INPS

Istituto Nazionale Previdenza Sociale

luglio 2020 – numero 21

WorkINPS *Papers*

What are the returns to apprenticeships?

Evidence from Italy

Luca Citino

Aggiornamento WorkINPS papers n.21 – Gennaio 2020

ISSN 2532 -8565

Lo scopo della serie WorkINPS papers è quello di promuovere la circolazione di documenti di lavoro prodotti da INPS o presentati da esperti indipendenti nel corso di seminari INPS, con l'obiettivo di stimolare commenti e suggerimenti.
Le opinioni espresse negli articoli sono quelle degli autori e non coinvolgono la responsabilità di INPS.

The purpose of the WorkINPS papers series is to promote the circulation of working papers prepared within INPS or presented in INPS seminars by outside experts with the aim of stimulating comments and suggestions.
The views expressed in the articles are those of the authors and do not involve the responsibility of INPS.

Responsabile Scientifico
Maurizio Franzini

Comitato Scientifico
Agar Brugiavini, Daniele Checchi, , Maurizio Franzini

In copertina: uno storico "Punto cliente" a Tuscania
INPS, Direzione generale, Archivio storico

I WORKINPS PAPER

Le basi dati amministrative dell'***INPS*** rappresentano una fonte statistica unica per studiare scientificamente temi cruciali per l'economia italiana, la società e la politica economica : non solo il mercato del lavoro e i sistemi di protezione sociale, ma anche i nodi strutturali che impediscono all'Italia di crescere in modo adeguato. All'interno dell'Istituto, questi temi vengono studiati sia dai funzionari impiegati in attività di ricerca, sia dai ***VisitInps Scholars***, ricercatori italiani e stranieri selezionati in base al loro curriculum vitae e al progetto di ricerca presentato.

I **WORKINPS** hanno lo scopo di diffondere i risultati delle ricerche svolte all'interno dell'Istituto a un più ampio numero possibile di ricercatori, studenti e policy markers.

Questi saggi di ricerca rappresentano un prodotto di avanzamento intermedio rispetto alla pubblicazione scientifica finale, un processo che nelle scienze sociali può chiedere anche diversi anni. Il processo di pubblicazione scientifica finale sarà gestito dai singoli autori.

Maurizio Franzini

What are the returns to apprenticeships?

Evidence from Italy

Aggiornamento WorkINPS papers n. 21* - gennaio 2020

Luca Citino

(†Bank of Italy and London School of Economics)

* la cui pubblicazione è stata autorizzata dal "Comitato Scientifico WorkINPS Papers" in data 1 luglio 2020 .

INPS.0072.25/06/2020.0000637 E

What are the returns to apprenticeships? Evidence from Italy*

Luca Citino†

January 2020

Abstract

What are the returns to apprenticeships? This paper tries to answer this question by leveraging novel administrative data from Italy on individual careers. We adopt a difference-in-difference methodology to compare the labor market outcomes of individuals starting an apprenticeship with those of similar individuals starting temporary contracts that, at least formally, do not provide training. We find apprenticeships to be a "double-edged sword". While they do guarantee a stronger labor market attachment during the first three years after the start of the contract, they produce ambiguous effects afterwards. Apprenticeships increase the probability of conversion to open-ended contracts, especially at the initial firm, but decrease the probability of obtaining further temporary jobs, especially at other firms. Quantitatively, this second effect prevails, generating a negative effect of the probability of having any job. These findings are consistent with a model where retention rates after the end of an apprenticeship convey stronger signals about workers' ability compared to retention after the end of a temporary contract.

JEL classification: J24, J62.
Keywords: Apprenticeships; Temporary contracts; Open-ended contracts; Careers; Firm size;

*I thank my advisors Jörn-Steffen Pischke and Stephen Machin for continuous guidance and support. The realization of this project was possible thanks to the sponsorships and donations in favour of the "VisitInps Scholars" program. I am very grateful to Massimo Antichi, Tito Boeri, Maria Cozzolino, Edoardo Di Porto, Pietro Garibaldi, Paolo Naticchioni and all the staff of Direzione Centrale Studi e Ricerche at INPS for making this programme possible and for their help with the data. This project has also benefited from discussions with Andrea Alati, Fabio Bertolotti, Giulia Bovini, Anna D'Ambrosio, Alessandra Fenizia, Simone Ferro, Nicola Fontana, Xavier Jaravel, Alessia Matano, Matteo Paradisi, Alessandra Pasquini, Kilian Russ, Viola Salvestrini, Vincenzo Scrutinio, Anna Vitali, Martina Zanella and seminar participants at INPS and at the LSE. The findings and conclusions expressed are solely those of the author and do not represent the views of INPS nor the Bank of Italy. All remaining errors are my own.

†Bank of Italy and London School of Economics; Contact: luca.citino@bancaditalia.it

1 Introduction

Apprenticeships are diffused in many European countries and constitute a middle-ground between high school and university education. Although there are differences across countries, they usually consist of job contracts where labour services are exchanged for certified training in an occupation and a salary (Snell, 1996; Ryan, 2012). In recent years apprenticeships acquired a prominent place in the policy discourse about youth unemployment and the NEET problem, with many governments offering reduced social security contributions or favorable taxation regimes to incentivize their use (Kuczera, 2017). Although in policy circles apprenticeships are often seen as a *panacea*, providing young people with good jobs and valuable skills, the economic reality may not be that simple. While it is true that apprentices ought to receive training by virtue of a contractual obligation, it is not a given that on-the-job training provided through apprenticeships has any real content. Firms may have scarce incentives to train if the human capital they need for production is general (Becker, 1962) and even more so if the labor market where they operate does not feature any frictions (Acemoglu and Pischke, 1999). In such cases, given the low enforcement level of apprenticeship contracts, firms may renege on the promise to provide training and the returns to apprenticeships would be close to zero (Dustmann and Schönberg, 2012). Conversely, firms will be more incentivized to provide training to young workers if the human capital they need is firm-specific, or if labor market frictions are substantial.

In this paper we empirically quantify the returns to apprenticeships by leveraging novel administrative matched employer employee data from the Italian Social Security Institute (INPS). We have access to the full working history for the universe of individuals born in Italy in 1980 and 1981, regardless of whether they have been employees in the private sector, dependent self-employed (*parasubordinati*) or self-employed. We define returns to apprenticeships as the extra gain coming from starting an apprenticeship compared to a temporary contract. Similarly to the former, temporary contracts also involve an employer-employee relationship but, at least formally, they don't require the firm to provide training. We perform this comparison in a difference-in-difference framework at the job spell level.

The comparison of apprenticeships with other temporary contracts is not completely new in the literature (Berton et al., 2011; Picchio and Staffolani, 2013) and is particularly relevant in the Italian setting. On the one hand the vast majority of apprenticeships happen when individuals have already left technical and vocational schools, and are not formally

linked to the education system.[1] Also, apprentices' training can take place entirely within the firm premises and trainees do not need to sit a formal examination at the end of the contract.[2] These characteristics make such contracts more similar to temporary training contracts than to a course of study. On the other hand the question of whether apprenticeships are really any different from other types of temporary contracts is recurrent in the Italian debate. Some scholars in other disciplines go as far as saying that "Although a number of legal provisions establish compulsory training during apprenticeship, reality is often very distant from the ideal apprenticeship model, and this tool becomes a mere instrument of exploitation of a flexible and cheaper labour force" (Tiraboschi, 2012). For these reasons we think that our focus on temporary contracts is indeed justified to evaluate the returns to apprenticeships in our setting. Given that temporary contracts are known to receive little or no firm-sponsored training (Booth et al., 2002; Albert et al., 2005), they are suited to gauge the magnitude of the returns to training at the *extensive margin*.

To preview our results, we find that apprenticeships are a "double-edged sword". They lead to higher conversion rates towards open-ended contracts, but have a negative effect on the probability of transitioning to other temporary contracts. Quantitatively, the second effect is stronger and produces a negative average treatment effect on the probability of having a job of any kind. We find that most conversions happen at the training firm, while the lack of job opportunities in other temporary contracts is explained by what happens in other firms. On the one hand, this indicates that training provided through apprenticeships is valuable and that training firms are able to appropriate some rents from it.[3] On the other hand it seems that "recalls" do not explain why apprentices spend less time churning between other temporary contracts.

Our findings can be rationalized in an asymmetric information model with adverse selection, where the absence of conversion to a permanent position for an apprentice conveys a stronger signal about ability than for a temporary contract. This can be the case if temporary contracts can fail to be converted because of reasons that are exogenous to the worker's ability with higher probability than apprentices (e.g. the task is temporary in nature ...). In this sense apprenticeships constitute a riskier investment compared to

[1] During the years 2007-2013 INPS data provide information on the type of apprenticeship contract. The share of apprenticeships linked to upper-secondary education was 19.8% in 2007 and steadily declined throughout the time window, reaching a low of 4.6% in 2013.

[2] *Cassazione, sent. 845/1988.*

[3] We leave the question open as to whether such training is firm-specific or if rent extraction is allowed by labour market imperfections (Acemoglu and Pischke, 1999)

a temporary contract, and its convenience may depend on a worker's ability level and preferences. Alternatively, apprentices may be acquiring firm-specific human capital that is not necessarily useful outside the training firm, and leads to a penalty in terms of future job offers from other firms.

When looking at earnings, we find that, conditional on working, apprenticeships pay off in the first three years after the start of the contract. However we fail to detect any long-run effects. Earnings effects are not significantly different from zero six years after having started the contract.

Our paper contributes to the literature on the returns to apprenticeships. Various other studies have measured the extent to which apprenticeships constitute a valid opportunity for the young, when compared to different alternative opportunities. The general consensus so far reached is that apprentices are better off in terms of wages if compared to low-skilled workers with no apprenticeship training, but not if compared to individuals completing full-time vocational education in the classroom (for a review see Samek Lodovici et al. (2013)). Also, while apprenticeships facilitate the school-to-work transition and pay off at an early stage of the working life, their effects may be more muted in the longer run (Samek Lodovici et al., 2013; Hanushek et al., 2017; Parey, 2016).

More in detail, Parey (2016) compares firm-sponsored training with school-based vocational education. He finds that the two tracks do not offer different returns, but that in the very short run firm-based apprenticeships provide stronger labor market attachment. He also finds no effects on wages. Similarly Albanese et al. (2017) compares two apprenticeship tracks that co-existed in Italy in the early 2000s, one of which emphasized firm-sponsored training rather than school-based vocational education. In line with Parey (2016), they find that firm-sponsored training improved the prospects of young workers, increasing their probability of transitioning to open-ended contracts but it also raised their wage levels, especially in bigger firms. Cavaglia et al. (2018) also find positive effects in the UK context. They find that apprentices yield substantial earning premia, especially for men. Fersterer et al. (2008) compare longer and shorter apprenticeships. For identification they exploit the unexpected closure of firms that employ apprentices at different tenure horizons. At such *intensive margin*, they find that an extra year into apprenticeship yields a 3.8% return in terms of higher earnings.

Due to a similar choice of a control group, the studies closest in spirit to ours are Picchio and Staffolani (2013) and Berton et al. (2011). The first paper exploits age limits in the Italian apprenticeship system and use a regression discontinuity design to compare

individuals who manage to get an apprenticeship just before age 30 and those who do not manage to do so. The authors find that, around age 30, individuals who start an apprenticeship are more likely to transition towards open-ended contracts, especially at the initial firm. The second paper uses a Multinomial Logit with individual fixed effects to study the transition matrices between different types of temporary contracts (including apprenticeships) and open-ended contracts. We extend these analyses in different ways: first we characterize the full time profile of returns to apprenticeships at the quarterly frequency and are able to look into the long run, up to six years after the start of the contract. Second, thanks to the matched employer-employee nature of the data we can look at how much of the conversion rate to open-ended contracts can be explained by the training firm or the other firms. Third we look at heterogeneous effects depending on firm size and are able to look at new outcomes that were unstudied before due to data limitations, such as the probability of entering self-employment.

The paper is structured as follows. In Section 2 we describe how apprenticeships are regulated in Italy and the data we employ for our analysis. In Section 3 we present our identification strategy and regression framework. In Section 4 and we present our main findings. In Section 5 we present some heterogeneity analysis along the firm size dimension. In Section 6 we present other results on the self-employment margin and on earnings. In Section 7 we discuss our results and in Section 8 we conclude.

2 Institutional framework and Data

2.1 Apprenticeships in Italy

The Italian apprenticeship system is made of three separate programmes, with different rules: (1) "right and duty" (*Apprendistato per l'espletamento del diritto/dovere di istruzione*), performed during upper secondary education for individuals aged 15-18 (2) "occupational" (*Apprendistato professionalizzante*), usually performed after the completion of secondary education, for individuals aged 18-29 and (3) "higher" (*Apprendistato di alta formazione e ricerca*), still oriented to individuals between 18 and 29, but who are enrolled in or have already earned a university degree and would like to carry out a thesis or a research project within a firm. In our analysis we require individuals to be at least 22 when doing their apprenticeships, so this excludes type (1) apprenticeships

by construction.[4] On the other hand in the data we do not have information needed to distinguish apprenticeships of type (2) from those of type (3) before 2007 or after 2013, so in what follows type (2) and type (3) are pooled together. Again, we stress the fact that the vast majority of apprenticeships in Italy are of type (2).

In terms of contractual obligations, apprenticeships are job contracts, limited to the private sector, in which worker and firm regularly pay social security contributions and work accidents insurance. The formal training content of apprenticeships is quite low. The minimum number of training hours that the firm must provide is 120 per year, split in the following way: 65% are dedicated to occupation-specific training and 35% are dedicated to general training (job safety, psychology of labor and team working). In exchange for training, firms obtain a reduction in social security contributions. The latter amounts to 10% of apprentices' gross earnings, compared to 27% for open-ended and temporary contracts. Also, firms can pay apprentices a lower wage, up to two levels below what a qualified worker would get, according to the corresponding collective bargaining agreement (CBA). At the end of the programme the workers receive a certification which is recognized by firms applying the same CBA. This implies a worker cannot be trained twice for the same occupation in the same CBA. Eligibility on the side of firms is linked to the presence of a *mentor*. The mentor must attend preparatory training and cannot train more than 5 individuals at each point in time. The law sets ceilings in apprenticeship use: they can never be more than the number of qualified workers in the firm (however if firm size is less than 3 the firm can hire up to 3 apprentices). Eligibility on the side of workers is exclusively age-dependent. Recent reforms raised the age limits (measured on the day of hiring).[5] A more complete description of the Italian apprenticeship contract and its recent reforms can be found in Albanese et al. (2017).

2.2 Data sources

We use administrative data on careers at the individual level made available by the Italian Social Security Institute (INPS) through the VisitINPS initiative. Below we present each source in detail:

[4] Our analysis excludes individuals younger than 22 at start of the contract in order to have sufficient information on the pre-event working history. This allows us to test whether individuals displayed parallel trends in the outcome variable before the onset of the contract.

[5] The 1997 (*Treu*) reform: from age 20 to age 24 (but 27 in regions entitled to EU structural funds - i.e. the South - and age 29 in artisan firms). The 2003 (*Biagi*) reform: from age 24 to age 29 in all firms in all regions.

Matched employer-employee data: our primary source is a matched employer-employee dataset covering all job spells in non-agricultural firms with at least one employee. The dataset spans the whole time period 1983-2017. The public sector and firms with no employees are not included. The data records the presence of job spells at the monthly frequency, which gives us the advantage tracing career dynamics at a very fine level. In each month we observe at which firm(s) the worker is employed, the type of contract(s) the worker has (open-ended, temporary), the type of work-time arrangement (full time or part time) and a coarse occupation code (apprentice, blue collar, white collar, supervisor or manager). Absent any change in the aforementioned characteristics, we observe one earning record per year for each worker. In case a worker has a contractual change during the year (e.g. becomes a white collar worker or changes firm) we see two separate earning records. This allows us to precisely separate earning records which belong to different contract characteristics, different firms and different years. For each individual we also observe a series of basic socio-demographic characteristics such as gender, year of birth and place of birth. Given the nature of the dataset, we are also able to build the total firm size in every year, and therefore check whether individuals starting apprenticeships in bigger firms obtain higher returns.

Dependent self-employment spells: starting from 1996, we also have information on dependent self-employment. The latter is a form of work where workers are formally self-employed but *de facto* employees (Williams and Lapeyre, 2017). This dataset also has a matched employer-employee structure. For each job spell we observe unique worker and firm identifiers, the beginning and end date of the spell, the type of contract and the overall compensation received for the job in every year. Given that firm and worker identifiers are the same across datasets we are able to merge this information with the matched employer-employee dataset.

Contribution Histories: for a subset of individuals in the matched employer-employee dataset we were able to obtain further information on their full contribution history, including spells as self-employed. This allows us to build more precise measures of labour market outcomes and investigate whether apprenticeships have an impact on the probability of entering self-employment. This dataset does not contain a firm identifier. We obtained such information for the universe of individuals born in Italy between 1980 and 1981, that is our main sample of interest.

2.3 Sample selection and variable construction

Our initial sample is made of all individuals born in Italy in 1980 and 1981. We focus on these two cohorts because information on whether an individual works in an open-ended or temporary contract is only available from 1998 onwards (approximately when our individuals leave upper-secondary education). On the other hand we don't choose cohorts younger than 1981 to have a long enough period to observe the evolution of the outcome variables. We restrict the sample only to those individuals who ever started an apprenticeship or a temporary contract between age 22 and age 29. We do not consider contracts starting before age 22 in order to have enough information on past working history, which is useful to check for the presence of underlying pre-trends.

In what follows, we refer to the start of either a temporary contract or an apprenticeship as an *event*. Apprenticeships are treatment events, while temporary contracts are control events. In our empirical strategy, we will look at the differential evolution of outcomes of interest around the event date, between these two types of events.

Among events we only consider first-time temporary contracts and first-time apprenticeships. Further apprenticeships or temporary contracts are not considered, although they contribute to the construction of the outcome variables. An individual may appear more than once (twice at most) in the sample if she starts a temporary contract and then starts an apprenticeship at a later age. In this case we include both events in the regression and study them separately.[6] To the contrary, if an individual starts an apprenticeship and then starts a temporary contract at a later age, only the apprenticeship is included as an event - the temporary contract is used for the construction of outcome variables. If apprenticeships indeed have dynamic effects, then including the latter type of temporary contracts in the set of events risks contaminating the control group and invalidating our design. For similar reasons we drop all individuals who do an apprenticeship and a temporary contract at exactly the same age. Our final sample consists of 285,422 events, either apprenticeships contracts (103,878) or a temporary contracts (181,544).

Although our data would allow us to construct employment outcomes at the monthly frequency, we collapse our dataset at the quarterly level for computational convenience. All employment outcomes are coded as dummy variables, taking value one when the condition is true for at least one month during the quarter. Due to workers changing jobs or holding multiple jobs within a quarter, employment outcomes are never mutually

[6] In order to treat this case we always include individual×event fixed effects, but cluster standard errors at the individual level.

exclusive.

2.4 Summary statistics

A description of our final sample can be found in Table 1. Apprentices and temporary contracts are not very dissimilar during the quarters leading to the start of the contract. While apprentices have slightly more work experience, they do not seem to have had higher probabilities to hold open-ended contracts before. Their previous wage levels (conditional on working) are also remarkably similar, indicating that apprentices are not particularly selected compared to workers obtaining temporary contracts. It is nonetheless true that apprentices start their contract approximately one year before. In our main specification we control for age fixed effects to account for these differences, although this makes little difference in the estimated coefficients.

3 Estimating returns to apprenticeships

In this paper we define returns to apprenticeship as the extra gain in labor market outcomes an individual obtains from starting an apprenticeship relative to another type of temporary contract that does not oblige the firm to provide training. We employ a dynamic difference-in-differences (DiD) strategy to compare the differential evolution of several labour market outcomes across individuals who start either type of contract. Our identification strategy is valid under a standard parallel-trend assumption i.e. individuals in apprenticeships would have followed the same *trend* as individuals in temporary contracts, had they started one. To corroborate the validity of this assumption, we check whether individuals starting apprenticeships were on different *trends* compared to individuals starting temporary contracts, in the quarters leading up to the start of the job. We find no evidence of underlying pre-trends, which reassures about the validity of our design.

Our unit of analysis is an individual i, whom we follow in the quarters k leading up to, and after an event j. Since the same individual may be present more than once in our data, we cluster standard errors at the individual level, but analyze each event separately and therefore include event-specific fixed effects. We run regressions of the form:

$$
\begin{aligned}
Y_{ijt} \;=\; & \alpha_j \;+\; \eta_t \;+\; \theta_a \;+\; \sum_{k=-4}^{23} \beta_k \;\times\; \mathbf{1}(\text{distance}_j = k) \\
& + \sum_{k=-4}^{23} \beta_k^T \;\times\; \mathbf{1}(\text{distance}_j = k) \;\times\; \text{Apprentice}_j \;+\; \epsilon_{ijt}.
\end{aligned}
\tag{1}
$$

where Y_{ijt} is a labor market outcome for individual i, around event j, measured in calendar year $\times$ quarter t; α_j are event fixed effects, which control for any time-invariant unobserved heterogeneity at the worker level when starting either her first apprenticeship or first temporary contract, and η_t are year $\times$ quarter fixed effects, which control for time-varying unobservables that are common across the two groups. We also include age fixed effects (θ_a), in quarters, to control for life-cycle patterns that are common across the two groups. Given that both our treatment and control group are assigned to a job contract at distance time $k = 0$, we include both a set of distance-to-event dummies that are common to both groups i.e. $\mathbf{1}(\text{distance}_j = k)$, and a set of distance-to-event dummies interacted with treatment i.e. $\mathbf{1}(\text{distance}_j = k) \times \text{Apprentice}_j$. This specification is very similar to Jaravel et al. (2018) and addresses the presence of dynamic effects around the start of the contract for both treatment and control group. The resulting coefficients may be interpreted as a tenure profile that is specific to each group.[7]

The coefficients of interest are the β_k^T, for $k \neq -1$. Due to multicollinearity issues we omit both $\mathbf{1}(\text{distance}_j = -1) \times \text{Apprentice}_j$ and $\mathbf{1}(\text{distance}_j = -1)$. All coefficients β_k^T must thus be interpreted as changes in the difference across the two groups relative to any pre-existing difference at distance $k = -1$ (one quarter before event). It follows that $\beta_k^T = \beta_k = 0 \, \forall k < 0$ implies the absence of differential trends in outcome variables before the start of the treatment.

[7] We are not including any other control that is time invariant such as firm characteristics in quarter $k = 0$, as these would be absorbed by the event fixed effects. On the other hand we do not condition on the covariates which vary after the start of the contract because these would constitute a bad control.

4 Main results

4.1 Graphical evidence on the returns to apprenticeships

As a first step in describing the kind of variation we exploit in the data, we turn our attention to Figure 1. The hollow markers represent the share of individuals who have an open-ended contract, as a function of event time k, for individuals who will start either an apprenticeship (circles) or a temporary contract (diamonds) in event time $k = 0$. The outcome can thus be interpreted as the probability of having an open-ended contract. The two curves evolve parallel in the quarters before the start of the contract, suggesting that our research design is valid. The solid circles instead are corresponding difference-in-differences estimates ($\hat{\beta}_k^T$) from specification 1. Associated 95% confidence intervals are also displayed. The graph displays an increase in labor market prospects following the start of either type of contract, as reflected in the higher probability of obtaining an open-ended contract in the quarters after $k = 0$. However, the dynamic evolution of the two paths clearly differs. Compared to temporary contracts, apprenticeships yield a negative short term effect, most likely due to the fact that individuals are locked-in their initial training contract (an "incapacitation effect"), but recover afterwards. The recovery from the negative effects follows a step function with more pronounced jumps at quarters 8, 12 and 16 after the start of the contract. This is reasonable because apprenticeships that are brought to completion have (in the majority of cases) fixed durations that are multiples of one year. We still see departures from the step function because apprenticeships may terminate before due to either of the two parties' willingness to stop.[8] After quarter 16 we see that apprenticeships have 8.5 p.p. higher probability of being converted to open-ended, an effect that remains stable up to six years after the start of the contract.

Given this framework, we now turn to the study of different outcomes. Together with the probability of being converted to open-ended contracts, in Figure 2 we overlay estimates for two other outcomes: the probability of having a temporary contract, and the probability of having either of the two, that is the probability of having any job that is not an apprenticeship.[9] When looking at the two other outcomes we see that starting

[8] By the law, apprenticeships have the same EPL coverage as open-ended contracts. They can only be dismissed under a "just cause" or "justified motive", because of economic or disciplinary reasons respectively. Temporary contracts can only be terminated under a "just cause". However firms can roll the latter over, generating more moments at which firms can terminate the working relationship.

[9] Individuals who have more than one job at the same time or transition from a job type to another within the same quarter will be recorded in the data as having *both* an open-ended and a temporary contract in the same quarter. For this reason the coefficient associated to "employee but not apprentice"

an apprenticeship instead of a temporary contract mechanically causes a sharp drop in both the probability of holding a temporary contract or having any job that is not an apprenticeship. Over time this effect is gradually reduced for both outcomes, as workers start new spells and transition towards different contractual forms. We see that by the end of the period, apprenticeships induce a decrease in the probability of having temporary contracts of around 13.1 p.p.. Quantitatively this effect is stronger than the positive effect on open-ended contracts first analyzed in Figure 1, which is reflected in coefficient associated with the probability of being in any job contract that is not an apprenticeship.

In sum, apprenticeships are indeed associated with higher probability of having an open-ended contract on average, but this comes at the expense of a much lower probability of having a temporary contract, with the second effect dominating. The combination of these forces implies a negative treatment effect of around 4 p.p on the probability of having any job that is not an apprenticeship after a six year period.

4.2 Decomposition according to firm mobility patterns

In the previous subsection we highlighted that apprenticeships confer to workers a higher probability of obtaining open-ended contracts and lower probabilities to have temporary ones. In this subsection we investigate where these gains or losses are accrued. It could be that apprenticeships lead to higher conversion rates to open-ended jobs at the training firms but lower probability of obtaining an open-ended contract elsewhere. Similarly, the lower probability of churning among other temporary jobs may be due to the fact that temporary contracts give workers the possibility to be periodically recalled by the same firm, a fact documented in Scrutinio (2019). In what follows we decompose both the probability of having an open-ended contract and the probability of having a temporary contract in spells at the initial firm and at other firms. Similarly to before, Figures 3 and 4 plot β_k^T coefficients and associated 95% confidence intervals.

Let us consider Figure 3 first. We see that apprenticeships have a positive impact on the probability of being employed under an open-ended contract at the initial firm but a negative effect on the same outcome in other firms. Although the overall effect is positive, the entirety of gains in terms of conversion to open-ended contracts are accrued at the training firm while the probability of obtaining open-ended contracts at other firms contributes negatively to the overall effect. This is consistent both with the accumulation of

is not necessarily equal to the sum of coefficients associated to "open-ended contracts" and "temporary contracts".

firm-specific human capital and a high degree of wage compression which limits poaching by competing firms in the post-training period. An apprenticeship increases on average the probability of conversion at the initial firm by 10.6 p.p.

Figure 4 has a similar structure and displays DiD estimates for the probability of having a temporary job (solid dots) and a decomposition thereof in the the probability of having it at the initial firm or in other firms. We see that apprenticeships do not miss out on the opportunity of obtaining other temporary contracts at the initial firm. However we see that the majority of the effect is explained by what happens in other firms. This goes against an explanation based on higher recall rates for temporary contracts. Rather, it seems that individuals in temporary contracts become more able to move across different firms with the same contractual form.

5 Heterogeneous effects

5.1 Effects by firm size

In this subsection we look at whether main results are different depending on the size of the firm where the individual starts the contract.[10] In order to do this we carry out the same analysis as before, separately for big and small firms. We classify a firm as being "big" if its average size is strictly greater than 15 in the solar year when the contract starts, and "small" otherwise.

To summarize results, we report β_{23}^{T} coefficients in bar charts and present the corresponding event study graphs in the Appendix. In Figure 5 we look at the probability of being employed under an open-ended contract, a temporary contract or either of the two 23 quarters (including 0) after the start of the contract. We see that the overall probability of having an open-ended contract is not different across the two groups. What differs is the probability of being employed in other temporary contracts. Big firms give a substantial disadvantage in this respect. As a consequence, the overall probability of having a job is negative only in big firms, but not in small firms.

When decomposing the rate of conversion to open-ended contracts in Figure 6 we notice two interesting facts. First, big firms convert apprenticeships to permanent positions at a much higher rate than small firms. The effect in small firms is 9 p.p. while the one

[10] We performed an heterogeneity analysis also based on gender and found identical results for men and women. Results are available upon request.

in big firms is 15 p.p, a 66% increase. Secondly, small firms produce higher rates of conversions to open-ended contracts in other firms. The same is not true for big firms, as they have a negative impact on the probability of obtaining permanent position in other firms. When looking at the overall effect, these two mechanism compensate each other: apprenticeships in both types of firms are associated an increase in the probability of having an open-ended contract by 12 p.p..

We perform a very similar exercise for the probability of being employed under temporary contracts. Results are displayed in Figure 7. We see that the qualitative pattern this time is very similar in both small and big firms. Apprenticeships unambiguously decrease the probability of churning in other temporary contracts, especially in firms other than the initial one. In small firms, the lack of other temporary contracts outside the initial firms accounts for about 90% of the overall impact, while the same figure is 94% for big firms.

6 Other results

6.1 Self-employment effects

Self-employment is very diffused in Italy and constitutes around 20% of the workforce, way above the European average (Istat, 2017).[11] It is therefore interesting to check whether apprenticeships contribute positively or negatively towards the individual choice of entering self-employment. From an economic standpoint, the direction of the effect is ambiguous. On the one hand apprenticeships increase the conversion rates at the initial firm, as firms train workers to keep them and extract rents from their accumulated human capital. On the other hand apprenticeships may want to learn a trade to establish their own entrepreneurial activity.

In Figure 8 we study three outcomes: the probability of working, the probability of being an employee and the probability of being self-employed. As described in previous sections, apprenticeships have a negative impact on the probability of being employees. Here we found that this is not compensated by the self-employment margin. To the contrary, apprenticeships have a negative impact on the probability of being self-employed. Despite being statistically significant, this effect is quantitatively small, in the order of magnitude of 1 p.p..

[11] Our definition of self-employed includes both freelancers (*libero professionista*), entrepreneurs (*titolare d'impresa*) and their collaborators (*coadiutore d'impresa*).

6.2 Wage effects

In Figure 9 we study the impact of apprenticeships on wages. Our dependent variable is now the log of quarterly earnings, conditional on working status. Our data does not record earnings at the quarterly frequency, but we still have information on the total amount of earnings received in a given year, separately by job characteristics and employer, in addition to detailed information on which exact months of the year these income flows refer to. In order to construct our measure of quarterly earnings we therefore apportion job-spell earnings to quarters based on the proportion of months accounted for by any given spell.[12]

Given that we established that apprenticeships have an impact on the overall probability of employment, our wage results ought to be interpreted with care. Conditional on having a job, we see that apprenticeships are associated with substantial wage gains. However the effects fade over time and are not statistically distinguishable from zero at the very last quarter of our observation period.

7 Discussion of the main findings

The main result in this paper is that on average apprenticeships can ease workers' transition towards open-ended contracts, but to the expense of fewer positions in other temporary contracts. The two effects do not mechanically cancel out: quantitatively, the second effect dominates, generating a negative impact on the probability of having any job. In this sense, apprenticeships seem to constitute a double-edged sword, because they allow workers to climb higher rungs on the job ladder but lead to higher penalties when conversion to open-ended does not happen.

There are different theoretical mechanisms that can rationalize these findings. The first possibility is that apprenticeships are more accurate screening devices for individual ability than are temporary contracts. Within the training firm, employers may learn workers' types precisely, thanks to higher monitoring and more frequent interactions. Other firms in the markets will then also have access to part of this private information, by observing apprentices' retention choice (or lack thereof). An apprenticeship that is not converted to an open-ended contract reveals the presence of a lower productivity type. To the contrary,

[12]Notice that our measure is imprecise only insofar a worker can receive a pay rise that is not also reflected in a job-title change. If instead a worker receives a pay rise but is also promoted from blue collar to white collar, we would observe two earning records

temporary contracts are not as precise screening devices as apprenticeships. While the initial firm may still learn a lot about worker types during this period, temporary contract may fail to be renewed because of exogenous reasons with higher probability, and therefore should lead to a weaker updating by the other firms in the market.

The second possibility is that apprentices acquire firm-specific skills that are not easily re-usable at other employers. To the contrary tasks performed in temporary contracts may be more standardized. Even here, dismissals after apprenticeships should lead to a penalty in the labor market, as time was "wasted" learning things not valued elsewhere. This would be consistent with recent evidence showing that apprenticeships may generate specific skills and scarce adaptability to new environments (Hanushek et al., 2017). The two stories are not necessarily mutually exclusive, and disentangling the two is left for future research.

8 Conclusions

In this paper we have analyzed the returns to apprenticeships by looking at a variety of labor market outcomes. In terms of conversion to open-ended contracts, apprenticeships are dominated by temporary contracts in the first three years after the start of the contract, but guarantee higher conversion rates afterwards, by about 8.5 p.p.. All of these extra conversions happen at the initial firm, while conversions to open-ended in other firms negatively contribute to the overall effect. While they increase the probability of accessing better jobs, they decrease the probability of obtaining further temporary contracts. This second effect is bigger (-13.1 p.p.) and negatively impacts the probability of having any job. We find transitions to self-employment not to be an important margin of adjustment in this context. Taken together, our results highlight a trade-off between the quality and the quantity of job offers that could result after starting an apprenticeship.

9 Tables

Table 1: Summary statistics

	Apprentices		**Temporary contracts**	
Variable (pre-event average)	**Mean**	**Std. Dev.**	**Mean**	**Std. Dev.**
Prob. of having any job	0.230	0.421	0.147	0.354
Prob. of being an employee	0.202	0.402	0.118	0.323
Prob. of being a blue collar	0.121	0.327	0.080	0.271
Prob. of being a white collar	0.084	0.278	0.039	0.193
Prob. having open-ended contract	0.096	0.295	0.114	0.318
Prob. having temporary contract	0.109	0.311	0	0
Age at start of spell (years)	24.098	1.978	25.01	2.139
Average monthly earnings (euros)	1250.82	610.614	1232.238	654.895
Number of spells	100,547		179,528	

Notes: This table provides descriptive statistics for our main sample. All variables are measured as an average of the four quarters before the start of the contract. All employment outcomes are dummy variables that take value one if the condition is true for at least one month during the quarter. As a consequence outcomes are never mutually exclusive. The probability of having any job includes both employment, dependent self-employment and self-employment. Average quarterly earnings is expressed in 2017 euros and winsorized at the 1st and 99th percentile. It includes all earnings from either employment and dependent self-employment. Earnings from self-employment are not included.

10 Figures

Figure 1: Probability of being an open-ended contract

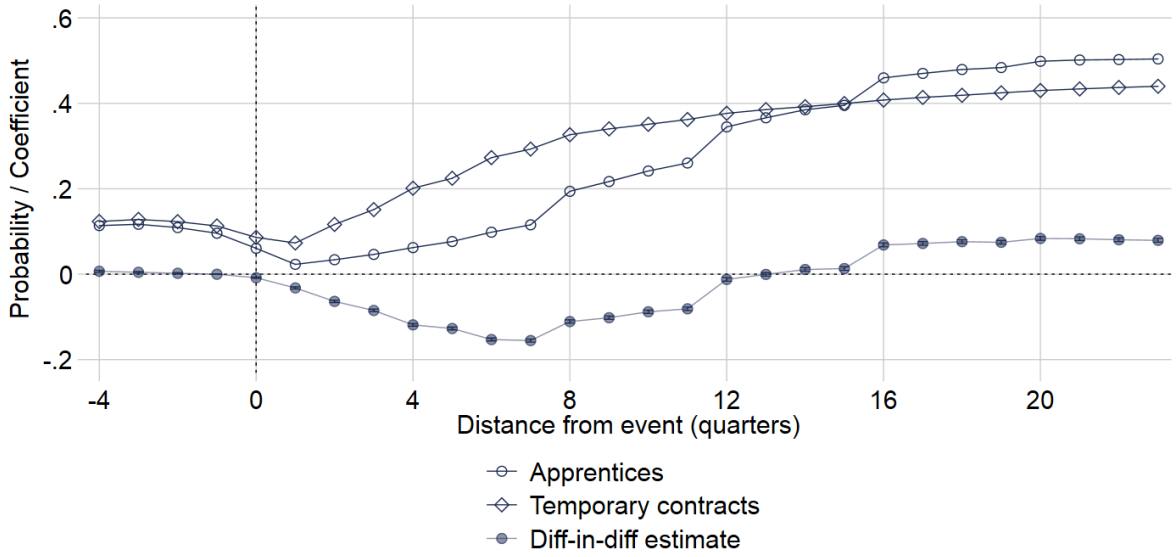

Note: The figure plots the dynamic evolution of the mean probability of being in an open-ended contract, for apprentices and individuals in temporary contracts (hollow circles and diamonds respectively). Solid blue circles indicate difference-in-differences estimates (β_k^T) from specification 1. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 2: Probability of being in temporary or open-ended contracts

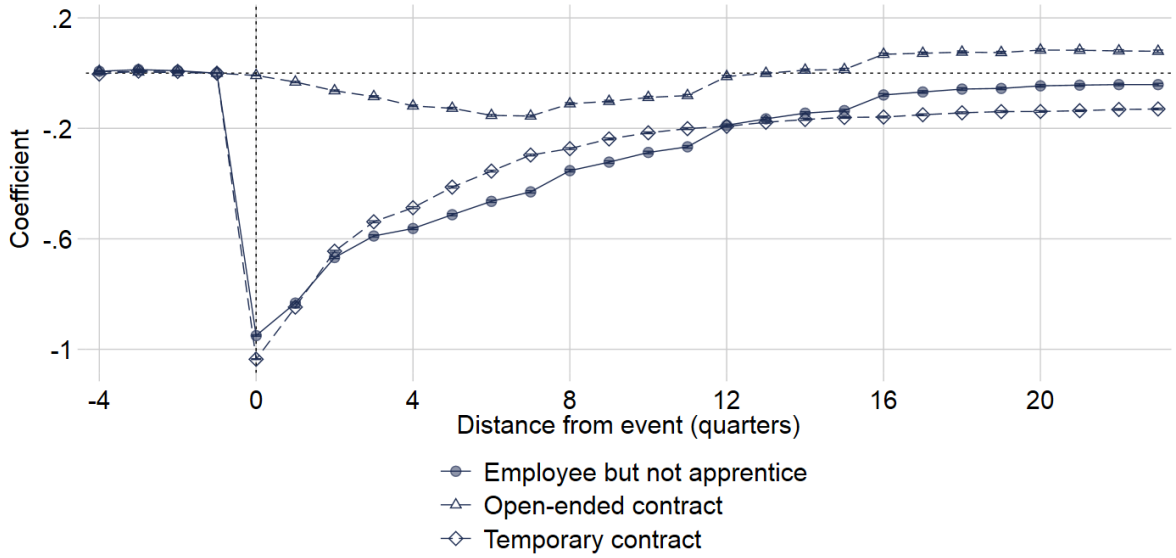

Note: The figure plots β_k^T coefficients from specification 1 for three outcomes: the probability of being employed under an open-ended contract, the probability of being employed under a temporary contract and the probability of being employed except for apprenticeship contracts. The latter constitutes the union of the former two events. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 3: Probability of being in open-ended contracts at initial or other firms

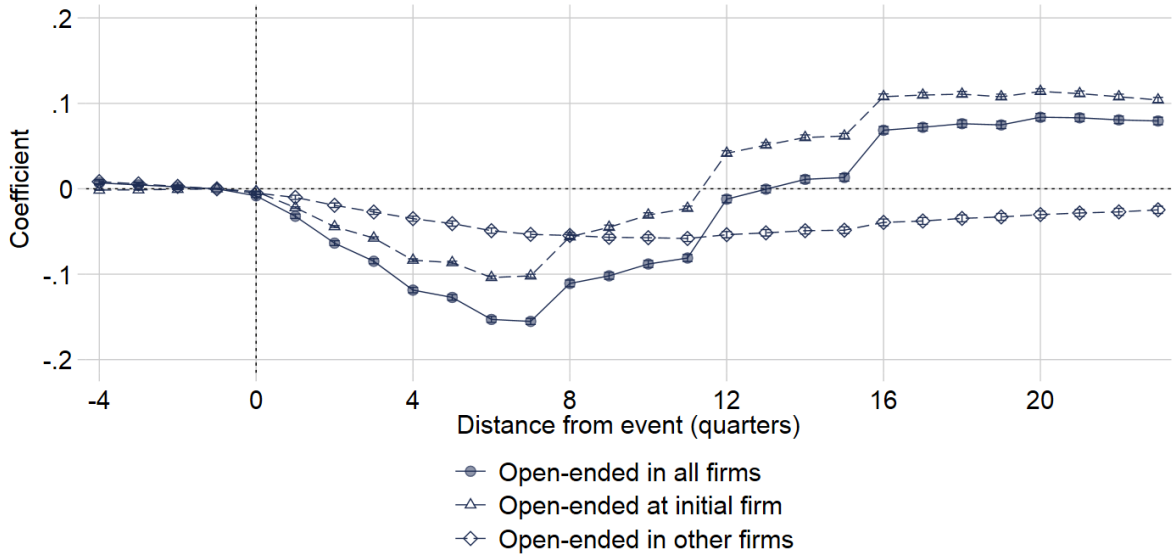

Note: The figure plots β_k^T coefficients from specification 1 for three outcomes: the probability of being employed under an open-ended contract, the probability of being employed under an open-ended contract at the same firm where the contract is started ($k = 0$) and the probability of being employed under an open-ended contract in firms other than the firm where the contract started. The latter constitutes the union of the former two events. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 4: Probability of being in temporary contracts at initial or other firms

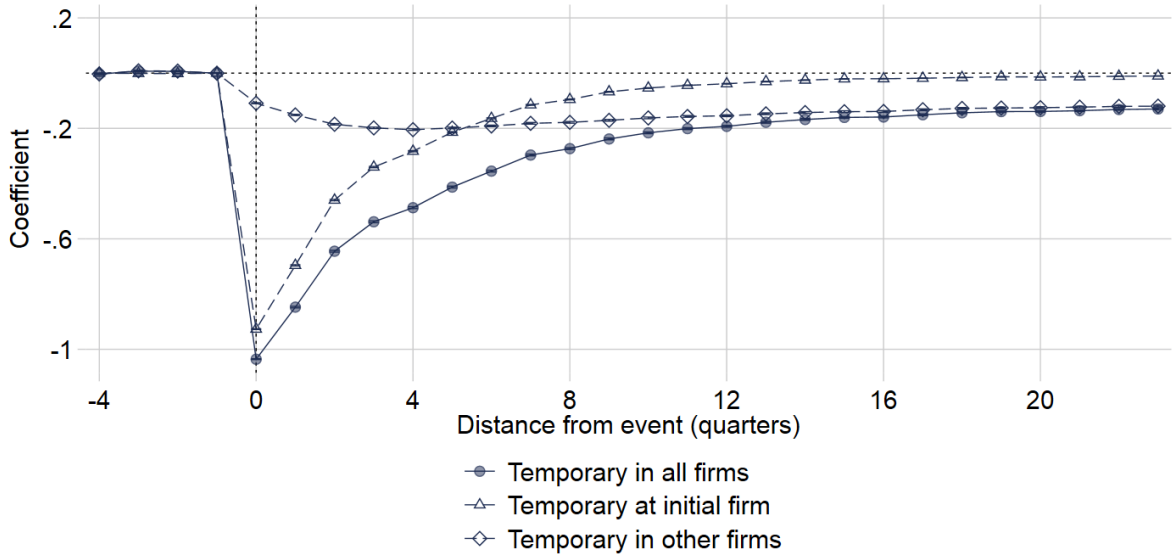

Note: The figure plots β_k^T coefficients from specification 1 for three outcomes: the probability of being employed under a temporary contract, the probability of being employed under a temporary contract contract at the same firm where the contract is started ($k = 0$) and the probability of being employed under a temporary contract in firms other than the firm where the contract started. The latter constitutes the union of the former two events. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 5: Probability of being in temporary or open-ended contracts by firm size

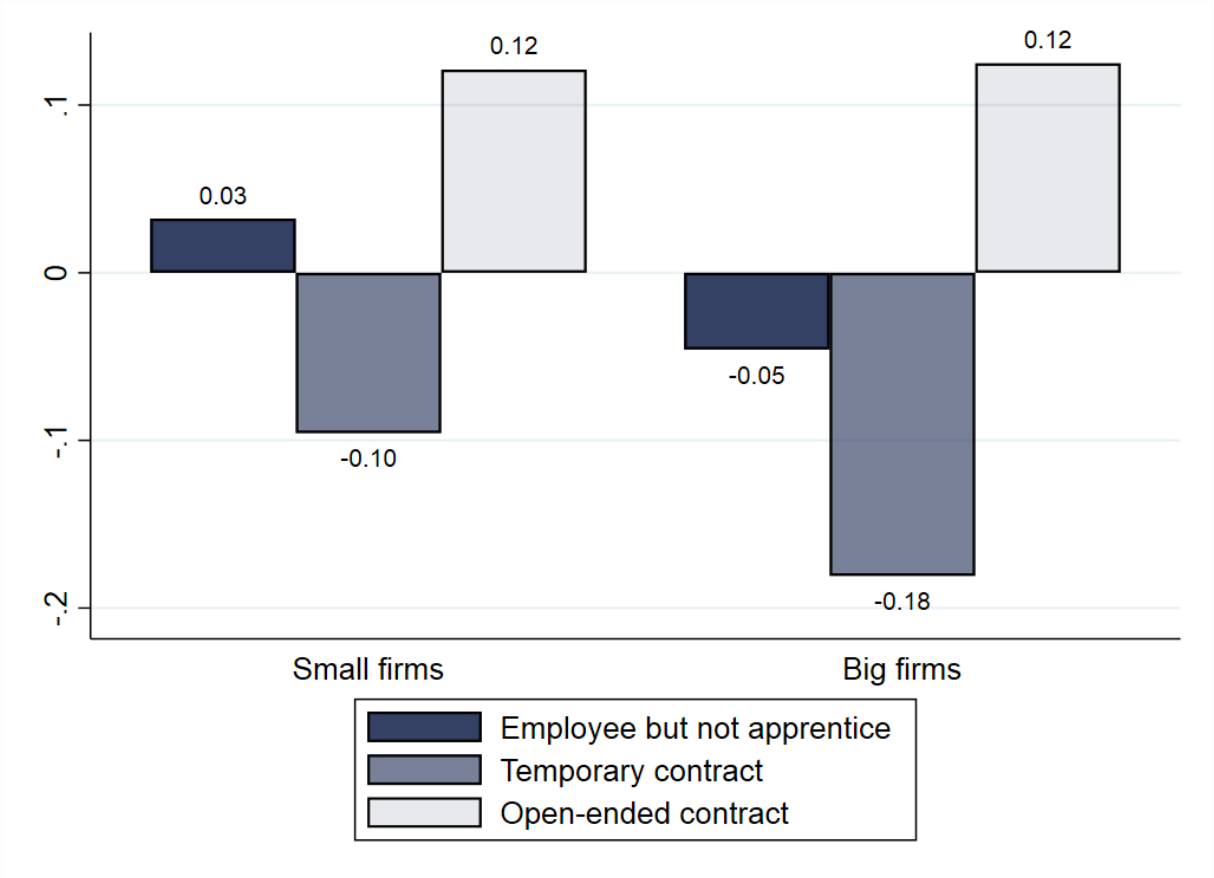

Note: The figure plots β_k^T coefficients from specification 1 for $k = 23$ only, run separately for contracts started in small firms and big firms. A firm is defined as big if its average size in the solar year when the contract starts is strictly greater than 15 and small otherwise. Three outcomes are displayed: the probability of being employed under an open-ended contract, the probability of being employed under a temporary contract and the probability of being employed except for apprenticeship contracts. The latter constitutes the union of the former two events. Standard errors for the difference-in-differences estimates are clustered at the individual level. Confidence intervals are not displayed, but estimates are always significant at the 1% level.

Figure 6: Probability of being in open-ended contracts at initial or other firms by firm size

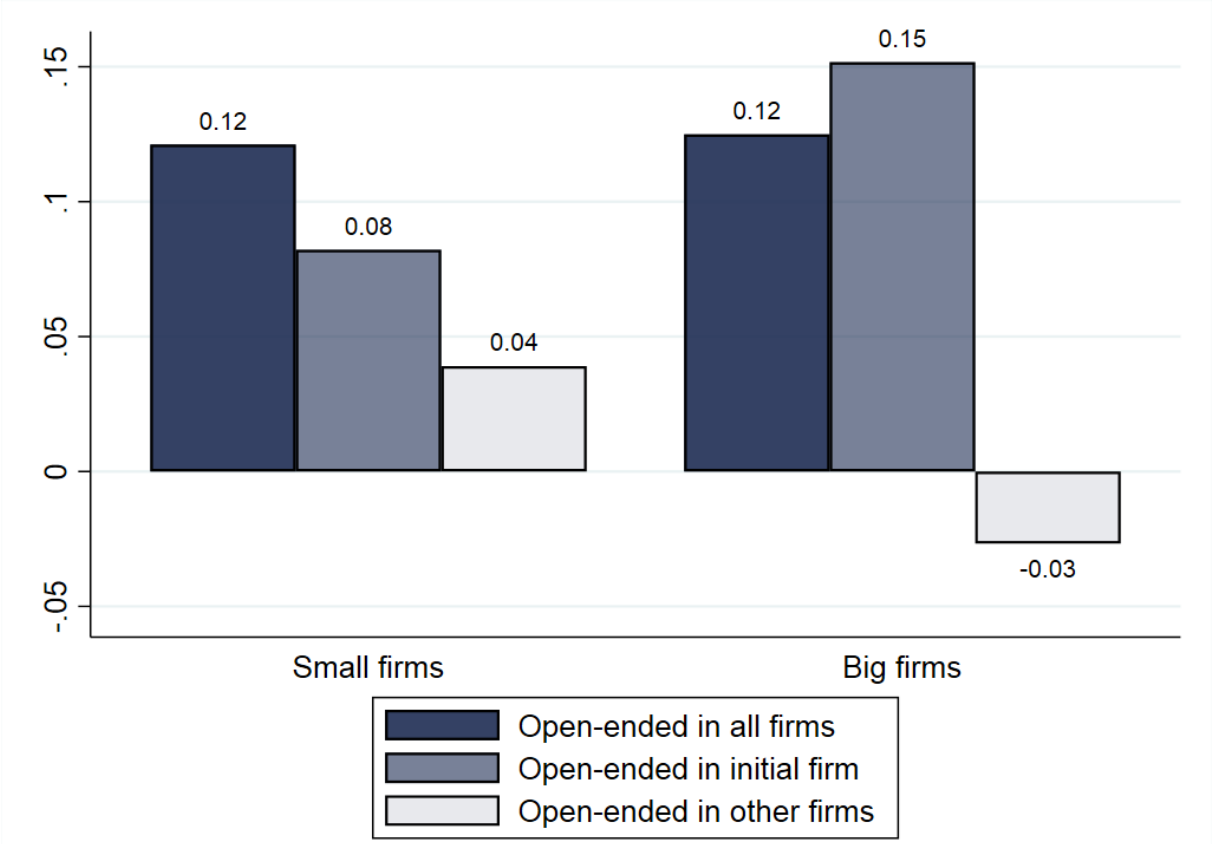

Note: The figure plots β_k^T coefficients from specification 1 for $k = 23$ only, run separately for contracts started in small firms and big firms. A firm is defined as big if its average size in the solar year when the contract starts is strictly greater than 15 and small otherwise. Three outcomes are displayed: the probability of being employed under an open-ended contract, the probability of being employed under an open-ended contract at the same firm where the contract is started ($k = 0$) and the probability of being employed under an open-ended contract in firms other than the firm where the contract started. Standard errors for the difference-in-differences estimates are clustered at the individual level. Confidence intervals are not displayed, but estimates are always significant at the 1% level.

Figure 7: Probability of being in temporary contracts at initial or other firms by firm size

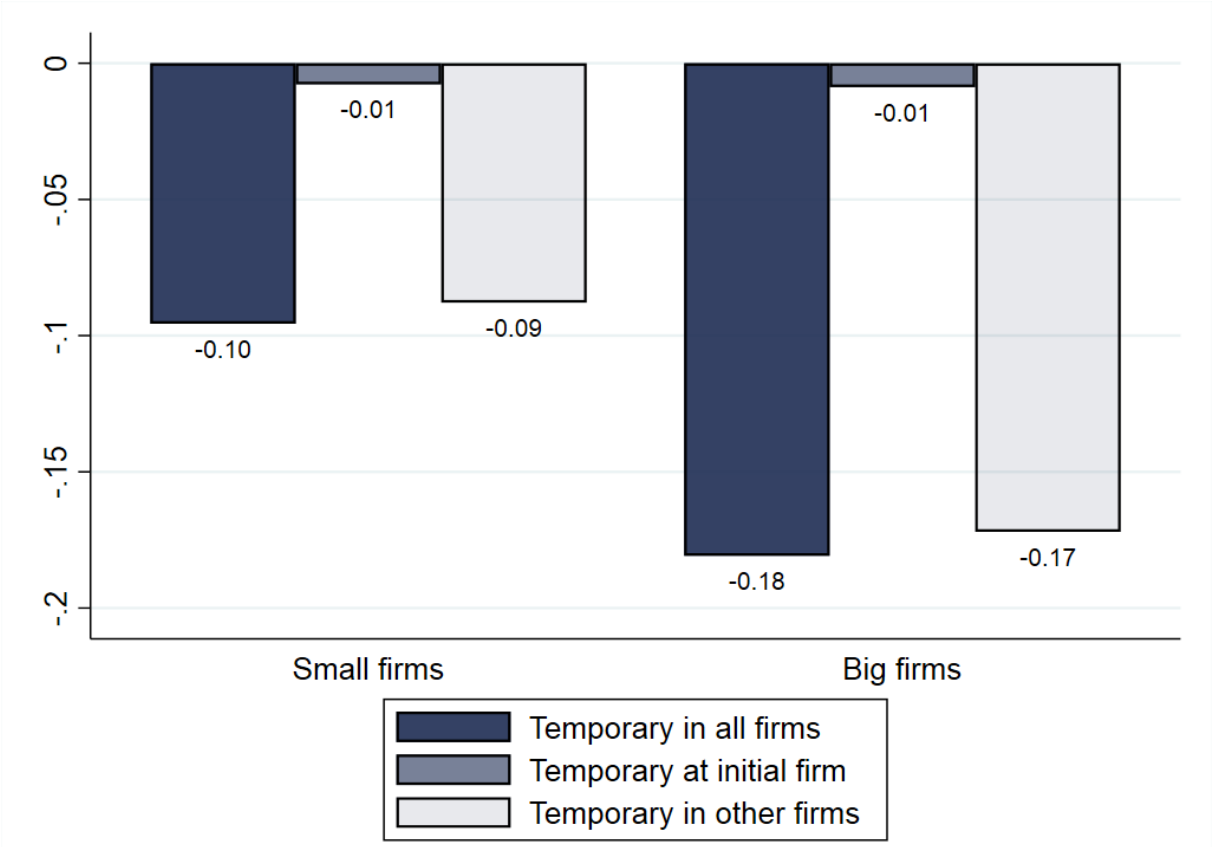

Note: The figure plots β_k^T coefficients from specification 1 for $k = 23$ only, run separately for contracts started in small firms and big firms. A firm is defined as big if its average size in the solar year when the contract starts is strictly greater than 15 and small otherwise. Three outcomes are displayed: the probability of being employed under a temporary contract, the probability of being employed under a temporary contract at the same firm where the contract is started ($k = 0$) and the probability of being employed under a temporary contract in firms other than the firm where the contract started. Standard errors for the difference-in-differences estimates are clustered at the individual level. Confidence intervals are not displayed, but estimates are always significant at the 1% level.

Figure 8: Employment and self-employment

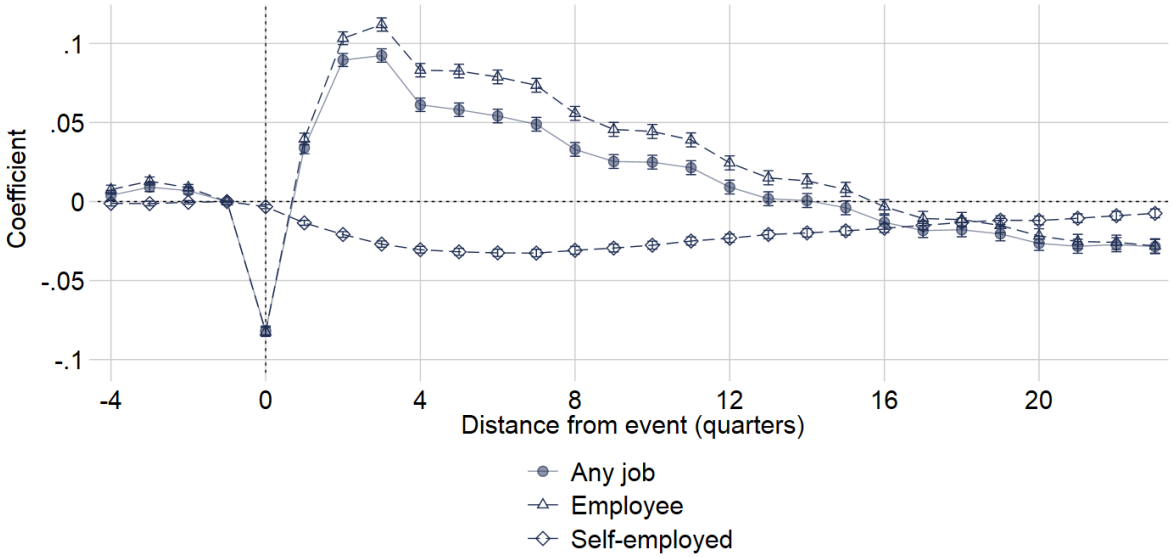

Note: The figure plots β_k^T coefficients from specification 1 for different outcomes. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Given that mechanically both groups have Pr(any job$= 1|k = 0) = 1$, the point estimate at $k = 0$ equals the level difference that exists between the two groups at $k = -1$. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 9: Log(quarterly earnings)

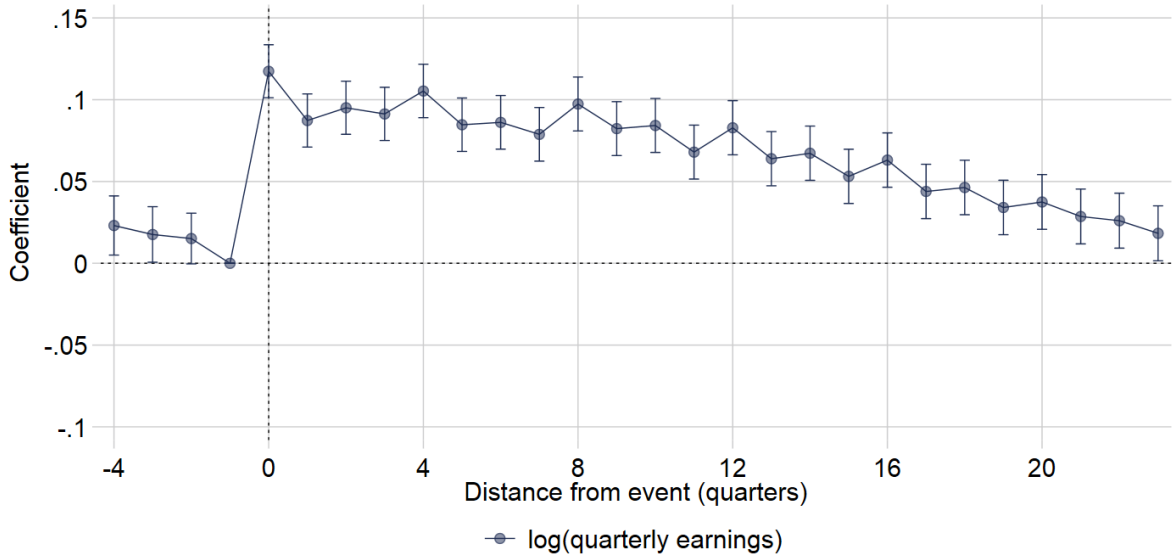

Note: The figure plots β_k^T coefficients from specification 1. The dependent variable is the natural logarithm of quarterly earnings, conditional on working status. Earnings include both labor income from employment and dependent self-employment. We have no reliable information on earnings as self-employed. Quarterly earnings are constructed by apportioning yearly earning amounts to quarters in proportion to the number of months spent in a given spell. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

References

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 Labor Market Reforms: Evidence from Italian Apprenticeships. Discussion Papers 10766, IZA.

Albert, C., García-Serrano, C., and Hernanz, V. (2005). Firm-provided training and temporary contracts. *Spanish Economic Review*, 7(1):67–88.

Becker, G. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy*, 70.

Berton, F., Devicienti, F., and Pacelli, L. (2011). Are temporary jobs a port of entry into permanent employment? evidence from matched employer-employee. *International Journal of Manpower*, 32(8):879–899.

Booth, A. L., Francesconi, M., and Frank, J. (2002). Temporary jobs: Stepping stones or dead ends? *The Economic Journal*, 112(480):F189–F213.

Cavaglia, C., McNally, S., and Ventura, G. (2018). Do apprenticeships pay? evidence for england. Discussion Paper 14, CVER.

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.

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.

Hanushek, E. A., Schwerdt, G., Woessmann, L., and Zhang, L. (2017). General education, vocational education, and labor-market outcomes over the lifecycle. *Journal of Human Resources*, 52(1):48–87.

Istat (2017). Focus statistiche: I lavoratori indipendenti. Technical report, Istat.

Jaravel, X., Petkova, N., and Bell, A. (2018). Team-specific capital and innovation. *American Economic Review*, 108(4-5):1034–73.

Kuczera, M. (2017). Incentives for apprenticeship. Education Working Paper 152, OECD.

Parey, M. (2016). Vocational schooling versus apprenticeship training. evidence from vacancy data. Working paper.

Picchio, M. and Staffolani, S. (2013). Does apprenticeship improve job opportunities? a regression discontinuity approach. *Empirical Economics*, pages 1–38.

Ryan, P. (2012). Apprenticeship: between theory and practice, school and workplace. In *The future of Vocational Education and Training in a changing world*, pages 402–432. Springer.

Samek Lodovici, M. et al. (2013). The effectiveness and costs-benefits of apprenticeships: Results of the quantitative analysis.

Scrutinio, V. (2019). The medium term effects of unemployment benefits. WorkINPS paper 18, INPS.

Snell, K. (1996). The apprenticeship system in british history: The fragmentation of a cultural institution. *History of Education*, 25(4):303–322.

Tiraboschi, M. (2012). Young workers in recessionary times: A caveat (to continental europe) to reconstruct its labour law? *E-Journal of International and Comparative Labour Studies*, 1(1-2).

Williams, C. and Lapeyre, F. (2017). Dependent self-employment: trends, challenges and policy responses in the eu. Employment Working Paper 228, ILO.

A Appendix: additional figures

Figure 10: Probability of being in temporary or open-ended contracts (small firms)

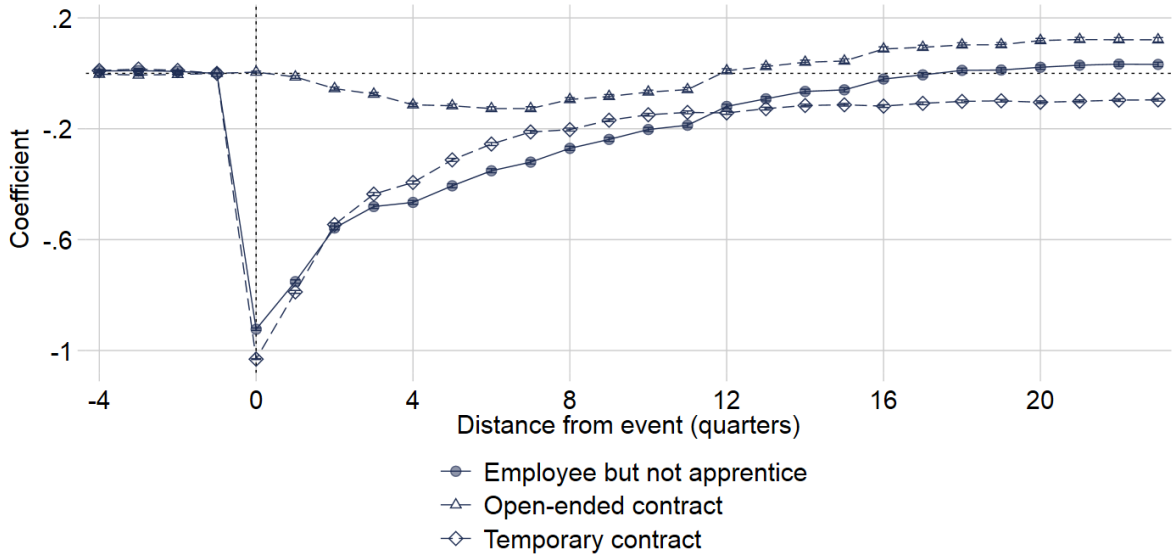

The figure plots β_k^T coefficients from specification 1 for three outcomes: the probability of being employed under an open-ended contract, the probability of being employed under a temporary contract and the probability of being employed except for apprenticeship contracts. The latter constitutes the union of the former two events. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 11: Probability of being in temporary or open-ended contracts (big firms)

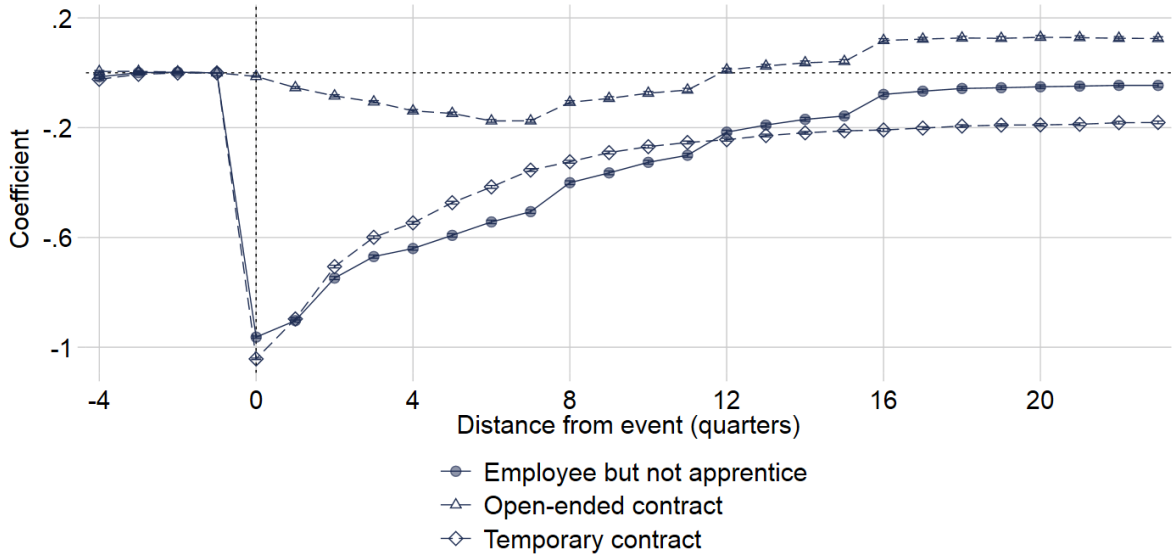

Note: The figure plots β_k^T coefficients from specification 1 for three outcomes: the probability of being employed under an open-ended contract, the probability of being employed under a temporary contract and the probability of being employed except for apprenticeship contracts. The latter constitutes the union of the former two events. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 12: Probability of being in open-ended contracts at initial or other firms (small firms)

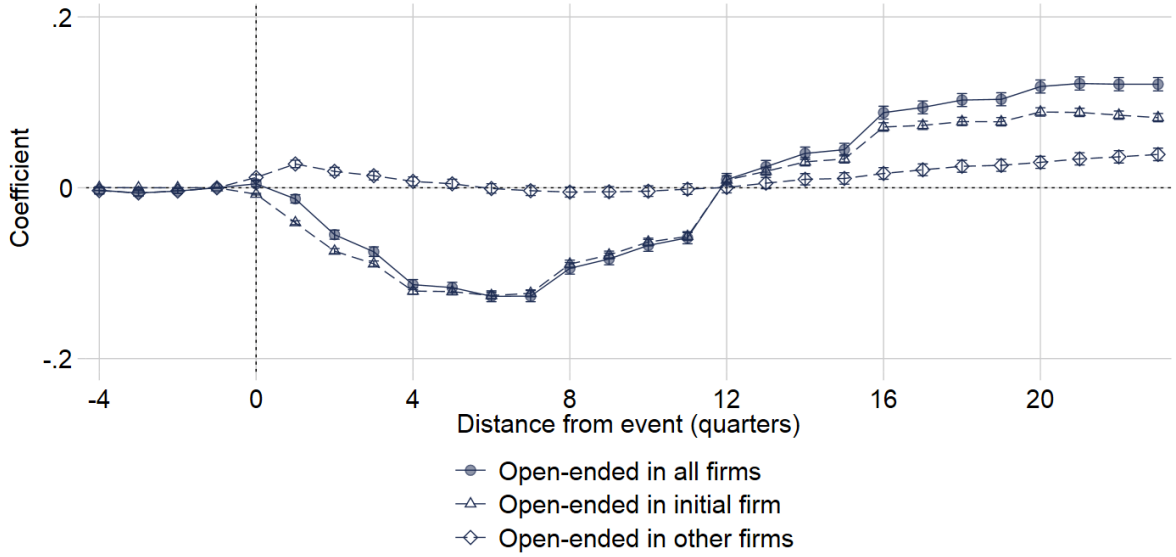

The figure plots β_k^T coefficients from specification 1 for three outcomes: the probability of being employed under an open-ended contract, the probability of being employed under an open-ended contract at the same firm where the contract is started ($k = 0$) and the probability of being employed under an open-ended contract in firms other than the firm where the contract started. The latter constitutes the union of the former two events. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 13: Probability of being in open-ended contracts at initial or other firms (big firms)

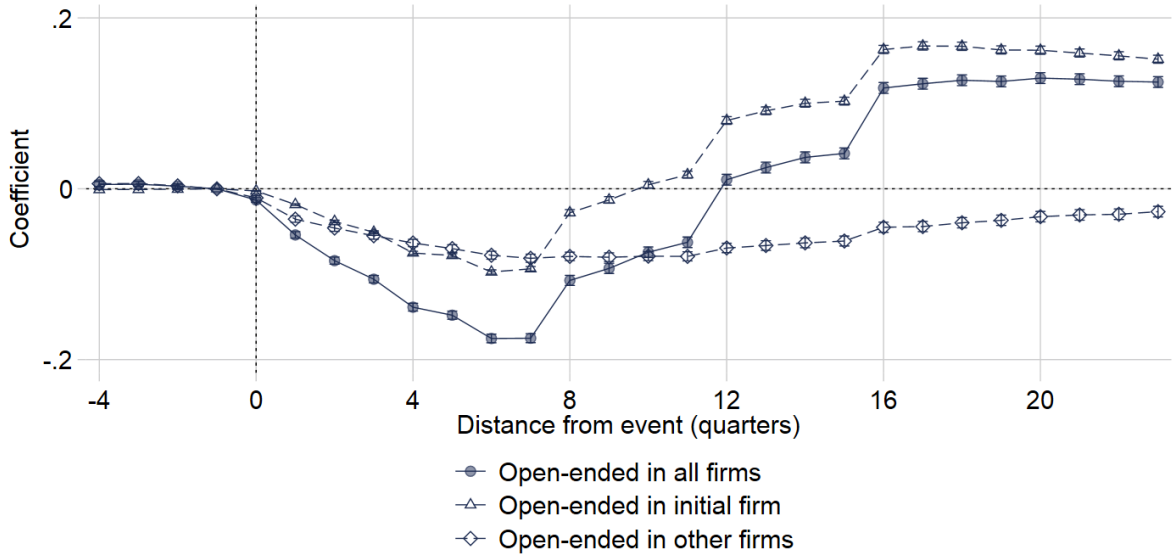

The figure plots β_k^T coefficients from specification 1 for three outcomes: the probability of being employed under an open-ended contract, the probability of being employed under an open-ended contract at the same firm where the contract is started ($k = 0$) and the probability of being employed under an open-ended contract in firms other than the firm where the contract started. The latter constitutes the union of the former two events. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 14: Probability of being in temporary contracts at initial or other firms (small firms)

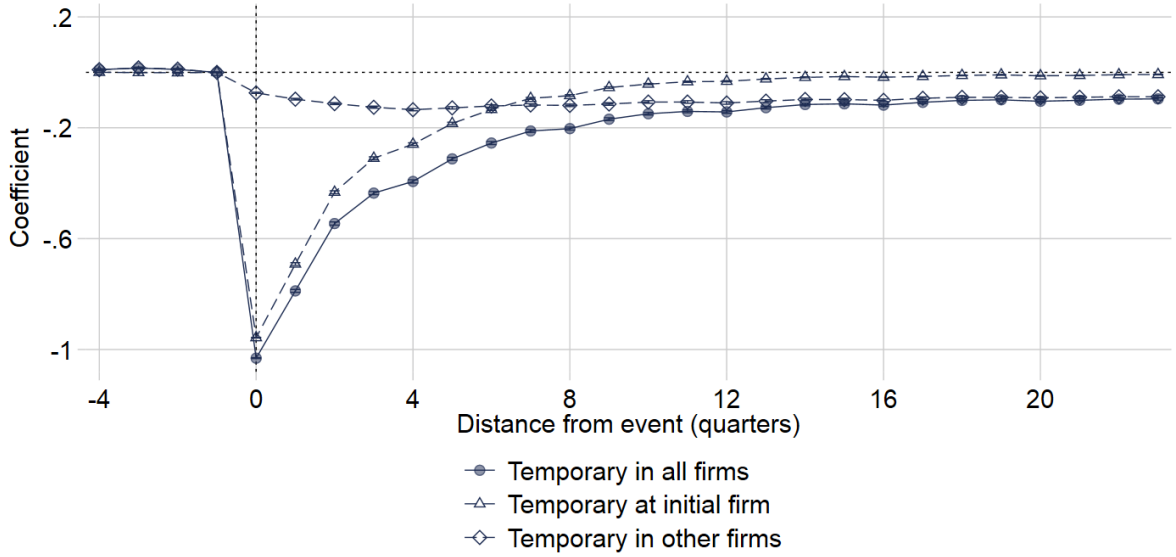

The figure plots β_k^T coefficients from specification 1 for three outcomes: the probability of being employed under a temporary contract, the probability of being employed under a temporary contract at the same firm where the contract is started ($k = 0$) and the probability of being employed under a temporary contract in firms other than the firm where the contract started. The latter constitutes the union of the former two events. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.

Figure 15: Probability of being in open-ended contracts at initial or other firms (big firms)

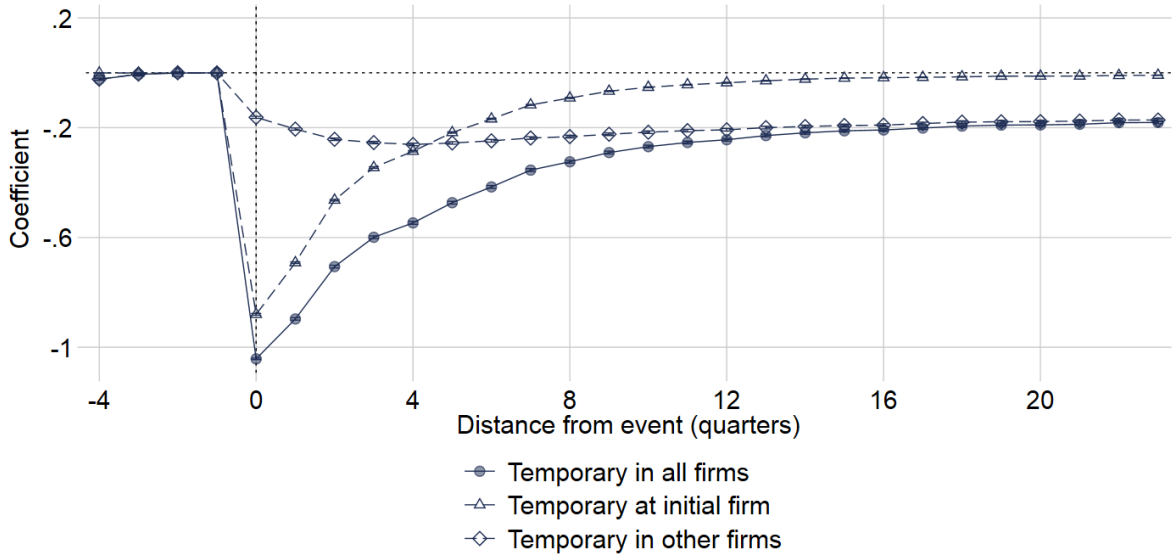

The figure plots β_k^T coefficients from specification 1 for three outcomes: the probability of being employed under a temporary contract, the probability of being employed under a temporary contract at the same firm where the contract is started ($k = 0$) and the probability of being employed under a temporary contract in firms other than the firm where the contract started. The latter constitutes the union of the former two events. The difference at event time $k = -1$ is normalized at zero. Event time $k = 0$ corresponds to the quarter when both the apprenticeship and the temporary contract start. Standard errors for the difference-in-differences estimates are clustered at the individual level and corresponding 95% confidence intervals are displayed.