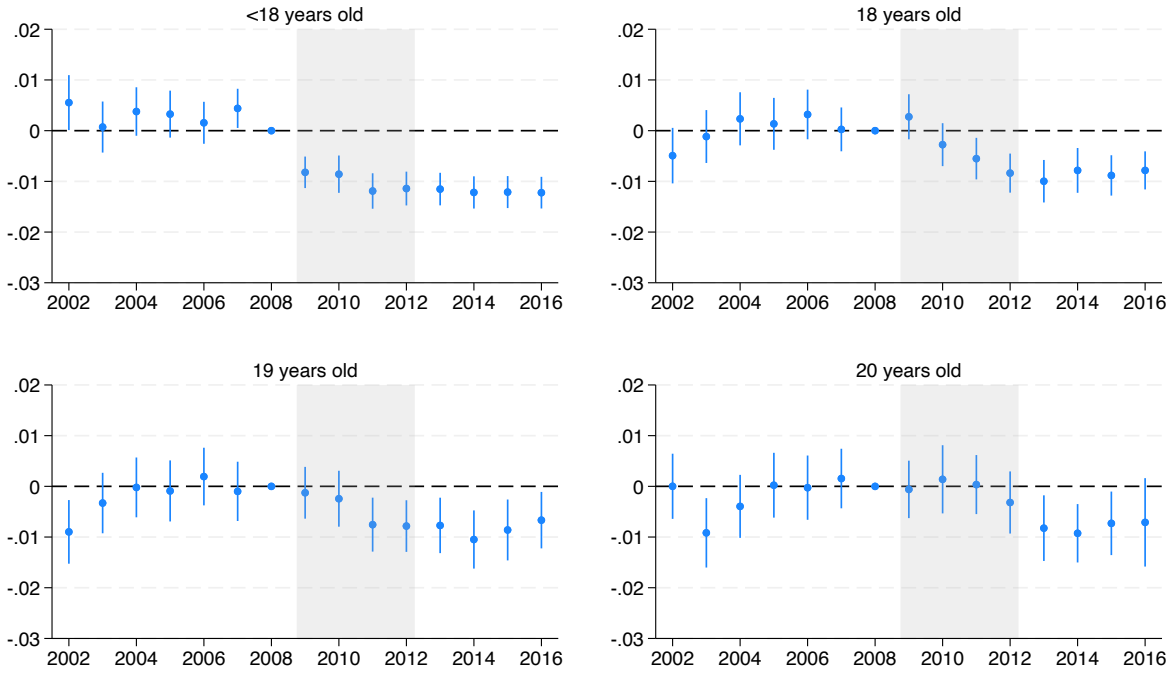

Panels A–D plot event-study difference-in-differences estimates from equation 1 for the share of workers by individual age groups in treated relative to control local labour markets. Panels correspond respectively to workers aged below 18, 18, 19, and 20 years old. Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around the point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.

Figure A4: Labour supply shock: older age groups

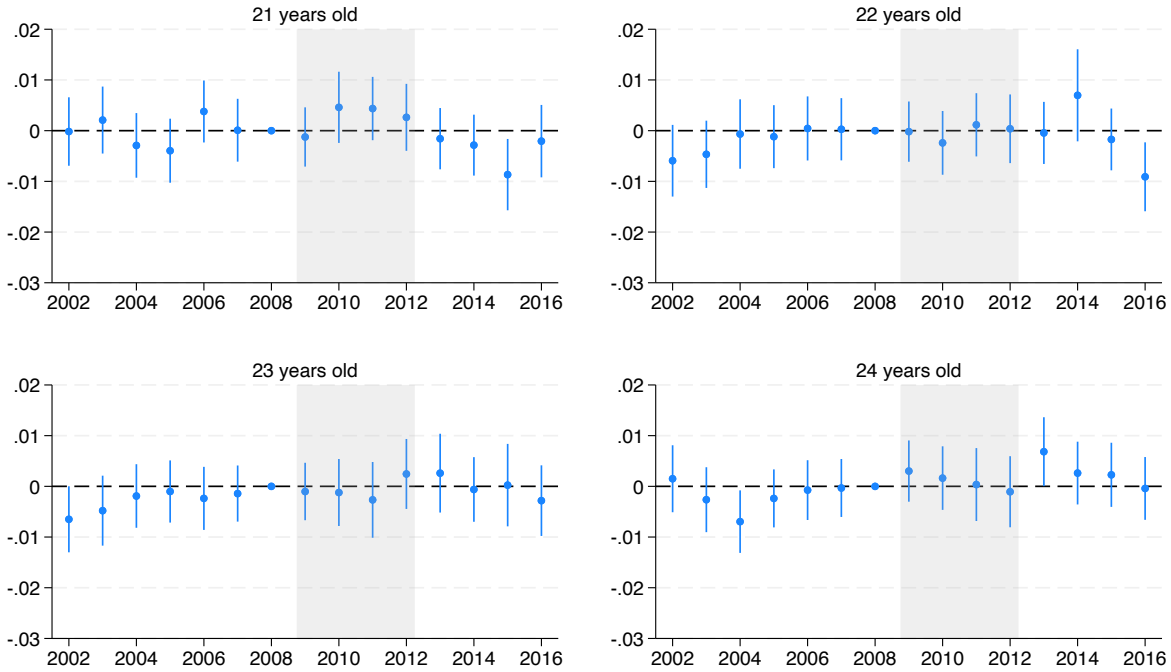

Panels A–D plot event-study difference-in-differences estimates from equation 1 for the share of workers by individual age groups in treated relative to control local labour markets. Panels correspond respectively to workers aged 21, 22, 23, and 24 years old. Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around the point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.

Figure A5: Educational composition effects: basic and university education levels

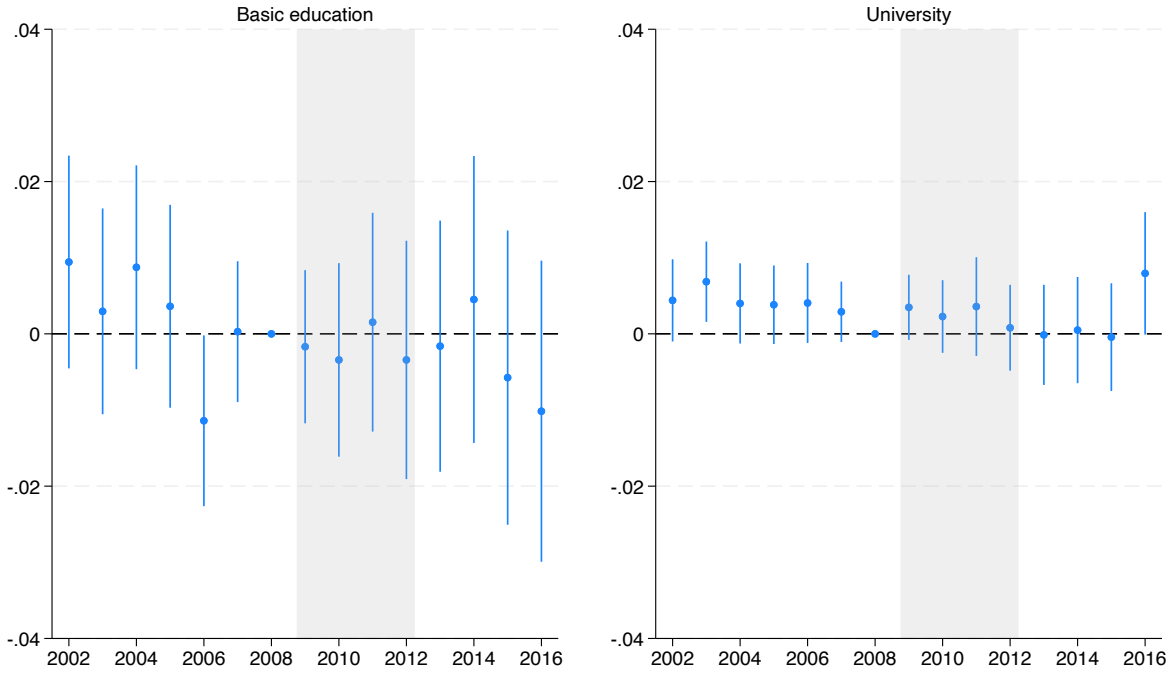

Panels A and B plot event-study difference-in-differences estimates from equation 1 for the share of workers with different education levels in treated relative to control local labour markets. Panel A shows effects for workers with at most basic education, and Panel B for those with a university degree. Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around the point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.

Figure A6: Robustness to Treatment Definition: Less Stringent Exposure Threshold

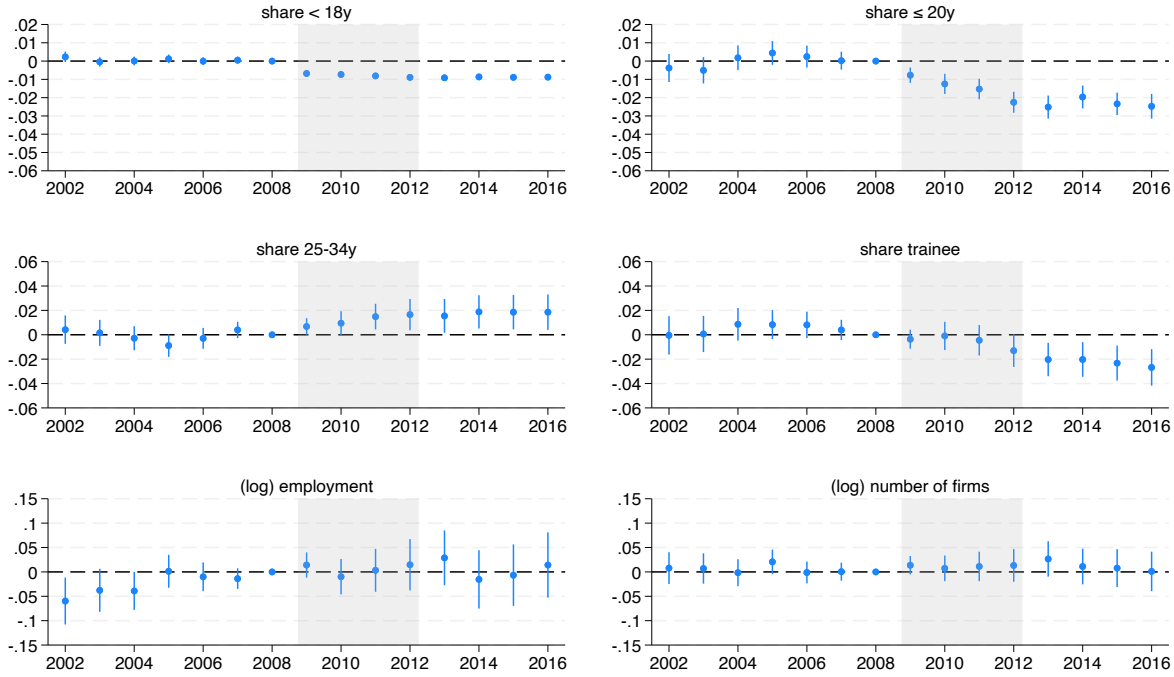

Event-study difference-in-differences estimates from equation 1 for the main outcomes using an alternative treatment definition. Treated local labour markets are defined as those appearing at least twice in the top quartile of minor-employment intensity in the pre-reform period, while control units appear in the top quartile only once. Panels display estimated effects for key outcomes: the share of minors (Panel A), the share of all teens aged ≤ 20 (Panel B), the share of young adults aged 25–34 (Panel C), the share of trainees (Panel D), total employment (Panel E), and the number of firms (Panel F). Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.

Figure A7: Robustness to Treatment Definition: Stricter Exposure Threshold

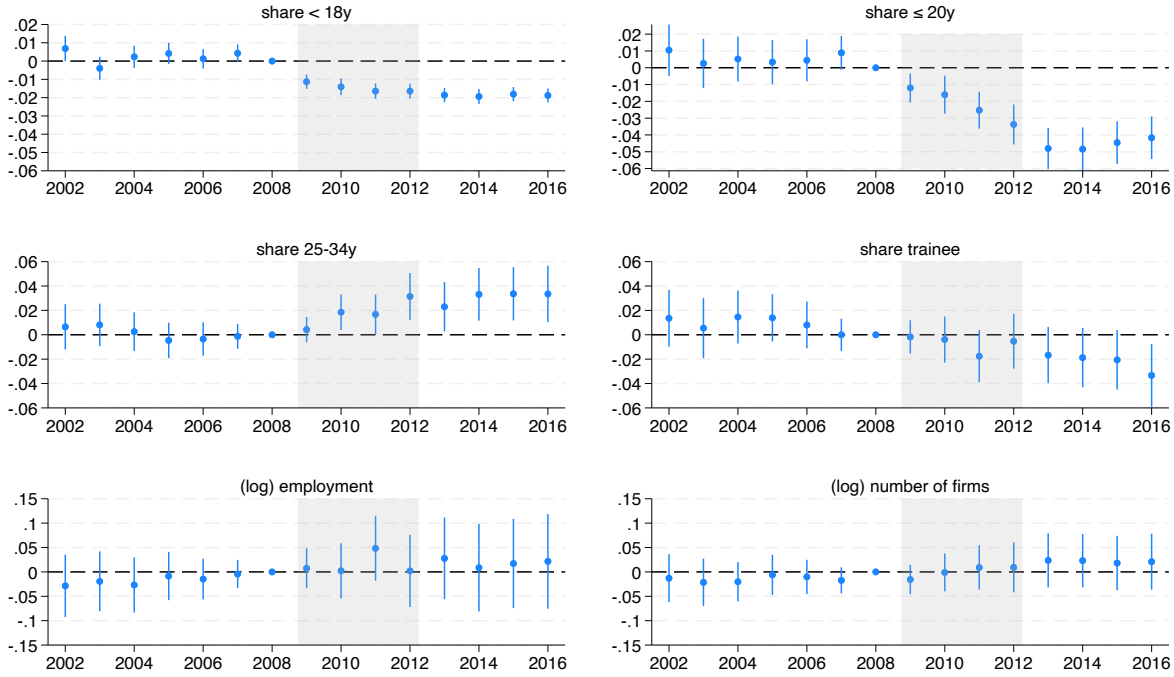

Event-study difference-in-differences estimates from equation 1 for the main outcomes using an alternative treatment definition. Treated local labour markets are defined as those appearing at least three times above the 90th percentile of minor-employment intensity in the pre-reform period, while control units appear at most twice. Panels show estimated effects for key outcomes: the share of minors (Panel A), the share of all teens aged ≤ 20 (Panel B), the share of young adults aged 25–34 (Panel C), the share of trainees (Panel D), total employment (Panel E), and the number of firms (Panel F). Shaded areas correspond to the transition period between the passing of the reform in 2009 and 2013, when its phased implementation culminated in a complete ban on full-time work by minors under 18. Vertical lines represent 95% confidence intervals around point estimates. Standard errors are clustered at the local labour market (sector–municipality) level.