

Are subsidies to apprenticeships a bad deal?

Evidence from a recent Italian reform *

Luca Citino^{†a} and Alessandra Fenizia^{‡b}

^aBank of Italy

^bGeorge Washington University

Preliminary

Abstract

This paper studies the effectiveness of a relative payroll tax reduction targeted to apprenticeship contracts. It uses rich administrative data and leverages an Italian reform that raised social security contributions for apprenticeship contracts but granted a substantial discount for firms with 9 employees or less. Although costs increase for all firms after the reform, the discount is effective at incentivizing the demand for apprentices among targeted firms. Relative to a counterfactual where costs increase equally for all firms, the reform is effective at incentivizing demand for apprentices by increasing the number of apprenticeship contracts five years after its implementation. Importantly, a non-negligible fraction of these apprenticeships is later transformed to open-ended contracts and we find no indication that the reform inadvertently incentivizes firms to misuse this contract. We also document that the reform is cost-effective from a fiscal standpoint. The cost of creating one apprenticeship is approximately €2,000, while an additional transformation to an open-ended contract costs €11,500 on average.

JEL classification: J23, J38

Keywords: apprenticeships, firms, labor costs, labor demand, incentives

*We are extremely thankful to our PhD supervisors David Card and Jörn-Steffen Pischke for their guidance and support during the early stages of this project. We also thank Fabrizio Balassone, Francesco D'Amuri, Steven Hamilton, Elira Kuka, Francesca Lotti, Matteo Paradisi, Vincenzo Scrutinio, Bryant Stuart, Roberto Torrini and Eliana Viviano for the useful discussion and comments. The realization of this project was possible thanks to the sponsorships and donations in favour of the "VisitInps Scholars" program. We are very grateful to Massimo Antichi, Elio Bellucci, Daniele Checchi, Mariella Cozzolino, Edoardo Di Porto, Paolo Naticchioni and all the staff of Direzione Centrale Studi e Ricerche for their invaluable support with the data and the institutional setting. The opinions expressed in this paper are those of the authors alone and do not necessarily reflect the views of the Bank of Italy or INPS. All errors are our own.

[†]Contact: luca.citino@bancaditalia.it

[‡]Contact: afenizia@gwu.edu

1 Introduction

Subsidies that incentivize firms to take on apprentices may constitute a double dividend for policy makers. On the one hand, they have the potential to correct some of the market failures associated with firm-sponsored training (Becker, 1962; Acemoglu and Pischke, 1998, 1999; Dustmann and Schönberg, 2012).¹ On the other hand, they can generate new and valuable employment opportunities for the youth, who are often seen as a vulnerable group in the labor market. For these reasons many governments offer reduced social security contributions or favorable taxation regimes to incentivize apprenticeship use (Kuczera, 2017).² Despite such policy interest, there is still limited empirical evidence on the effectiveness of these subsidies.

In this paper we attempt to fill this gap by exploiting an Italian reform that provides financial incentives to small firms relative to large firms for the use of apprenticeship contracts (Law n.296/2006). While the reform increased the overall social security contributions (SSC) firms were required to pay for apprentices, it provided a substantial discounts to firms with 9 employees or less during the first two years of the contract. Before 2007 firms used to contribute a weekly fee of 2.85 euros for each apprenticeship contract. After 2007, the same social contributions were raised to 10% of the apprentice's wage in firms with more than 9 employees, but only to 1.5% and 3% in firms with 9 employees or less, during the first and second year of the contract, respectively. The eligibility for the SSC discount was not conditional on firms having net employment growth over the period nor on converting a given share of apprentices to open-ended contracts. As a consequence, our results do not embed mechanical effects which may be present in other reforms.

We base our analyses on the confidential matched employer-employee dataset collected by the Italian Social Security Agency (INPS), covering the universe of firms with at least one employee during the period 1983-2018. These administrative data contain rich information on job spells including detailed contract type and social security contributions. We also combine the matched employer-employee dataset with firm-level balance sheet data from *Cerved*.

¹In a competitive labor market firms will not have any incentive to provide general training to workers unless the worker pays for such training herself. While this is not a market failure *per se*, financial markets incompleteness may prevent workers to take substantial wage cuts (or accept negative wages) in exchange for valuable training. As a consequence training provision would remain below the first best allocation. In this context public subsidization of apprenticeships can act as a grant to the worker.

²Historically, these policies were popular in developed countries only. However, they caught the attention of governments of developing countries who recently started experimenting with it (see e.g. Crépon and Premand (2019)).

Our empirical strategy exploits the differential change in SSC for firms above and below the 9-employee threshold in a difference-in-differences design. This strategy compares the evolution of the outcomes of eligible (between 5 and 9 employees at baseline) and ineligible firms (between 10 and 14 employees at baseline). The identifying assumption is that eligibility status is not predictive of potential changes in the outcomes of interest (parallel trends assumption). While we cannot directly test this assumption, we provide evidence on the absence of pre-trends in the outcomes, corroborating the validity of the design. We further address some concerns related to our design in a series of robustness checks where we show that our results are not confounded by differences in baseline firm characteristics, industry-specific time trends, firm size and exposure to the Great Recession.

We begin our empirical analysis by documenting some interesting descriptive facts associated with this reform. First, very few firms take up the policy. In 2007, at the onset of the reform, only 17.5% of firms in our sample claim the discount. On the one hand, this is consistent with a general unwillingness of firms to use this type of contract, something that has also been documented in other settings (Caicedo et al., 2020; Alfonsi et al., 2020). On the other hand, between 50% and 80% of eligible firms *that hire new apprentices* – and thus are interested in using this contract – do not take up the policy, thus leaving money on the table.³ Furthermore, we provide evidence consistent with non-compliance with policy rules by firms. Many firms with more than 9 employees take up the policy, even if they are not eligible to do so. Also, we do not find any bunching of firms at the threshold, either in 2007 or in subsequent years. This mitigates concerns that the reform generates costly firm-size distortions, as it is often the case with size-dependent policies (Garicano et al., 2016; Caicedo et al., 2020).

Our results indicate that the SSC discount targeted to small firms is effective at increasing the demand for apprenticeship contracts in eligible relative to ineligible firms. While the number of new apprenticeships hired at both firm types decreases during the study period, eligible firms hire more apprentices throughout and increase the stock of apprenticeship contracts 5 years after the reform relative to ineligible firms. Perhaps unsurprisingly, this policy does not appear to impact workers' *net* salaries, as wage rigidities imposed by the collective bargaining agreements likely impede this margin of adjustment.⁴

One may be concerned that firms may opportunistically hire apprentices to take advantage of the SSC discount and then let them go once their contract expires. Reassuringly, we find

³This is consistent with two non-mutually exclusive explanations: (i) firms may be not be aware of this policy (i.e., inattention); (ii) the bureaucratic costs associated with claiming the discount (and possibly hiring an apprentice for the first time) may be too high relative to the value of the discount in SSC.

⁴Firms can pay apprentices a lower wage, up to two levels below what a qualified worker would get, according to the corresponding collective bargaining agreement (Albanese et al., 2017).

that a non-negligible fraction of apprenticeships is transformed to open-ended contracts at the same firm.⁵ A related concern that is often brought up in relation to apprenticeships is that firms may view these contracts as source of “cheap labor” and hire apprentices without providing the appropriate level of training (Tiraboschi, 2014). As we cannot measure whether apprentices are trained appropriately, we can not investigate the extent to which apprenticeship contracts are used as a training contract. However, we can evaluate several margins that are indicative of whether firms engage in strategic behavior and take up the subsidy with the only purpose of cutting down costs (i.e. subsidy misuse): strategic separations from expensive apprentices, substitutability with other contract types, shorter contract duration, few transformations to open-ended contracts, and changes in the quality of new hires. Although we cannot rule out that apprentices may not be adequately trained, we find no indication that firms take up the subsidy to merely reduce costs.

Finally, we estimate the number of jobs created thanks to the SSC discount and the direct fiscal cost per job created relative to a scenario where payroll taxes increased for all firms equally. We do so by instrumenting the endogenous take up (or the firm-level SSC discount) with the size threshold interacted with year dummies in a 2SLS framework. We find that – over a five year period – the average complier firm hires 2.86 more apprentices and transforms 0.5 more (of them) into open-ended contracts than it would have done in the absence of the SSC discount. This implies that roughly only 1 in 6 apprentices is eventually offered an open-ended contract at the same firm.⁶ When comparing these figures with the actual amount of money spent on this policy, we find the discount to be quite effective at the margin. A €1,000 expenditure generates an average of 0.482 extra apprenticeship contracts. This implies that €2,075 of public transfers are on average sufficient to create an apprenticeship contract. Similarly, an additional €1000 generates 0.087 conversions to open-ended contracts, which translates into €11,494 per conversion. Importantly, these figures do not represent the *economic* cost of creating such jobs, but only the direct government *transfer* that is necessary to incentivize firms.⁷ They also exclude indirect savings on other social programs that workers

⁵Of course an increase in transformations after the reform is not sufficient to conclude that apprenticeships are effective in providing workers with valuable training. Alternatively, firms may be using apprenticeship contracts as a screening device to select the most productive workers. Unfortunately, we do not directly observe training provision in our data and cannot disentangle these two competing explanations.

⁶In this setting complier firms are those that hire at least an apprentice because of the SSC discount but would have not done so in the absence of the discount.

⁷The economic cost is the private resource cost associated with job creation e.g. other workers’ time, physical capital used by workers, etc. plus the welfare loss from raising distortionary taxes needed to finance the transfer to firms. The latter parameter is known in the public finance literature as the *marginal cost of public funds*. While an exact assessment for the case at hand would require knowledge of the type of taxes needed to finance this reform, some estimates for Italy can be found in Kleven and Kreiner (2006) and are between 0.3 and 0.8. These numbers imply that raising distortionary taxes to collect €1 of revenue leads to a welfare cost between 30 and

may have received if unemployed and taxes collected on labor while employed. If anything, these two latter channels decrease the actual fiscal cost per job created, making the reform even more cost-effective.

While we cannot provide a full assessment as to why the cost per transformation is low at this stage, we provide a suggestive explanation and leave a more comprehensive treatment to future research.⁸ The SSC discount does not *directly* subsidize permanent employment. It just incentivizes firms to do more training and screening of potentially productive matches. It follows that, at the point of transformation, offering an open-ended contract must be privately optimal for the firm even absent the SSC discount. In other words, the cost per transformation is purely driven by the transformation rate of apprentices and not by the cost associated to permanent employment. The tilted schedule of the subsidy further contributes to the cost-effectiveness. This is because the transfer is higher when labor demand is more elastic (due to uncertainty around workers' types or training success) and lower when labor demand is less elastic (due to match-specific investments).

This paper relates to three strands of literature. First, it fits into the nascent literature that investigate *firm-level* responses to apprenticeship regulation ([Cappellari et al., 2012](#); [Caicedo et al., 2020](#); [Alfonsi et al., 2020](#); [Crépon and Premand, 2019](#)). While a large body of work has studied the returns to apprenticeships for *workers* ([Fersterer et al., 2008](#); [Lodovici et al., 2013](#); [Albanese et al., 2017](#); [Picchio and Staffolani, 2019](#); [Cavaglia et al., 2020](#); [Alfonsi et al., 2020](#)), much less is known about how firms respond to these type of policies. Understanding firm responses is key to design policies that are both cost-effective and that minimize unintended consequences (e.g. misuse of apprenticeship contract and firm-size distortions). [Caicedo et al. \(2020\)](#) analyze the introduction of minimum and maximum apprenticeship quotas (dispensable upon payment of a fee) and decreases apprentices' minimum wages in Colombia. They find that the reform considerably increased the number of apprentices but at the cost of sizable firm-size distortions. [Alfonsi et al. \(2020\)](#) and [Crépon and Premand \(2019\)](#) design a field experiment in Uganda and Côte d'Ivoire respectively, to study both worker and firm responses to subsidies to vocational training and apprenticeships. On the side of firms [Alfonsi et al. \(2020\)](#) find that few apprentices kept working for the same employer once they completed their subsidized training and that their returns faded quickly. The authors argue that this may be due not to the lack of training, but rather the lack of *certified*

80 cents.

⁸A full quantitative assessment of the employment effects of the reform and its cost would need a general equilibrium model taking into account – among other things – how worker flows in our sample crowd out worker flows outside our sample (for example at very big firms). We therefore cannot rule out that the open-ended contracts being formed in our sample would not have materialized outside our sample.

training, stressing the importance of asymmetric information. They also find no evidence of firms substituting other types of workers with apprentices. [Crépon and Premand \(2019\)](#) provide wage subsidies to workers enrolling onto dual apprenticeships, rather than traditional ones. They document that firms substitute towards dual apprenticeships and that the crowding out of traditional ones is small in magnitude. We contribute to this literature in three ways. First, we show that modest targeted firm subsidies successfully induce firms to employ more apprentices without generating firm-size distortions. Second, to the best of our knowledge, we are the first to analyze the impact of such policies on firms' choices using rich administrative data in a developed country setting.⁹ Third, we investigate whether hiring credits incentivize firms to use apprenticeship contract as a mere source of cheap labor. We document no evidence of this phenomenon.

Second, this paper fits in the literature that examines the impact of payroll deductions on employment and wages ([Bohm and Lind, 1993](#); [Bennmarker et al., 2009](#); [Egebark and Kaunitz, 2013](#); [Korkeamaki and Uusitalo, 2006](#); [Huttunen et al., 2013](#); [Saez et al., 2019](#)). These studies examine the effects of permanent payroll reductions in Nordic European countries and document mixed results on employment and wages. This paper provides novel empirical evidence on the effectiveness of a temporary reduction in payroll taxes in a Southern European country characterized by wage rigidities.

Finally, we relate to the literature on the effectiveness of hiring credits ([Sestito and Viviano, 2018](#); [Cahuc et al., 2019](#); [Neumark, 2013](#)). Although hiring credits have often been used to sustain employment during periods of economic crisis, there is very little empirical evidence on the effectiveness of these measures during a recession ([Cahuc et al., 2019](#)). Moreover, there are some concerns that their effectiveness is strongly diminished when they target disadvantaged groups ([Neumark, 2013](#)). We shed light on these topics by showing that the SSC effectiveness is not hampered by the crisis and that the payroll reduction we study is effective even though it targets a disadvantaged group. Unlike previous studies, we evaluate both the short- and the medium-term impacts of the reform and show that the effectiveness of the SSC discount does not fade over time. This is especially important as temporary subsidies, like the one studied by [Cahuc et al. \(2019\)](#), may induce a shift in the timing of hiring without affecting the long-term employment level. We also contribute to this literature by documenting that the light touch intervention we study appears to generate long-term worker-firm matches using a limited amount of government resources. Overall these results suggest that wage rigidities play a key role in our setting and that the effectiveness of hiring

⁹Other studies in the European context have looked at the utilization of apprenticeships by firms but do not rely on quasi-experimental designs to tease out causal impacts ([Wolter et al., 2006](#); [Mohrenweiser and Zwick, 2009](#)).

credits and payroll subsidies may hinge on the degree of wage adjustments.

The paper is structured as follows. Section 2 describes the apprenticeship system in Italy and the 2007 reform. Section 3 discusses our data sources and documents some interesting stylized facts. Section 4 develops the empirical strategy and discusses the identifying assumption. Sections 5 and 6 report the main results. Section 7 discusses the robustness checks. Section 8 presents an additional analysis that helps us quantify the cost effectiveness of the policy. Section 9 concludes.

2 Institutional setting

2.a Apprenticeships in Italy

The Italian apprenticeship system is composed of three distinct types of contracts, and each one of them has different rules. In practice, as of 2017, around 95% of apprenticeship contracts fall under the “Occupational” apprenticeship category (*apprendistato professionalizzante*) (D’Arcangelo et al., 2019). The latter is a job contract with a maximum duration of 6 years, at the end of which the worker is awarded a professional qualification in a given occupation. Such contract is limited to the private sector and to individuals aged 18-29. The content of training courses is regulated by regions and collective bargaining agreements (CBA), which keeps it separate from the schooling system. The law imposes a minimum of 120 training hours per year. 80 hours must be dedicated to occupation-specific training while the remaining 40 relate to general training (job safety regulation, psychology of labor and team working). The other two – quite marginal – forms of apprenticeship are “right-duty” apprenticeship (*apprendistato per l’espletamento del diritto/dovere di istruzione*) and “higher” apprenticeship (*apprendistato di alta formazione e ricerca*). While the former is partly tied to the education system and targets individuals aged 15-18, the latter is meant for individuals aged 18-29 who are enrolled in or have earned a college degree and want to carry out a specific research project at the firm premises.

For all of the three apprenticeship types, firms are granted substantial reductions in social security contributions (SSC) in return for the training they provide. A favourable SSC regime also applies during the first year of transformation to open-ended contract (Law 56/1987). We provide more details about SSC regimes in Section 2.b, as they are the key features of the reform we are studying. In addition to SSC discounts, employers are allowed to pay apprentices lower salaries compared to what a trained worker would get at the end of the same apprenticeship, according to the relevant CBA.

The law imposes limitations on the use of apprenticeship contracts. A firm cannot employ more apprentices than regular workers. The provision does not apply to firms that have at most three employees; these firms can employ up to three apprentices. As for the transformation of apprenticeship contracts into open-ended ones, during the period under study no provision in the law forced firms to transform a minimum number of apprentices. Nonetheless, many CBAs set transformation quotas, whereby the firm needs to transform a certain fraction of the apprenticeship contracts used over a certain period of time (e.g. the last two years). Quotas vary by CBA and sometimes are as high as 70%. (D’Agostino et al., 2010).

2.b The 2007 Budget Bill

The reform of social security contributions (SSC) for apprenticeship contracts that we study in this paper was introduced with the 2007 Budget Bill (Law n.296/2006), passed on December 27, 2006 and effective since January 1, 2007. Among various other measures concerning labor markets and social protection the 2007 Budget Bill increased the SSC charged on firms for their apprenticeship contracts. This was needed to finance the simultaneous introduction for apprentices’ paid sick leave, up to 180 days a year.¹⁰ Importantly for us the increase in SSC was differentiated depending on whether the firm had more or less than 9 employees. Figure 1 illustrates how SSC changed in response to the 2007 Budget Bill. Before 2007 employers used to pay a fixed weekly fee of 2.85 euros for each apprenticeship contract (green triangles in Figure 1).¹¹ Starting on January 1, 2007 firms with more than 9 employees (labeled “large” firms in the graph) were required to pay 10% of the apprentice’s wage in social contributions (hollow circles in Figure 1). Firms with at most 9 employees (labeled “small” firms in the graph) were required to pay 1.5% of the apprentice’s wage in the first year of the contract, 3% in the second year, and 10% in all the following years (orange circles in Figure 1). The increase in SSC applied to both existing apprenticeship contracts and those signed after January 1, 2007. The eligibility for the SSC discount was not conditional on firms having net employment growth over the period nor on converting a given share of apprentices to open-ended contracts. As a consequence, our results do not embed mechanical effects which may be present in other reforms. We stress that no other pre-existing or concurrent policy was discontinuous at nine employees.¹²

¹⁰With regard to these issues, the main measures concerned the contrast of informality in the labor market, modifications to unemployment insurance and short-time work and reductions in the tax wedge on labor. The full text of the law can be accessed at <http://www.parlamento.it/parlam/leggi/062961.htm>

¹¹The weekly fixed contribution was equal to 2.95 euros for apprentices eligible for occupational injury insurance. A 2.85 (2.95) weekly fee translates into $2.85 \times 52 = 148.2$ ($2.95 \times 52 = 153.4$) euros per year.

¹²Reassuringly, Figures 4a and 4b report the distribution of firm size in the two years leading up to the reform and show no discontinuity in firm size at 9 employees. As the government changed the apprenticeship

To quantify the changes in SSC for an average apprenticeship contract (€1,000 per month), Table 1 discusses a numerical example. Columns 1 and 2 report the SSC for firms below and above the 9 employee threshold under the old regime; column 3 and 4 report the analogous SSC under the new regime; finally, columns 5 and 6 report the change in SSC in the new regime relative to the old one. Firms with more than 9 employees experience an increase in SSC of €1,052 per apprentice per year (column 5), while firms with 9 employees or less experience an increase of €32, €212, and €1052 per apprentice in the first, the second, and the third year of the contract, respectively (column 6). Overall, firms experience an increase in the cost of the average apprenticeship contract but this effect is greater for firms above the size threshold.

The firm size that determines the eligibility for the SSC discount is full-time equivalent employment and it excludes apprentices, temporary agency workers, workers who are on leave (unless the firm hires a substitute), and workers who have been hired with an on-the-job training contract.¹³ More details on how we construct firm size are provided in Appendix A. For contracts signed after January 1, 2007 the firm size at the time of hiring determines the eligibility for the discount. For pre-existing contracts, eligibility is determined based on the 2006 average firm size. The 2007 Budget Bill also introduces minor changes to the SSC for other contract types but none of these changes exhibits a discontinuity at 9 employees.

Next, we document some stylized facts that illustrate some interesting features of this reform. First, only a modest share of firms takes up the policy. In 2007, at the onset of the reform, only 17.5% of firms in our sample both use apprenticeships and claim the discount. This is consistent with a general unwillingness of firms to use this type of contract, which has also been documented in other settings (Caicedo et al., 2020; Alfonsi et al., 2020)

Second, we document evidence consistent with non-compliance with policy rules. Figure 2 illustrates the relationship between the share of firms that claim the subsidy in January 2007 and small bins of the 2006 policy-relevant average firm size. We calculate that on average 14.5% of firms below the size threshold claim the SSC discount in January 2007 while only 3.4% firms do so above the threshold. Importantly, the probability of claiming the subsidy tapers off around the 9-employee threshold. The absence of a sharp discontinuity in the take-up rate prevents us from using a regression discontinuity design, we discuss this point further in Section 4. One might expect the share of firms who claim the subsidy to drop to zero as we move past the threshold. This is not the case due to two reasons. First, the policy-

contribution schedule in 2012, we limit our sample period to 2004-2011 to avoid picking up the effect of the 2012 reform.

¹³The definition of firm size is contained in the provision issued by INPS, which is in charge of implementing this policy (*circolare n. 22, 2007*). In our empirical analysis, we follow this definition closely.

relevant firm size to claim the subsidy is the size at the time of hire, while the horizontal axis reports the average 2006 firm size (i.e., measurement error). Second, a non-negligible number of ineligible firms claim the discount, which is suggestive of non-compliance with policy rules. As we discuss in Section 4, this type of non-compliance biases our reduced-form results downward and our point estimates can be interpreted as a lower bound on the true effect.¹⁴

Third, we discuss three factors that are likely to contribute to these modest take-up rates: (i) firms may not be aware of this policy; (ii) the bureaucratic costs associated claiming the discount (and possibly hiring an apprentice for the first time) may be too high relative to the value of the discount in SSC; (iii) the apprenticeship contract is an underutilized contractual arrangement in the Italian labor market. While we can not separately identify the first two mechanisms, we investigate the extent to which they are jointly relevant in our setting. To this end, we construct the firm-level “attention rate” as in Cahuc et al. (2019). This is defined as the probability of claiming the SSC discount conditional on hiring at least an apprentice. Figure 3 depicts the relationship between the attention rate and small bins of the policy-relevant firm size in 2007. The attention rate displays a similar pattern as the take-up rate: eligible firms display higher attention rates than ineligible ones and the attention rate tapers off around the 9-employee threshold. In particular, roughly 80% of firms with five employees claim the subsidy when they hire new apprentices, while only 20% of firms with 14 employees claim the subsidy conditional on signing new apprenticeship contracts. This Figure implies that between 20% and 50% of eligible firms leave money on the table. One may expect the attention rate of the firms with more than 9 employees to be close to zero as these firms do not qualify for the discount. This is not the case and the positive attention rates for these firms are consistent with the evidence documented in Figure 2.

Finally, we examine the distribution of firm-size in the years after the reform and look for evidence of bunching below the threshold. Figures 4c through 4f report the distribution of firm size by year and show no indication of firms manipulating their size to become eligible for the discount, mitigating concerns that the reform generated firm-size distortions. This finding is consistent with two non-mutually exclusive explanations. First, the discount may not be sufficiently large to induce firms to take costly actions to manipulate their size (e.g. let some workers go). Second, firms may not need to resort to these costly actions if eligibility requirements are not strictly enforced (i.e., non-compliance).

¹⁴As detailed in Appendix A, we do not observe flags for one particular type of on-the-job training contract which ought to be excluded from the policy-relevant firm size calculation. This makes us overstate the policy-relevant firm size for firms that use such contracts, and classify them as non-eligible when they are in fact eligible. As long as these contracts are not too common, such source of measurement error should be small.

3 Data and measurement

3.a Data

In this paper we use confidential administrative data from Italy provided by the Italian Social Security Institute (INPS). We use a matched employer-employee dataset based on monthly compulsory communications that firms are required to send to INPS for administrative purposes (*UNIEMENS* module). This covers the universe of all private non-agricultural firms with at least one employee and spans the period 1983-2018. We also match the INPS data with balance sheet data from Cerved. In what follows we describe our data sources and discuss the sample selection.

Matched employer-employee data: the raw data is at the job-spell level i.e. each record identifies a unique combination of contractual characteristics observed for a given firm-worker match in a given year. If a worker changes a contractual characteristic during the year we observe two records, one for each configuration. As for characteristics, we observe the contract type (open-ended or temporary), work-time arrangement (part-time or full-time) and a coarse job ladder code (apprentice, blue collar, white collar, supervisor or manager). For each job spell we also observe the duration in days and accrued earnings. If the worker does not experience any change in these characteristics in a given year, we only observe one earning record. If the worker changes any of these (including where she works) we are able to observe two separate earning records.

For each worker-firm match we also observe the start and end dates of the contract, which is crucial for us to construct hires and separations. Thanks to the matched employer-employee nature of the dataset, for each worker hired by the firm we are able to reconstruct his whole working history during the years 1983-2018. This allows us not only to look at the workers' previous experience, but also at the realized contract duration and future contracts. We are also able to observe a variety of indicators for particular types of workers which enter (or are excluded from) the formula for the policy-relevant firm size. Equally importantly for us, the INPS data contains a flag for whether the firm claimed the SSC discount and the bracket applied (1.5%, 3% or 10%). We provide a detailed definition of the variables used in the analysis in Appendix B.

Since 2005 (two years before the reform is enacted) we can also rely on a more disaggregated dataset where earnings are reported at the much finer monthly frequency. This is important for us as it allows us to precisely distinguish salaries that apprentices accumulate during their first year of tenure against subsequent years, as well as computing the exact monetary

amounts that the firm receives because of SSC discounts.

Cerved data: the balance sheet information comes from *Cerved*, administered by the *Cerved Group*. This contains information on balance sheets and income statements of all Italian incorporated companies, regardless of whether they have employees or not. We use this data to construct financial indicators such as the labor share, liquid assets over assets, investment over assets and cash flow over assets, which we use in our robustness tests.

Sample selection: our period of analysis starts in 2004, three years before the reform, and ends in 2011, when the 2012 Budget Bill changed the SSC schedule once again. The pre-reform data allows us to conduct tests on the presence of differential pre-trends between firms above or below the size threshold. We restrict our main sample to all firms whose policy-relevant firm size is between 5 and 14 in 2006, and who have been active during all years from 2004 to 2006. Our selection procedure yields a sample of 193,297 firms. The matched INPS-Cerved sample, used in some of the robustness tests, is considerably smaller and only made of 98,084 firms.

3.b Descriptive Statistics

Table 2 displays firm characteristics at baseline. Panel A reports the descriptive statistics on worker composition and firm age, Panel B and C illustrate the industry composition and the geographic location of firms, respectively. Column 1 reports the characteristics for the full sample; columns 2 and 3 display the statistics for eligible and ineligible firms respectively. Firms are classified as eligible vs ineligible based on their average policy-relevant firm size in 2006. Eligible firms include all incumbent firms between 5 and 9 employees, while ineligible firms include all incumbent firms with at least 9 but less than 14 employees. Panel A shows that, although worker composition does not appear to differ dramatically between eligible and ineligible firms, eligible firms tend to employ a slightly younger workforce (27% of their employees have not yet turned 30 vs 25% of workers in large firms) and a larger share of female workers (36% vs. 34%). Both eligible and ineligible firms are well established and they have been around for more than 14 years on average. As one might have expected, ineligible firms are on average older (and this difference is about one year). Panel B investigates whether industry composition differs between these two groups. The industry composition appears to be fairly well balanced. The only exception is the manufacturing sector: the average ineligible firm is 8 percentage points more likely to be in manufacturing compared to the average eligible firm in our sample. Finally, Panel C explores the geographic location of firms in our setting and shows that eligible firms are not disproportionately represented in any of the twenty Italian regions.

Next, we discuss the summary statistics of our main outcomes of interest. Table 3 has the same structure as Table 2 and reports the descriptive statistics computed over the full sample period. The average firm in our sample has 9.03 full-time equivalent employees (FTE) and eligible firms employ by construction fewer employees than ineligible ones (7.41 vs 12.16 workers). The fact that full-time equivalent employment is only slightly smaller than the average firm size suggests that most of the workers are employed full-time. Given the relatively small size of firms in our sample one may be concerned about attrition. Having selected our sample such that firms are active in all years between 2004 and 2006 implies that we selected a sample of relatively stable firms. As a result, the survival rate is extremely high (92%) and there is no evidence of differential attrition of eligible vs ineligible firms (survival rate of 92% vs 93%). The average firm employs only 0.41 apprentices in any given period. While this figure may appear fairly low, this is consistent with apprenticeship being a somewhat underutilized contractual arrangement in the Italian labor market. Anecdotally, while some firms use this contract type over and over again, most firms do not use it at all. The average apprenticeship contract lasts for about 20 months. Firms replace a substantial share of apprentices in every period as they hire 0.31 and separate from 0.34 apprentices every year. Despite the high turnover, the average yearly number of transformations into open-ended contracts is approximately a fourth of the apprenticeship stock. Not surprisingly, ineligible firms employ more apprentices than eligible firms on average, but apprenticeship contracts are more popular among eligible firms as the higher apprentice-to-size ratios attest.

3.c Wage effects and pass-through

In a perfectly competitive and spot labor market, worker mobility ensures that a targeted reduction in social security contributions has no differential impact on the net wages paid by firms above or below the size threshold. Indeed, as long as workers view such firms as perfect substitutes, they will apply for jobs at firms paying the highest net wages. This puts downward pressure on wages until they are again equalized. To the contrary, labor costs will differ between eligible and ineligible firms, and this determines an increase in the quantity of labor demanded by eligible firms. The perfectly competitive model therefore predicts zero pass-through of a targeted discount of SSC.

Our setting departs from the perfectly competitive model in a few ways. For example, apprentices may not be able to switch firm midway through an apprenticeship to take advantage of higher wages, perhaps because of some firm-specific human capital that is sunk. Alternatively, search frictions or idiosyncratic preferences for some firms may also hinder reallocation of workers in response to wage changes. In such a setting it is likely that

part of the discount that treated firm received is transferred to workers.

In Figure 5 we plot the evolution over time of average real *net* monthly earnings for apprenticeships in firms whose 2006 average size was above or below the 9 employee threshold (“ineligible” and “eligible” firms, respectively). Earnings display a marked increase in real terms, which may be driven by firm growth over time. However, there is virtually no difference in the earnings of apprentices employed at eligible and ineligible firms before the reform and no difference emerges after 2007, consistent with the competitive labor market model. While contracts for incumbents were signed before 2007 and may be hard to renegotiate, those for new hires could in principle reflect the changes in SSC introduced by the 2007 budget bill. To investigate whether this pattern is driven by incumbent apprentices only, we replicate the same analysis using only the monthly earnings of new hires, and find very similar results (which are reported in Appendix D).

4 Empirical strategy

Our identification strategy leverages the discontinuous change in the costs of apprenticeship contracts introduced by the 2007 Budget Bill. As documented in Section 3.b the take-up does not exhibit a discontinuity at the threshold, thus we cannot rely on a regression discontinuity design. We build on the work of Cahuc et al. (2019) – who faces a similar challenge as we do – and we opt for a difference-in-differences (DiD) design. Our estimating equation is

$$y_{it} = \alpha_i + \tau_t + \sum_{k \neq 2006} \beta_k \cdot \mathbf{1}(\text{year}_t = k) \cdot T_i + \epsilon_{it}, \quad (1)$$

where y_{it} represents the outcome of interest (e.g., the number of apprentice hires) at firm i in year i . We control for firms (α_i) and calendar year fixed effects (τ_t). The firm fixed effects account for time-invariant differences across firms, while time fixed effects control shocks that are common across firms. T_i is our eligibility indicator and takes value one for firms whose average size in 2006 was at most 9 employees. In our baseline estimates standard errors are clustered at the firm level to account for serial correlation in the outcome of interest (Bertrand et al., 2004). Results are virtually unchanged when we cluster at the 2-digit sector level or at the local labor market level (see Appendix D). We are interested in coefficients β_k which capture the average difference between eligible and non-eligible firms in year k relative to the same difference in year 2006 (our reference year).

The identifying assumption of our DiD design is that eligibility status is not predictive of

potential *changes* in the outcomes of interest (parallel trends assumption). While we cannot directly test this assumption, we can probe the validity of our design by testing whether the β_k coefficients associated to the years prior to the reform are significantly different from zero.

Given that our eligibility measure does not capture the actual receipt of the SSC discount, our point estimates only measure a reduced-form (also called intention-to-treat) effect. As shown in Figure 2, some ineligible firms claim the SSC discount in our context. This type of non-compliance biases our results toward zero and our point estimates can be interpreted as a lower bound on the true effects. In Section 6 we provide a thorough description of an additional analysis where we adopt a 2SLS strategy to rescale our reduced-form effects by the extra share of firms actually claiming the discount among the eligible ones.

Next, we discuss some threats to our empirical strategy and foreshadow some robustness tests we present in Section 7. First, aside from effects related directly to firm size, it may be the case that firms above and below the 9 employee threshold differ along characteristics that predict the potential evolution of apprenticeship use. To control for this potential source of omitted variable bias, we augment specification 1 and include baseline firm-level covariates interacted with year fixed effects. We also propose an additional robustness check where we augment our baseline specification with region \times sector linear time trends.¹⁵

One may be concerned that eligible and ineligible firms may react differently to the same subsidy only because they have different sizes, and that our estimates pick up a combination of the reform impact and this size effect. The relevance of this confounding factor fades as the difference in firm size between the two groups shrinks. Therefore, we mitigate this concern by repeating our analysis on the sub-sample of firms who employ between 8 and 11 employees.

In the same vein, we also check that our regression results are not confounded by firms being differentially exposed to the Great Recession. In order to address this concern, we evaluate whether firms located in local labor markets that are differentially affected by the crisis exhibit a differential response to the reform. We split our sample according to whether the 2007-2010 change in the local unemployment rate was above or below the median.¹⁶

¹⁵The inclusion of *linear* time trends in this latter specification is done just for computational convenience, due to the high number of fixed effects.

¹⁶Local unemployment rates are measured at the local labor market level and are retrieved from the National Statistical Institute (Istat).

5 Reduced-Form Evidence

This section presents the reduced-form results of our analysis, which capture both within-firm changes in employment conditional on policy take up and take up itself. We first look at worker flows i.e., hires, separations, and transformations of apprenticeship contracts. We then move on to examine the impact of the discount on the overall number of apprentices working at the firm (the “stock” of apprenticeships) and contract duration. Finally, we discuss whether the discount may incentivize firms to hire apprentices as a form of cheap labor without training them appropriately. While we do not observe actual training, we can evaluate several margins that are indicative of whether firms engage in strategic behavior and take up the subsidy with the only purpose of cutting down costs (i.e. subsidy misuse): re-labeling of existing contract as apprenticeship contracts, substitutability between apprenticeships and temporary contracts, strategic separations, transformations, and apprentices’ selection.

5.a Impacts on Apprentices Flows

As for the analysis on apprentice flows we look at three outcomes: hires of apprentices, separations from apprentices and transformations of apprenticeships into open-ended contracts.

Figure 6 reports the point estimates and the associated 95% confidence intervals for the β_k coefficients from equation 1, for the three outcomes of interest. In all cases the coefficients leading up to the year of the reform are not significantly different from zero, indicating that the policy-relevant size threshold is not predictive of differences in trends in the pre-reform years. This provides evidence in favor of the parallel trend assumption that underlies our DiD design and that we discussed in Section 4.

In 2007 the impact of the discount on the number of new hires is positive but small and not statistically significant (blue circles). Starting from 2008, it increases gradually, peaks in 2009 and then stabilizes around 0.06 in 2011. Compared to the average yearly number of apprenticeship hires in our sample (0.31) this corresponds to a 20% increase. This point estimates would indicate that the discount induces on average one in approximately 17 firms ($1 \div 0.06$) to hire an apprentice in each year.

When looking at separations (red diamonds), the discount does not appear to have an impact in the first two years suggesting that firms are not terminating incumbent apprentices’ contracts to substitute them with new apprentices at a lower cost. The impact on separations starts increasing in 2009 and remains positive and statistically significant for the rest of the sample period. This result is consistent with the fact that the increase in hiring results into

a *lagged* increase in separations. Interestingly, the number of transformations increases on impact (albeit it is only marginally statistically significant in 2007) and remains positive and statistically significant throughout the period (black triangles). This finding is likely to be driven by the floors on conversion rates of apprentices to open-ended contracts and the caps on the total number of apprentices working at the firm mandated by the CBAs and the law, respectively. As described in Section 2.b, during the period under study firms were restricted by law to have fewer apprentices than qualified workers. Also, several collective bargaining agreements required firms to convert a certain share of incumbent apprentices to open-ended positions before being able to hire new apprentices. In other words, firms that are close to the cap or that have not recently converted any of their apprentices to an open-ended position may have an incentive to do so if this allows them to hire new apprentices in the future. While the increase in the number of transformations may be reflective of the discount inducing firms to shift to shorter contracts, we show in Section 5.d that this is not the case. We probe the robustness of our results with several checks, which we present in Section 7.

5.b Impacts on the Stock of Apprenticeships

Next, we examine how the SSC discount impacts the stock of apprenticeships. This measure nicely aggregates the results on workers flows and proxies for the popularity of this contract type.

We measure the stock of apprentices *pro rata temporis* and in full-time equivalent units. Specifically, the stock of apprentices is not measured as the number of workers ever employed in a given year but it is adjusted both for the number of months an individual works at the firm during a given year and for part-time work.¹⁷

Figure 7 reports the point estimates and the associated 95% confidence interval for the β_k from equation 1. The coefficients in the years leading up to the reform are not statistically different from zero and the pattern looks rather flat, corroborating our identifying assumption. The stock of apprentices decreases in 2007 and 2008, and the effect is statistically significant. From 2009 onward we have a trend inversion, which continues until 2011. Both coefficients in 2010 and 2011 are positive and statistically significant. The moderate decline in the stock of apprentices in the first two years after the introduction of the SSC discount is consistent

¹⁷For example, if an individual is hired to work full time from March to December, she would count as 0.75 units towards the firm full-time equivalent employment in the year. Likewise, if she worked part-time from January to December, she would count as 0.5 full-time equivalent units. One implication is that, while the month when a worker is hired or let go does not impact the analysis of the worker flows discussed in Section 6, it will be reflected the analysis presented in this section.

with two non-mutually exclusive explanations. First, this results lines up with the evidence presented in Figure 6: the number of new hires in the first two years is positive but not sufficiently large to counteract the increase in transformations and separations. Second, if the conversions of apprentices to open-ended contracts occur earlier in the year than hiring, this can cause a temporary decrease in the stock of apprentices. By the end of our sample period, the SSC discount induces one in 23 firms ($1 \div 0.044$) to have one additional full-time employment unit (namely a full-time apprentice for 12 months). We postpone the discussion of the effects size to Section 8 where we re-scale our reduced form effects by the take-up of the policy.

5.c Impacts on Contract Duration

After having examined the discount impact on apprentice flows and the stock of apprentices, we investigate whether firms strategically shift towards shorter apprenticeships to reduce labor costs. As explained in Section 2.b, the 2007 reform introduces a tilted schedule for apprentice social contributions for firms below 9 employees. In the absence of valuable training, if apprentices are perfectly substitutable and there are no search frictions, small firms have an incentive to churn through many short apprenticeship contracts to take advantage of the lower contribution in the first two years. To evaluate whether firms are strategically shortening contract duration, we examine the number of new apprenticeships hires by apprenticeship contract duration. More specifically, we construct a set of mutually-exclusive outcomes as the number of apprenticeship hires in each given year that last v months (where $v \in \{[0, 12], [13, 24], [25, 36], [37, 48], [49, \max]\}$). Figure 8 reports this analysis and shows that a large share of new contracts signed after 2007 last at most 12 months. This Figure displays a very clear gradient where for each year the number of newly signed contracts declines as contract duration increases. To study whether the discount incentivizes shorter contracts, Table 4 compares the distribution of the duration of newly signed contracts generated by the discount (column 1) with the baseline distribution of contract duration (column 2).¹⁸ These two distributions look fairly similar suggesting that the discount did not shorten contract duration.¹⁹

To corroborate our results we also estimate equation 1 using average apprentice contract

¹⁸We construct the share of newly signed contracts of duration v generated by the discount as the number of contracts of duration v generated by the discount over the total number of new contracts generated by the discount.

¹⁹Since apprenticeship duration is only observed *conditional* on hiring, we cannot reject the alternative possibility that our effects are driven by firms with relatively longer apprenticeship durations shortening their contracts so that it *matches* the baseline distribution of contract duration.

duration as the dependent variable. We assign to each firm i in year t the average (forward looking) duration of all the apprenticeship contracts started by the firm in year t .²⁰ While this regression fits naturally in the framework we developed so far, the dependent variable is defined only for firms that hire apprentices in a given year. This analysis thus exploits the variation in contract durations for firms that hire apprentices both before and after the reform. Figure 9 reports the results and shows that the discount does not impact average apprenticeship contract duration. We conclude that we find no evidence that firms strategically shift toward shorter contracts as a response to the discount.

5.d Misuse of the Subsidy

One concern that is often brought up in relation to apprenticeships is that they may not constitute a valuable investment for firms or workers. To the contrary, firms may just view them as source of “cheap labor” (Tiraboschi, 2014).

As we can not measure whether apprentices are trained appropriately, we can not investigate the extent to which they are used as a training contract. However, we can evaluate several margins that are indicative of whether firms engage in strategic behaviour and take up the subsidy with the only purpose of cutting down costs (i.e. subsidy misuse): contract “re-labeling”, substitutability between apprenticeships and temporary contracts, strategic separations from more expensive apprentices, shorter contract duration, transformations to open-ended contracts, and changes in the quality of new hires.

Substitutability with other contract types: We also study whether the discount induces firms to increase the number of apprenticeship contracts at the expense of other contract types. To this end, we investigate whether the number of young workers’ hires changes in response to the discount. If firms were hiring apprentices at the expense of other contract types, we would likely not see an increase in the *total* number of young workers hired. We define a young hire as any hiring of workers who are at most 29 years old. Figure 10 reports the estimated coefficients and the 95% confidence intervals. Although we observe a significant increase in the number of newly hired young workers, we do not view this finding as conclusive. Indeed, the pattern of the estimated coefficients in the periods leading up to the reform shows that the parallel pre-trend assumption is unlikely to hold for this outcome.²¹

²⁰An explicit formula can be found in Appendix B, together with the definition of the other variables used in the analysis.

²¹In future versions of this paper we plan to examine the substitutability of apprentices contracts and temporary contracts more closely.

Strategic separations: If apprentices did not receive any training and were perfectly substitutable with new untrained workers and in the absence of search costs, we would expect eligible firms to let their incumbent apprentices go and substitute them with newly hired apprentices at lower cost. In Figure 6 – which we have thoroughly described in Section 5.a – we see that separations do not react on impact the discount. Rather, they pick up only two years after, likely as a consequence of the increase in hiring. Indeed, as seen in Section 3.b, the average apprenticeship lasts 20 months and there is no evidence that firms are shortening their apprenticeship contracts (see Section 5.c). It thus seems that new hires do not come at the expense of incumbent workers leaving the firm. This is evidence that the discount did not induce firms to substitute incumbent apprentices with cheaper newly hired ones. This can either be rationalized with the presence of valuable match-specific training or with the presence of relatively high search costs.

Shorter durations for new apprenticeships: The reform we are studying provides a SSC discount only during the first two years of an apprenticeship contract. Assuming that firms provide at least *some* training to apprentices, the discount should induce some firms to accelerate the training process or to shift their apprenticeships towards occupations that require shorter training periods altogether.²² Again, the evidence shown in Section 9 is not consistent with firms strategically manipulating contract duration.

Transformation into open-ended contracts: Whenever the ability of the trainee is not too low for the job at hand, apprenticeship contracts should lead to transformations into open-ended contracts in the short- or medium-term. Firms may hire young workers as apprentices only to take advantage of the discount in SSC and plan to let them go at the end of their contract. The increase in the number of transformations to open-ended contracts documented in Figure 6 is somewhat encouraging in so far as the discount appear to generate at least some long-lasting relationship. An increase in transformations after the discount is consistent with two non-mutually exclusive explanations. Firms may choose to convert some apprentices to open-ended contracts to rip the benefits of having trained them. Alternatively, firms may choose to retain some apprentices even in the absence of training because they represent a good match for the firm. Unfortunately we cannot distinguish between these two explanations.

Apprentices' selection: Finally, if the discount induces small firms to use apprenticeship contracts as a mere source of cheap labor, firms may choose to invest less in the search of talented apprentices and compromise on the quality of new hires. We test whether the

²²Garicano and Rayo (2017) show that, if training is general, firms have an incentive to provide inefficiently slow training in order to extract rents from trainees. It is therefore a possibility that accelerating training is not detrimental to workers' accumulation of human capital

quality of new hires changes as a response to the discount along two dimensions: previous salary and previous experience of newly hired apprentices. Figure 11 reports the point estimates and the associated 95% confidence intervals. We find no evidence that firms hire a different type of apprentices along these two dimensions.²³

We conclude that, while we cannot evaluate whether apprentices are trained appropriately, we find no indication that firms take up the subsidy to merely reduce costs.

6 Number of jobs generated among compliers

The reduced form effects presented in previous sections capture both firm-level changes in employment conditional on policy take-up and changes in policy take-up itself. In what follows we employ an instrumental variable strategy to rescale the reduced form effects by the difference in take-up between the eligible and the ineligible group in each year. This allows us to obtain measures of the number of jobs generated among complier firms. We estimate the following model:

$$Y_{it} = \alpha_i + \gamma_t + \sum_{k=2007}^{2011} \lambda_k \cdot \mathbf{1}(\text{take-up}_{it} = 1) \cdot \mathbf{1}(\text{year}_t = k) + v_{it}, \quad (2)$$

where $\mathbf{1}(\text{take-up}_{it} = 1)$ is a dummy that takes value 1 if the firm is receiving *some* SSC credit on its apprenticeship contracts during year t . This includes the discount on both incumbents and new hires. Since there is no policy before 2007, $\mathbf{1}(\text{take-up}_{it} = 1)$ will be equal to zero for all firms in the sample in the years 2004 to 2006. We still run the regression on all years but exclude $\mathbf{1}(\text{take-up}_{it} = 1) \cdot \mathbf{1}(\text{year}_t = k)$ related to $k = 2004, 2005, 2006$. The remaining λ_k coefficients can thus be interpreted as average differences between treatment and control group relative to average pre-existing differences in years 2004-2006.

We estimate 2 by means of two-stage least squares (2SLS) and instrument $\mathbf{1}(\text{take-up}_{it} = 1) \cdot \mathbf{1}(\text{year}_t = k)$ with $\mathbf{1}(\text{size}_{i,2006} \leq 9) \cdot \mathbf{1}(\text{year} = k)$ (for every k). Our 2SLS estimator captures the change in Y for firms who take up the discount.

We first consider the number of apprenticeship hires as an outcome and report 2SLS estimates in Figure 12. Omitted dummies are set to zero. Consistent with our reduced-form results, we find that the effect unfolds gradually and reaches its peak in 2009, staying constant afterwards. Three years after the implementation of the policy, complier firms hire between

²³We reckon that ability may not be perfectly reflected in wages or experience for this type of workers, and that firms may be screening on other dimensions that we cannot observe in the data.

0.79 and 0.96 more apprentices per year, totalling 2.86 apprentices over five years. We now turn to Figure 13, where we show results on the stock of apprentices. Due to high turnover rates among apprentices, changes in this outcome are less stark. Similarly to the reduced form pattern, the stock of apprentices slightly decreases in 2007 and 2008 and then rises steadily in the following years. At the end of our sample period the policy compliers employ 0.8 more apprentices on average. Since apprenticeship contracts are thought as ports of entry into stable employment for the youth, it is important to look at how many of these contracts are transformed into open-ended ones. We show results on transformations in Figure 14. As expected, here we find considerably smaller point estimates, around 0.1-0.15 depending on the year. Cumulatively the number of transformations totals 0.52 over a five-year period, that is one extra open-ended contract every two firms who take up the policy. While the number of transformations may appear low, these numbers should be interpreted in light of the (rather low) monetary amounts at stake. In the following subsection we compare these effects with the euro amounts that firms received, in order to compute the actual cost per job generated, which we think is a relevant parameter for the welfare analyses of these policies.

7 Robustness Checks

In this section we address some concerns relative to our empirical strategy. First, we show that our results are not driven by differences in baseline firm characteristics and industry-specific trends. Second, we discuss whether firm size may be driving our findings and provide a robustness check that addresses this concern. Third, we document that the differential exposure to the Great Recession is not confounding our results.

7.a Firm-Level Controls

Aside from effects related directly to firm size, it may be that firms above and below the 9 employee threshold differ along characteristics that predict the potential evolution of apprenticeship use. We mitigate these concern by running additional specifications where we add as controls a rich set of baseline firm covariates (\mathbf{x}'_i) interacted with year fixed effects. Our augmented specification reads:

$$y_{it} = \alpha_i + \tau_t + \sum_{k \neq 2006} \gamma_k \cdot \mathbf{1}(\text{year}_t = k) \cdot T_i + \sum_{k \neq 2006} \mathbf{x}'_i \delta_k \cdot \mathbf{1}(\text{year}_t = k) + \epsilon_{it}. \quad (3)$$

where the set of controls \mathbf{x}'_i includes the share of workers aged 29 or less, the share of workers

aged 30-49, the share of workers aged 50 or more, the share of female workers, the share of apprentices, the share of blue-collar workers, the share of white-collar workers, the share of managers, the labor share, liquid assets over total assets, investment over assets, and cash flow over total assets (all measured in 2006). For this set of results the sample size is 98,084 instead of 193,297 as it just includes the firms for which we have balance sheet information.²⁴ The results are presented in Figure 17. We also present robustness tests where instead of firm-level covariates we include linear trends for sector and region. Our specification in this case is:

$$y_{it} = \alpha_i + \tau_t + \sum_{k \neq 2006} \eta_k \cdot \mathbf{1}(\text{year}_t = k) \cdot T_i + \sum_{k \neq 2006} \mathbf{s}'_i \zeta_k \cdot t + \epsilon_{it}. \quad (4)$$

where \mathbf{s}'_i is a set of aggregate controls. One at the time we use: (i) two-digit sector dummies (ii) region dummies (iii) two-digit sector \times region dummies.²⁵ Results are reported in Figure 18. In all cases we see that the qualitative patterns and magnitudes are very similar to our baseline results.

7.b Firm Size

One may be concerned that small and large firms may react differently to the discount (i.e., “size effect”) and that our reduced-form estimates may pick up a combination of the reform impact and this size effect. In order to test for this we run alternative specifications that compares firms that are closer to the threshold and therefore more similar in terms of size.²⁶ More specifically, we replicate our main analysis on the subsample of incumbent firms that have between 8 and 11 employees in 2006.

Figure 19 reports the point estimates and associated 95% confidence interval for the β_k from equation 1 estimated on main analysis sample (blue circles) and on the 8-11 subsample (red circles). The point estimates are exceptionally stable across samples and their difference is not statistically significant. The number of hires is the only notable exception; in this case the point estimates are slightly smaller in magnitude for the 8-11 subsample than those for the main sample. Yet, the confidence intervals overlap for the vast majority of point estimates.

²⁴The baseline results are not altered by just using the *Cerved* rather than the full sample. Results are available upon request.

²⁵We use *linear* trends instead of non-parametric ones purely for computational convenience due to the high number of fixed effects involved.

²⁶The ideal strategy would involve using a regression discontinuity design to evaluate the impact of the reform at the threshold. As we discussed in Section 4, the absence of a sharp discontinuity in the take-up rate documented in Figure 2 prevents us from pursuing this strategy.

Overall, we find no evidence that size effects may be confounding our estimates.

7.c The Great Recession

One may worry that the sharp increase in the estimated impact on hiring and separations that take place in 2009 may be driven by the Great Recession. As long as aggregate shocks, such as the Great Recession, yield common effects on the two groups, these will be absorbed by the year fixed effects. Yet, one may worry that the Great Recession may have a disproportionate impact on firms below the 9-employee threshold in our sample.

To address this concern, we evaluate whether our reduced-form effects are heterogeneous across labor markets exposed more or less severely to the Great Recession. If the crisis was a main driver of our results, then we would expect our estimated effects to be larger in labor markets that were hit more severely. To test this mechanism, we run the following specification

$$y_{it} = \alpha_i + \tau_t + \sum_{k \neq 2006} \beta_k^L \mathbf{1}(\text{year}_t = k) \cdot T_i \cdot (1 - H_{i(l)}) + \sum_{k \neq 2006} \beta_k^H \mathbf{1}(\text{year}_t = k) \cdot T_i \cdot H_{i(l)} + u_{it}, \quad (5)$$

Where all the variables are defined as in equation 1 and $H_{i(l)}$ is an indicator that takes value one when the 2007-2010 change in the unemployment rate of local labor market l (where firm i operates) is above the median change in unemployment rate.²⁷ β_k^H and β_k^L are the main coefficients of interest and they identify the estimated impact for markets above and below median exposure.

Figure 20 reports these coefficients with their 95% confidence intervals. Red dots are associated with β_k^H and blue dots with β_k^L ones. This analysis shows that for all the outcomes of interest the estimated impact of the reform is not heterogeneous across local labor markets that have been affected more or less severely by the Great Recession.

While the unemployment rate is our preferred proxy for the Great Recession, we repeat this analysis using other proxies of financial constraints at the firm level. Firms that were financially more vulnerable in our baseline year may suffer bigger employment losses during the crisis. We follow Saez et al. (2019) and use three proxies for financial constraints: (i) liquid assets over total assets, (ii) cash-flow over total assets, and (iii) total revenues. We estimate model 5 substitute $H_{i(l)}$ with a dummy for whether the proxy for firm i is above the median of

²⁷We follow the 2001 definition by the National Institute of Statistics (Istat), which has grouped Italian municipalities in 686 local labor markets, according to commuting patterns, similarly to US commuting zones.

the distribution of the proxy at baseline. The overall pattern of our results remains unchanged and we conclude that we do not find evidence that the Great Recession acts as a confounder in this setting (see Appendix D, Figures D.3,D.4,D.5).

8 Cost per job created

The SSC discount studied in this paper is a temporary measure offering cost reductions to firms employing apprentices during the first two years of the contract. Such cost reductions are greater during the first year than in the second year. In Section 5 and 6 we showed that firms respond to the policy by increasing both the number of apprenticeship hires and eventually the number of transformations to open-ended contracts. In this section we seek to understand how much money the government had to transfer to firms in order to generate these jobs relative to a scenario where payroll taxes increased for all firms equally.

To this end, we construct a monetary measure (M_{it}) of the SSC discount that each firm receives in a given year, relative to a world where it was not implemented (and both groups of firms are subject to the same SSC rates). The expression for M_{it} then is:

$$M_{it} = \begin{cases} 0, & \text{if } \mathbf{1}(\text{take-up}_{it} = 0) \\ 0.085 \cdot W_{it}^1 + 0.07 \cdot W_{it}^2, & \text{if } \mathbf{1}(\text{take-up}_{it} = 1) \end{cases} \quad (6)$$

where W_{it}^j is the total *net-of-SSC* wage bill (in thousands €) for apprentices employed by firm i in year t during their j -th year of tenure. Indeed the reform grants a 8.5 p.p. discount on SSC during the first year of the contract and a 7 p.p. discount during the second year.²⁸ This amount corresponds to the reduction in costs that firms experience if they take up the discount. One assumption we make here is that labor choices are static and all that is relevant for the firm decision is the money that it receives in each year. We deem this assumption to be reasonable for relatively small firms who may be cash constrained.²⁹

Similarly to subsection 6 we run 2SLS specifications of the form

²⁸In the workers file (UNIEMENS) we observe monthly flags for whether a given apprenticeship contract receives SSC discounts and the exact bracket (8.5% or 7%). This level of detail allows us to separately identify whether an apprentice has been employed for one or two years.

²⁹In Figure D.6 in Appendix D we show that the results are very similar when constructing a measure of the net present value of such transfers for each year t . This alternative measure assumes a discount factor of zero, so that all transfers are simply summed.

$$H_{it}^j = \alpha_i + \gamma_t + \sum_{k=2007}^{2011} \theta_k^j \cdot M_{it} \cdot \mathbf{1}(\text{year} = k) + v_{it}, \quad (8)$$

where H_{it}^j is the number of apprentice hires with duration j (in years). These are the same outcomes that we have used in Section 5.c. For every year k we instrument $M_{it} \cdot \mathbf{1}(\text{year} = k)$ with $\mathbf{1}(\text{size}_{i,2006} \leq 9) \cdot \mathbf{1}(\text{year} = k)$. In this way our 2SLS coefficients θ_k^j can be interpreted as the number of new hires – lasting j years – per thousand Euro spent in year k that are generated in year k . These parameters are informative of the number of jobs and job-years created per Euro spent on the discount. In the same spirit, we also consider an alternative specification:

$$\text{Transf}_{it} = \alpha_i + \gamma_t + \sum_{k=2007}^{2011} \rho_k \cdot M_{it} \cdot \mathbf{1}(\text{year} = k) + v_{it}, \quad (9)$$

where Transf_{it} is the number of apprenticeship transformations to open-ended contracts. The ρ_k can be interpreted very similarly to the θ_k^j coefficients, except that they are not indexed by j .

The θ_k^j coefficients are reported in Figure 15. Similarly to the reduced form effects (Figure 8), this figure shows that, even keeping the monetary transfer constant, most new hires of apprentices tend to be relatively short and the frequency of hires decreases with apprenticeship duration. When summing the θ coefficients across years for each duration j ($\sum_{k=2007}^{2011} \theta_k^j$) and dividing through by the overall sum of coefficients ($\sum_{j=1}^5 \sum_{k=2007}^{2011} \theta_k^j$) we obtain the fraction of jobs of duration j created per €1,000 spent. Per €1,000 spent, 49.7% of these new hires last less than a year and another 47% lasts between 1 and 3 years. Only a marginal fraction lasts more than 3 years. The ρ_k coefficients are displayed in Figure 16. Similarly to what we had found in Section 6, their magnitude is much smaller than coefficients related to hires, because there are fewer transformations than hires occurring at any given firm.

In order to compute the cost per apprenticeship job created, we divide €1,000 by the number of apprenticeship jobs created in the average year per €1,000 spent, that is:

$$\frac{\text{€1,000}}{\frac{1}{5} \sum_{j=1}^5 \sum_{k=2007}^{2011} \theta_k^j} = \frac{\text{€1,000}}{0.482} = \text{€2,075} \quad (10)$$

Thus we find that on average firms in our sample need an incentive of €2,075 to generate

an apprenticeship contract. We perform an analogous computation for transformations to open-ended contracts in the average year:

$$\frac{\text{€}1,000}{\frac{1}{5} \sum_{k=2007}^{2011} \rho_k} = \frac{\text{€}1,000}{0.087} = \text{€}11,494 \quad (11)$$

where we find that on average firms require an incentive of €11,494 to generate a transformation to open-ended contracts. While these figures may seem small, they are in line with recent estimates in the literature. Similarly to us, [Cahuc et al. \(2019\)](#) find that the cost per job is around €8,400 per year. The subsidy they study is also temporary and targeted to low wage workers, like the one in our setting.

Although we cannot disentangle the exact mechanisms driving these results, we try to provide some qualitative considerations. We leave a fuller treatment of these topics to further research. The SSC discount is not *directly* subsidizing permanent employment. It just incentivizes firms to do more training and screening of workers for a limited period of time. Once the SSC discount expires, some of these subsidized apprenticeships will have turned out to be productive matches for the firm, while some of them will have not. Given that the subsidy is small and only affects the relative price of training and screening, it must be that any decision to offer an open-ended contract at the end of training must be privately optimal for the firm, even absent the SSC discount.³⁰ In other words the government generates transformations only *indirectly*. The effectiveness of the policy should be only driven by the transformation rate of apprentices and not by the costs associated to permanent employment. Because of this, the economic incentive that the government needs to provide to indirectly to these transformations is low. A second channel through which the policy is presumably cost effective is the tilted schedule of the SSC discount. The firm is plausibly receiving a higher transfer when the demand for apprenticeships is more elastic, as there is a lot of uncertainty around workers' types and training success, and decreases in the following years, when match-specific investments have been made.

To conclude this section, we propose a simple back-of-the-envelope calculation that produces a rough estimate of the total number of apprenticeship jobs and job years saved thanks to the SSC discount, for the firms in our sample.³¹ We proceed as follows: each coefficient θ_k^j

³⁰If the SSC was bigger in size, then the firm could use part of the public transfers to finance training and part to sustain the costs of a transformation into an open-ended contract. Since the measure we study here is modest in size, we abstract from these considerations.

³¹We are working on an extension that looks at the number of jobs-years after the end of the apprenticeship contract, by taking into account subsequent years of tenure at the training firm. Preliminary results indicate that there is a minority of matches that still exist in 2018, nine years after the initial hire. Given that the SSC

corresponds to the additional number of apprenticeship hires of duration j that is generated by spending €1,000 in year k on this policy. Since we know the amount of money spent on the policy in each year, we can back out the total number of apprenticeship jobs generated by the policy. For each contract duration j we have that:

$$\Delta\text{Jobs}^j = \sum_{k=2007}^{2011} \theta_k^j \cdot \mathbf{M}_k \quad (12)$$

where \mathbf{M}_k is the total amount spent for the policy in year k on *all* firms belonging to our sample, both eligible and not eligible. We can then easily recover the number of jobs and job years as follows:

$$\Delta\text{Jobs} = \sum_{j=1}^5 \Delta\text{Jobs}^j \quad (13)$$

$$\Delta\text{Job-years} = \sum_{j=1}^5 j \cdot \Delta\text{Jobs}^j \quad (14)$$

The key assumption behind this exercise is that we can extrapolate the average treatment effects θ_k^j to the whole sample and that relative differences across firms reflect absolute employment changes. This would be violated if there were spillovers between eligible and ineligible units or if there were general equilibrium effects.³²

We report the monetary amounts (average per firm and total) in Table 5, and the results of this exercise in Table 6. We find that the reform generated approximately 54,000 jobs and 87,000 job years. An analogous exercise for open-ended contracts yields a total of 11,281 transformations.

8.a Back of the envelope calculation

In Section 2.b we documented that the social security contributions increased on average for all firms and that this increase was more pronounced for firm above the 9-employee

discount is offered just during the first two years of the apprenticeship contract, these jobs are quite cheap from a fiscal perspective and most likely pay for themselves thanks to subsequent taxes collected on labor. All in all our estimates provide a lower bound on the number of job years that is generated thanks to this reform.

³²This is similar in spirit to Autor et al. (2013) who want to have an estimate of the total number of jobs lost in the US due to Chinese competition and can only compare commuting zones that are more or less exposed to such competition.

threshold. In the rest of the paper we exploited the plausibly random reduction in social security contributions paid by firms eligible for the subsidy relative to ineligible firms. In this section, we perform a simple back of the envelope calculation to estimate the reduction in the number of apprentice hires and open-ended contracts in our sample driven by the overall increase in the cost of social security contributions.

We do so by computing the elasticity of the number of hires to the change in social security contributions as the ratio of the percentage change in hiring estimated in Subsection 5.a and the percentage change in social security contribution costs for the average apprenticeship contract (i.e., a 20 month contract paying 1,000 euros per month). Under the assumptions of constant elasticity and no of general equilibrium effects, we estimate that the overall increase in social security contributions after 2007 resulted in a reduction of 34.54% apprenticeship hires at firms between 9 and 14 employees, which corresponds to 25,800 fewer apprenticeship contracts (for details on how we construct these figures, please refer to Appendix C).³³ Similarly, we estimate that these higher costs caused a reduction of the number of apprenticeship transformations to open-ended contracts of 24.79% in our sample (roughly corresponding to 5,000 open-ended contracts).

9 Conclusions

Although apprenticeships do not represent a very popular contractual arrangement, several countries offer lower social security contributions or favorable taxation regimes to incentivize their use. Yet, little is known about the firm-level responses to these type of measures and their cost-effectiveness.

In this paper we have studied the effectiveness of hiring credits for apprenticeship contracts. We have exploited a reform that increased the SSC contributions for apprenticeships but granted a discount to firms below 9 employees. We have found that lower SSC rates stimulate demand for this contract type and contained the overall decrease in the number of apprenticeship contracts. Overall the reform generated 2.96 apprenticeship contracts and 0.52 transformations per complier firm over a period of five years. On average the government spent €2,075 to shield each additional apprenticeship contract and €11,400 to generate an open-ended contract. These figures represent an upper bound for the cost of this policy as they do not incorporate the social transfers that would have been directed to the apprentices who would have otherwise been unemployed and the tax revenues that these contracts generate.

³³By “constant elasticity” we mean that we assume the elasticity to be constant over time and not to differ between eligible and non eligible firms.

We conclude that light touch interventions, like the policy we study, may be effective at increasing the utilization of apprenticeship contracts without having unintended consequences such as firm-size distortions and strategic firm behavior.

In line with our findings, the policies that proved to be successful at stimulating labor demand were implemented in context characterized by wage rigidities ([Cahuc et al., 2019](#); [Saez et al., 2019](#)). We speculate that the extent to which wages can adjust in the aftermath of the policy is a key determinant of the effectiveness of payroll reductions. Our results are unlikely to generalize to settings where wages adjust flexibly.

References

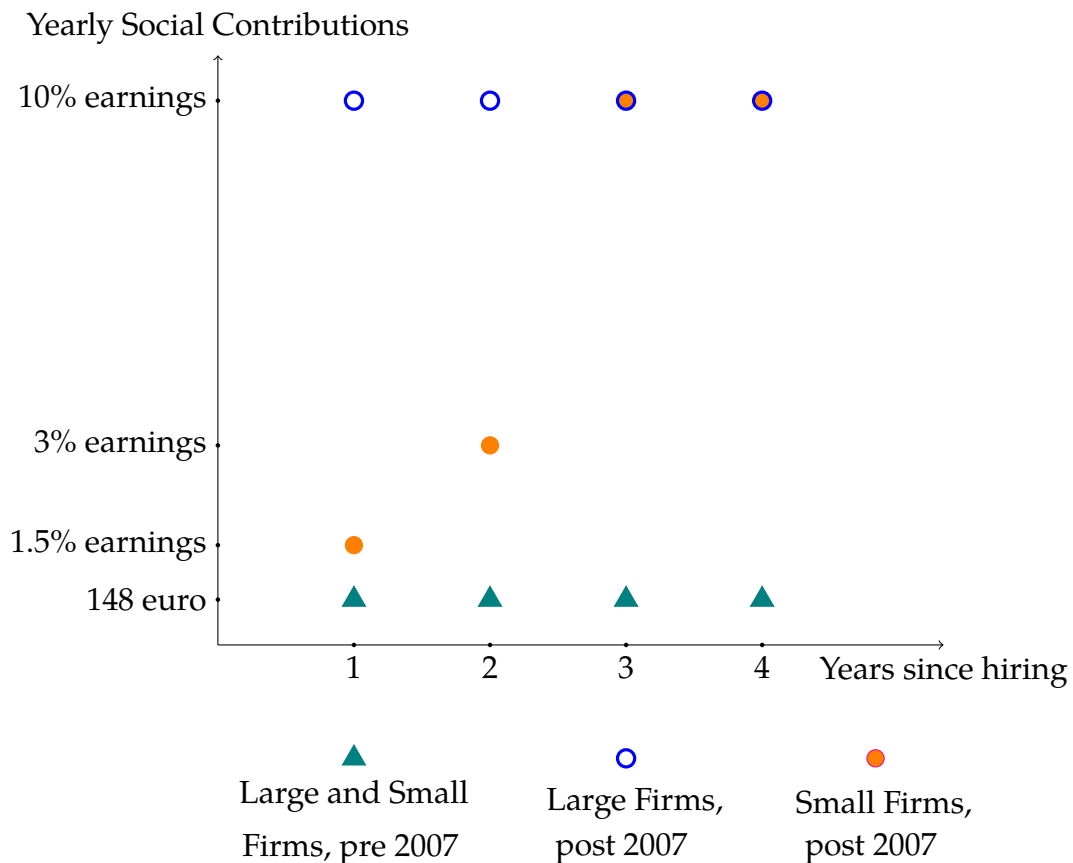
- Acemoglu, D. and Pischke, J.-S. (1998). Why do firms train? theory and evidence. *The Quarterly journal of economics*, 113(1):79–119.
- Acemoglu, D. and Pischke, J.-S. (1999). The structure of wages and investment in general training. *Journal of political economy*, 107(3):539–572.
- Albanese, A., Cappellari, L., and Leonardi, M. (2017). The effects of youth labour market reforms: evidence from italian apprenticeships. *Oxford Economic Papers*.
- Alfonsi, L., Bandiera, O., Bassi, V., Burgess, R., Rasul, I., Sulaiman, M., and Vitali, A. (2020). Tackling youth unemployment: Evidence from a labor market experiment in uganda. *Econometrica*, 88(6):2369–2414.
- Autor, D., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *The American Economic Review*, 103(6):2121–2168.
- Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of political economy*, 70(5, Part 2):9–49.
- Bennmarker, H., Mellander, E., and Öckert, B. (2009). Do regional payroll tax reductions boost employment? *Labour Economics*, 16(5):480–489.
- Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*, 119(1):249–275.
- Bohm, P. and Lind, H. (1993). Policy evaluation quality: A quasi-experimental study of regional employment subsidies in sweden. *Regional Science and Urban Economics*, 23(1):51–65.
- Cahuc, P., Carcillo, S., and Le Barbanchon, T. (2019). The effectiveness of hiring credits. *The Review of Economic Studies*, 86(2):593–626.
- Caicedo, S., Espinosa, M., and Seibold, A. (2020). Unwilling to train? firm responses to the colombian apprenticeship regulation. Technical report, CESifo Working Paper.
- Cappellari, L., Dell’Arlinga, C., and Leonardi, M. (2012). Temporary employment, job flows and productivity: A tale of two reforms*. *The Economic Journal*, 122(562):F188–F215.
- Cavaglia, C., McNally, S., and Ventura, G. (2020). Do apprenticeships pay? evidence for england. *Oxford Bulletin of Economics and Statistics*, 82(5):1094–1134.

- Crépon, B. and Premand, P. (2019). Direct and indirect effects of subsidized dual apprenticeships.
- D'Agostino, S. et al. (2010). *Apprendistato: un sistema plurale: 10. rapporto di monitoraggio*. Isfol.
- D'Arcangelo, A., Romito, A., et al. (2019). *L'apprendistato tra continuità e innovazione: 18. rapporto di monitoraggio*. Technical report.
- Dustmann, C. and Schönberg, U. (2012). What makes firm-based vocational training schemes successful? the role of commitment. *American Economic Journal: Applied Economics*, 4(2):36–61.
- Egebark, J. and Kaunitz, N. (2013). Do payroll tax cuts raise youth employment? Working Paper Series 2013:27, IFAU - Institute for Evaluation of Labour Market and Education Policy.
- Fersterer, J., Pischke, J.-S., and Winter-Ebmer, R. (2008). Returns to apprenticeship training in austria: Evidence from failed firms. *Scandinavian journal of economics*, 110(4):733–753.
- Garicano, L., Lelarge, C., and Van Reenen, J. (2016). Firm size distortions and the productivity distribution: Evidence from france. *American Economic Review*, 106(11):3439–79.
- Garicano, L. and Rayo, L. (2017). Relational knowledge transfers. *American Economic Review*, 107(9):2695–2730.
- Huttunen, K., Pirttilä, J., and Uusitalo, R. (2013). The employment effects of low-wage subsidies. *Journal of Public Economics*, 97:49–60.
- Kleven, H. J. and Kreiner, C. T. (2006). The marginal cost of public funds: Hours of work versus labor force participation. *Journal of Public Economics*, 90(10-11):1955–1973.
- Korkeamäki, O. and Uusitalo, R. (2006). Employment Effects of a Payroll-Tax Cut: Evidence from a Regional Tax Exemption Experiment. Discussion Papers 407, VATT Institute for Economic Research.
- Kuczera, M. (2017). Incentives for apprenticeship.
- Lodovici, M. S., Comi, S., Origo, F., Patrizio, M., Torchio, N., Speckesser, S., and Montalt, J. V.-B. (2013). The effectiveness and costs-benefits of apprenticeships: Results of the quantitative analysis.
- Mohrenweiser, J. and Zwick, T. (2009). Why do firms train apprentices? the net cost puzzle reconsidered. *Labour Economics*, 16(6):631–637.

- Neumark, D. (2013). Spurring job creation in response to severe recessions: Reconsidering hiring credits. *Journal of Policy Analysis and Management*, 32(1):142–171.
- Picchio, M. and Staffolani, S. (2019). Does apprenticeship improve job opportunities? a regression discontinuity approach. *Empirical Economics*, 56(1):23–60.
- Saez, E., Schoefer, B., and Seim, D. (2019). Payroll taxes, firm behavior, and rent sharing: Evidence from a young workers’ tax cut in sweden. *American Economic Review*, 109(5):1717–63.
- Sestito, P. and Viviano, E. (2018). Firing costs and firm hiring: evidence from an italian reform. *Economic Policy*, 33(93):101–130.
- Tiraboschi, M. (2014). Young workers in recessionary times: A caveat (to continental europe) to reconstruct its labour law? GundersenM. FazioF.(Eds.), *Tackling youth unemployment*, pages 3–26.
- Wolter, S. C., Mühlemann, S., and Schweri, J. (2006). Why some firms train apprentices and many others do not. *German Economic Review*, 7(3):249–264.

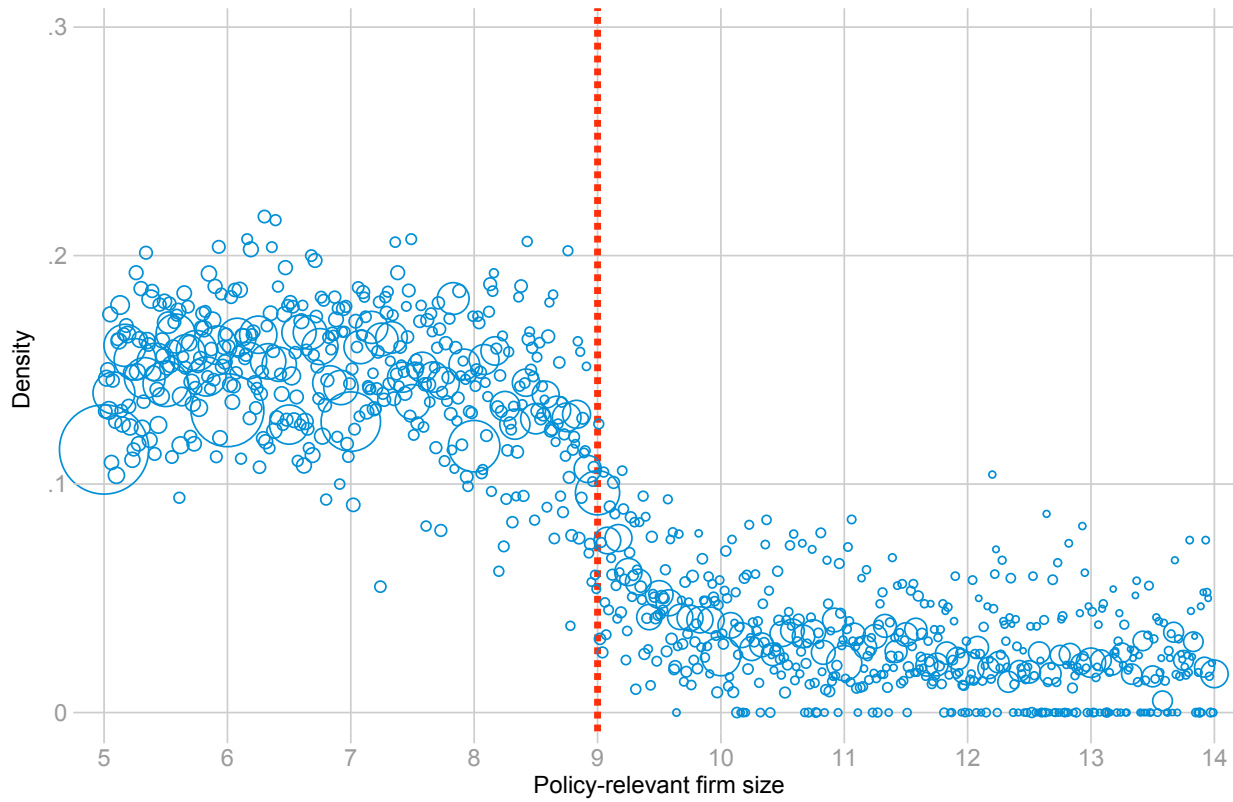
Figures

Figure 1: Social Contributions for Apprenticeship Contracts




Note: This Figure illustrates how yearly social security contributions for apprenticeship contracts changed in response to the 2007 Budget Bill. Before 2007 all employers paid a fixed weekly fee of 2.85 euros per apprenticeship contract. The yearly social contributions are computed as $2.85 \times 52 = 148.2$ euros (green triangles). Yearly social contributions for the period after January 1, 2007 are computed as a percentage of yearly earnings (*imponibile previdenziale*) and the schedule differs between firms below or above the 9-employee threshold (labeled “small” and “large” firms, respectively). Social contributions amount to 10% of the apprentice’s earnings for large firms (blue hollow circles). Small firms pay 1.5% of the apprentice’s earnings in the first year of the contract, 3% in the second year, and 10% in the third year and all the following ones (orange circles).

Figure 2: Fraction of firms taking up the SSC discount by firm size



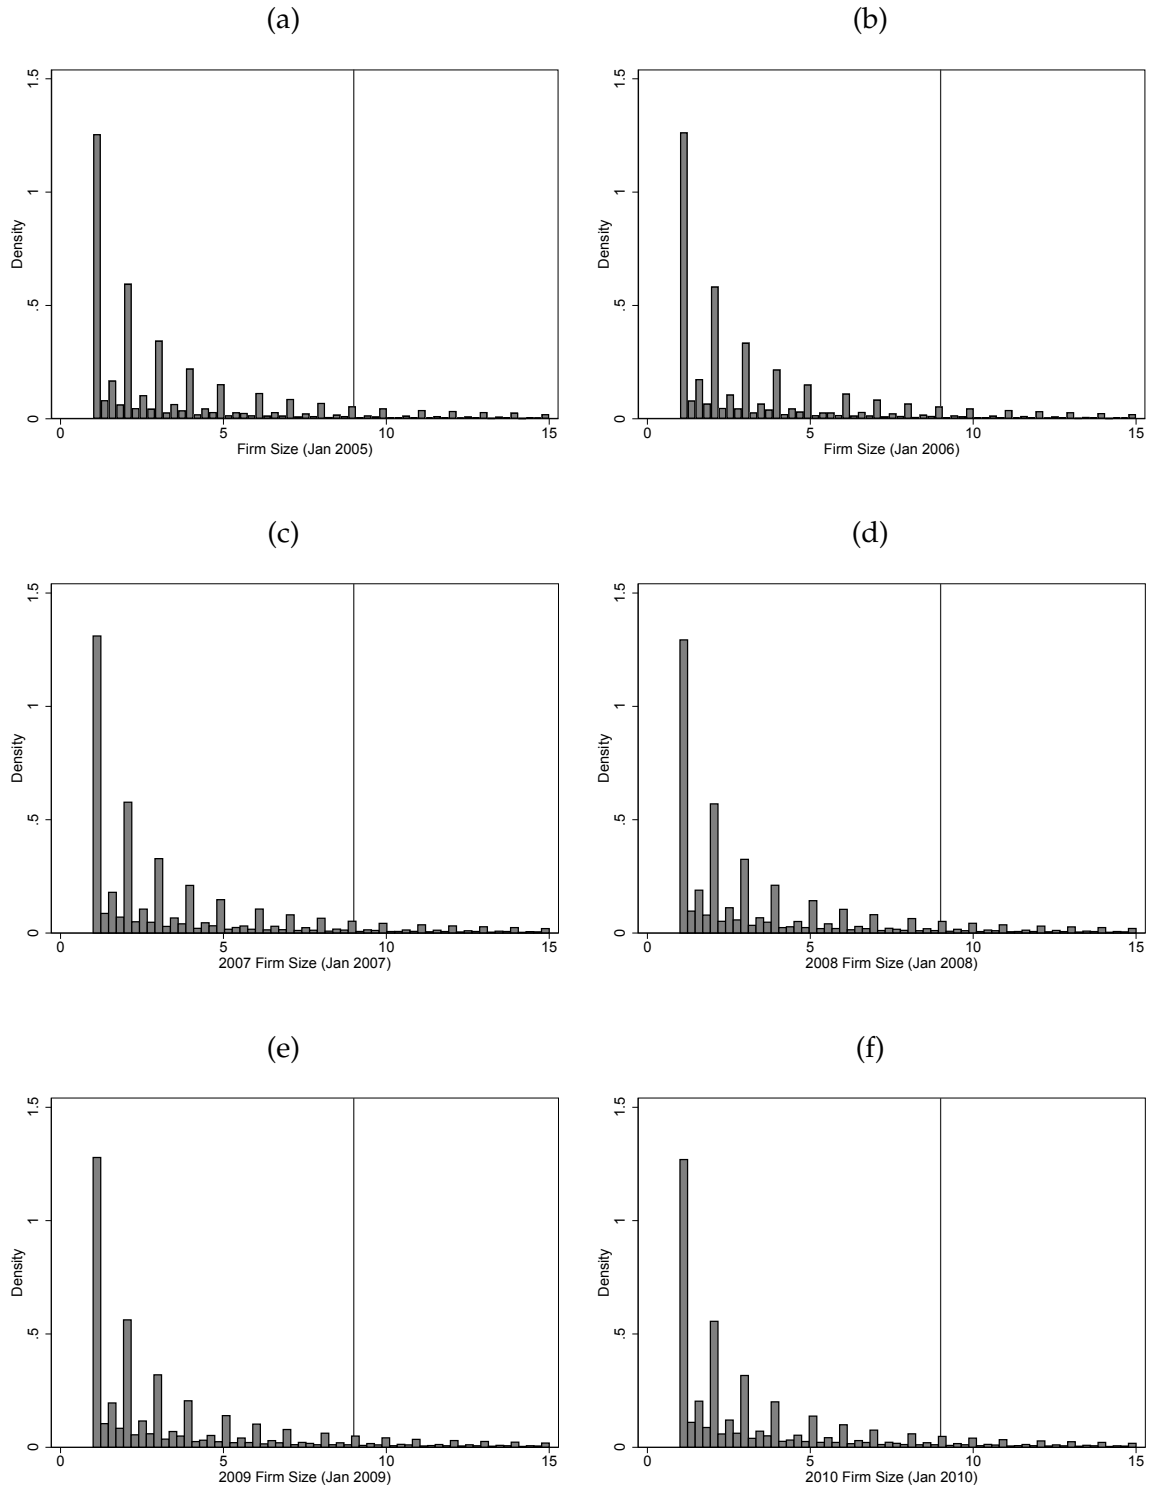
Note: This figure illustrates the relationship between the share of firms that take up the SSC discount in January 2007 and bins of 2006 (policy-relevant) average firm size. The size of the bubble represents the number of firms in each bin. The red dashed line indicates the 9-employee threshold.

Figure 3: Attention rate by firm size




Note: This figure illustrates the relationship between 2007 attention rate and bins of 2006 (policy-relevant) average firm size. The attention rate is constructed as the share of firms that claim the subsidy relative to the number of firms that hire at least one apprentice. The red dashed line indicates the 9-employee threshold.

Figure 4: Distribution of Firm Size by Year



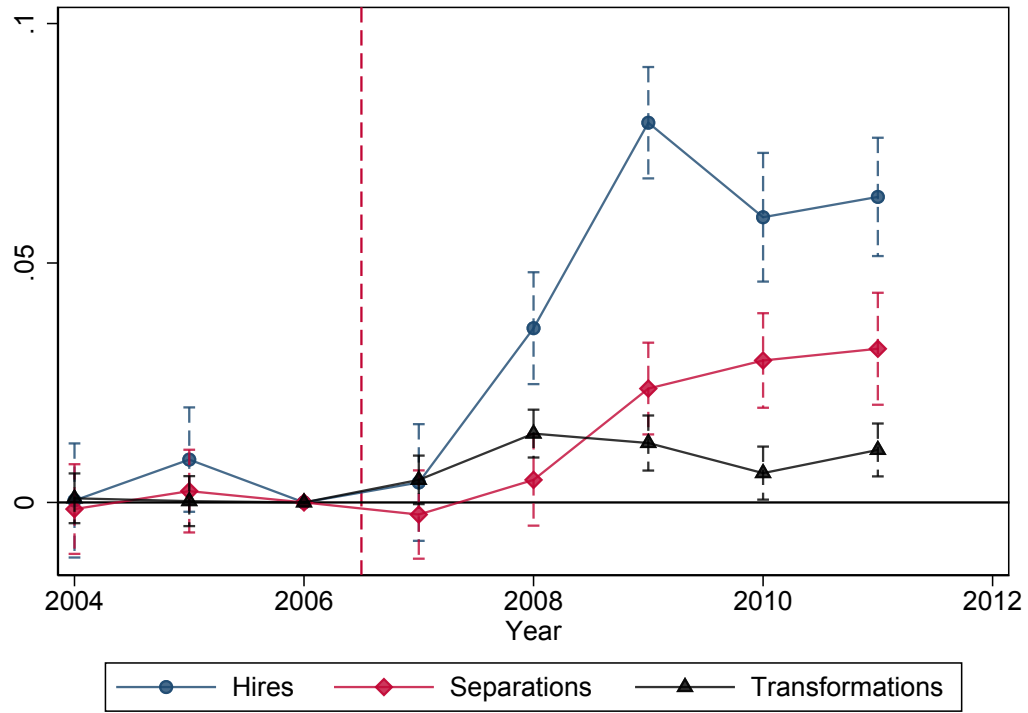
Note: This figure reports the distribution of the policy-relevant firm size by year. Panel A through F report the distribution from 2005 to 2010, respectively.

Figure 5: Changes in the monthly earnings of apprentices by firm eligibility status




Note: This figure reports average monthly earnings, net of SSC for the stock of apprenticeship contracts in each year for the firms in our sample. Firms are divided according to whether their policy-relevant firm size was between 5 and 9 (blue circles) or between 10 and 14 (red circles) in 2006.

Figure 6: Reduced-form impact of the reform on apprentice flows - OLS estimates




Note: This figure reports the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, i.e., $\hat{\beta}_k$ from equation 1. The coefficients β_{2006} are normalized to zero. The dependent variables are number of apprenticeship hires (blue circles), number of apprenticeship separations (red diamonds), and number of transformations of apprenticeship contracts to open ended contracts (black triangles). Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 7: Reduced-form impact of the reform on the stock of apprenticeships - OLS estimates




Note: This figure reports the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, i.e., $\hat{\beta}_k$ from equation 1. The coefficients β_{2006} are normalized to zero. The dependent variable is the firm-level stock of apprenticeship contracts in year t , measured *pro rata temporis* and in full-time equivalent terms. Each apprenticeship contract accounts for $\frac{m}{12}$ towards the firm-level count, where m is the number of months the apprenticeship contract last during a given year. The contribution of a given apprenticeship contract in a given month is reduced proportionally to the number of effective hours if the apprentice works part-time. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 8: Reduced-form impact of the reform on new hires of apprentices by their contract duration - OLS estimates




Note: This figure reports the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, i.e., $\hat{\beta}_k$ from equation 1. The coefficients β_{2006} are normalized to zero. These dependent variables are the number of newly hired apprentices by apprenticeship contract duration, classified in five mutually exclusive and jointly exhaustive bins. More specifically, this set of outcomes is constructed as the number of apprenticeship hires in each given year that last v months (where $v \in \{[1, 12], [13, 24], [25, 36], [37, 48], [49, \max]\}$). Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 9: Reduced-form impact of the reform on apprenticeship contract duration - OLS estimates




Note: This figure reports the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, i.e., $\hat{\beta}_k$ from (1). The coefficients β_{2006} are normalized to zero. The dependent variable is the average forward-looking duration of all apprenticeship contracts started in year t (in months). It is defined only for firms starting new apprenticeship contracts in year t . Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 10: Reduced-form impact of the reform on the number of young hires - OLS estimates




Note: This figure reports the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, i.e., $\hat{\beta}_k$ from equation 1. The coefficients β_{2006} are normalized to zero. The dependent variable is the firm-level number of young hires in year t . Young hires are at most 29 years old at the moment of hire. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 11: Reduced-form impact of the reform on the characteristics of new hires - OLS estimates




Note: This figure reports the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, i.e., $\hat{\beta}_k$ from (1). The coefficients β_{2006} are normalized to zero. The two dependent variables are the average labor market experience (in months) of apprentices hired in year t (blue diamonds), and the average weekly wage (in €) that new apprentices were earning before starting the apprenticeship. Zeros are included if the newly hired apprentice has no previous experience. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 12: Effects of the reform on new hires of apprentices - 2SLS Estimates




Note: The figure reports the λ_k regression coefficients and the associated 95% confidence interval that identify the impact of the reform on complier firms (equation 2). The coefficients $\lambda_{2004}, \lambda_{2005}, \lambda_{2006}$ are normalized to zero. The outcome is the number of new apprentices hires. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 13: Effects of the reform on the stock of apprenticeship contracts - 2SLS Estimates




Note: The figure reports the λ_k regression coefficients and the associated 95% confidence interval that identify the impact of the reform on complier firms (equation 2). The coefficients $\lambda_{2004}, \lambda_{2005}, \lambda_{2006}$ are normalized to zero. The outcome is the stock of apprenticeship contracts in year t , measured *pro rata temporis* and in full-time equivalent terms. Each apprenticeship contract accounts for $\frac{m}{12}$ towards the firm-level count, where m is the number of months the apprenticeship contract last during a given year. The contribution of a given apprenticeship contract in a given month is reduced proportionally to the number of effective hours if the apprentice works part-time. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 14: Effects of the reform on the number of transformations of apprenticeship contracts to open-ended contracts - 2SLS Estimates




Note: The figure reports the λ_k regression coefficients and the associated 95% confidence interval that identify the impact of the reform on complier firms (equation 2). The coefficients $\lambda_{2004}, \lambda_{2005}, \lambda_{2006}$ are normalized to zero. The outcome is the number of transformations of apprenticeship contracts to open ended contracts. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 15: Effects of € 1000 reduced social security contributions on new hires of apprentices by duration - 2SLS estimates



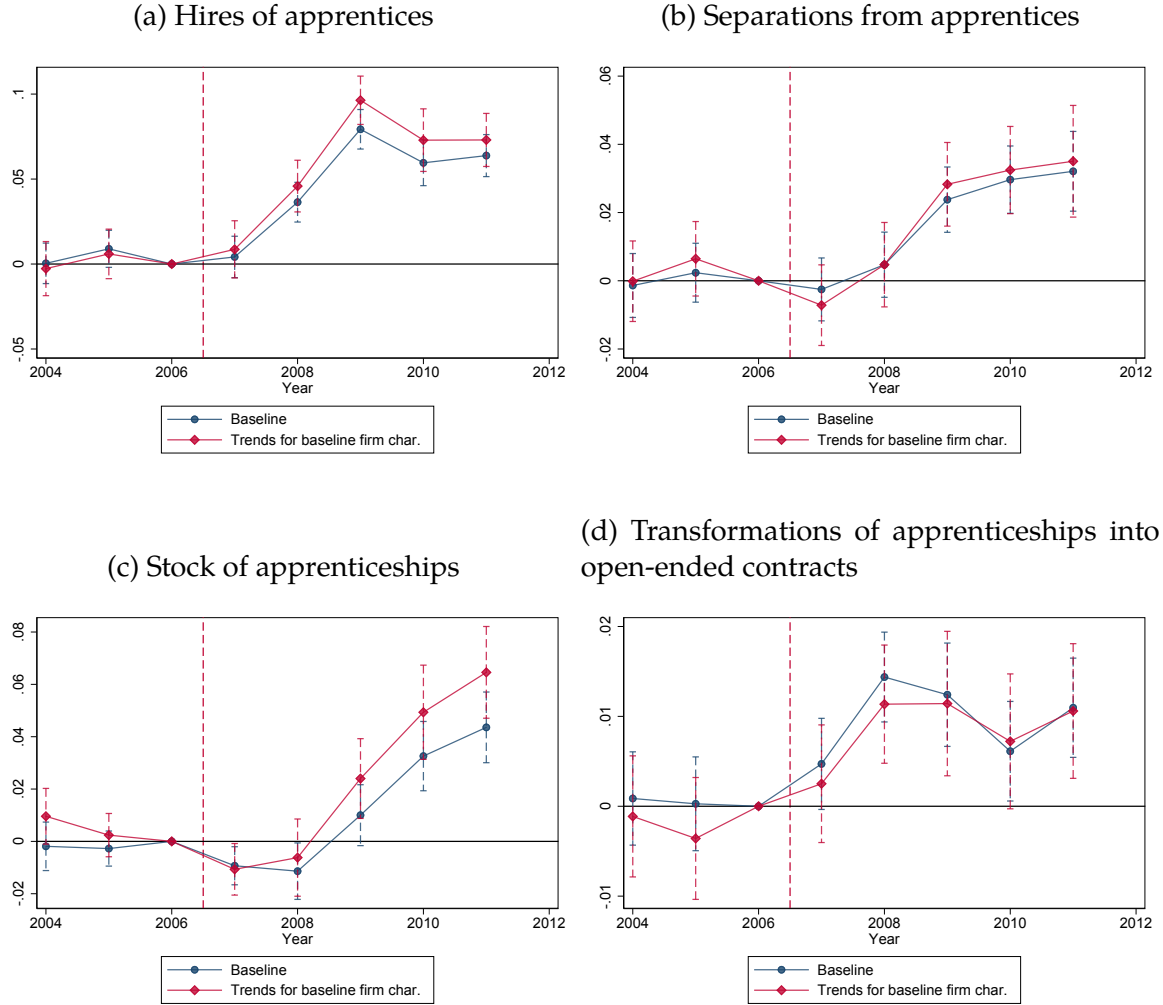
Note: The figure reports the θ_k^j regression coefficients and the associated 95% confidence interval that identify the impact of a €1000 cost reduction for apprenticeship contracts on new hires of apprentices among complier firms (equation 8). The coefficients $\theta_{2004}^j, \theta_{2005}^j, \theta_{2006}^j$ are normalized to zero. These dependent variables are the number of newly hired apprentices by apprenticeship contract duration, classified in five mutually exclusive and jointly exhaustive bins. More specifically, this set of outcomes is constructed as the number of apprenticeship hires in each given year that last v months (where $v \in \{[1, 12], [13, 24], [25, 36], [37, 48], [49, \max]\}$). Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 16: Effects of €1000 reduced social security contributions on the number of transformations of apprenticeships into open-ended contracts - 2SLS estimates



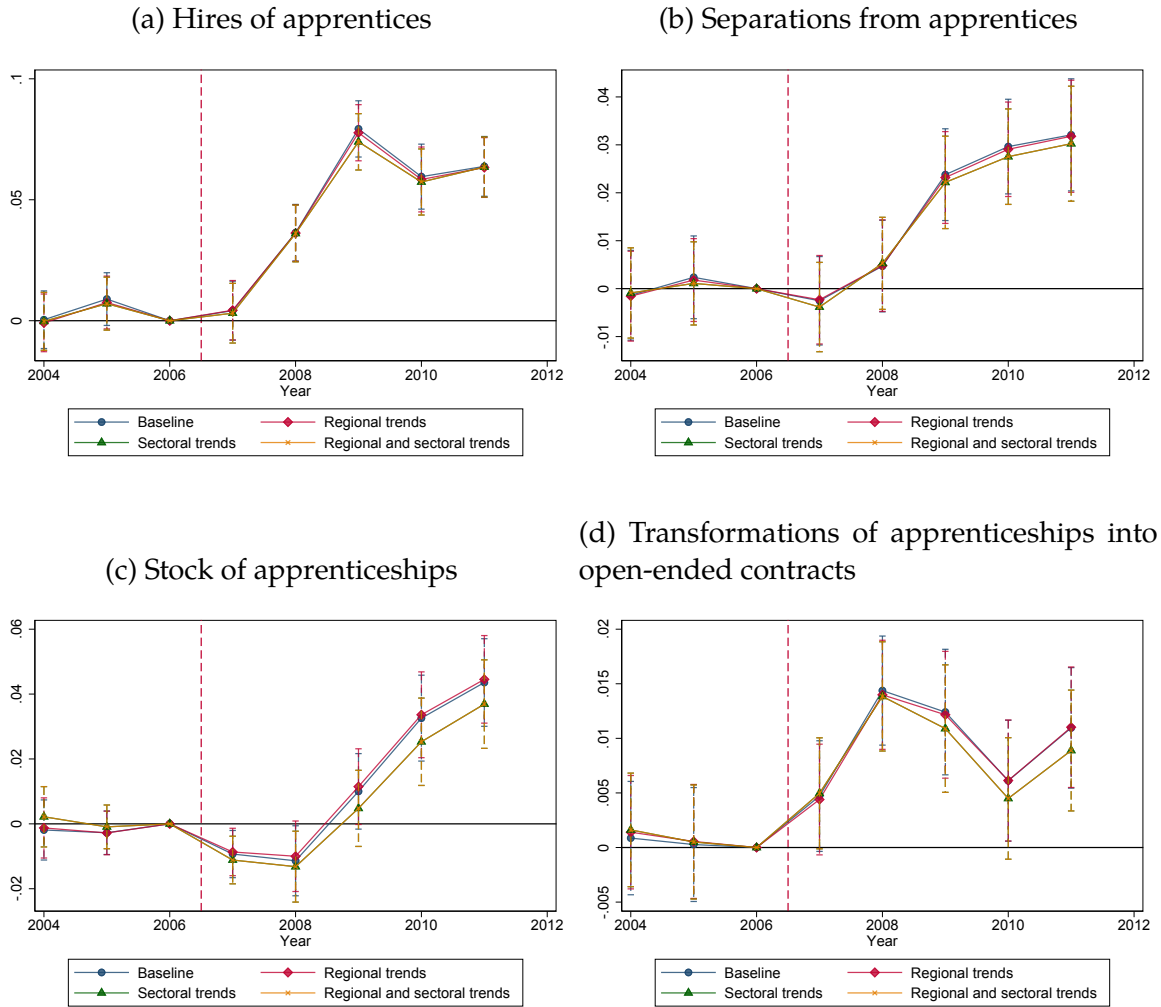
Note: The figure reports the ρ_k regression coefficients and the associated 95% confidence interval that identify the impact of a €1000 cost reduction for apprenticeship contracts on new hires of apprentices among complier firms (equation 8). The coefficients $\rho_{2004}, \rho_{2005}, \rho_{2006}$ are normalized to zero. The outcome is the number of transformations of apprenticeship contracts to open ended contracts. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 17: Reduced-form impact of the reform: robustness to the inclusion of trends for baseline firm characteristics - OLS estimates



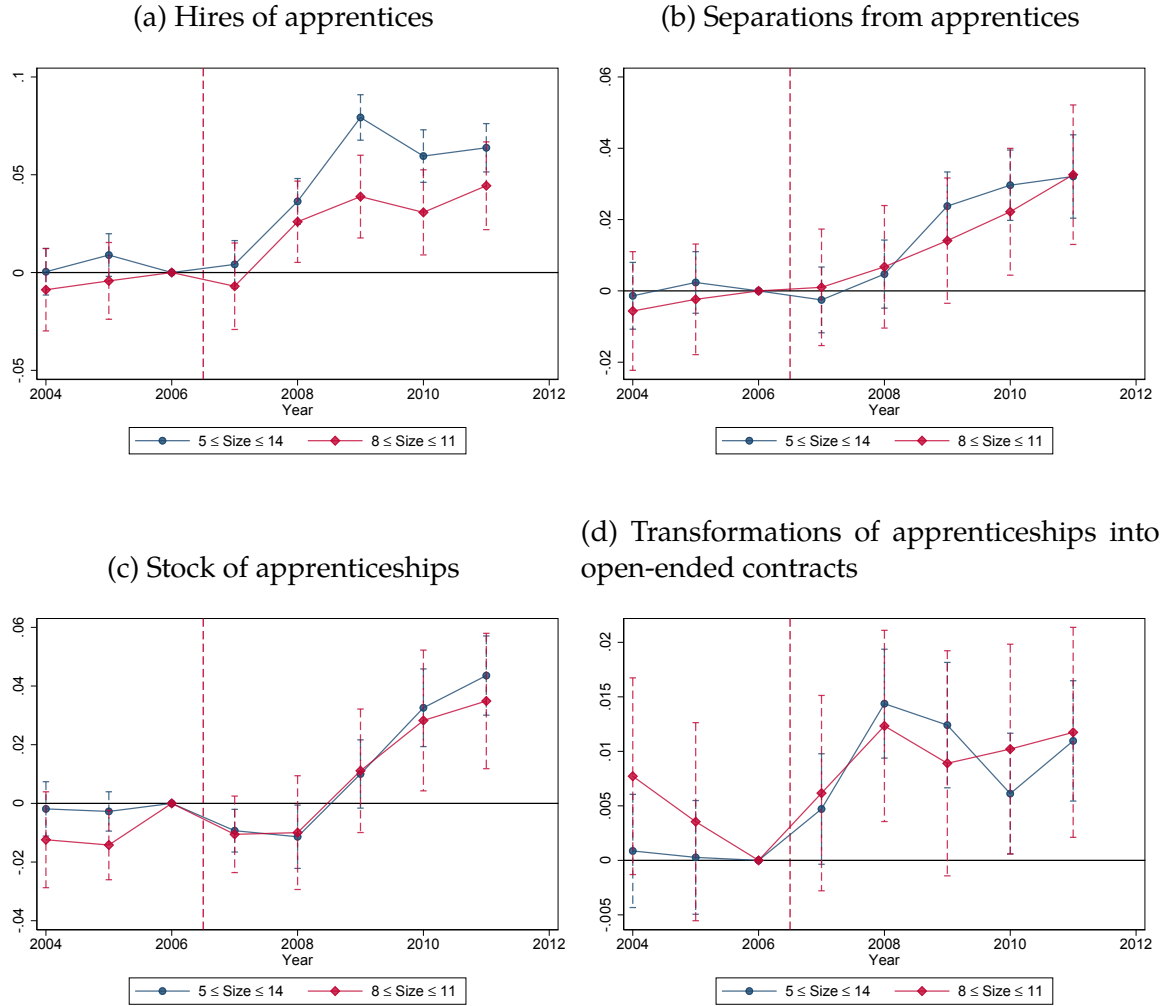
Note: These figures report the regression coefficients and the associated 95% confidence intervals that identify the reduced form impacts of the reform, with and without the inclusion of baseline controls interacted with year fixed effects. These correspond to parameters β_k in equation 1 and γ_k in equation 3. $\beta_{2006}, \gamma_{2006}$ are normalized to zero. Baseline controls include the share of workers aged 29 or less, the share of workers aged 30-49, the share of workers aged 50 or more, the share of female workers, the share of apprentices, the share of blue-collar workers, the share of white-collar workers, the share of managers, the labor share, liquid assets over total assets, investment over assets, and cash flow over total assets (all measured in 2006). The regressions with the additional controls are run on the restricted sample for which we observe all the covariates. This corresponds to 98,084 firms. The baseline regressions are run on the full sample that is 193,297 firms. The names of the outcomes are reported on top of the relative subfigure. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 18: Reduced-form impact of the reform: robustness to the inclusion of sector and province-specific trends - OLS estimates



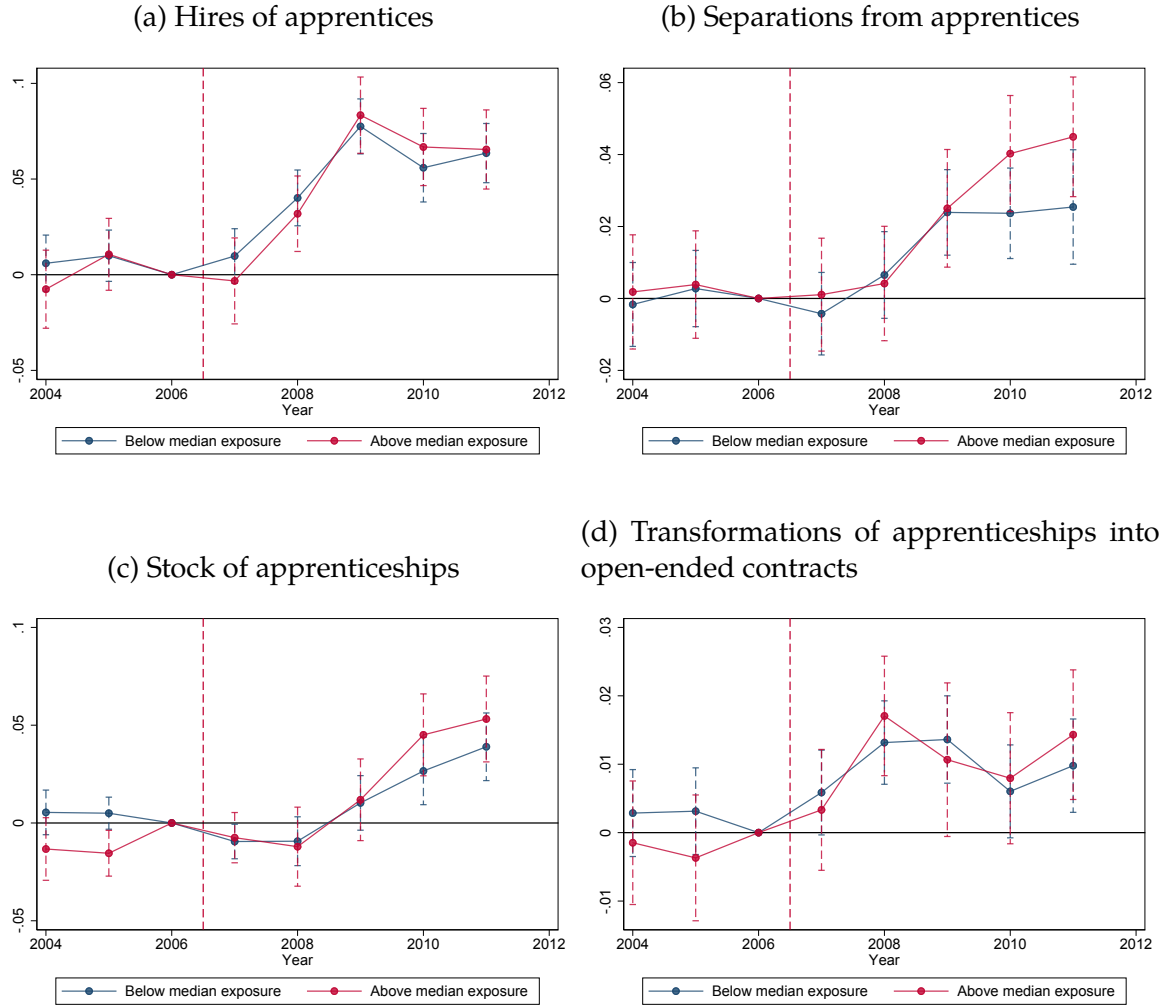
Note: These figures report the regression coefficients and the associated 95% confidence intervals that identify the reduced form impacts of the reform, with and without the inclusion of (i) two-digit sector specific linear trends (green triangles) (ii) region specific linear trends (red circles) (ii) two digit sector \times region specific linear trends (yellow crosses). These correspond to η_k coefficients from equation 4. The names of the outcomes are reported on top of the relative subfigure. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 19: Heterogeneous effects by window size - OLS estimates



Note: These figure reports the β_k regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform (equation 1), depending on whether we only consider firms whose 2006 policy-relevant size is $\in [5, 14]$ (blue circles) or $\in [8, 11]$ (red circles). The names of the outcomes are reported on top of the relative subfigure. Standard errors are clustered at the firm level. The x-axis indexes time.

Figure 20: Heterogeneous effect by local exposure to the Great Recession - OLS estimates



Note: These figure reports the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, depending on the level of exposure to the Great Recession. These correspond to coefficients $\hat{\beta}_k^L$ (blue circles) and $\hat{\beta}_k^H$ (red circles) from equation 5. The coefficients relative to $k = 2006$ are normalized to zero. Exposure is defined as the 2007-2010 change (in p.p.) in the unemployment rate of the local labor market where the firm operates. The names of the outcomes are reported on top of the relative subfigure. Standard errors are clustered at the firm level. The x-axis indexes time.

10 Tables

Table 1: Yearly Social Contributions for the Average Apprenticeship Contract

Years since hiring	Before Jan 1, 2007		After Jan 1, 2007		$\Delta_{After-Before}$	
	Large (1)	Small (2)	Large (3)	Small (4)	Large (5)	Small (6)
1	148	148	1200	180	1052	32
2	148	148	1200	360	1052	212
3	148	148	1200	1200	1052	1052

Note: This Table illustrates how yearly social security contributions for the average apprenticeship contract changed in response to the 2007 Budget Bill. Before 2007 all employers paid a fixed weekly fee of 2.85 euros per apprenticeship contract. The yearly social contributions are computed as $2.85 \times 52 = 148.2$ euros. Yearly social contributions for the period after January 1, 2007 are computed as a percentage of yearly wages and the schedule differs between large and small firms. Social contributions amount to 10% of the apprentice's wage for large firms. Small firms pay 1.5% of the apprentice's wage in the first year of the contract, 3% in the second year, and 10% in the third year and all the following ones. To compute the change in social contributions implied by this policy, we use the yearly average 2006 apprentice wage which is equal to 12,000 euros.

Table 2: Firm Characteristics at Baseline (2006)

	Firm Size $\in [5, 14]$	Size $\in [5, 9)$	Size $\in [9, 14]$
<i>Panel A: Worker Composition and Firm Age</i>			
Share of Workers Aged 0-29	0.262	0.269	0.25
Share of Workers Aged 30-49	0.598	0.593	0.607
Share of Workers Aged 50-100	0.14	0.139	0.143
Share Female	0.353	0.358	0.342
Firm age	14.876	14.566	15.483
<i>Panel B: Industry Composition</i>			
Agric., Hunt, and Forestry	0.005	0.005	0.004
Fishery	0.002	0.001	0.002
Extraction	0.005	0.004	0.006
Manufacturing	0.345	0.318	0.398
Energy	0.001	0.001	0.001
Construction	0.163	0.171	0.148
Commerce	0.194	0.203	0.178
Hospitality	0.072	0.078	0.06
Transport and Comm.	0.053	0.053	0.053
Finance	0.01	0.01	0.008
Real Estate and Other Prof. Activ.	0.092	0.096	0.084
Other	0.059	0.059	0.058
<i>Panel C: Geography</i>			
Abruzzo	0.02	0.02	0.021
Basilicata	0.007	0.007	0.006
Calabria	0.019	0.02	0.017
Campania	0.067	0.068	0.065
Emilia Romagna	0.09	0.089	0.092
Friuli Venezia Giulia	0.023	0.023	0.023
Lazio	0.076	0.076	0.076
Liguria	0.024	0.025	0.024
Lombardy	0.217	0.213	0.224
Marche	0.032	0.032	0.033
Molise	0.004	0.004	0.004
Piedmont	0.073	0.073	0.074
Apulia	0.047	0.048	0.044
Sardinia	0.024	0.025	0.022
Sicily	0.05	0.052	0.046
Tuscany	0.076	0.076	0.077
Trentino Alto Adige	0.024	0.025	0.023
Umbria	0.016	0.016	0.124
Valle d'Aosta	0.002	0.003	0.002
Veneto	0.108	0.107	0.111
Observations	193296	127857	65439

Notes: The full sample includes all incumbent firms between 5 and 14 employees in 2006. Firms are classified as small vs large based on their policy-relevant firm size in 2006. Small firms include all incumbent firms between 5 and 9 employees, while large firms include all incumbent firms with more than 9 but less than 14 employees. All statistics are calculated across firm-year observations at baseline (i.e., 2006). The number of observations for Panel B is 193018, 127641, and 65377, respectively.

Table 3: Firm Outcomes (2004-2011)

	Firm Size $\in [5, 14]$	Size $\in [5, 9)$	Size $\in [9, 14]$
Avg. N Full-time eq. employees	9.03	7.41	12.16
Avg. N Employees	9.41	7.75	12.61
N Apprentices	0.41	0.36	0.49
N Transformations	0.09	0.08	0.12
N Separations	0.34	0.31	0.40
N Hires	0.31	0.28	0.36
Duration	20.13	19.89	20.55
N	1,422,023	937,184	484,839

Notes: The full sample includes all firms which employed between 5 and 14 employees in 2006 and were also active in years 2004 and 2005. Average outcomes are computed over the period 2004-2011. Firms size is the policy-relevant firm size measured in 2006. Details on firm size calculations are described in Appendix A.

Table 4: Distribution of Contract Duration

Contract Duration	Distribution of Newly Signed Contracts (1)	Baseline Contract Distribution (2)
0-12 months	50.7	48.7
13-24 months	22.5	19.4
25-36 months	15.7	16.1
37-48 months	7.8	9.7
49 months or more	3.3	6.5

Note: In column (1) we report the fraction of apprenticeship contracts generated because of the reform by duration (in months). This is obtained by summing regression coefficients β_k in Figure 8, for every k , separately by contract duration, and dividing by the overall sum of all β_k coefficients from the same graph. In column (2) we report the same fraction, which we compute as an average of the same outcome variables across all firms in our sample for the years 2004-2011.

Table 5: Public spending on the SSC discount

	<i>Apprenticeship contracts</i>		
	Avg. spending per firm (€)	N. of firms	Total spending (€ mln)
2007	1,036	36,111	37,411
2008	1,143	30,674	35,060
2009	1,105	23,542	26,014
2010	1,069	19,486	20,831
2011	1,185	16,856	19,974
Total		126,669	139,290

Notes: The first column reports the average spending per firm taking up the policy in a given year. The second column reports the total number of firms who reported taking up the policy in a given year. The third column is obtained by multiplying the second and third column and reports the total amount spent on the policy in each year. These descriptives are computed from our final sample of 193,297 firms.

Table 6: Cost per job and job-year generated

Apprenticeship duration	<i>Apprenticeship contracts</i>	
	Jobs	Job-years
Less than 1 year	28,404	28,404
Between 1 and 2 years	15,461	30,922
Between 2 and 3 years	11,538	34,615
Between 3 and 4 years	1,681	6,723
More than 4 years	-2,739	-13,695
Total	54,345	86,969

Notes: The second columns reports the number of new apprenticeship contracts associated to the reform, cumulatively over the period 2007-2011, by realized duration. The third column reports the number of job-years associated to such apprenticeships and is obtained - through an approximation - by multiplying the figures in the second column by the upper bound of the duration (e.g. $30,922 = 15,461 \times 2$).

Appendix

Appendix A The Policy-Relevant Firm Size

The 2007 Budget Bill does not define how to compute the policy-relevant firm size and delegates this task to INPS. INPS details how to compute firm size in a provision issued in January 2007 (*circolare n. 22, 2007*). We follow this definition closely.

The firm size that determines the eligibility for the SSC discount is full-time equivalent employment and it excludes apprentices, temporary agency workers, workers who are on leave (unless the firm hires a substitute), and workers who have been hired with an on-the-job training contract. The types of job training contracts that are excluded from the computation of firm size are those created under the following provisions: exD.lgs.251/2004, D.lgs.n.276/2003, law n.223/1991.

Our rich administrative data contains detailed information on workers' contracts and allows us to construct a fairly accurate proxy for the policy-relevant firm size. In this context there are two sources of potential measurement error. First, INPS data does not contain a flag for the on-the-job training contracts created under the exD.lgs.251/2004. Anecdotally, this contractual arrangement is very rare and it is unlikely to generate substantial measurement error. Second, our proxy does not account for workers who are on temporary leave (e.g., sick leave or maternity leave).

Appendix B Variables Definition

In this section, we define the variables we use in the empirical analysis and provide further details about the institutional background related to these variables.

Eligibility indicator: the definition of our eligibility indicator is time invariant and is based on the average policy-relevant firm size in 2006. The assignment rule therefore is:

$$T_i = \begin{cases} 0 & \text{if } \mathbf{1}(\text{size}_{i,2006} > 9) \\ 1 & \text{if } \mathbf{1}(\text{size}_{i,2006} \leq 9) \end{cases} \quad (15)$$

This definition gives us 65,439 control firms and 127,858 treated firms.

Hires, separation and transformations: We define the variable "hires" as the number of newly established apprenticeship contracts at firm i in year t . When the contract comes

to an end, the worker can either leave the firm, which we term “separation”, or she may see her contract converted into an open-ended contract at the same firm, which we call “transformation”.

More specifically an individual is considered hired as an apprentice in year t at a given firm A if she appears with an apprenticeship contract in such firm in year t but did not hold any apprenticeship contract at the same firm during year $t - 1$. Our definition includes workers who had already had a spell at the same firm in the past, just with a different contract type, and excludes individuals who perform consecutive apprenticeships at the same firm. We have checked that the vast majority of new hires concern individuals who were not working at the same firm in the previous year, so alternative definitions are unlikely to change our results. One drawback with our hiring measure is that we cannot observe if the firm rolls apprenticeship contract over, as we do not observe the exact level the worker is training for. If a firm uses two consecutive apprenticeship contracts with the same worker, we are bound to classify this as a unique contract. An apprentice is considered to be separated from a given training firm A in year t if she was holding an apprenticeship contract in such firm in year t but does not appear in firm A in year $t + 1$, with *any* contract. Transformations to open-ended contracts are instead easier to classify because we can observe a specific flag that identifies such an event. After having constructed the relevant variables we aggregate those up to the firm level in each year.

Apprenticeship duration: our measure of duration is based on the number of months between the hiring date of an apprentice, and the last month when she is observed with an apprenticeship contract at the firm.³⁴ Realized duration can be different from *ex-ante* contractual duration if either the apprentice decides to quit midway through the contract or the firm decides to lay her off. In order to aggregate durations up to the level of firms-year (it) observations, we construct the following measure:

$$D_{it} = \frac{1}{N_{it}} \sum_{v=1}^{N_{it}} D_{vit} \quad (16)$$

where N_{it} is the number of apprenticeship contracts *started* by firm i in year t and D_{vit} is the duration of contract v started by firm i in year t . Notice that this measure is defined only for years when the firm starts new apprenticeship contracts.

Wages: the measure of wages we can observe in our data is monthly earnings *net* of SSC

³⁴Similarly to before, such last date requires that the worker is not observed at the same firm with any contract during the whole of year $t + 1$.

(*imponibile previdenziale*) which in Italy constitutes the basis on which such contributions are computed. Such measure includes all compensation that the firm pays the worker, but excludes all compensation that the worker receives from INPS in case there is an event leading to a temporary reduction in working hours or absence (e.g. sickness leave or maternity leave). To the contrary, if there is a fraction of compensation that the employer needs to pay when these events occur, that will be included in our measure. All of this implies that this wage measure is not a contractual wage but is tightly linked to how many hours are effectively worked during the year.

Thanks to the monthly frequency of the data, and a flag that indicates whether the firm claimed the SSC discount in that month, we reconstruct a measure of earnings that is *gross* of social security contributions, and which we will call labor cost for brevity.³⁵ Contrary to the rest of the paper, our analysis on wages and labor costs only considers the years 2005-2011. This is because information on earnings at the monthly frequency is available only from 2005 onwards.

Appendix C Back of the Envelope Calculation

In this Appendix we describe in detail the back of the envelope calculation presented in Subsection 8.a. As discussed above, we compute the elasticity of the number of apprenticeship hires with respect to the change in social security contributions rate and use it to provide a back of the envelope estimate of the aggregate effects of the policy in our sample. Our elasticity formula is:

$$\epsilon_{h,t} = \frac{dh/h}{dt/t} \quad (17)$$

where dh/h is the *average* percentage change in cumulative hires between 2006 (our pre-reform baseline) and 2011 and dt/t is the percentage change in the SSC rate for the *average* apprenticeship contract. dh is obtained from estimated coefficients of equation 1, while for h we use the mean of the outcome in Table 3. We obtain $dh/h = 0.77$ ($0.24/0.31$). Correspondingly we construct the average increase in the SSC rate for the average apprenticeship contract at firms above 9 employees relative to firms with less than 9 employees and standardize it by the social security contribution rate in 2006. In this setting the average apprenticeship contract lasts 20 months and pays 1,000 euros per month. Hence, the average monthly social security contribution rate for a firm that pays 1.5% for the first 12 months and 3% for the re-

³⁵This measure excludes other forms of labor costs that are related to taxes.

maintaining 8 months is 2.1%. The implicit monthly social security contribution rate at baseline is 1.25% ($150/(12 \times 1,000)$). As a consequence we find $dt/t = 632\%$ ($((10\% - 2.1\%)/1.25\%)$).


The elasticity of the number of hires to the change in social security contributions is constructed as the ratio of the percentage change in hiring and the percentage change in social security contribution rates for the average apprenticeship contract. We find an elasticity $\epsilon_{h,t} = -0.12$ ($0.77/-6.32$).

Using the distribution of firms in our sample, we estimate that the average increase in social security contribution rates is approximately 3.5%. This represents an increase of 281.96% relative to the baseline social contribution rate ($281.96\% = 3.5\%/1.25\%$). Finally, we estimate that the average reduction in hires driven by the increase in social contributions is equal to 34.54% ($34.54\% = 281.96\% \times (-0.12)$) and use it to obtain the reduction in the number of apprenticeship hires in our sample which amounts to 25,854 ($25,854 = 74,852 \times 34.54\%$).

We follow the same procedure to construct the elasticity in the number of transformations to the change in social security contributions. In this case, the average increase in the number of transformations amounts to 0.56 and the elasticity is equal to -0.09 ($0.56/-6.32$). This leads to an expected average reduction in transformations of 24.79% and a reduction in the number of hires in our sample of 4,960 contracts.

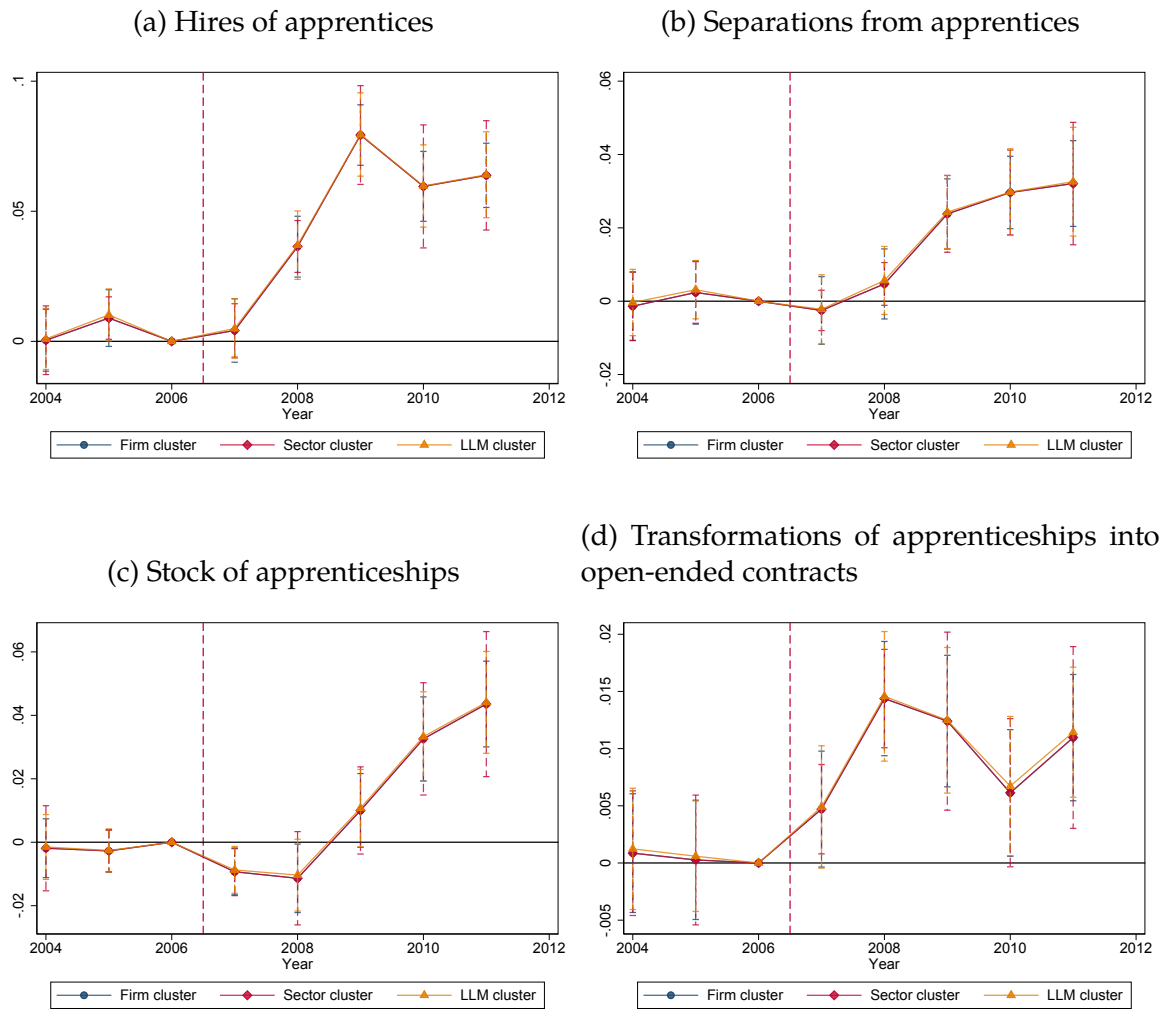
Appendix D Additional Figures

Figure D.1: Changes in the monthly wages of apprentices by eligibility status - new hires only



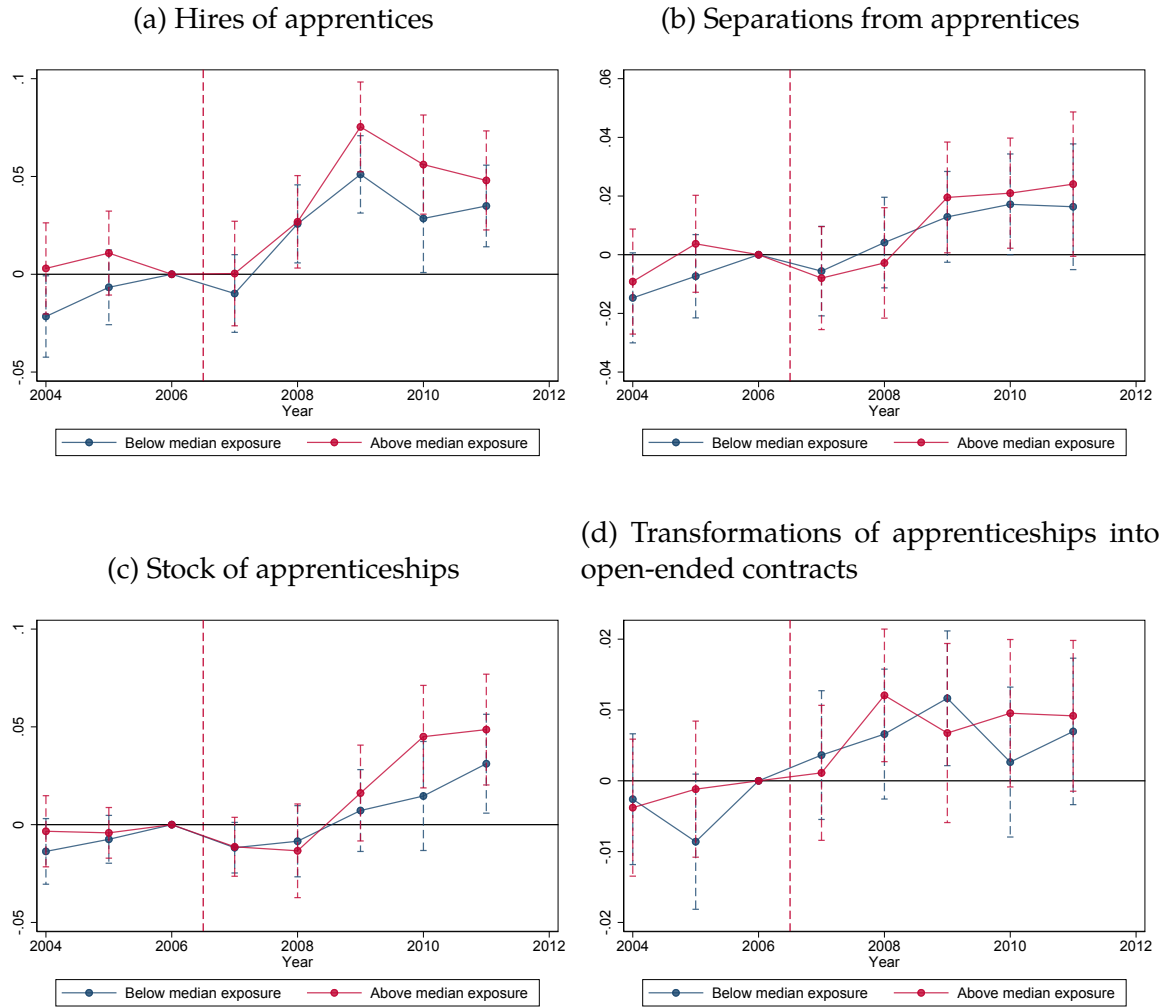
Note: This figure reports average monthly earnings, net of SSC for the new apprenticeship hires in each year t for the firms in our sample. Firms are divided according to whether their policy-relevant firm size was between 5 and 9 (blue circles) or between 10 and 14 (red circles) in 2006.

Figure D.2: Reduced-form impact of the reform: alternative clustering schemes for standard errors - OLS estimates



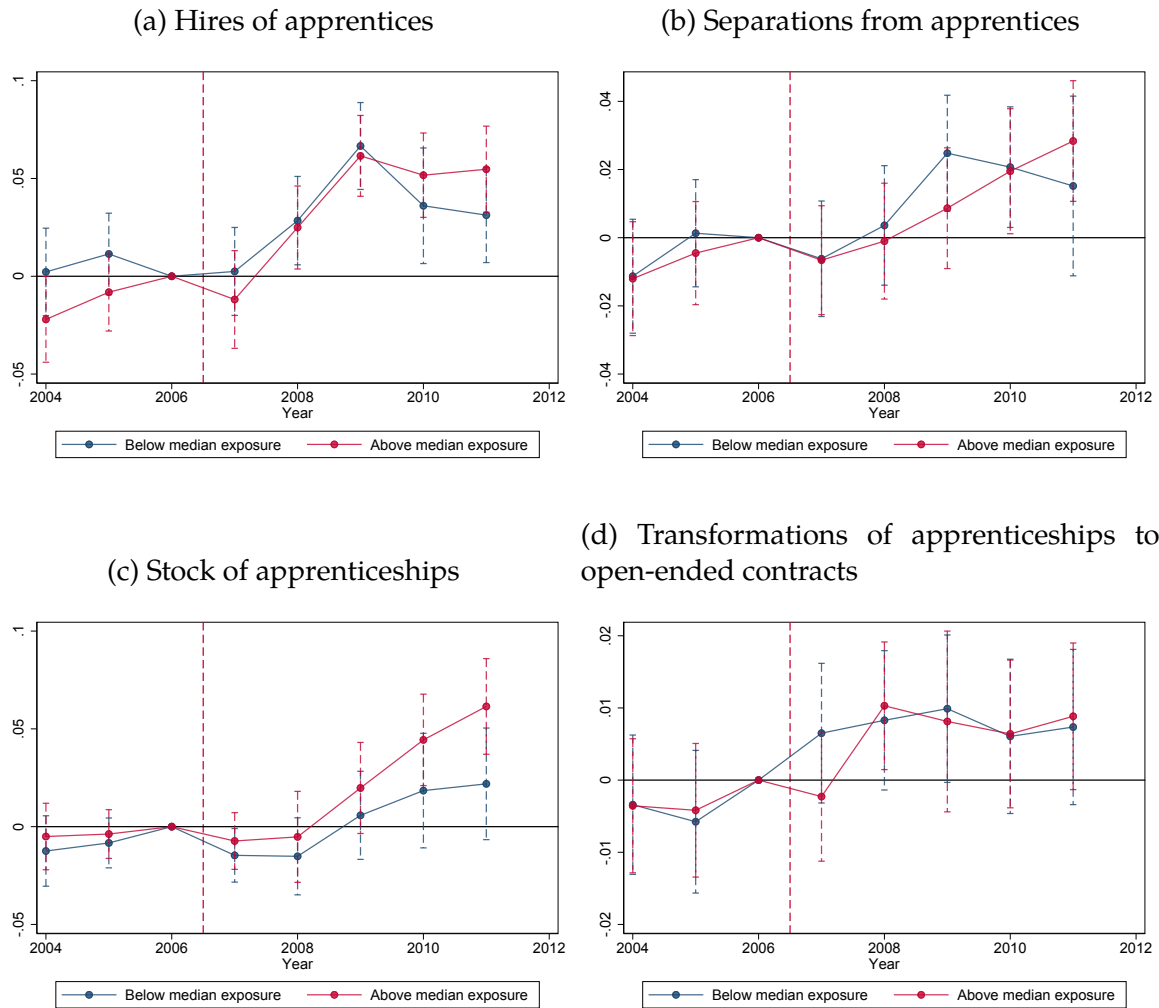
Note: These figures report the regression coefficients and the associated 95% confidence intervals that identify the reduced form impact of the reform, under different clustering schemes for the standard errors. Sector cluster are two-digit NACE sectors, LLM cluster are 686 local labor markets.

Figure D.3: Heterogeneous effects by exposure to financial constraints (liquid assets over total assets) - OLS estimates



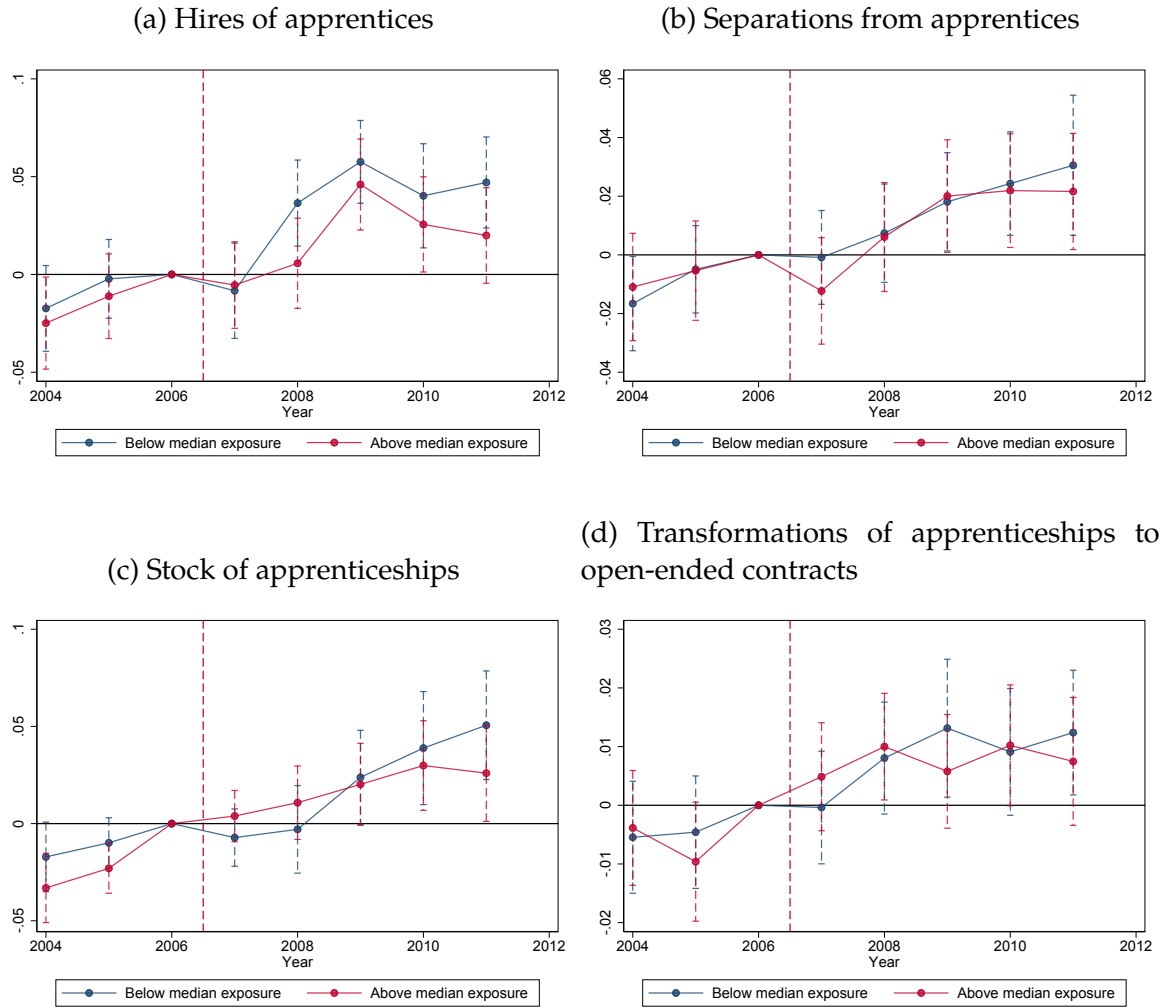
Note: These figures report the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, depending on the level of exposure to financial constraints. The indicator is the ratio of liquid assets over total assets and is defined for 98,084 firms. Standard errors are clustered at the firm level.

Figure D.4: Heterogeneous effects by exposure to financial constraints (cash flow over assets)
- OLS estimates



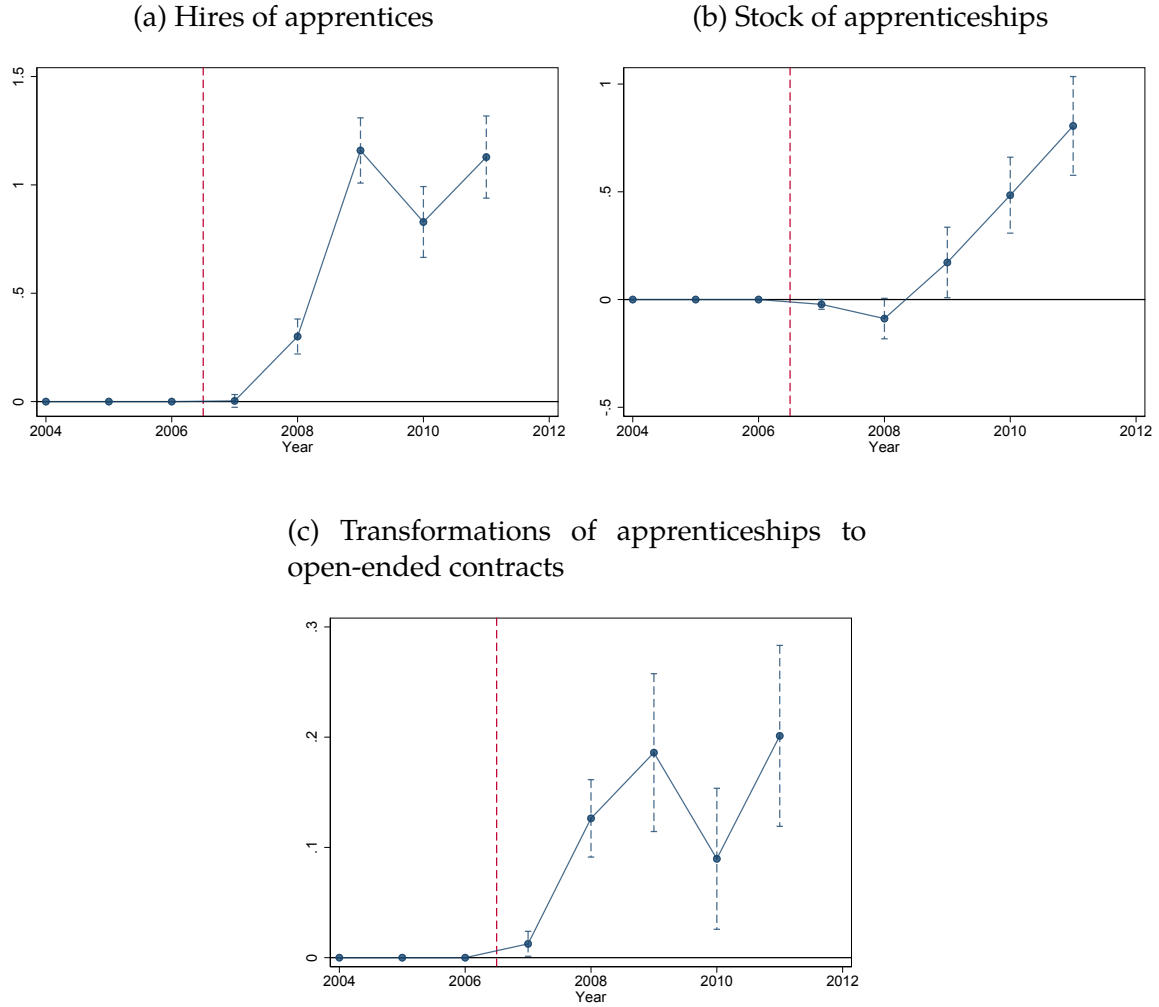
Note: These figures report the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, depending on the level of exposure to financial constraints. The indicator is the ratio of cash flow over assets and is defined for 98,084 firms. Standard errors are clustered at the firm level.

Figure D.5: Heterogeneous effects by exposure to financial constraints (total revenues) - OLS estimates



Note: These figures report the regression coefficients and the associated 95% confidence intervals that identify the reduced-form impacts of the reform, depending on the level of exposure to financial constraints. The indicator is total revenues and is defined for 98,084 firms. Standard errors are clustered at the firm level.

Figure D.6: Jobs created per €1,000 spent: alternative definition of M_{it}



Note: The figure reports the θ_k^j and ρ_k regression coefficients together with the associated 95% confidence interval that identify the impact of a €1000 cost reduction for apprenticeship contracts on new hires of apprentices, the stock of apprenticeships and transformations of apprenticeships to open-ended contracts among complier firms (equation 8 and equation 9). The definition of M_{it} is the forward-looking sum of all transfers that firm i receives for apprenticeship contracts started in year t . All coefficients for years 2004, 2005 and 2006 are normalized to zero. Standard errors are clustered at the firm level. The x-axis indexes time.