No Teens, No Tech: How Shortages of New Entrants Hinder Firm Technology Investments

Cäcilia vom Baur*

June 17, 2026

Abstract

What drives firm investments in new technologies, a key source of economic growth? I argue that young labor market entrants, i.e. trainees, are a central determinant because of their low opportunity costs and high returns from acquiring new tech skills. I provide causal evidence for this hypothesis by leveraging a large and exogenous trainee supply shock caused by an education reform in Germany in 2001. Lower trainee supply decreases firm capital investments, in particular in industries undergoing strong technological change, and deteriorates the technical status of firms' machinery, indicating complementarity between new entrants and new technologies. Consistent with the skill-acquisition mechanism, this investment-depressing effect is stronger in firms characterized by high human capital investments in their trainees. These findings cast doubt on counteracting shortages of young entrants by substituting labor with capital.

Keywords: Technology Adoption, Labor Shortages, Firm Investments, Vintage-Specific Skills

JEL: D22, D24, J21, J24, O33

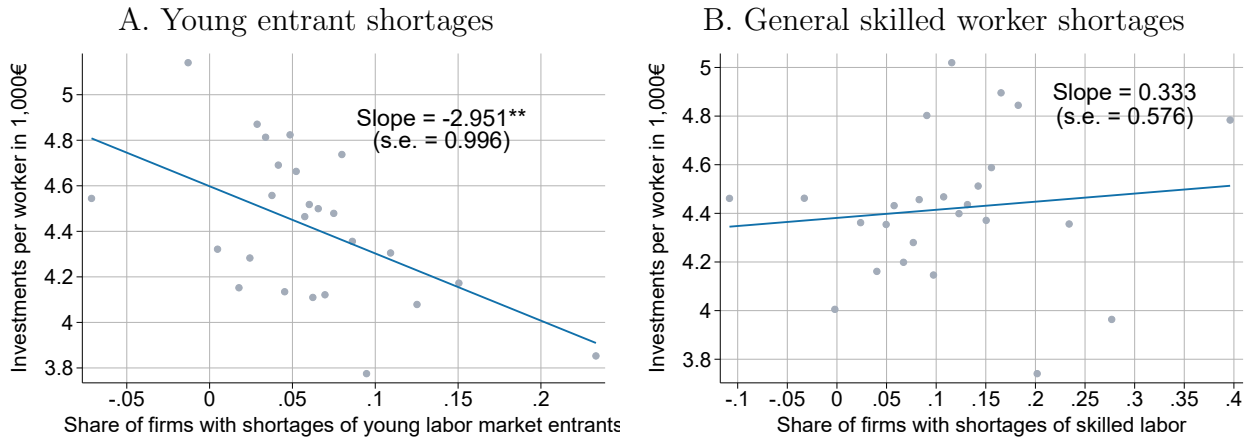
*Goethe University Frankfurt, vombaur@econ.uni-frankfurt.de

This paper was previously circulated under my maiden name, Cäcilia Lipowski. I am greatly indebted to Anna Salomons and Ulrich Zierahn-Weilage for their advice and support. For helpful comments and conversations, I thank Daron Acemoglu, Melanie Arntz, David Autor, Eduard Brüll, Christian Dustmann, Guido Friebel, Katja Görlitz, Maarten Goos, Simon Jäger, Morten Olsen, Harald Pfeifer, Pascual Restrepo, Johannes Schmieder, Fabian Waldinger, Anna Waldman-Brown, and Nicolas Ziebarth. I thank conference and seminar participants at the EEA (Barcelona), EALE (Prague), TPRI (Boston), Skills-for-the-Future Conference (LISER), CESifo Summer Institute on Re-skilling and Skill Shortages (Venice), Leading House Conference on the Economics of Educational and Vocational Training (Zurich), at seminars at University of Bonn, Frankfurt School of Finance and Management, Goethe University Frankfurt, ifo Munich, IWH Halle, IZA, LMU Munich, University of Mannheim, Rockwool Foundation Berlin, University of Strasbourg, Utrecht University and ZEW Mannheim, and numerous other conferences and seminars. The project was mainly written at the ZEW Mannheim and was financially supported by the Leibniz Association through the Leibniz Professorship for Applied Labor Economics at the University of Heidelberg (P56/2017) and ZEW Mannheim.

1 Introduction

Economic growth depends not only on technology invention but also on diffusion across firms. The supply of young labor market entrants is potentially an important determinant of firms' technology adoption: entrants' low opportunity costs and high returns to skill acquisition may make them crucial complements to new technologies that require new skills. German regions with more severe shortages of young entrants indeed exhibit significantly less firm capital investments (Figure 1, Panel A), controlling for region and year fixed effects. In contrast, shortages of skilled workers in general are not associated with lower investments (Panel B), indicating that young entrants play a distinct role from incumbents.[1] Also, younger and lower-tenured workers are more likely to use AI (Humlum & Vestergaard, 2025; Pulito et al., 2026; Bick et al., 2026). Establishing a causal link between entry-level workers and firms' technology adoption is difficult, however: Labor supply shortages, while being widely discussed in the policy debate (e.g. Fuest & Jäger, 2023), are rarely sharp and exogenous. Effects of technology adoption on younger versus older workers, in turn, may be confounded with differential hiring and firing patterns by age groups (Autor & Dorn, 2009).

Figure 1: Labor shortages and firm investments at the regional level



Notes: Binscatter at the county level. Based on the IAB establishment panel survey, 1997–2006; years chosen for data availability reasons. N=1,961. Controlling for region and year fixed effects; labor shortages lagged by one year compared to investments. For more information and other specifications, see Appendix A.1. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

In this paper, I identify the causal role of the supply of young labor market entrants in firms' technology adoption by exploiting an education reform that induced a negative supply shock of vocational trainees in Germany. I show that reduced entrant supply significantly hinders firms' technology adoption, indicating complementarity between new technologies and entrants, explained by their comparative advantage in skill acquisition.

I start by establishing two stylized facts consistent with this complementarity. First, in a representative survey, a substantial share of firms states reports offering vocational training

[1] Equivalent patterns are found for other specifications, see Appendix A.

because it helps adapting to technological change and ensures the constant supply of new skills. Second, young workers are significantly more likely to work with new technologies than workers aged 30 and above.

Next, I sketch a stylized framework that produces complementarity between new entrants and new technologies through what I call the "new-skills" mechanism. The complementarity hinges on the arguably plausible assumption that new entrants are initially less productive than incumbent workers. Consider a task-based vintage capital model where operating a new technology creates a new task requiring a new skill. Take for example office technologies. These technologies automate labor in data entry tasks, augment labor in decision-making tasks, and create new tasks requiring new skills related to system navigation. New entrants face lower opportunity costs and greater productivity gains from learning such new skills compared to incumbents, since their initial productivity is lower. Firms thus choose to complement new technologies with new entrants. The framework delivers three testable predictions: First, firms reduce investments in new technologies when new entrants are scarce. The reason is that both retraining incumbents to use new technologies and attracting additional entrants via higher wages raise the cost of technology adoption relative to the no-shortage scenario. Second, this investment reduction is driven by firms characterized by heavily investing in trainees' new skills and in industries undergoing technological change (third).

To identify the causal effect of a supply reduction of new entrants on firms' technology adoption—the main focus of this paper—I exploit a natural experiment arising from an education reform. In 2001, two East German federal states (henceforth "treated states") increased the length of schooling required for the university entrance qualification by one year.[2] This reform created a one-time missing cohort of school graduates from the upper track. About one third of them usually pursue vocational training, accounting for 20% of training firms' entry-level hires. The reform hence results in a missing cohort of trainees from the upper school track. There was no comparable shock in the other four East German states, henceforth "control states". This shock has a number of beneficial features: First, its exogenous and temporary nature allows me to estimate the direct causal relationship between new entrants and technology adoption, and to pin down the underlying mechanism. Other labor supply shocks usually combine a variety of aspects—a smaller number of total workers, reduced consumption, increased savings etc.—impeding such conclusions. Second, the identification is sharp because trainees tend to stay close to their hometowns, unlike university students who often relocate for their studies and jobs. Third, the shock is large, making the study high-powered to detect effects, and market-wide, making substitution effects between firms almost impossible. It therefore mimics effects of other market-wide shocks, such as population aging.

In Germany, vocational training is a full-time educational program following high-school,

[2] Among others, Büttner & Thomsen (2015); Morin (2015); Muehlemann et al. (2022); Marcus & Zambre (2019) and Dorner et al. (2024) exploit this and the opposite reform to study the effect on school grades, university enrollment, trainee employment and trainee wages. So far, no study has looked at effects on firms.

taking place in firms and vocational schools simultaneously. Trainees often remain at their training firm after completing the program. Trainees from the upper school track studied in this paper, i.e. those with 12 or 13 years of schooling, henceforth "highly educated trainees", make up almost 30% of all trainees[3] while the majority have 9 or 10 years of schooling, henceforth "low-educated trainees." In general, vocationally trained workers are considered middle-skilled professionals, rather than low-skilled laborers (Fitzenberger et al., 2024), and make up the largest share of the workforce with two-thirds.[4]

To estimate a causal parameter, I compare investments of establishments in treated East German states to those in control East German states in a difference-in-differences (DiD) event study design. In the following, I use the terms "firm" and "establishment" interchangeably for simplicity. I focus on training firms, defined as firms that employed at least one highly educated trainee prior to the reform. The identification strategy rests on two main assumptions. First, absent the reform, firm outcomes in treated and control states would have followed parallel trends. This assumption is supported by pre-treatment data. To rule out that concomitant industry-specific shocks drive the results, I match treated firms to comparable control firms within industries. Second, I assume that firms in control states are not affected by the reform. This assumption is likely to hold since vocational trainees are highly immobile (Muehlemann et al., 2022).[5] Accordingly, cross-state commuting does not increase after the reform.

Using a large and representative firm panel survey (IAB[6] Establishment Panel) linked with social security records (LIAB), I directly observe firms' trainee employment, and total capital investments (my main outcome). I employ two strategies to link capital investments to technology adoption: First, I distinguish firms by the sector they operate in, assuming that sectors with large productivity gains in this time period experience strong technological change. Second, I study firms' self-assessed technical status of machinery.

I provide four key empirical findings. First, the education reform causes a temporary trainee shortage. Firms' employment stock of highly educated trainees decreases by a quarter for four years, consistent with supply being reduced by one entire cohort and vocational training usually lasting three years plus one year for delays due to military/social service. There is no increased trainee hiring prior to the shock as trainee supply almost exclusively consists of fixed-sized school graduation cohorts, making it highly inelastic. Training wages do not increase and firms do not compensate for the absence of highly educated trainees by hiring more low-educated trainees or workers with completed vocational training. These rigidities may reflect the temporary nature of the shock. However, the economic framework suggests that both uncompensated trainee supply shocks and those partly offset by wage and worker

[3] Ministry of Education & Research (2025). In 2000, the share was 16% (Federal Statistical Office, Genesis-Online, 2022a).

[4] Sample of Integrated Labor Market Biographies (SIAB), own calculations.

[5] Only 2.2% of trainees move across federal states for their vocational training (Socio-Economic Panel (SOEP), own calculations) and only 5% commute between federal states (LIAB, own calculations).

[6] Institute for Employment Research

entry responses yield qualitatively similar effects on technology adoption.

The second key finding is that the trainee shortage depresses firm investments. Investments decrease sharply in treated compared to control training firms during the years of the trainee shortage. There is no adjustment in firms' investments in anticipation of the trainee shortage. This can be rationalized by firms either being unaware of the impending shortage—indeed, I find no evidence that treated training firms expected a shortage of new entrants—or by opposing firm responses that cancel each other out: some firms reduce technology adoption prior to the shortage, knowing they need to train several trainee cohorts on the new technology, but only a few would arrive before the shortage, while other firms invest more in anticipation of the shortage. Firms substantially reduce investments when trainees are missing: investments drop by 8% at the median (20% at the mean), or €1,140 per worker (€2,740). This is equivalent to a reduction by 0.4% (0.9%) of the capital stock. Hence, the effect is larger than what would result from a "mechanical" reduction of capital in proportion to employment, in line with the argument that young entrants are disproportionally relevant for technology adoption. Firms do not adjust investments for all workers equally but instead forgo investment spikes: the probability of undertaking large investments decreases by 15 percentage points (pp). Investment levels of treated firms return to those of control firms after the shock but do not exceed them. The absence of delayed compensating investments may reflect perfect worker-capital complementarity (normal-sized post-treatment cohorts only trigger normal-sized investments), investment frictions (firms cannot invest in twice as large amounts at the same time), or missing follow-up investments (causing treated firms to always lag behind control firms).

The finding that firms time their investments to align with the availability of trainees informs the main take-away of the paper: Trainees are key complements to capital investments. Incumbent workers, not being short in supply, do not compensate for the role of new entrants in firm technology adoption. This central finding survives a battery of robustness checks, for example including the West German states as control states, or an estimation at the region-industry level.

I confirm the causal link to trainee supply using two auxiliary identification strategies: First, I compare the difference in investment trends between *training* firms in treated and control states to the same difference between *non-training* firms in treated and control states in a triple DiD design. Confirming the link to the trainee shortage, training firms reduce investments compared to non-training firms in treated states. Second, I exploit firms' pre-reform exposure to highly educated trainees with a Bartik-style IV, finding that more exposed firms decrease investments to a greater extent.

As the third key finding, the complementarity between new entrants and investments is linked to skill acquisition. Comparing firms that heavily invest in their trainees' human capital—inferred from high pre-treatment trainee retention rates, as is standard in the literature—with firms that primarily employ trainees as low-cost labor, I find that the investment-depressing effect is driven by the former. This suggests that it is the trainees' role in acquiring new skills

rather than their role as low-cost labor that makes them complementary to new technologies.

Fourth, the investment decline is linked to the reduced adoption of new technologies: only firms operating in sectors characterized by rapid technological progress decrease their investments during the trainee shortage, while firm investments in tech-stagnant sectors are completely unaffected. I also use a different outcome variable to confirm the link between investments and technology adoption: firms' self-assessed technical status of machinery depreciates in treated compared to control firms. I also find indicative evidence that firm output is reduced in response to foregone technology adoption.

The paper is most closely related to a strand of literature that studies how technology invention and adoption respond to factor prices determined by their relative abundance (e.g. Zeira, 1998; Acemoglu, 1998, 2002). For example, an increased supply of skilled labor intensifies the adoption of skill-complementing technologies (Beaudry et al., 2010; Carneiro et al., 2023; Terry et al., 2026), while a decrease in the supply of (low-)skilled labor increases wages and incentivizes firms to patent and adopt labor-saving technologies (Lewis, 2011; Hornbeck & Naidu, 2014; Clemens et al., 2018; Hémous et al., 2025; Danzer et al., 2024; Andersson et al., 2022; San, 2023, vice versa for an increase). This paper contributes to the literature in two dimensions. First, it highlights a novel channel through which labor supply affects technology adoption which counteracts the standard labor-saving channel: while reduced labor supply creates incentives to adopt labor-saving technologies (the standard channel), the lack of the most complementary workers constrains adoption by increasing adoption costs. Angelini (2023) incorporates both channels in a macro model of population aging and finds that, for high levels of young worker shortages, the constraining effect dominates. Second, this paper focuses on young labor market entrants—a decisive but previously overlooked worker group—and finds that a decrease in their supply harms technology diffusion. My findings hereby also inform the debate around effects of labor shortages and population aging (Abeliansky & Prettner, 2017; Acemoglu & Restrepo, 2022; Le Barbanchon et al., 2023).

I also contribute to the literature on new technologies and new skills. The literature documents that new technologies create new tasks (Autor et al., 2024) that require vintage-specific skills (Chari & Hopenhayn, 1991; Salomons et al., 2025), and worker training (Bartel & Sicherman, 1998; Bresnahan et al., 2002; Battisti et al., 2023), while rendering skills of incumbents obsolete (Deming & Noray, 2020). In macro settings, vintage-specific skills have been put forward as the reason why technological change takes place through the entry of young workers, rather than upskilling incumbents (MacDonald & Weisbach, 2004; Cavounidis & Lang, 2020; Adão et al., 2024; Hassan et al., 2025). They might also be the reason why younger workers use AI more than experienced workers (Humlum & Vestergaard, 2025; Pulito et al., 2026; Bick et al., 2026). The contribution of this paper is to provide causal micro-level evidence for vintage effects and how they make firms forego technology adoption when new entrants are scarce.

In most general terms, this paper contributes to the extensive literature on the relationship between labor and technology. Canonical frameworks consider labor-technology complemen-

tarity based on workers' education level or job tasks (e.g. Autor et al., 2003; Acemoglu & Autor, 2011). Most empirical studies examine the relationship by analyzing the labor demand effects of technological change (e.g. Acemoglu et al., 2020; Aghion et al., 2023; Bessen et al., 2023; Bonfiglioli et al., 2024). A subset of this literature focuses on older workers, finding that they are more likely to retire early, and to not receive rents when firms invent or adopt new technologies (e.g. Aubert et al., 2006; Ahituv & Zeira, 2011; Battisti et al., 2023; Aghion et al., 2024). Conceptually, this paper identifies a novel factor determining labor-technology complementarity beyond their education level or job tasks: workers' propensity to learn and adapt, which makes young labor market entrants a key complement to new technologies. Empirically, I identify the causal relationship in the reverse way, leveraging a clean identification strategy free from endogeneity concerns in firm technology adoption and from confounding age-specific hiring and firing patterns Autor & Dorn (2009).

This paper is structured as follows. The next section presents stylized facts that motivate the following simple economic framework. Section 3 provides an overview of the German vocational training system and the education reform. Section 4 describes the data. I present the DiD event study approach in Section 5, followed by the results of the reform on trainee employment (Section 6) and firm technology investments (Section 7). Section 8 concludes.

2 Young labor market entrants and new technologies

2.1 Stylized facts

I provide two stylized facts that inform the simple economic framework.

Firm statements. The first stylized fact is that firms explicitly report offering vocational training to acquire new skills and to adapt to technological advances. I use the representative BIBB-Cost-Benefit firm survey, which asks firms engaged in vocational training about their reasons for offering it. Among East German firms in 2000 in industries exposed to technological change[7] (corresponding to the setting of the natural experiment), more than half report using vocational training to substantially improve their adaptability to technological change and to enhance their innovative capabilities, while only below 20% report not to do so, see Table 1. Similar shares are observed for firms using vocational training to ensure a constant supply of new skills and knowledge. For more details, see Appendix C.

Technology use by age groups. The second stylized fact is that young workers are more likely to use new technologies than older workers. I use a large, representative employee survey in Germany in 1999, 2006 and 2012 that asks respondents about their main working tools

[7]defined based on EU-KLEMS data. For more information, see 7.4.

Table 1: Reasons for offering vocational training according to firm statements

	Applies	Does not apply
Substantially improves adaptability to technological change	51.5%	18.2%
Enhances innovative capabilities	55.5%	17.1%
Ensures constant supply of new skills and knowledge	52.3%	18.7%

Notes: Based on the BIBB-Cost-Benefit Survey 2000. Firms in East Germany in industries with above-median exposure to technological change. Responses range on a scale from 1 ("Does not apply at all") to 5 ("Fully applies"). Applies: categories 4+5. Does not apply: Categories 1+2. Using representative survey weights. N=521.

(IAB/BIBB/BAuA Qualification and Career Survey). I construct a binary variable capturing the use of new technologies—computers and computer-controlled machines— as the main working tool and regress this variable on an age dummy, controlling for industry, occupation, year, education, and gender. Workers still in vocational training are unfortunately not included in the survey. Table 2, column 1 shows the results. Strikingly, workers below the age of 30 are 4.2pp (12%) more likely to mainly work with new technologies than workers in the same industry, same detailed occupation, and same education group, but aged 30 and above. Similar results are found for other European countries based on computer use at work reported in the European Working Conditions Survey (see Appendix C), and in studies on the adoption of AI (Humlum & Vestergaard, 2025; Bick et al., 2026; Pulito et al., 2026) suggesting that cohort effects in technology use are a widespread phenomenon.

I next perform sample splits by educational degree, see column 2 to 4. Workers with vocational training are not only those driving the differential technology use by age groups, they also have the highest overall usage level of new technologies. This highlights the importance of studying middle-skilled professionals in technology *adoption* or usage, in contrast to tertiary educated workers likely being more relevant for technology *invention*. For more information, see Appendix C.

Table 2: Technology use by age

	Overall (1)	No education (2)	Completed vocational training (3)	Tertiary educated (4)
$<$ 30 years	4.17***	3.42	4.96***	2.50
	(0.70)	(2.13)	(0.83)	(1.58)
Mean dep. variables	34.15	27.58	37.63	24.43
N	51,976	4,083	37,309	10,584

Notes: Outcome: Use of computers and computer-controlled machines as main working tool (0/100). Based on the BIBB-BAuA Qualification and Career Survey 1999, 2006 and 2012. All regressions control for dummies for survey wave, gender, education level, occupations (353), industries (17). Heteroscedasticity-robust standard errors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

2.2 Stylized economic framework

Guided by these stylized facts, I present a simple economic framework that produces complementarity between new entrants and firm investments in new technologies. The set-up is based on the vintage human capital model à la Chari & Hopenhayn (1991) and the tasks framework à la Acemoglu & Autor (2011): each new technology vintage introduces a new task requiring a new skill specific to this vintage. Firms incur capital adjustment costs of worker training in this new skill and assign the new task to workers with a comparative advantage in skill acquisition. Under the plausible assumption that new entrants are initially less productive than incumbent workers, firms assign the new technology tasks to entrants, making technology adoption endogenous to their supply.

Set-up. Firms operate with an infinite horizon and employees work in overlapping generations for two periods, such that in each period, firms have access to two worker cohorts: entrants—untrained workers with productivity A_0—, and incumbents. In each period, firms produce one final good Y by combining perfectly substitutable tasks y under the following production function:[8]

$$Y = \sum_{\nu=0}^{\tau+1} y_\nu = \sum_{\nu=0}^{\tau+1} A_\nu L_\nu \tag{1}$$

with $\tau + 1$ the latest technology vintage currently available. At the beginning of each period, a new technology vintage ν with productivity A_ν creating a task y_ν becomes exogenously available. Compared to the previous technology vintage, the new vintage increases worker productivity by $\Delta A_\nu \sim Pois(1)$. Hence, technological progress is always weakly positive, but rarely large. Only workers trained for a vintage ν—denoted L_ν—can use this vintage. Training takes the fraction $\kappa \in (0; 1]$ of one period for all technologies and workers. Adoption costs of the new vintage are borne by the firm and consist of foregone output (at base productivity) during training.

The price for the final product is fixed to one. Workers are identical in all aspects other than their cohort such that worker cohorts and worker productivity/technology types coincide. Wages w_ν are in proportion to, but below worker productivity, $w_\nu = (1 - \theta)A_\nu$—with an exogenous wage wedge $\theta \in [0, 1)$—such that benefits from technology-induced productivity increases are not completely passed on to workers.[9] A worker remains with her firm in the next period with the probability $\rho \in [0, 1]$.

[8] This production function zooms in on the labor reinstatement channel of new technologies, i.e. new technologies creating new tasks performed by humans, while abstracting from automation or factor-augmenting forms which may happen simultaneously and have been studied extensively.

[9] The renunciation of the assumption that wages are equal to marginal productivity is well backed up in the literature, in particular in the context of firm training (e.g. Konings & Vanormelingen, 2015).

Firm maximization problem. Given perfect substitutability between tasks, firms maximize profits by independently deciding for each cohort ν whether or not to assign it the new task and provide the necessary training. Assuming no discount for future periods for simplicity, for each cohort, adopting the new vintage increases net output by

$$\underbrace{(1-\kappa)(y_{\tau+1}-y_{\nu})}_{\text{Net output surplus in period 1}} + \underbrace{\rho[(y_{\tau+1}-y_{\nu})-(w_{\tau+1}-w_{\nu})L_{\nu}]}_{\text{Net output surplus in period 2}} \tag{2}$$

at the cost of foregone output $C = \kappa y_{\nu}$. Hence, a firm adopts the new vintage for a cohort if and only if the net output surplus of adoption is larger than the adoption costs:

$$[1-\kappa+\rho\theta](A_{\tau+1}-A_{\nu}) > \kappa A_{\nu} \qquad . \tag{3}$$

Given that the net output surplus of adoption decreases and that the adoption costs increase in workers' initial productivity A_{ν}, it is more often profitable for firms to assign new technology tasks to entrants than it is to incumbents.[10]

Figure 2 visualizes the trade-off. New technologies below the productivity threshold A' are not profitable to adopt, not even for entrants. New technologies above A' but below A'' are assigned to entrants only. New technologies above the threshold A'' are assigned to both incumbents and entrants. However, these highly productive technologies only arrive rarely.[11]

Testable predictions. The framework provides three testable predictions.

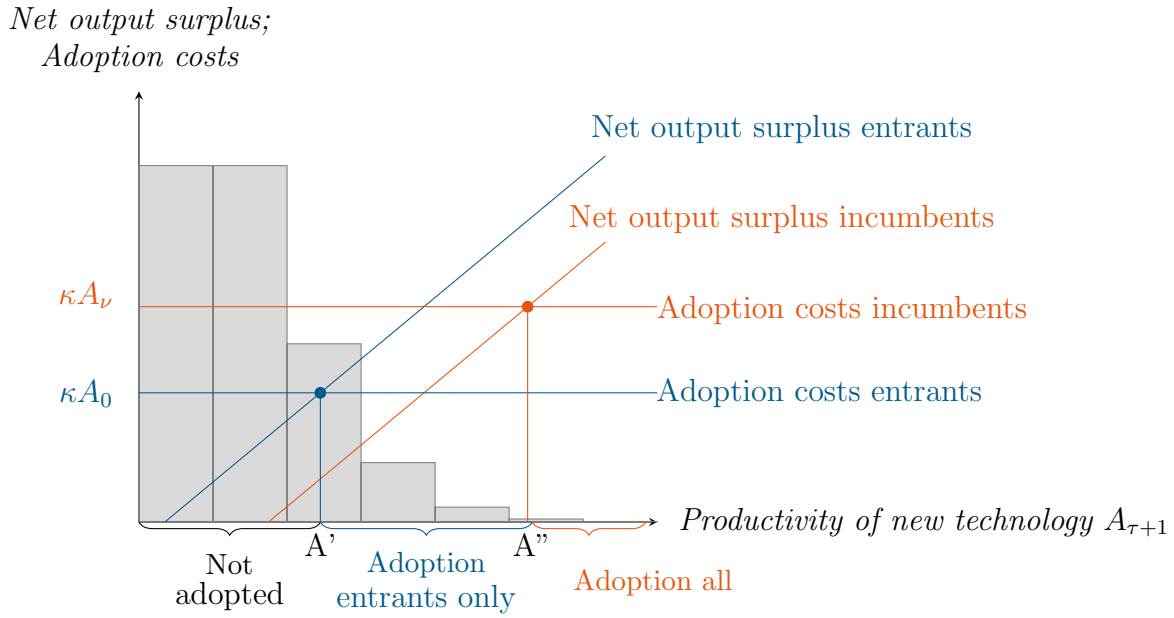
(1) Firm investments in new technology vintages decrease in periods without supply of young labor market entrants and reach standard levels thereafter.

When the entrant cohort is absent, firms invest in the new technology vintage if and only if the new vintage is productive enough to make retraining incumbents profitable. For productivity levels $A' \leq A_{\tau+1} < A''$, firms reduce their capital investments in this period. When entrants are available again, firms face the standard maximization problem, evaluating whether or not to adopt the following vintage. Hence, there is no overshooting of investments following the missing entrant cohort. This prediction remains qualitatively unchanged when allowing for wage responses to the entrant supply shock because entrants' wage increases always raise the

[10] Incumbents have an additional disadvantage over entrants, the so-called horizon channel: they will not remain at the firm since they only live for two periods, hence $\rho = 0$. The point of the framework is to show that even absent the horizon channel, firms choose to complement new technologies with entrants over incumbents. This opportunity-cost-productivity channel is the more empirically applicable one.

[11] In the case of increasing and convex capital adjustment (training) costs, as the literature typically assumes, highly productive technologies are also only adopted by entrants because training costs are prohibitively large for incumbents, see Appendix F.

Figure 2: Firms' costs and benefits of technology adoption



Notes: Profitability of assigning new technology-using task to entrants versus incumbent workers. The histogram shows the distribution of the productivity increase related to the new technology vintage.

productivity threshold above which a new technology is adopted.[12,13]

(2) Only firms with sufficiently high trainee retention rates decrease investments when young labor market entrants are missing.

The higher a firm's retention probability ρ, the more profitable it is to train and adopt a new technology. Firms with retention rates below a certain threshold decide not to train their entrants but utilize them for low-productivity tasks. Consequently, only firms with a retention rate high enough to train their entrants reduce technology adoption if entrants are missing.

(3) Investments decrease when young labor market entrants are missing only when the new technology vintage is sufficiently more productive than previous vintages.

For productivity levels of the new vintage $A_{\tau+1} < A'$, firms do not adopt the new technology even in presence of entrants.

[12] An increase in entrants' market wage corresponds to a reduction in θ, i.e. the productivity share captured by the firm. With decreasing θ, firms' returns to training and adoption decrease, corresponding to a lower LHS of equation (4), i.e. a downward shift of the entrants' net output surplus in Figure 2. Consequently, the threshold productivity above which technology adoption is profitable for entrants, A', increases and technology adoption decreases when entrant supply is reduced. If entrant supply is inelastic, for example because it is fixed by birth cohort sizes, raising wages may only attract experienced workers with higher baseline productivity A_τ. In this case, the opportunity costs and net output surplus of retraining these replacement workers is higher than in the no-shortage scenario, further raising the technology adoption threshold A'.

[13] The purpose of this paper is to focus on the task creation effect of new technologies. Allowing for technology-induced automation creates incentives for firms to adopt labor-saving technologies when trainees are scarce and their wages high. Angelini (2023) incorporates both channels in a macro model and finds that, for high levels of young worker shortages, the technology-depressing effect dominates.

3 The German vocational training system and the education reform

Below, I describe the German vocational training system, frequently studied in economic research (e.g. Acemoglu & Pischke, 1998; Dustmann & Schönberg, 2009, 2012), and the education reform.

3.1 The German vocational training system

Vocational training is a key component of both the German education system and labor market, with approximately 60% of the working population having undergone such training (Sample of Integrated Labour Market Biographies, own calculations). Vocational trainees are regarded as future middle-skilled professionals and work in occupations that typically require bachelor's or associate's degrees in other countries, such as the U.S.

Adolescents usually start vocational training after graduating from one of the following three high-school tracks: the basic track (*Hauptschule*, 9 years of schooling), which qualifies students for vocational training in blue-collar occupations; the intermediate track (*Realschule*, 10 years), which prepares students for any vocational training, including training in white-collar occupations (media, financial services, or retail occupations, in addition to manufacturing and technical occupations); or the upper-track (*Gymnasium*, 12 or 13 years) which is required for university studies. At the time, approximately a third of the upper-track school graduates choose to undergo vocational training,[14] such that in 2000, 16% of trainees had a university entrance qualification (*Abitur;* Federal Institute for Vocational Education and Training, 2009). Nowadays, these shares are higher (Ministry of Education & Research, 2025). Trainees rarely move or commute to their workplace: only 2% of vocational trainees move across states for their vocational training (SOEP, own calculations) and only 5% commute across states (LIAB, own calculations). After completing vocational training, which typically lasts three years, 43% of trainees in East Germany in 2001 remained at their training company, with large variation across industries, e.g. 75% in the banking and insurance sector (Ministry of Education & Research, 2004, for East Germany, 2001). Again, these shares are significantly higher nowadays.

Vocational training in Germany is commonly provided within the dual system, which combines on-the-job training at a firm (3-4 days per week) with vocational schooling (1-2 days per week). This paper exclusively focuses on the on-the-job training part. Trainees are hired by their training company, and receive a wage that is usually subject to collective bargaining agreements and low.[15]

[14]There were approximately 200,000 university entrants and 100,000 vocational training entrants with university qualification in 2000 (Federal Statistical Office, Genesis-Online, 2022c; Federal Institute for Vocational Education & Training, 2002). Similarly, Heine et al. (2005) report that 28% of upper-track graduates from 1999 had enrolled in university studies six months after graduation, while 21% had started vocational training. 32% were in civil or military service, hence pursuing vocational training or higher education with one year delay.

[15]The average monthly gross compensation agreed by collective bargaining was €555 in 2000 (Federal Insti-

Vocational training is comparable to on-the-job training in other countries with two notable exceptions: First, in addition to on-the-job training at the firm, trainees receive vocational schooling, which equips them with skills that may be external to the firm. Second, regularly updated training curricula ensure that the training content includes up-to-date technical skills (Salomons et al., 2025).

3.2 The reform

Prior to German reunification in 1990, upper-track school graduates underwent 12 years of schooling in East Germany and 13 years in West Germany. After reunification, in an effort to align the two education systems, Brandenburg switched to 13 years in 1994,[16] while Saxony and Thuringia retained the 12-year system. Saxony-Anhalt and Mecklenburg-Western Pomerania transitioned from 12 to 13 years with the graduation cohort of 2001: the reform that I exploit in this paper. I assign Saxony-Anhalt and Mecklenburg-Western Pomerania as treated states and the other four East German states (Brandenburg, Saxony, Thuringia, East Berlin) as control states. The education reform was decided in May 1996 in Mecklenburg-Western Pomerania and in January 1998 in Saxony-Anhalt.[17] By lengthening the years of schooling, the reform increased the level of education. More importantly, because the last cohort completing 12 years graduated in 2000 and the first cohort completing 13 years graduated in 2002, the reform resulted in a one-time missing upper-track school graduation cohort in spring 2001. Figure 3 depicts the sharp drop in the number of upper-track school graduates in 2001 – in Mecklenburg-Western Pomerania from 6,400 to 300, and in Saxony-Anhalt from 9,400 to 400 – while the figures remain relatively constant in the control states.

The education reform was initiated by the Social Democratic Party (SPD), which entered the government in both treated states in 1994.[18] The SPD advocated for 13 years of schooling to promote equal opportunity, while its main opponent, the Christian Democratic Union (CDU), pushed for a 12-year system to enhance the international competitiveness of German school graduates. I provide evidence that the governance of the SPD, or related policy or socio-economic changes, are unlikely to confound the effect of the education reform by comparing several state metrics between treated and control states before and after the reform, as well as between state-periods governed and not governed by the SPD, in Appendix B. I also discuss why concomitant investment tax programs are unlikely to have confounded the trainee supply shock.

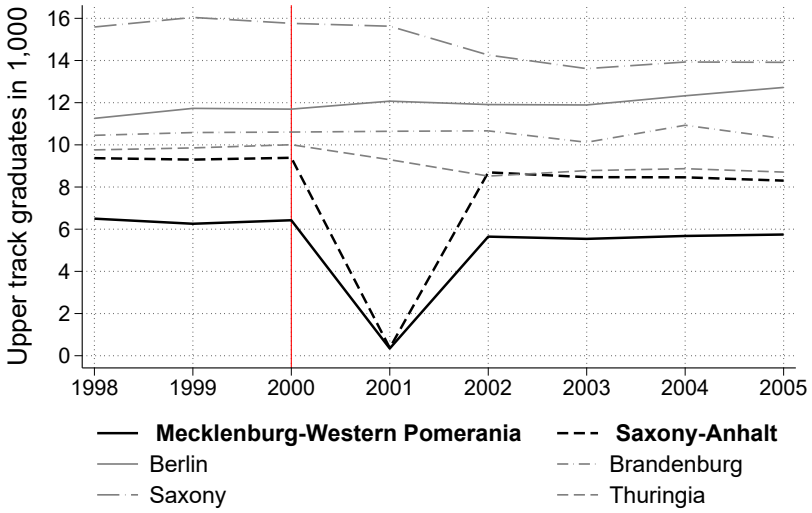
tute for Vocational Education & Training, 2022).

[16] Results are robust to excluding Brandenburg from the set of control states.

[17] For more information on the education reforms, see Kühn et al. (2013) and Helbig & Nikolai (2015). Between 2007 and 2013, all German federal states adopted to the 12-year system, with Saxony-Anhalt making the change in 2007 and Mecklenburg-Western Pomerania in 2008. To avoid potential confounding effects from these changes, this study ends in 2006.

[18] Brandenburg was always governed by the SPD since reunification which is why they changed from 12 to 13 years immediately after reunification.

Figure 3: The missing school graduation cohort



Notes: Source: Federal Ministry of Education & Research (2022).

How does the missing school graduation cohort translate into the labor market? Usually, one-third of the upper-track school graduates start vocational training. The missing school graduates of spring 2001 are hence expected to result in a missing entry cohort of highly educated trainees in fall 2001, and to reduce the stock of highly educated trainees for three consecutive years, 2001–2003, given that vocational training typically lasts three years. At that time, males in Germany had to do military service of 10 months when reaching the age of 18, partly postponing the missing entry and prolonging the reduction in the supply of trainees by one year, i.e. until 2004.

Dorner et al. (2024) analyze how the education reform affects trainee employment, finding sizable and significant reductions. I use the temporary shock in trainee supply to study subsequent effects on firm investments. I focus on those upper-track school graduates who subsequently start vocational training instead of those going to university because vocational trainees are less likely to postpone their labor market entry and to move or commute across federal states, endorsing the identification strategy. Note that the labor supply shock is unlikely to be confounded by a labor demand shock: the overall number of consumers remains unchanged, only the composition adjusts. Since trainees earn low wages, consumption patterns of trainees versus students is unlikely to have caused relevant demand changes.

Perhaps improving the credibility of the research design, both treated states are economically fairly different: Mecklenburg-Western Pomerania, located in the northeast of Germany along the Baltic Sea, is a predominantly rural and sparsely populated federal state. Its economy is defined by small and medium-sized enterprises in agriculture, maritime industries, mechanical engineering, and tourism. Saxony-Anhalt, situated in central Germany features a more urban environment. It is characterized by the chemical industry, mechanical engineering, and automotive supply. Both states, as well as control states, are characterized by excess trainee

supply and high unemployment rates during this period, namely 17.8% in Mecklenburg-Western Pomerania and 20.2% in Saxony-Anhalt in 2000, see Figure A3, Panel B. Nonetheless, 16% of the firms at the time report to experience severe skill shortages (LIAB), in particular in ICT skills (e.g. Mitteldeutsche Zeitung, 2001).

4 Firm panel data

The analysis is based on the Linked-Employer-Employee-Data of the Institute for Employment Research (IAB), the LIAB-QM, which combines the annual representative IAB Establishment Panel survey with administrative employment information of all employees at surveyed establishments. The number of surveyed establishments has risen from 4,000 in 1993 to 15,000 in 2023. The survey is conducted at the workplace level, enabling the distinction between treated and untreated establishments based on their location.[19] I use the terms "firm" and "establishment" interchangeably for simplicity.

Employment information is based on administrative records reported to social security insurance. While employment information is reported in the cross-sectional data as of June 30 each year,[20] most vocational training programs start in the fall, such that trainees starting in fall each year appear in the data with a lag of one year. To account for this, I use employment figures from June of the following year to represent employment levels in fall of the current year. The data provides a distinction between trainees and workers who have completed their training program, in addition to wages and employment status. Also, information on schooling allows me to distinguish "highly educated" from "low-educated" trainees, i.e. trainees with a university entrance qualification and those with a lower schooling degree, respectively. This is important since the education reform directly affects highly educated trainees only.[21]

I restrict the data in the following way. First, I limit the data to firms in East Germany including Berlin, since firms in East Germany are potentially not comparable to firms in West Germany. Results remain very similar when including firms in West Germany as control firms. Second, I exclude firms in the health/education/social service sectors, because vocational training in most related occupations is entirely school-based; as well as firms in agriculture and arts/sports based on their extremely low share of highly educated trainees. Third, I drop very small firms, defined as those with fewer than five full-time employees subject to social security contributions in 2001, because small firms tend to exhibit volatile investment behavior. The results remain robust when including these small firms. Fourth, I only keep firms that invest at least once throughout the entire time period out of the fear that never-investing firms may

[19] The data does not allow to assign establishments to parent companies, precluding a within-company cross-establishment design.

[20] I use the cross-sectional data which comprises employment spells that encompass June 30 of each year. The longitudinal data, encompassing all spells, is not suitable because it is only available from 2009 onward.

[21] I use the harmonized version of the schooling variable based on the imputation procedure by Thomsen et al. (2018) and Fitzenberger et al. (2006).

misreport their investments. Again, results are robust to this restriction. The resulting firm panel 1997–2006 is unbalanced: in the final dataset, 47% of firms observed in 1997 are still present in 2006. In Section 7.5, I show that the results remain robust to restricting to a balanced set of firms for the years 1997–2003.

The final sample comprises 2,246 distinct firms, of which 744 are treated (416 in Saxony-Anhalt and 328 in Mecklenburg-Western Pomerania). Table 3 shows summary statistics. In sum, all firms cover almost 270,000 workers per year, amounting to approximately 3.6% of the East German workforce in a year.[22] I observe 15,010 trainees on average across years, of which 2,456 (16%) are highly educated. In 78% of the firm-by-year observations, no highly educated trainee is employed, and 62% of firms never employ a highly educated trainee over the entire time window 1997–2006.

Table 3: Summary statistics – Full data sample

	Mean	SD	Min	Max	Yearly sum
# workers	163.75	452.38	1	12,133	267,389
# trainees	9.19	52.10	0	3,181	15,010
# highly educated trainees	1.50	8.84	0	461	2,456
No highly educated trainee	.78	.41	0	1	1,276
No highly educated trainee ever	.62	.49	0	1	1,013

Notes: Summary statistics in the full data sample (training and non-training firms, 1997–2006). SD: standard deviation. Yearly sum across all observed firms. Entry-level hires are defined as workers below age 30 who are either undergoing vocational training or entering social-security-covered employment for the first time.

Training versus non-training firms. Since the reform predominantly affects highly educated trainees, I focus on training firms, defined as firms with at least one highly educated trainee in fall 1995, 1996, or 1997 (based on employment information from June 1996, June 1997, and June 1998). These years are chosen to minimize the chance that a firm's training status is affected by firms' anticipation of the trainee shortage. This divides the sample into 571 training firms and 1,675 non-training firms. Training and non-training firms are fundamentally different. Training firms are more than four times as large in employment, operate more often in the business service and public administration sector, and less often in construction, retail/motor vehicles, and hospitality and other services.

Descriptive statistics based on the final matched sample of training firms (for a description of the matching procedure, see Section 5) are given in Table 4. On average, firms employ 294 workers. Vocationally trained workers are with an absolute number of 222 (76%) the majority. A training firm employs on average 4.13 highly educated trainees. While the employment share

[22]The average yearly working population in East Germany from 1997 to 2006 was 7.43 million according to Statistisches Landesamt (2023).

of highly educated trainees is thus fairly small (1.4%), highly educated trainees make up 20.1% of firm entry level hires.

Investments. Each year, firms are asked whether they invested in four investment types in the last year: (1) production facilities, plant and equipment, furniture and fixtures (henceforth "production equipment"), (2) communication technology, electronic data processing (henceforth "ICT"), (3) real estate and buildings (henceforth "real estate"), and (4) means of transport, transportation systems (henceforth " transport"). Unfortunately, information on investment types only distinguishes between "no" or "any" investments in this category. Since the subsequent analyses reveal that the intensive investment margin is much more affected than the extensive investment margin, these variables will not be used to study technology adoption. Instead, if the firm reports any of the above investments, the firm is surveyed on the total amount of annual capital investments.

I scale investments by firm employment to account for the fact that large firms tend to invest more. To curtail the impact of extremely large investments, I winsorize values that exceed the third quartile plus four times the interquartile range.[23] Across firm-year observations, 92% show investments; the mean investment is €5.4million, or €14,500 per worker, while the median investment is €1.4million, or €7,600 per worker. The investment distribution is still highly right-skewed and has a large standard deviation of €14,460 per worker.

Table 4: Descriptive statistics – Training firms, matched sample

	p5	Median	Mean	p95	SD
# workers	21	199	294	828	353
# workers with+in vocational training	18	155	222	650	262
# highly educated trainees	0.00	1.00	4.13	17.00	6.84
% highly educated trainee hires in entry-level hires	0.00	8.33	20.14	75.00	28.14
Any investments (1/0)	0.00	1.00	0.92	1.00	0.27
Total investments in €1,000	0	1,425	5,410	19,388	11,174
Investments per worker in €1,000	0.00	7.60	14.54	37.50	14.46
Projected technical status of machinery	3.00	4.00	3.92	5.00	0.63

Notes: Averages across 1997–2000. p5 – 5th percentile; p95 – 95th percentile; SD – Standard deviation.

While investments certainly include non-technology-related items, I use them as main outcome because they strongly correlate with several indicators of technological progress: First, in 2000, firms report to allocate on average 25% of their investments to ICT. Since this information is only available from 2000 onward, I do not use it as a main outcome of interest. Second, firms with higher investment levels have a significantly better self-assessed technical status of machinery, see Table 5, column 1. Investment increases over time within firms are

[23]The factor four is chosen to reflect the fact that investment distributions tend to be highly right-skewed. Results are robust to using other values or trimming the uppest percentile, see Section 7.5.

also significantly positively correlated with improvements in the technical status (column 2). Third, investments are positively associated with the incidence of firm product innovation,[24] both in the cross section and over time (columns 3 and 4). Last, investments are associated with less routine employment (columns 5 and 6), indicating that investments partly belong to the group of routine-replacing computer-based technologies that were heavily implemented in this decade.[25]

Table 5: Correlations between firm investments and other firm indicators

	Technical status		Product innov. (1/0)		Routine employment	
	(1)	(2)	(3)	(4)	(5)	(6)
Inv. per worker in €1,000	0.006***	0.004***	0.259***	0.245***	-0.227***	-0.028**
	(0.001)	(0.000)	(0.039)	(0.056)	(0.029)	(0.010)
Year FE + Firm FE		✓		✓		✓
N	21,591		17,026		43,681	

Notes: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors clustered at the firm level. Technical status on a scale from 1 to 5; Routine employment on a scale from -1 to 1.

Technical status of machinery. To better get at firm technology adoption, I employ two strategies. First, I infer firm exposure to technological change based on firms' sectors. Second, I use the technical status of a firm's machinery as additional outcome. Firms are asked to assess the overall technical status of their machinery compared to other establishments in the same industry on a scale from 1 ('completely out-of-date') to 5 ('state-of-the-art'). Reported values are plausible: in the cross section, larger firms and firms with higher AKM fixed effects report a significantly better technical status, also within broad industry groups (Table A1).

In principle, a better technical status of machinery could reflect the good condition of existing technologies, while the goal is to study the adoption of *new* technologies. Convincingly, I find that increases in firms' technical status are accompanied by more organizational changes,[26] see Table A1. This suggests that changes in the technical status indeed capture the

[24]Product innovation is a binary variable equal to one if a firm introduced a completely new service or product ("Have you added a completely new service or product [...]"), or substantially improved it in a given year ("Has your company improved or further developed a service or product that you already offered [...]?"). This variable is available for the years 1997, 2000, 2003, and 2006. Process innovations are only captured from 2007 onward and are therefore not included.

[25]Routine employment is computed as the firm's average share of routine tasks minus the share of non-routine tasks, based on the occupations held by its workers and time-constant occupation-specific task shares. To measure routine employment, I follow Spitz-Oener (2006) and use the 1999 BIBB-BAuA Qualification and Career Survey, which asks workers about the tasks performed in their jobs. By aggregating this information within detailed occupations and five task categories, I compute for each occupation and task category the share of tasks performed among all tasks in this category. For each firm, I then compute the average task share across all workers in each year, and compute routine employment as the share of routine cognitive and routine manual tasks minus the share of non-routine analytic, non-routine interactive and non-routine manual tasks. Changes in firms routine employment can therefore only stem from changes in occupational employment.

[26]I follow Battisti et al. (2023) and define organizational change on a scale from 0 to 4 by adding up the following binary indicators: restructuring of departments or areas of activities, downward shifting of responsibilities

introduction of new technology.

The technical status is a stock variable, unlike investments, such that current investments affect firms' future technical status. To purge the analyses from this time lag and convert the stock into a flow variable, I construct the variable "projected technical status" which attributes a firm's technical status back to the years in which the investments influencing that status were made. To get at the structure of the time lag, I regress a firms' technical status on the sequence of current and past investments, controlling for year and firm fixed effects, see Table A2. The investments of the past three years mainly affect a firms' current technical status. I therefore compute the projected technical status as the weighted sum of the firm's technical status over the next three years, with weights equal to the estimated regression coefficients.[27]

5 Event study approach

The identification strategy exploits the quasi-random assignment of the education reform to firms, inducing exogenous variation in the supply of upper-track school graduates across firms and years. I compare training firms in treated and control states before and after the reform in a difference-in-differences (DiD) event study design by estimating the following specification:

$$Y_{jt} = \sum_{\substack{t=1997,\\ t\neq 2000}}^{2006} \beta_t(\text{Treat}_{s(j)} \times \text{Year}_t) + \psi_t + \phi_j + \epsilon_{jt} \tag{4}$$

where Y is one of several outcomes such as investments, j denotes the firm, $s(j)$ the firm's federal state, and t the calendar year. I study firms up to 2006 because a different education reform affects trainee supply from 2007/2008 onward. Treat is a binary variable with Treat = 1 if the firm is located in a state undergoing the reform and zero otherwise. ψ_t captures calendar-year fixed effects. Firm fixed effects ϕ_j capture time-constant level differences between firms. The vector $\beta_t, t \geq 2001$, includes the coefficients of interest, namely the difference in firm outcomes in treated states compared to control states following the reform in 2001 purged from the baseline difference in 2000. The event study thus identifies the causal effect of a firm facing a state-wide negative trainee supply shock.[28] I follow Roth et al. (2023) for the most recent suggestions for DiD estimations. Since treatment is not staggered, potential biases common to two-way fixed effects estimators in a staggered setting (e.g. Goodman-Bacon, 2021) are irrelevant here.

and decisions, introduction of team work/working groups with their own responsibilities, and introduction of units/departments carrying out their own cost-benefit calculations.

[27] In particular, I compute Projected technical status$_t$ = $\beta_{t-1}/(\beta_{t-1}+\beta_{t-2}+\beta_{t-3})$Technical status$_{t-1}$ + $\beta_{t-2}/(\beta_{t-1}+\beta_{t-2}+\beta_{t-3})$Technical status$_{t-2}$ + $\beta_{t-3}/(\beta_{t-1}+\beta_{t-2}+\beta_{t-3})$Technical status$_{t-3}$ with β_{t-1} to β_{t-3} the regression coefficients from Table A2. A projected technical status of 4.5 in 2001 may, for example, by the mix of a technical status of 4 in 2002 (weighted with an importance of 50%), of 5 in 2003 (weighted with an importance of 35%) and 5 in 2004 (weighted with an importance of 15%).

[28] Note that this is different to the causal estimate of a firm employing x fewer trainees.

For brevity, I also estimate the equivalent DiD specification, assigning the years 1997–2000 as pre-period and the years 2001–2003 as post-period (dropping the years 2004–2006):

$$Y_{jt} = \delta(\text{Treat}_{s(j)} \times \text{Post}_t) + \xi \text{Post}_t + \lambda_j + u_{jt} \tag{5}$$

where the coefficient of interest is δ, the difference in the average post-period outcomes compared to the average pre-period outcomes for treated compared to control firms.

Matching. Treated training firms may differ from control training firms in aspects that expose them to different potentially confounding factors. To ensure that treated and control firms are comparable, and therefore exposed to similar potential confounders, I match firms based on their pre-treatment characteristics in two steps. First, I match firms within nine industry groups. Results are robust to matching on more detailed industry classes. By matching within industries, the reform effects are devoid of distortions arising from industry-specific shocks. Second, within industries, I perform a Mahalanobis distance matching with replacement. The matching variables include average pre-treatment log overall employment, average pre-treatment employment shares of highly educated trainees, and the increase in pre-treatment log employment between 1997 and 2000.[29] To avoid excessively limiting the sample size while ensuring good compatibility, I keep the three control firms with the smallest Mahalanobis distance for each treated firm and subsequently discard the worst 10% of all matches. Results are robust with respect to both aspects. I also report the results without matching in Section 7.5 which are very comparable.

Table 6 shows characteristics of treated compared to control firms for the unmatched and matched sample. Prior to matching, treated firms are significantly smaller than control firms. The matching works well in eliminating differences in observable firm characteristics, both targeted and non-targeted ones.

The identification of the causal effect relies on three main assumptions.

Assumption 1 – Parallel trends. First, I assume that firm outcomes in treated states would have evolved in parallel to those in control states in the absence of the reform. Parallel trends prior to the shock reveal that this assumption is likely to hold, see Sections 6 and 7.

In Appendix B, I show that key state metrics such as unemployment, population size, education expenditure, public debt and public investments do not change differentially in treated compared to control states post 2000. Likewise, I argue that the concomitant investment tax programs studied in Lerche (2025) and Siegloch et al. (2025) are unlikely to confound the effect of the education reform. Moreover, one might be concerned that the introduction of the euro in 2002, the German Hartz reforms over 2003–2005, the bust of the dot-com bubble in

[29] This requires firms to be present both in 1997 and 2000, reducing the extent of panel attrition: 47% of matched training firms present in 1997 are still present in 2006.

Table 6: Matched training firms – Balancing table and descriptive statistics

	Unmatched		Matched	
	Δ Mean	SE	Δ Mean	SE
	A. Targeted variables			
Log employment	-0.37	0.13**	-0.13	0.12
Δ log employment	-0.09	0.05	-0.04	0.03
% highly educ. trainees in total empl.	-2.04	3.49	0.23	0.28
	B. Non-targeted variables			
# workers	-220	74***	-28	38
# highly educated trainees	-3.51	1.84	-0.29	0.72
Any investments (1/0)	-0.02	0.02	-0.00	0.02
Investments per worker in €1,000	1.13	1.19	0.53	1.47
Capital stock per worker in €1,000	48.77	28.78	35.81	34.42
Projected technical status	0.04	0.06	0.04	0.07
Number of firms	571		445	

Notes: Averages across 1997–2000. Δ log employment refers to the change in log employment between 1997 and 2000. SE: Standard error. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

2000, or China's accession to the World Trade Organization in 2001/2002 might confound the reform effect. However, my design accounts for this since these shocks likely affected treated and control East German states similarly, especially within industries. In addition, I perform several robustness checks addressing this concern, all of which successfully pass.

Assumption 2 – No anticipation. The second assumption is that firms did not change their behavior prior to the reform. Even though the reforms were decided in 1996 and 1998, the event study estimates show no evidence of anticipatory behavior.

Students may also have anticipated the reform. There was, however, very little scope for them to react: When the reform was decided, students of the missing graduation cohort were in grade 7 in Mecklenburg-Western Pomerania and grade 9 in Saxony-Anhalt. Since the school track was chosen after grade 6, it was not impacted by the reform. Also, school graduates may delay or accelerate the start of their vocational training in response to the shock. This would bias the estimates toward zero.

Assumption 3 – No spill-overs/SUTVA. Third, I assume that control states are not affected by the reform, and treated states are not affected by the absence of the reform in control states. This assumption is violated if trainees move or commute across states. The data allows me to identify cross-state commuting. Trainees in the affected states rarely commute across states (2.7% prior to 2001; 2.9% for workers with completed vocational training) compared to workers with a university degree (5.3%), and this share does not change in response to the

reform, see Section 6. To investigate whether school graduates move for their apprenticeship, I turn to the Socio-Economic Panel (SOEP) which tracks individuals from childhood onward. The cross-state trainee mobility rate is extremely low at 2.2%. Further, there is no instance of a highly educated trainee relocating to one of the treated federal states in the post-reform years 2001, 2002 or 2003 in the SOEP. However, if trainees moved or commuted from control states to treated states in response to the reform, this would bias the estimates toward zero.

Interpreting the reform as trainee supply shock. Beyond identifying the causal impact of the reform, I aim to attribute the effects on firm investments to the temporary drop in trainee supply. This requires that no other aspect of the reform affects investments.

Another aspect of the reform is the increase in the years of schooling which increased the education level of highly educated trainees, their age, and potentially the share of upper-track graduates who start vocational training. However, these two factors changed permanently, such that effect dynamics will help distinguish these permanent adjustments from the temporary trainee shortage. In addition, higher levels of education would, if anything, induce more, not less, investments, and therefore provide a lower bound of the effect.

One potential concern is that the supply shock may be accompanied by demand changes. However, the postponement in the start of vocational training is unlikely to meaningfully affect consumer demand. First, because the overall population size remains constant. Second, since trainees earn low wages, students' counterfactual demand–had they started vocational training instead of staying in school for an additional year–is unlikely to differ significantly. Turning to firm demand, low trainee wages also prevent a meaningful decrease in the firm wage bill when firms reduce trainee employment, making it unlikely to present a confounding channel.

Trainee distribution across firms. Even if the estimated parameters of interest, $\hat{\beta}_t$, identify the unbiased effect of facing a trainee shortage, they are subject to the realized distribution of trainees across firms. In particular, for the investment outcome, $\hat{\beta}_t$ are small if trainees are primarily missing in firms that would not have invested in the absence of the shock, and $\hat{\beta}_t$ are large if trainees are primarily missing in firms that would have invested in the absence of the shock. I address this concern in two ways. First, I additionally run the DiD analysis at the region-industry level and find very comparable investment effects. Second, in order to identify the effect on investments independent of the realized trainee distribution, I propose a complementary identification strategy in Appendix D: I predict the trainee distribution across firms based on a Bartik-style instrument of firms' pre-reform use of trainees and the state-level shift in trainee employment induced by the reform. This allows me to identify a different causal parameter, namely the effect of employing one fewer trainee. This analysis is, however, more demanding and subject to further assumptions, which is why my preferred identification strategy is the DiD event study design.

Inference. Standard errors are commonly clustered at the treatment level to account for cluster-level shocks (e.g. Abadie et al., 2023). Here, this would result in a small number of clusters, i.e. federal states. For valid inference with a small number of clusters, I follow Roth et al. (2023), and assume that state-specific shocks are small relative to idiosyncratic firm-level error terms, justifying clustering at the firm level, potentially at the expense of a small violation of parallel trends. The data strongly support this assumption. For the main outcome, investments per worker, the within-state standard deviation of the error term is 91.4% of the overall standard deviation, indicating little within-state correlation (see Table A3). During the shock years 2001–2003, this share rises further to 93.0–98.4%, suggesting even less state-level clustering. I therefore cluster standard errors at the firm level. Additionally, I perform permutation (Fisher randomization) tests, comparing the t-statistic of the treatment effect for the actual treatment assignment and all permuted treatment assignments across federal states. Finally, I also report alternative confidence intervals based on wild t-bootstraps clustered at the state level (Cameron et al., 2008).

6 Bite of the reform

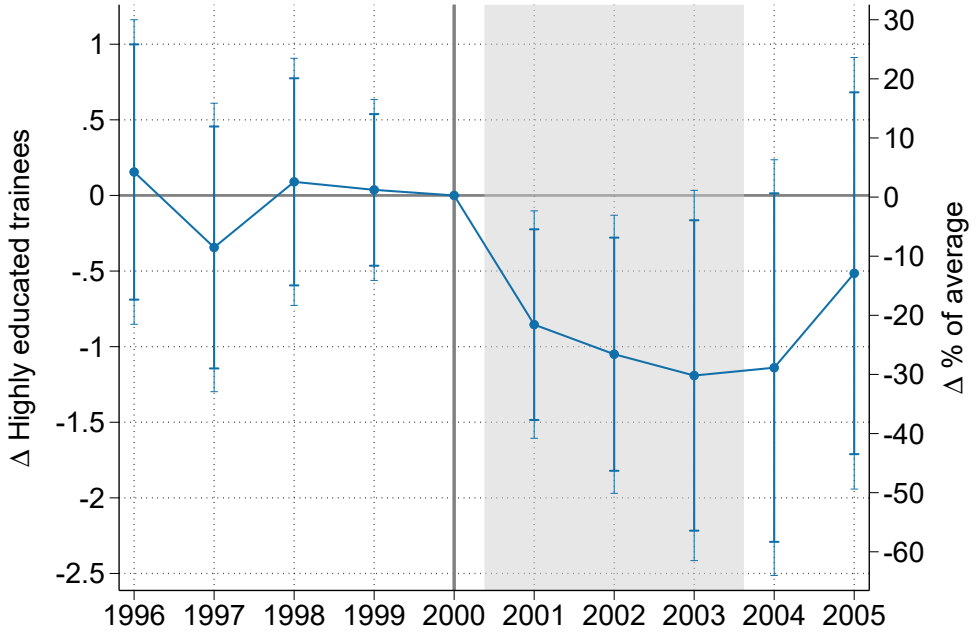
6.1 Effect on trainee employment

Figure 4 displays the difference in the employment stock of highly educated trainees in treated compared to control firms based on results of the DiD event study regression outlined in equation (4). Endorsing the identification strategy, firms' highly educated trainee employment evolves in parallel in control and treated states prior to 2001. This is expected, as trainee supply almost exclusively consists of fixed-sized school graduation cohorts, making it highly inelastic. During the years 2001–2004, employment of highly educated trainees is significantly lower in treated compared to control firms. Considering the typical training duration of three years plus one year delay for those going to military or social service, these are precisely the years the majority of the missing school graduates would have undergone vocational training. Absolute employment of highly educated trainees decreases by 1.01 on average in affected years or 27% of the pre-shock average. The effect corresponds to a reduction in the stock of trainees by a quarter, i.e. one out of four cohorts for four years, or, put differently, to a decline by one entire cohort. While a reduction of one worker may appear small, it represents a 20% reduction in the firm's entry level hiring. Consistent with the timeline of the shock, the employment gap shrinks in 2005 and becomes statistically insignificant.

6.2 Wage and worker substitution effects

The administrative labor market data allow me to study firms' adaptation strategies, such as changes in trainee wages, or worker substitution patterns. To investigate such effects, I employ

Figure 4: Effect on employment of highly educated trainees



Notes: Event study coefficients of the interaction terms Treat × Year plus 90% and 95% confidence bands. Based on equation (4). Based on employment data from June 1997 to June 2006 to capture employment as of fall 1996 to fall 2005. Left axis: Difference in the absolute employment stock of highly educated trainees per firm between treated and control firms. Right axis: Difference expressed in percent of the average pre-treatment absolute employment stock of highly educated trainees among treated firms. Standard errors clustered at the firm level. N=3,609.

the corresponding DiD specification given in equation (5). Results are given in Table 7.

There is no evidence of an increase in the wages of highly educated trainees in response to the negative supply shock (column 2). This is in contrast to what standard economic theory predicts. To understand the absence of any wage effects, it is important to keep in mind that trainee wages are very rigid —often set by collective bargaining agreements—, that the shock was only temporary, and that the supply of highly educated school graduates is fixed by cohort size, giving very little scope for wage increases to induce more employment.[30]

There is no statistically significant substitution for missing highly educated trainees. Prominent candidates as substitutes are low-educated trainees, highly educated trainees from other federal states, and workers who have already completed vocational training. However, I find no statistically significant effect on the number of low-educated trainees (column 3). In con-

[30]In fact, there are a multitude of reasons that potentially explain the lack of a wage adjustment. First, firms likely shy away from increasing wages in response to a temporary shock because downward rigid wages will impede a subsequent wage decline once the supply shock dissipates. Second, trainee wages in Germany are set at a very low level and are paid only throughout the three-year vocational training period. Hence, even a hypothetical doubling of training wages would result in negligible changes in absolute lifetime income. Instead, trainee supply responds to anticipated post-training wages (Neuber-Pohl et al., 2023) that remain unchanged in the present case. Third, the vast majority of training wages are set by collective bargaining agreements, and even firms that are not part of those agreements tend to base their wages on such agreements. Of course, firms could deviate upwards. In that case, works councils, which would have to approve training wages in large firms, would likely oppose unequal treatment of trainees. Finally, this finding is in line with the results by Muehlemann et al. (2022) in the case of the opposite, positive supply shock of trainees.

Table 7: DiD Results – Wage and worker substitution effects

	Shock	Absence of wage and entry responses						
	# highly educated trainees	Log wage highly educ. trainees	# low-educated trainees	% highly educ. commuting trainees	Log highly educ. VT employment	Internal retraining	Log employment	Log wage
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treat × Post	-1.01*	0.01	-0.94	2.08	0.07	-0.04	-0.06	-0.00
	(0.47)	(0.03)	(1.42)	(2.66)	(0.07)	(0.06)	(0.06)	(0.01)
Mean dep. variable	4.14	2.99	11.65	5.19	2.25	0.47	5.03	4.21
N	3,133	2,025	3,133	1,364	2,930	1,552	3,133	3,133

Notes: DiD coefficients based on equation (5). Pre: 1996–2000 (based on data from June 1997–June 2001). Post: 2001–2003 (based on data from June 2002–June 2004). Standard errors clustered at the firm level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Mean dep. variable: Average outcome in 1997–2000. Column 4: a commuter is defined as a person living and working in two distinct federal states. This variable is available from 1999 onward. Column 6: Internal retraining is the number of retraining incidences per worker (number of workers fixed to pre-treatment levels). VT: completed vocational training. Log employment and log wages refer to all workers excluding trainees.

sequence, overall trainee employment also drops. The low substitutability between low- and highly educated trainees, in line with Muehlemann et al. (2022), is likely related to distinct skill sets, the specialization in different occupations, and the unchanged demand for low-educated trainees against a fixed supply of school graduates. Also, there is no statistically significant increase in cross-state commuting of highly educated trainees from a different federal state (column 4),[31] supporting the SUTVA assumption. Column 5 shows that the employment of highly educated workers who have completed their training program does not increase in response to the trainee shortage either, indicating that already trained workers are not suitable substitutes for trainees. Likewise, none of the unreported coefficients for the outcomes of log number of workers below the age of 30, workers with vocational training below the age of 30, or highly educated workers with vocational training below the age of 30 are significant. Firms may also increase retraining of incumbent workers to overcome skill shortages caused by the negative trainee supply shock. Again, I find no evidence for this (column 6). Overall firm size and log wages remain unaffected as well, see columns 7 and 8.

I confirm the employment effects in a different dataset, SIEED, consisting of a representative 1% establishment sample with information from social security. The data source is the same as for the main dataset, but the sample is different and larger, and survey outcomes such as investments are not observed. Table A4 reveals that DiD estimates are equivalent in this sample: employment of highly educated trainees drops by 29%, and coefficients for trainee wages and other workers types remain insignificant. This larger dataset enables me to estimate the effects on highly educated trainee employment by occupation groups, see Table A5. I find significant employment declines in production (-62%) and business service occupations (-37%).

To sum up, the reform leads to a large decline in the employment of highly educated trainees that is not accompanied by higher trainee wages, and not compensated for with other worker

[31] The coefficient of interest captures potential increases in commuting into treated states in addition to potentially reduced commuting into control states, and thus provides an upper bound of the true effect.

types. According to the economic framework, the absence of entrants reduces the adoption of technologies for which retraining incumbents is not profitable. Even if the shock triggered wage and entry responses, adoption would remain limited, as retraining of replacement workers at higher wages would also be unprofitable.

7 Effect on firm technology investments

I now turn to the effect on firm total capital investments, which will be linked to foregone technology adoption below, interpreting them as the causal effect of a negative trainee supply shock. Figure 5 shows the difference in investments per worker between treated and control firms over time. Convincingly, there is no pre-trend, supporting the assumption that investments in treated states would have evolved in parallel with investments in control states absent the shock. The fact that firms do not adjust their investment behavior in anticipation of the trainee shortage may be due to several reasons: First, firms are unaware of the impending shortage: According to the IAB establishment survey, in 2000, treated training firms did not expect significantly elevated shortages of young workers for 2001 and 2002 compared to control training firms. Second, if trainees and investments are perfect complements, it is optimal for firms to not adjust investments prior to the shock given that trainee employment was stable before 2001. Third, opposing anticipatory behaviors across firms may cancel each other out: in the years before the shortage, some firms may have already foregone technology adoption of new technologies, which require a critical mass of trainee cohorts to be trained on, knowing that only a few trainee cohorts would arrive before the trainee shortage. Other firms, however, may have invested more in anticipation, leading to net zero effects.

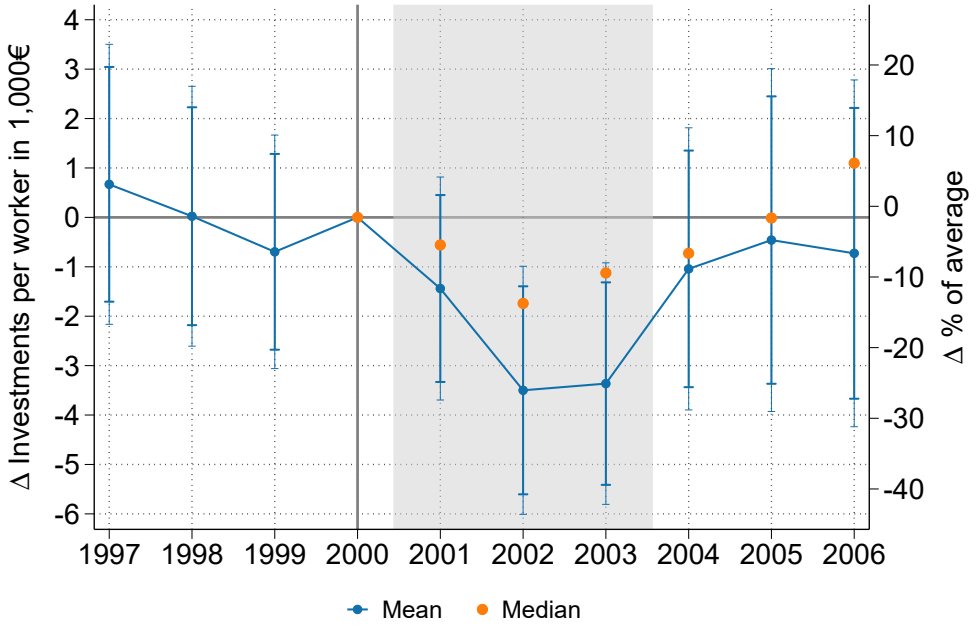
Firms time their investment to coincide with trainee supply: investments per worker significantly drop during the time of the trainee shortage in treated compared to control firms. This is the key finding of the paper and confirms the first prediction of the economic framework. The decrease is temporary: investments of treated firms return to those of control firms in 2004 to 2006, suggesting a link to the short-term supply shock, rather than a response to a new long-term reality in which trainees might be more highly skilled, older, or otherwise different.

Firms reduce investments by €2,740 per worker in 2001 to 2003, corresponding to a decline of (2,740/13,380=) 20.5% of pre-treatment average investments among treated firms, or 0.9% of the capital stock per worker.[32] The effect on investments per worker is not an artifact of changes in the denominator: the estimated effect is similar when fixing the denominator to firm size in 1997 and when using a time-variant denominator, see Table 8, columns 1 and 2.

Investment levels of treated firms reach those of control firms after the shock but do not

[32]To impute the capital stock, I follow Müller (2008, 2017) and exploit information on total investments, the proportion of net investments, dummy variables representing the four investment types, and industry. In particular, I establish a starting value for the capital stock using investments in the first three observed years (1996, 1997 and 1998 at the earliest) and project the capital stock for subsequent years using investment information and sector-specific depreciation rates.

Figure 5: Effect on investments per worker in €1,000



Notes: Event study coefficients of the interaction terms Treat × Year plus 90% and 95% confidence bands. Based on equation (4). Standard errors clustered at the firm level. Left axis: Difference in investments in €1,000 divided by total employment in 1997 between treated and control firms. Right axis: Difference expressed in percent of the average pre-treatment investments per worker among treated firms. N=3,609. Mean: OLS estimates. Median: Median of the treatment effect distribution.

exceed them. Why do treated firms not invest more in the years after the shock? Several reasons could explain this. If trainees and capital are perfect complements, the normal-sized post-treatment trainee cohorts only lead to normal-sized investments. Investment frictions—organizational or financial—may inhibit treated firms to carry out the sum of forgone investments and current investments simultaneously, leading to treated firms always lagging behind control firms. Similarly, if investments trigger follow-up investments, a delay in the initial investment naturally causes delays in subsequent ones. Also, treated firms may decide to skip one technology vintage altogether, corresponding to the formulation of the economic framework.

7.1 Size of the effect

The estimated investment decline is larger than the "mechanical" effect of reducing investments in proportion to trainee employment. The BIBB-Cost-Benefit-Survey allows me to explicitly determine the "mechanical" costs: In 2000, East German firms spent an average of €487 per trainee per year on equipment and material (Beicht et al., 2004). This implies a negligible "mechanical" reduction of $(1.01 \times 487/294 =)$ €1.67 per worker.[33]

A reason for the large effect may be that technology investments are lumpy: not all firms

[33] In addition, firms reported €1,530 of "other costs" per trainee per year, including costs for teaching material, fees, and training administration. If a firm interpreted all of these costs as capital investments, the "mechanical" reduction in investments would still be as small as €6.93 per worker.

constantly adopt new technologies, but if they do, they make costly, indivisible investments (e.g. Cooper et al., 1999; Bessen et al., 2020). Hence, firms not planning to invest remain unaffected by the trainee shortage while the other firms forego large investments, leading to large average treatment effects. I test this hypothesis in two ways: First, I compute the median instead of average treatment effect. In particular, I compute the average difference in investments between each treated firm and its three control firms for each year, and then plot the median of these distributions within years in Figure 5. The median decrease is €1,140 per worker in 2001 to 2003, equivalent to (1,140/13,380=) 8.5% of the pre-treatment average, or 0.4% of the capital stock per worker. The negative but smaller median decline indicates that the effect is not driven by outliers but reflects a broad-based investment decline, with the mean effect inflated by large effects for a subset of firms. Second, I study the extensive and intensive investment margin separately, see Table 8, columns 3 to 5. The extensive margin—whether firms invest or not—is unaffected, which is expected since 92% of observations show strictly positive investments. The intensive margin—measured as log investments—adjusts significantly. Most importantly, treated firms are 15pp less likely to make large investments (investments in the upper tercile of the investment per worker distribution of at least €15,940 per worker and €30,190 per worker on average) than control firms when trainees are scarce.

To sum up, average effects on investments are large, likely because firms who would have undertaken large, lumpy investments in the counterfactual scenario forgo (infinitely postpone or skip) technology vintages. Given that highly educated trainees account for 20.14% of entry-level hires and that their absence reduces investments by 8.5% at the median, a simple back-of-the-envelope calculation assuming homogeneous entry-level workers suggests that approximately ($8.5 \times 100/20.14 =$) 42.2% of firm investments are induced by entry-level hiring.

Table 8: DiD Results – Additional investment effects

	Inv. per worker in €1,000		Any inv. (1/0)	Log(Inv.)	Large inv. (1/0)
	constant workers	time-variant workers			
	(1)	(2)	(3)	(4)	(5)
Treat × Post	-2.74*	-2.64*	-0.01	-0.27*	-0.15**
	(1.07)	(1.06)	(0.03)	(0.13)	(0.05)
Mean dep. variable	13.40	14.56	0.92	14.08	0.38
N	2,853	2,853	2,850	2,550	2,550

Notes: DiD coefficients based on equation (5). Pre: 1997–2000. Post: 2001–2003. Standard errors clustered at the firm level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Mean dep. variable: Average outcome in 1997–2000. Large inv.: Investments in the upper tercile of the distribution of strictly positive investments per worker assigned as one, and zero otherwise.

7.2 Triple DiD

To increase confidence that the investment decline is linked to the trainee shortage and not to confounding events in the treated states, I incorporate non-training firms into the design: non-training firms in the same industries, states and years should be subject to the same confounders as training firms but much less affected by the trainee shortage. I run a triple DiD regression, comparing the difference in the investment decline between training and non-training firms in treated states to the difference in control states over time:

$$Y_{jt} = \sum_{\substack{t=1997,\\ t\neq 2000}}^{2006} \beta_t(\text{Train}_j \times \text{Treat}_{s(j)} \times \text{Year}_t) + \sum_{\substack{t=1997,\\ t\neq 2000}}^{2006} \gamma_t(\text{Train}_j \times \text{Year}_t) + \nu(\text{Train}_j \times \text{Treat}_{s(j)}) + \rho_{skt} + \psi_t + \phi_j + \epsilon_{jt} \quad (6)$$

I additionally control for all corresponding two-way-interaction terms as well as state-industry-specific year fixed effects ρ_{skt}, such that the coefficients of interest are identified off the variation between training and non-training firms operating in the same state-year-industry cell.[34]

Figure 6 shows the estimated regression coefficients of interest β_t. Due to the more demanding model, the post-treatment coefficients are not statistically significant at conventional levels but exhibit a very similar pattern and coefficient magnitude as in the baseline DiD regression. This provides further evidence that the investment drop is indeed related to the negative trainee supply shock. I prefer the baseline DiD because the triple DiD obscures treatment effects on non-training firms, which may arise if these firms would start training in the counterfactual scenario without a trainee shortage, as well as through spillover effects.

In Appendix D, I run a Bartik-style IV regression to analyze whether firms that are more exposed to the trainee shock reduce investments more. To do so, I instrument firms' trainee employment with firms' initial employment of highly educated trainees (share) and the reform (shift). Just as the triple DiD regression, this analysis helps tie the investment decline to the trainee shortage. Also, it identifies the effect independent of the realized distribution of trainees across training firms. The analysis reveals that more exposed firms indeed reduce investments significantly more.

7.3 The new-skills mechanism

According to the economic framework, firms with high trainee retention rates view trainees as a human capital investment for future production (the so-called 'investment model' in the training literature, see Stevens, 1994), whereas firms with low retention rates employ trainees

[34] In the baseline DiD event study, industry balance was achieved by matching on industry.

Figure 6: Triple difference specification – Effect on investments per worker in €1,000



Notes: Event study coefficients of the interaction terms Train × Treat × Year plus 90% and 95% confidence bands. Based on equation (6). Standard errors clustered at the firm level. N=14,657.

as inexpensive labor for current production without training them (the so-called 'production model', see Lindley, 1975). If the investment decrease is driven by the role of trainees in skill acquisition rather than their role as inexpensive labor, the framework predicts that investments should drop more in firms with higher retention rates.

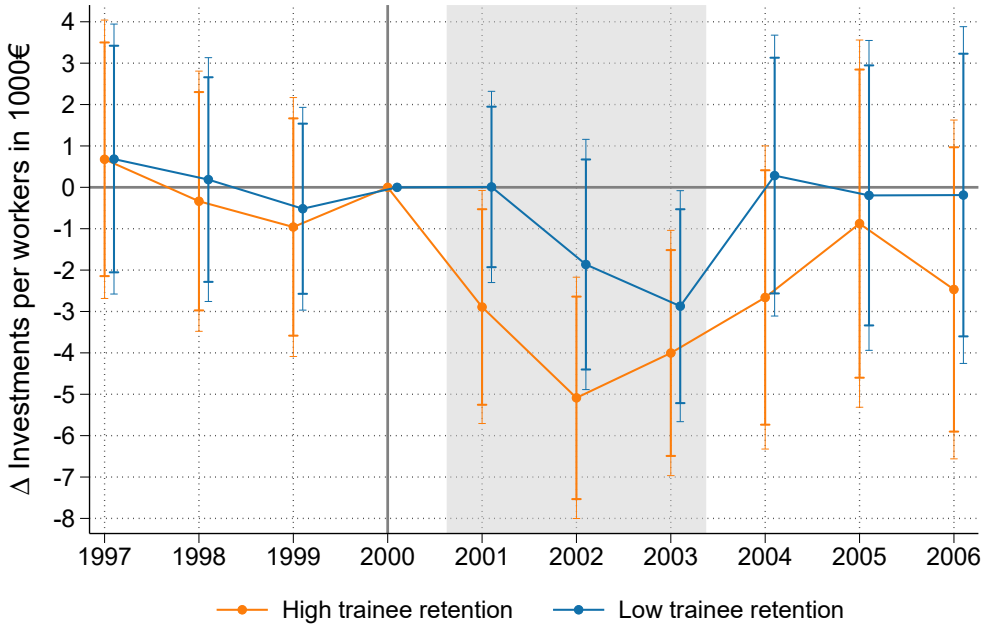
To test this, I run the DiD event study regression including the interaction term Treat × Year × Trainee retention rate and all corresponding two-way interaction terms. The retention rate is defined as the pre-treatment fraction of trainees staying at the firm upon training completion.[35] I additionally include the triple interaction term Treat × Year × Shock exposure plus all corresponding two-way interactions to compare investment drops between two treated firms with the same exposure to the reform but with different training strategies. Shock exposure is defined as pre-reform employment of highly educated trainees. Figure 7 shows the predicted investment changes for firms with a high (low) trainee retention rate, evaluated at the 75th (25th) percentile of the distribution, corresponding to a retention rate of 68% (33%). Firms with high retention rates reduce investments by, on average, €4,000 over 2001–2003, equivalent to (4,000/13,560=) 29% of their pre-treatment levels. In contrast, treated firms with low retention rates reduce investments by only €1,580 (1,580/12,390=13%). The pattern is robust to controlling for Treat × Year × Industry and Treat × Year × Firm size interactions, along with all corresponding lower-order terms (see Table A6). This effectively

[35] This information is based on two questions from the establishment panel on the number of trainees retained by the firm and the number of successfully completed vocational trainings. If the ratio of these two variables is not available, I construct the ratio based on the social security data.

compares treated firms with identical reform exposure, industry, and size, but different training strategies.

To sum up, I find evidence for the second prediction of the economic framework: firms investing in their entrants' new skills are those reducing investments when entrants are missing. This suggests that the underlying mechanism behind the complementarity between new entrants and new technology is indeed entrants' comparative advantage in skill acquisition.

Figure 7: Effects by trainee retention rate



Notes: Based on the triple interaction effect Treat × Year × Trainee retention rate, evaluated at the 25th and 75th percentiles of the trainee retention distribution.

Alternative mechanisms. There are at least two alternative mechanisms that produce complementarity between young labor market entrants and technology adoption other than trainees' low opportunity costs and high returns to skill acquisition. First, expected payoffs of skill acquisition are higher for younger workers (the "horizon" channel in Cavounidis & Lang, 2020). Second, younger workers may generally possess more up-to-date tech skills. While both channels likely play a role, they are unlikely to fully explain the investment decline: Both mechanisms imply that the returns to training decrease *gradually* with worker age. However, workers from the previous entrant cohort, i.e. trainees in their second year of vocational training, are only one year older and have therefore only marginally lower returns to training or only marginally worse technology skills than new entrants. To explain why the previous cohort cannot compensate for the missing entrant cohort, there must be a *discontinuous*, or at least very large, jump in the profitability of adoption across cohorts (see also Cavounidis & Lang, 2020). Jumps in productivity levels between entrants and the previous cohort provide such a discontinuity: according to the BIBB-Cost-Benefit-Survey, firm revenues from skilled labor

activities of second-year (third-year) trainees are a staggering 134% (254%) higher than of first-year trainees (Schönfeld et al., 2016, Table 18).

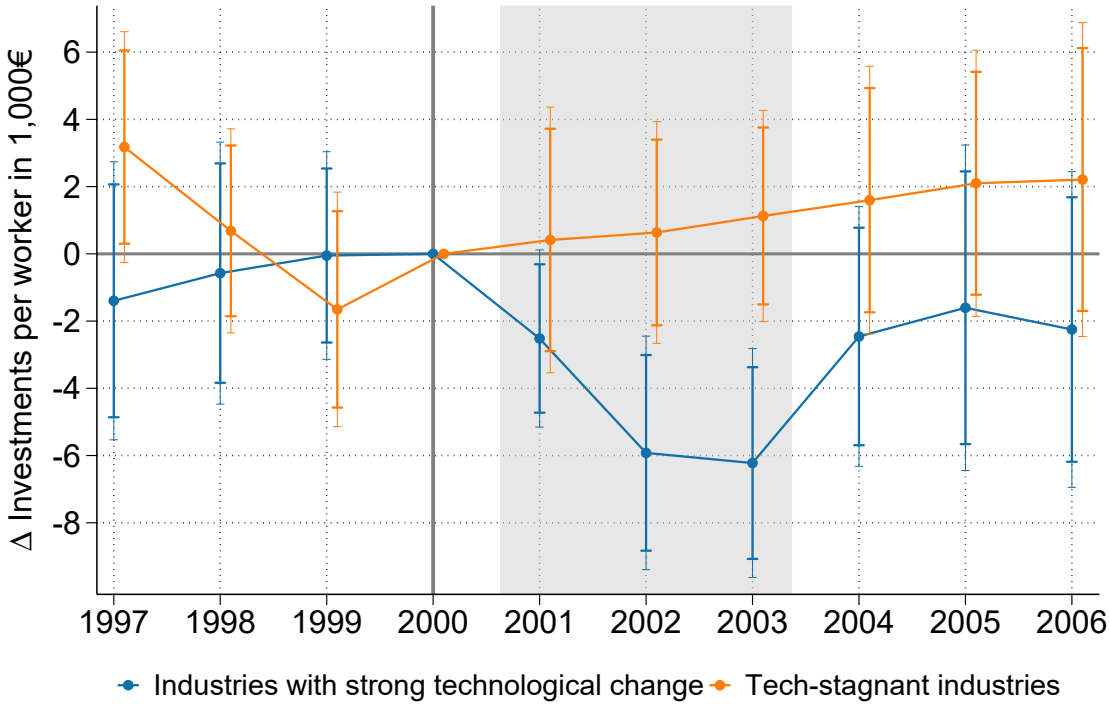
7.4 Effect on firm technology adoption

I next investigate whether the investment decline is linked to the foregone adoption of new technology. I employ two strategies: First, I conduct a heterogeneity analysis by industry exposure to technological progress. Second, I run the DiD event study using firms' self-assessed technical status of machinery as auxiliary outcome.

Industry exposure to technological progress. For the investment decline to be driven by foregone technology adoption, the effect needs to be concentrated in industries experiencing rapid technological change. To test this, I compute log labor productivity for 40 distinct sectors based on EU-KLEMS data (Stehrer, 2024), defined as value added over employment, and its increase over the five years around the shock, i.e. 1998-2003. Sectors with above-median labor productivity growth are categorized as experiencing strong technological change, and as tech-stagnant otherwise. Sectors with strong technological change include manufacturing industries (e.g. Manufacturing of chemicals and chemical products with an increase of log labor productivity by 0.19) as well as service industries (e.g. Telecommunications with an increase of 0.43, or real estate with an increase of 0.14). Figure 8 shows the DiD event study coefficients for this sample split. Firms in industries with strong technological change clearly drive the investment decline: investments decrease by, on average, €4,890 per worker over 2001–2003, corresponding to (4,890/16,230=) 30% of their pre-treatment levels. Firms in tech-stagnant industries, in contrast, do not significantly adjust their investments. The absence of any investment decrease in tech-stagnant industries provides clear evidence that a trainee shortage hinders the adoption of new technologies, whereas other capital investments are unaffected.

Firms' technical status of machinery. Turning to my second strategy, I study the effect of the trainee shortage on firms projected technical status, i.e. the self-assessed technical status of a firm's machinery on a scale from 1 ('completely out-of-date') to 5 ('state-of-the-art'), projected back to the years when investments affecting the technical status took place, see again Section 4. As shown in Figure 9, Panel A, treated firms report an outdated projected technical status of their machinery compared to control firms starting in 2002. The depreciation is statistically significant and meaningful in magnitude: a depreciation by 0.26 in 2004 corresponds to 26% of firms reporting a deterioration by one category, or a decrease by a third of a standard deviation. Treated firms' technical status starts to converge back to control firms' technical status by 2005/2006. The result survives in the triple DiD specification, albeit with reduced statistical significance, see Panel B: within treated states, the decrease is driven by training compared to non-training firms.

Figure 8: Effects by industry-level technological change



Notes: Event study coefficients of the interaction terms Treat × Year plus 90% and 95% confidence bands. Based on equation (4). Standard errors clustered at the firm level. Sample split by industry, with industries experiencing strong technological change defined as those with above median growth in log labor productivity over 1998–2003 according to EU-KLEMS data, and as tech-stagnant otherwise. N=2,052 for firms in industries with strong technological change and N=1,557 for firms in tech-stagnant industries.

I conclude that at least part of the investment decline is due to the reduced adoption of new technology and that, hence, young labor market entrants are important complements to firms' technology adoption. In Appendix E, I investigate whether the foregone technology adoption translates into foregone firm output and find indicative evidence of sizable output declines.

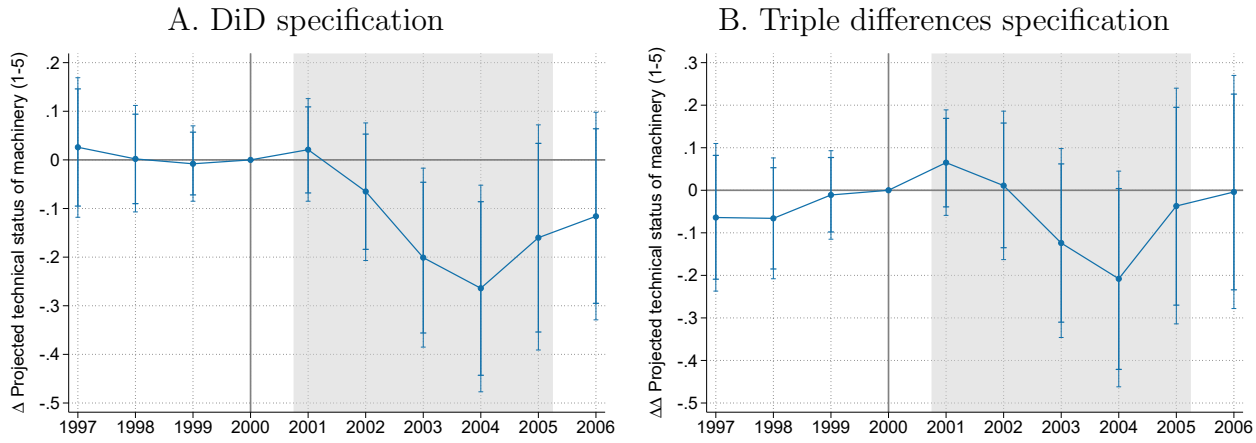
7.5 Robustness

The negative reform effect is robust to a wide range of specifications. To see this, I present the DiD estimates for the main outcome, investments per worker, in Figure 10.

Next, I expand the set of control firms to include West German firms, which were initially excluded because they may differ from East German firms and may be exposed to different shocks. As an advantage, West German firms are less subject to spill-overs such as trade or competition than East German firms. Convincingly, including them in the control group yields very similar estimates. When excluding Berlin or Saxony from the set of control states, due to their slightly different demographic and economic trends and due to their potentially different exposure to the dot-com bubble, the results also remain robust.

Firms may leave the sample over time. The negative estimate remains robust when restricting the sample to firms observed in every year between 1997 and 2003.

Figure 9: Effect on technology adoption



Notes: Panel A: Event study coefficients of the interaction terms Treat × Year plus 90% and 95% confidence bands. Based on equation (4). N=3,025. Panel B: Event study coefficients of the interaction terms Train × Treat × Year plus 90% and 95% confidence bands. Based on equation (6). N=12,104. Standard errors clustered at the firm level. Projected technical status: Technical status of a firm's machinery on a scale from 1 ('completely out-of-date') to 5 ('state-of-the-art'), projected to the year in which the investments influencing that status were made.

The effect is negative within both treated states separately, despite their different industry structures and geography. However, the effects are less precise due to the smaller sample size.

Firms at federal state borders might be less affected by the reform because they may attract trainees from control states. I use firms' cross-state commuting shares as a proxy for worker supply from other states. Excluding firms with a pre-treatment commuter share in the highest decile does not meaningfully affect the results, confirming that trainee mobility is irrelevant.

In the main specification, I restrict the sample to firms with at least five employees and with positive investments at least once during the time period of interest. The estimate is completely unaffected by this restriction.

First, the result is qualitatively robust to the trimming procedure of the outcome: trimming the uppest percentile instead of values that exceed the third quartile plus four times the interquartile range leads to larger but noisier estimates. This is consistent with the problem of a highly right-skewed variable.

Results are robust to weighting observations by firms' pre-treatment employment size.

A concern for causality is that the results may reflect shocks unrelated to the trainee shortage. Acknowledging that states may be on different time trends, and controlling for them linearly does not affect the results.[36] Next, I directly control for economic state-level log GDP and log public investments, both interacted with the Post dummy. The estimate of interest remains robust.

The trainee supply shock is a market-wide shock, arguably allowing for little redistribution of trainees across firms in the time of the shortage. Therefore, the effect at the firm should be

[36] I predict state-specific time trends based on the period 1997–2000 and residualize investments with respect to this (predicted) time trend in order to no confound the time trend with the shock.

similarly sized as the effect at the market level. I repeat the analysis at the county-industry level. Convincingly, the effect is extremely similar in size.

The results are also robust to not matching treated and control firms, as well as different specifications of the matching procedure. When using the one nearest neighbor instead of the three nearest neighbors, the estimated coefficient remains very similar, despite a significant reduction in sample size. The results are also robust to the inclusion of the 10% most distant matches that are discarded in the main specification.

7.6 Inference with few clusters

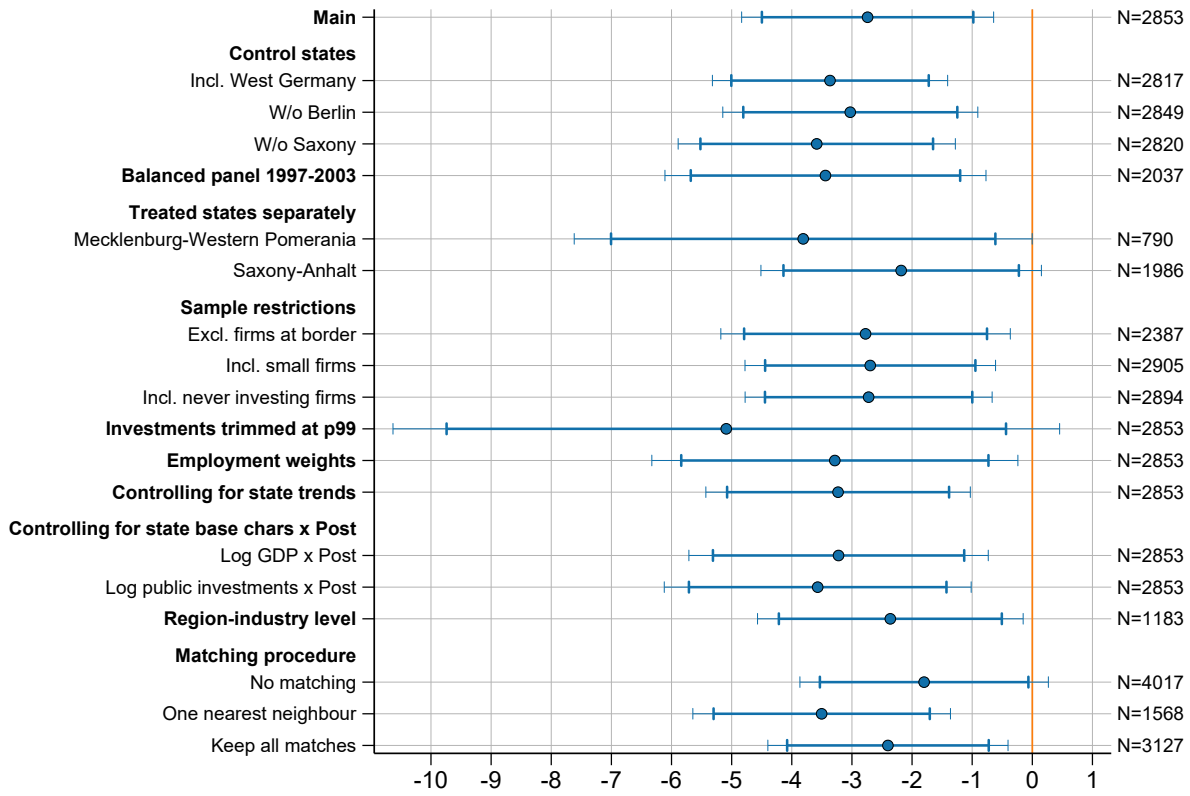
Until now, I have assumed that state-specific shocks are small compared to the idiosyncratic error terms at the firm level, justifying to cluster standard errors at the firm level. I next perform permutation tests which have been suggested as a valid method for inference when the number of clusters is small (e.g. Roth et al., 2023). Figure 11, Panel A, shows the t-statistics for the DiD event study estimates based on the actual treatment assignment and for all permuted treatment assignments. The t-statistics are based on standard errors clustered at the firm level that account for sampling errors of firms within states. Following the 2001 reform, the t-statistics based on the actual treatment assignment are by far more negative than any t-statistic based on a permuted treatment assignment. Hence, it is very unlikely that cluster-level shocks are large enough to cause the observed investment decline.

Since the number of possible permutations within East Germany is limited to 15,[37] I repeat the permutation test across the 10 West German federal states, yielding $(\sum_{k=2}^{8} \binom{10}{k} =)$ 1,002 permutations with two to eight treated states. Again, the t-statistic of actual treatment assignment is much more negative than the 5% and 2.5% most negative t-statistics under permuted treatment assignment, see Panel C.

Finally, I report confidence intervals based on wild cluster bootstraps at the state level (Cameron et al., 2008). With only six clusters, the wild cluster bootstrap likely produces unreliable confidence intervals in the East German sample. I nonetheless report the result for completeness; the estimates for the years 2002 and 2003 are significant at the 10% level (Panel B), though this should be interpreted with caution. Turning to the sample of both East and West German states, relying on 16 clusters, I find significance at the 5% level for the estimates of the years 2002 and 2003.

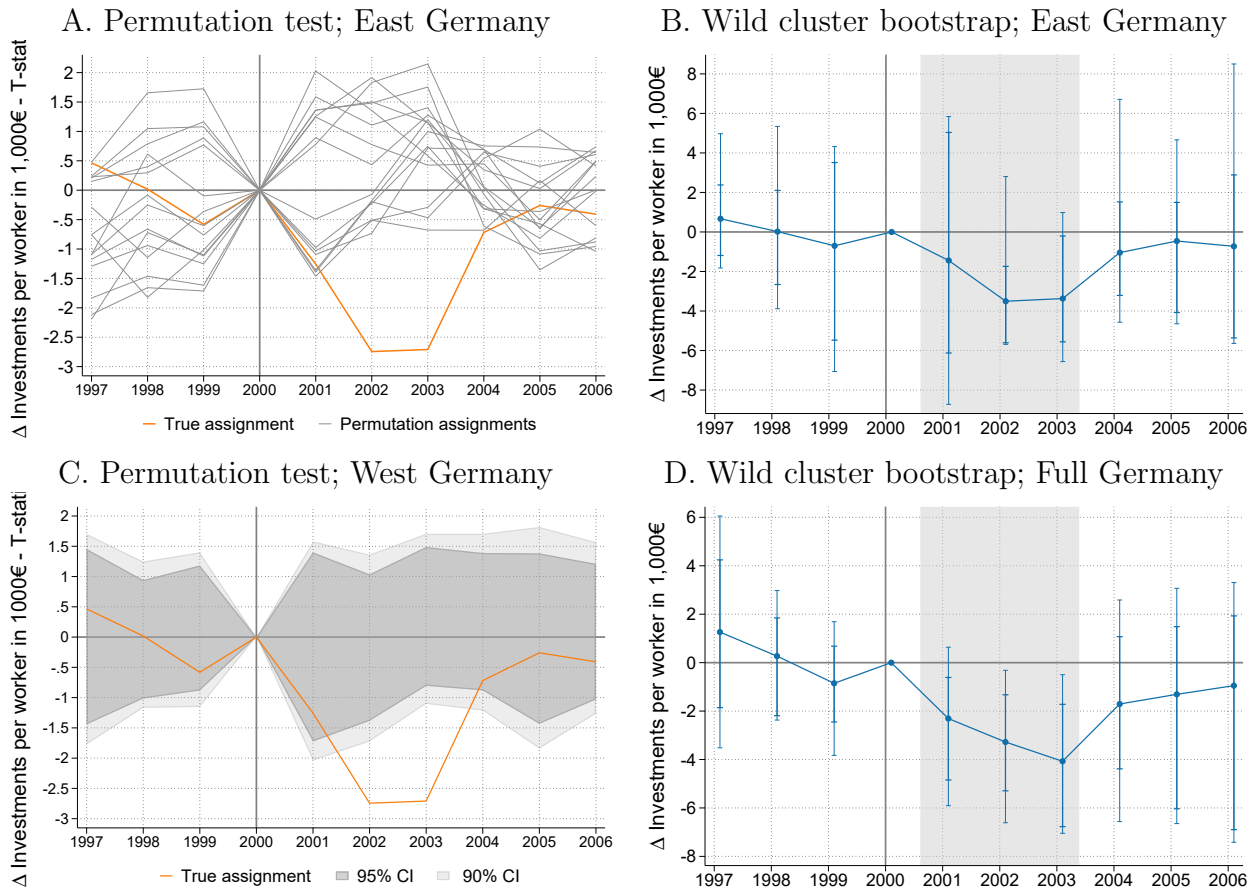
[37] 4 permutations with one treated state, 6 permutations with two treated states, 4 permutations with three treated states, and one permutation with four treated states.

Figure 10: Robustness



Notes: DiD coefficients and 90% and 95% confidence bands of the term Treat × Post. N indicates the number of observations in the respective estimation. *Control states:* Additionally including all West German training firms as control firms; dropping Berlin or Saxony from the set of control firms. *Balanced panel 1997-2003:* Sample restricted to firms observed in each year between 1997 and 2003. *Treated states separately:* Only using treated firms from one treated state and dropping firms from the other. *Excl. firms at border:* Dropping those 10% of firms with the highest 1999 cross-state commuter share of workers with vocational training. *Employment weights:* Weighting firms by their total employment count in 1997. *Controlling for state trends:* Additionally controlling for predicted linear state-specific time trends based on the pre-period. *Controlling for state base characteristics × Post:* Additionally controlling for 1997 baseline characteristics at the state level times a Post dummy. *Region-industry level:* Analysis at the level of local labor markets (69 counties) times 8 broad industries, resulting in 182 distinct labor-market-industry cells. Cells are weighted by their number of underlying firms in 1997 to increase precision. *Matching procedure:* Using only the nearest neighbor instead of the three nearest neighbors as control firms; keeping all matches instead of discarding the furthest 10% of all matches.

Figure 11: Effect on investments – Inference diagnostics



Notes: Panel A and C: T-statistics of the DiD event study regression coefficients using the actual treatment assignment (orange line) and all possible permutation assignments within East Germany (Panel A; gray lines) within West Germany (Panel C; gray areas). Panel B and D: Coefficients and 90% and 95% confidence intervals using wild bootstraps clustered at the state level.

8 Conclusion

In this paper, I provide causal evidence that a reduction in the supply of young labor market entrants substantially reduces firm investments, linked to a decrease in the adoption of new technologies. The empirical strategy exploits an education reform in Germany that creates a missing school graduation—and therefore trainee—cohort in some states, and compares firm outcomes with control states. I hereby provide evidence that trainees are complements to firm technology adoption. This is due to their low opportunity costs and/or high expected returns from skill acquisition: firms investing in their trainees' skill acquisition drive the investment decline, in contrast to firms viewing trainees as inexpensive labor. Supplemental descriptive evidence beyond the reform setting reinforces this finding: young workers more often work with new technologies than older workers across education levels and countries, and firm investments are generally negatively correlated with regional shortages of young labor market entrants, unlike staff shortages or shortages of skilled labor.

This paper contributes several novel insights. First, the availability of young labor market entrants is a key determinant of firm technology adoption. Second, one important factor determining how well workers complement new technologies is workers' comparative advantage in skill acquisition—beyond the factors previously emphasized, namely education and job tasks. As a consequence, shortages of complementary workers may not only incentivize firms to replace them by adopting labor saving technologies—a channel studied extensively in the literature—but may also constrain adoption because the workers supposed to implement and use this technology are unavailable. Third, retraining incumbents is costly compared to training young labor market entrants. This leads to strong vintage effects, where different worker cohorts posses distinct, vintage-specific skills, which slows down technological change (Adão et al., 2024).

By establishing complementarity between new entrants and new technologies, this paper also informs on the macro debate on endogenous technological change and demographic change: countries with lower population growth or shortages of middle-aged workers are found to adopt more (labor-saving) robots (Abeliansky & Prettner, 2017; Acemoglu & Restrepo, 2022). The "new-skills" channel proposed in this paper works in the opposite direction. In line with this argument, Angelini (2023) finds that above a certain tipping point, population aging *reduces* investments in labor-saving and skill-intensive new technology.

From a policy perspective, the findings stress the importance of attracting young labor market entrants or subsidizing retraining of incumbent workers to foster technology adoption. The results also have implications for the design of education systems: the finding that firms shy away from retraining incumbent workers trained a few years ago indicates that skills acquired through vocational training may be overly specific (Hanushek et al., 2017).

References

Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. M. (2023). When should you adjust standard errors for clustering? The Quarterly Journal of Economics, 138(1), 1–35.

Abeliansky, A., & Prettner, K. (2017). Automation and demographic change. Available at SSRN 2959977.

Acemoglu, D. (1998). Why do new technologies complement skills? Directed technical change and wage inequality. The Quarterly Journal of Economics, 113(4), 1055–1089.

Acemoglu, D. (2002). Directed technical change. The Review of Economic Studies, 69(4), 781–809.

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Handbook of Labor Economics (Vol. 4, pp. 1043–1171). Elsevier.

Acemoglu, D., Lelarge, C., & Restrepo, P. (2020). Competing with robots: Firm-level evidence from France. In AEA Papers and Proceedings (Vol. 110, pp. 383–88).

Acemoglu, D., & Pischke, J.-S. (1998). Why do firms train? Theory and evidence. The Quarterly Journal of Economics, 113(1), 79–119.

Acemoglu, D., & Restrepo, P. (2022). Demographics and automation. The Review of Economic Studies, 89(1), 1–44.

Adão, R., Beraja, M., & Pandalai-Nayar, N. (2024). Fast and slow technological transitions. Journal of Political Economy Macroeconomics, 2(2), 000–000.

Aghion, P., Akcigit, U., Hyytinen, A., & Toivanen, O. (2024). A year older, a year wiser (and farther from frontier): Invention rents and human capital depreciation. Review of Economics and Statistics, 106(4), 974–982.

Aghion, P., Antonin, C., Bunel, S., & Jaravel, X. (2023). Modern manufacturing capital, labor demand, and product market dynamics: Evidence from France.

Ahituv, A., & Zeira, J. (2011). Technical progress and early retirement. The Economic Journal, 121(551), 171–193.

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

Angelini, D. (2023). Aging population and technology adoption (Tech. Rep.). Department of Economics, University of Konstanz.

Arntz, M., Lipowski, C., Neidhöfer, G., & Zierahn-Weilage, U. (2025). Computers as stepping stones? technological change and equality of labor market opportunities. Journal of Labor Economics, 43(2), 503–543.

Aubert, P., Caroli, E., & Roger, M. (2006). New technologies, organisation and age: firm-level evidence. The Economic Journal, 116(509), F73–F93.

Autor, D., Chin, C., Salomons, A., & Seegmiller, B. (2024). New frontiers: The origins and content of new work, 1940–2018. The Quarterly Journal of Economics, 139(3), 1399–1465.

Autor, D., & Dorn, D. (2009). This job is "getting old": measuring changes in job opportunities using occupational age structure. American Economic Review, 99(2), 45–51.

Autor, D., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. The Quarterly Journal of Economics, 118(4), 1279–1333.

Barsi, M. (2001). Betriebsbesuch: Arbeitskräfte reichen nicht – landtechnik gmbh sucht noch fachleute. Mitteldeutsche Zeitung. (Ausgabe: Weissenfels)

Bartel, A. P., & Sicherman, N. (1998). Technological change and the skill acquisition of young workers. Journal of Labor Economics, 16(4), 718–755.

Battisti, M., Dustmann, C., & Schönberg, U. (2023). Technological and organizational change and the careers of workers. Journal of the European Economic Association, 21(4), 1551—1594.

Beaudry, P., Doms, M., & Lewis, E. (2010). Should the personal computer be considered a technological revolution? Evidence from US metropolitan areas. Journal of Political Economy, 118(5), 988–1036.

Beicht, U., Walden, G., & Herget, H. (2004). Kosten und Nutzen der betrieblichen Berufsausbildung in Deutschland. Bielefeld: W. Bertelsmann Verlag.

Bessen, J., Goos, M., Salomons, A., & van den Berge, W. (2020). Firm-level automation: Evidence from the Netherlands. AEA Papers and Proceedings, 110, 389–393.

Bessen, J., Goos, M., Salomons, A., & Van den Berge, W. (2023). What happens to workers at firms that automate? The Review of Economics and Statistics, 1–45.

Bick, A., Blandin, A., Deming, D. J., Fuchs-Schündeln, N., & Jessen, J. (2026). Mind the gap: Ai adoption in europe and the us (Tech. Rep.). National Bureau of Economic Research.

Bonfiglioli, A., Crino, R., Fadinger, H., & Gancia, G. (2024). Robot imports and firm-level outcomes. The Economic Journal, 134(664), 3428–3444.

Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. The Quarterly Journal of Economics, 117(1), 339–376.

Büttner, B., & Thomsen, S. L. (2015). Are we spending too many years in school? Causal evidence of the impact of shortening secondary school duration. German Economic Review, 16(1), 65–86.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. The Review of Economics and Statistics, 90(3), 414–427.

Carneiro, P., Liu, K., & Salvanes, K. G. (2023). The supply of skill and endogenous technical change: evidence from a college expansion reform. Journal of the European Economic

Association, 21(1), 48–92.

Cavounidis, C., & Lang, K. (2020). Ben-Porath meets Lazear: Microfoundations for dynamic skill formation. Journal of Political Economy, 128(4), 1405–1435.

Chari, V. V., & Hopenhayn, H. (1991). Vintage human capital, growth, and the diffusion of new technology. Journal of Political Economy, 99(6), 1142–1165.

Clemens, M. A., Lewis, E. G., & Postel, H. M. (2018). Immigration restrictions as active labor market policy: Evidence from the mexican bracero exclusion. American Economic Review, 108(6), 1468–87.

Cooper, R., Haltiwanger, J., & Power, L. (1999). Machine replacement and the business cycle: lumps and bumps. American Economic Review, 89(4), 921–946.

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

Deming, D. J., & Noray, K. (2020). Earnings dynamics, changing job skills, and stem careers. The Quarterly Journal of Economics, 135(4), 1965–2005.

Dorner, M., Görlitz, K., & Jahn, E. J. (2024). The impact of a missing school graduation cohort on the training market. Economics of Education Review, 103, 102580.

Dustmann, C., & Schönberg, U. (2009). Training and union wages. The Review of Economics and Statistics, 91(2), 363–376.

Dustmann, C., & Schönberg, U. (2012). What makes firm-based vocational training schemes successful? The role of commitment. American Economic Journal: Applied Economics, 4(2), 36–61.

Ellguth, P., Kohaut, S., & Möller, I. (2014). The IAB Establishment Panel – Methodological essentials and data quality. Journal for Labour Market Research, 47(1–2), 27–41.

Federal Institute for Vocational Education and Training (Ed.). (2009). Datenreport zum Berufsbildungsbericht 2009. informationen und Analysen zur Entwicklung der beruflichen Bildung. (Tech. Rep.). (Table A5.4.2-1)

Federal Statistical Office, G.-O. (2022). 12411-0010: Bevölkerung: Bundesländer, Stichtag. (January 18, 2024; data license by-2-0; own calculation/own presentation.)

Federal Statistical Office, G.-O. (2023a). 13211-0007: Arbeitslose, Arbeitslosenquoten, gemeldete Arbeitsstellen: Bundesländer, Jahre. (January 18, 2024; data license by-2-0; own calculation/own presentation.)

Federal Statistical Office, G.-O. (2023b). 21711-0010: Ausgaben der öffentlichen Haushalte für Bildung: Bundesländer, Jahre, Körperschaftsgruppen, Aufgabenbereiche der öffentlichen Haushalte. (January 18, 2024; data license by-2-0; own calculation/own presentation.)

Federal Statistical Office, G.-O. (2023c). 71141-0006: Investitionsausgaben der öffentlichen Haushalte: Bundesländer, Jahre, Körperschaftsgruppen, Art der Investitionsausgaben [ohne

Krankenhäuser und Hochschulkliniken mit kaufmännischem Rechnungswesen.]. (January 18, 2024; data license by-2-0; own calculation/own presentation.)

Federal Statistical Office, G.-O. (2023d). 71321-0002: Schulden der öffentlichen Gesamthaushalte: Bundesländer, Stichtag (31.12.1992-31.12.2005), Körperschaftsgruppen, Schuldenarten. (January 18, 2024; data license by-2-0; own calculation/own presentation.)

Federal Statistical Office, G.-O. (2023e). 82111-0001: Vgr der Länder (Entstehungsrechnung) – Bruttoinlandsprodukt zu Marktpreisen (nominal): Bundesländer, Jahre; i. jew. Preisen (Mill. Euro). (January 18, 2024; data license by-2-0; own calculation/own presentation.)

Federal Statistical Office, Genesis-Online. (2022a). Ausbildungsverträge: Deutschland, Jahre, Geschlecht, (November 07, 2022; data license by-2-0; own calculation/own presentation.)

Federal Statistical Office, Genesis-Online. (2022b). Bevölkerung: Bundesländer, Stichtag, Altersjahre. (November 07, 2022; data license by-2-0; own calculation/own presentation.)

Federal Statistical Office, Genesis-Online. (2022c). Studienanfänger: Deutschland, Semester, Nationalität, Geschlecht. Statistik der Studenten. (November 07, 2022; data license by-2-0; own calculation/own presentation.)

Federal Institute for Vocational Education, & Training (Eds.). (2002). Berufsbildungsbericht 2002. (Tech. Rep.). (Table 17)

Federal Institute for Vocational Education, & Training. (2022). Collectively agreed training allowances database.

Federal Ministry of Education, & Research. (2022). Tab 2.3.15. School-graduates and school-leavers, by type of school-leaving certificate, Länder and sex. (October 18, 2022; https://www.datenportal.bmbf.de/portal/en/K233.html; own calculation/own presentation.)

Fischer, G., Janik, F., Müller, D., & Schmucker, A. (2009). The IAB Establishment Panel — Things users should know. Schmollers Jahrbuch. Zeitschrift für Wirtschafts- und Sozialwissenschaften, 129(1), 133–148. doi: 10.5164/IAB.FDZD.2103.en.v1

Fitzenberger, B., Holleitner, J., & Kagerl, C. (2024). Herausforderungen für die arbeitsmärkte der zukunft am beispiel deutschland. Wirtschaftsdienst, 104(8), 519–523.

Fitzenberger, B., Osikominu, A., & Völter, R. (2006). Imputation rules to improve the education variable in the iab employment subsample. Journal of Contextual Economics–Schmollers Jahrbuch(3), 405–436.

Fuest, C., & Jäger, S. (2023). Können höhere löhne zur überwindung des fachkräftemangels beitragen. Wirtschaftsdienst, 103(4), 253–258.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Journal of Econometrics, 225(2), 254–277.

Hanushek, E. A., Schwerdt, G., Woessmann, L., & Zhang, L. (2017). General education, voca-

tional education, and labor-market outcomes over the lifecycle. Journal of Human Resources, 52(1), 48–87.

Hassan, T. A., Kalyani, A., & Restrepo, P. (2025). New technologies and the college premium (Tech. Rep.).

Heine, C., Spangenberg, H., & Sommer, D. (2005). Studienberechtigte 2004: erste Schritte in Studium und Berufsausbildung; Vorauswertung der Befragung der Studienberechtigten 2004 ein halbes Jahr nach Schulabgang im Zeitvergleich. HIS.

Helbig, M., & Nikolai, R. (2015). Die Unvergleichbaren: Der Wandel der Schulsysteme in den deutschen Bundesländern seit 1949. Julius Klinkhardt.

Hémous, D., Olsen, M., Zanella, C., & Dechezleprêtre, A. (2025). Induced automation innovation: evidence from firm-level patent data. Journal of Political Economy, 133(6), 1975–2028.

Hornbeck, R., & Naidu, S. (2014). When the levee breaks: black migration and economic development in the American South. American Economic Review, 104(3), 963–90.

Humlum, A., & Vestergaard, E. (2025). The unequal adoption of chatgpt exacerbates existing inequalities among workers. Proceedings of the National Academy of Sciences, 122(1), e2414972121.

IAB establishment panel (IAB BP) — Version 9317 v1. (2019). doi: 10.5164/IAB.IABBP9317.de.en.v1

Konings, J., & Vanormelingen, S. (2015). The impact of training on productivity and wages: firm-level evidence. Review of Economics and Statistics, 97(2), 485–497.

Kühn, S. M., Ackeren, I. v., Bellenberg, G., Reintjes, C., & Brahm, G. i. (2013). Wie viele Schuljahre bis zum Abitur? Eine multiperspektivische Standortbestimmung im Kontext der aktuellen Schulzeitdebatte. Zeitschrift für Erziehungswissenschaft, 16(1), 115–136.

Le Barbanchon, T., Ronchi, M., & Sauvagnat, J. (2023). Hiring frictions and firms' growth. Available at SSRN 4105264.

Lerche, A. (2025). Direct and indirect effects of investment tax incentives. American Economic Review, 115(8), 2781–2818.

Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. The Quarterly Journal of Economics, 126(2), 1029–1069.

Lindley, R. M. (1975). The demand for apprentice recruits by the engineering industry, 1951-71. Scottish Journal of Political Economy, 22(1), 1–24.

Linked-employer-employee-data from the IAB (LIAB): LIAB-cross-sectional model 2 (LIAB QM2) 1993–2019, version 1. (2021). doi: 10.5164/IAB.LIABQM29319.de.en.v1

MacDonald, G., & Weisbach, M. S. (2004). The economics of has-beens. Journal of Political Economy, 112(S1), S289–S310.

Marcus, J., & Zambre, V. (2019). The effect of increasing education efficiency on university

enrollment evidence from administrative data and an unusual schooling reform in Germany. Journal of Human Resources, 54(2), 468–502.

Ministry of Education, F., & Research (Eds.). (2004). Berufsbildungsbericht 2003. (Tech. Rep.).

Ministry of Education, F., & Research (Eds.). (2025). Berufsbildungsbericht 2025. (Tech. Rep.).

Morin, L.-P. (2015). Cohort size and youth earnings: evidence from a quasi-experiment. Labour Economics, 32, 99–111.

Muehlemann, S., Dietrich, H., Pfann, G., & Pfeifer, H. (2022). Supply shocks in the market for apprenticeship training. Economics of Education Review, 86, 102197.

Müller, S. (2008). Capital stock approximation using firm level panel data: A modified perpetual inventory approach. Jahrbücher für Nationalökonomie und Statistik, 228(4), 357–371.

Müller, S. (2017). Capital stock approximation with the perpetual inventory method: An update. FDZ-Methodenreport, 5, 2017.

Neuber-Pohl, C., Pregaldini, D., Backes-Gellner, U., Dummert, S., & Pfeifer, H. (2023). How negative labor supply shocks affect training in firms: Lessons from opening the Swiss-German border. Swiss Leading House "Economics of Education" Working Paper, 203.

Pulito, G., Pytlikova, M., Schroeder, S., & Lodefalk, M. (2026). Who adopts AI? Evidence on firms, technologies and workers (Tech. Rep.). GLO Discussion Paper.

Rohrbach-Schmidt, D., & Tiemann, M. (2013). Changes in workplace tasks in Germany — evaluating skill and task measures. Journal for Labour Market Research, 46(3), 215–237.

Roth, J., Sant'Anna, P. H., Bilinski, A., & Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. Journal of Econometrics, 235(2), 2218–2244.

Ruf, K., Schmidtlein, L., Seth, S., Stüber, H., & Umkehrer, M. (2021). Linked-employer-employee-data from the IAB (LIAB): LIAB-cross-sectional model 2 (LIAB QM2) 1993–2019. FDZ-Datenreport, 03/2021 (en). (Tech. Rep.). doi: 10.5164/IAB.FDZD.2103.en.v1

Salomons, A., Zierahn-Weilage, U., & vom Baur, C. (2025). Expertise at work: New technologies, new skills, and worker impacts. IZA Discussion Paper, 18248.

San, S. (2023). Labor supply and directed technical change: Evidence from the termination of the bracero program in 1964. American Economic Journal: Applied Economics, 15(1), 136–63.

Schönfeld, G., Jansen, A., Wenzelmann, F., & Pfeifer, H. (2016). Kosten und Nutzen der dualen Ausbildung aus Sicht der Betriebe. Ergebnisse der fünften BIBB-Kosten-Nutzen-Erhebung

(Tech. Rep.).

Siegloch, S., Wehrhöfer, N., & Etzel, T. (2025). Spillover, efficiency, and equity effects of regional firm subsidies. American Economic Journal: Economic Policy, 17(1), 144–180.

Socio-Economic Panel. (SOEP, 2019). Socio-Economic Panel (SOEP), Version 34, Data for years 1984-2017 (SOEP-Core v34). (DOI: 10.5684/soep.v34)

Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands: Looking outside the wage structure. Journal of Labor Economics, 24(2), 235–270.

Statistisches Landesamt, H. (Ed.). (2023). Erwerbstätigenrechnung. Erwerbstätige in den Ländern der Bundesrepublik Deutschland 1991 bis 2022. Berechnungsstand: Mai 2023 (Tech. Rep.).

Stehrer, R. (2024). wiiw growth and productivity database.

Stevens, M. (1994). An investment model for the supply of training by employers. The Economic Journal, 104(424), 556–570.

Terry, S. J., Chaney, T., Burchardi, K. B., Tarquinio, L., & Hassan, T. A. (2026). Immigration, innovation, and growth. American Economic Review, 116(3), 828–861.

Thomsen, U., Ludsteck, J., Schmucker, A., et al. (2018). Skilled or unskilled-improving the information on qualification for employee data in the IAB employee biography. FDZ-Methodenreport, 9(2018), 22.

Walden, G., Beicht, U., & Herget, H. (2009). BIBB-Cost-Benefit-Survey 2000. doi: 10.7803/370.00.1.2.10

Zeira, J. (1998). Workers, machines, and economic growth. The Quarterly Journal of Economics, 113(4), 1091–1117.

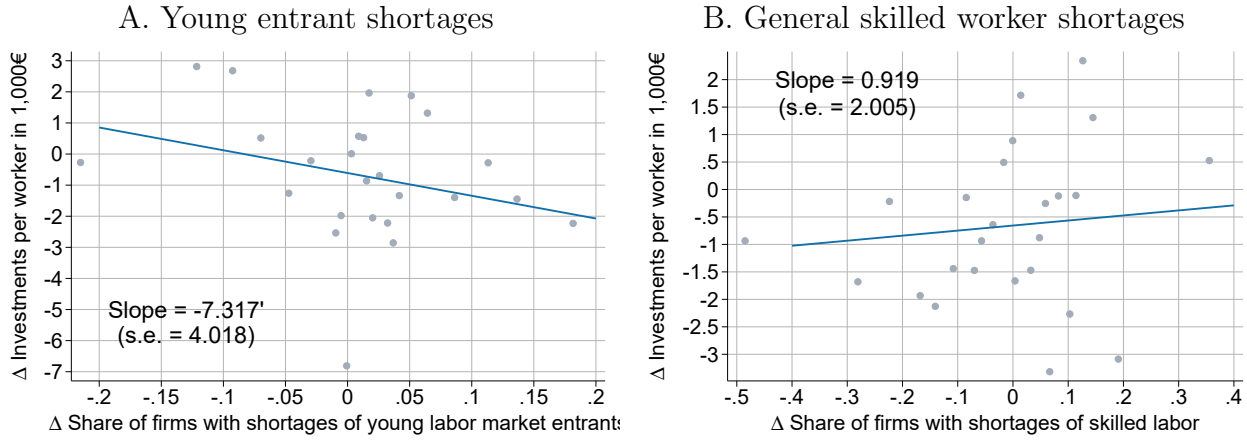
A Data and additional results

A.1 Regional correlations between firm investments and labor shortages

The analysis is based on the IAB Establishment Panel. I use firms' reasons for staffing problems, including shortages of young labor market entrants, shortages of skilled labor, and general staff shortages.[38] I aggregate this information, available for 1997, 1999, 2000, 2004, and 2006, and firm investments per worker to the level of local labor markets, i.e., 401 counties in both East and West Germany. I use regional rather than firm-level shortage indicators to mitigate the concern that firms experiencing labor shortages may differ systematically from those without in ways that confound the investment effect. I drop county-year cells with less than 10 observations and weight observations by the sum of their firms survey weights, with results robust to using other weights.

In the main Figure 1, I exploit yearly fluctuations: I use all county-year observations, controlling for county and year fixed effects. In Figure A1, I analyze long differences over a 10 year time window. The result is comparable: regions with intensifying shortages of young workers have decreasing investments, while this is not the case for general skilled worker shortages.

Figure A1: 10-year differences in firm investments and labor shortages at the regional level



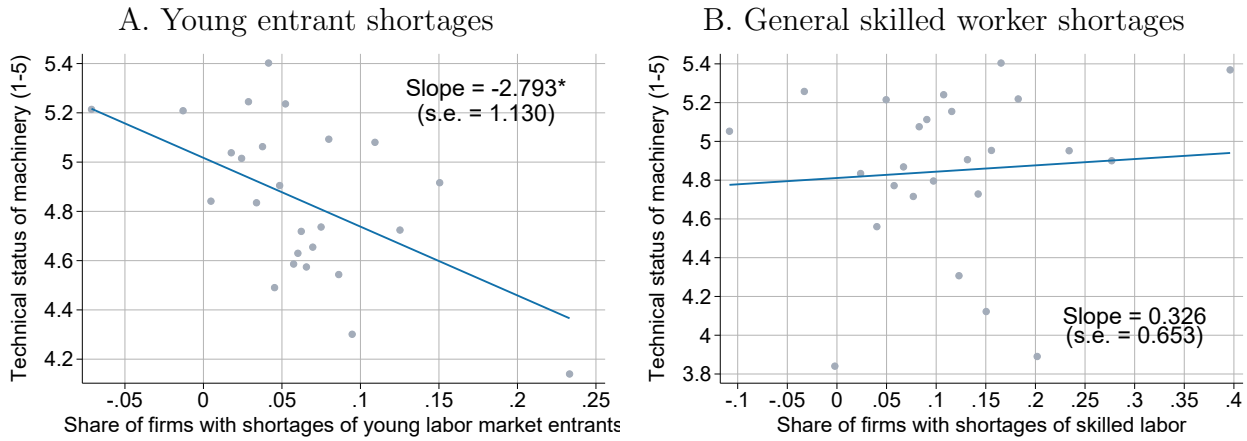
Notes: Binscatter at the county level. Based on the IAB establishment panel survey. Long-differences over 1997–2006; years chosen for data availability reasons. N=181. ' $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, ** $p < 0.001$.

Figure A2 reproduces the main Figure, but looking at firms' self-reported technical status of machinery instead of investments. As described in Section 4, firms are asked to assess the overall technical status of their production equipment compared to other establishments in the same industry on a scale from 1 ('completely out-of-date') to 5 ('state-of-the-art'). Firms in

[38]The exact expression for "shortage of young labor market entrants" in the survey is "Nachwuchsmangel", with "Nachwuchs" referring to the next generation in demographic terms. The expression used for skilled labor shortages is "Difficulties in finding the necessary skilled workers on the labor market".

regions with more severe shortages of young workers report worse technical status' of machinery. Again, this is not the case for general shortages of skilled workers.

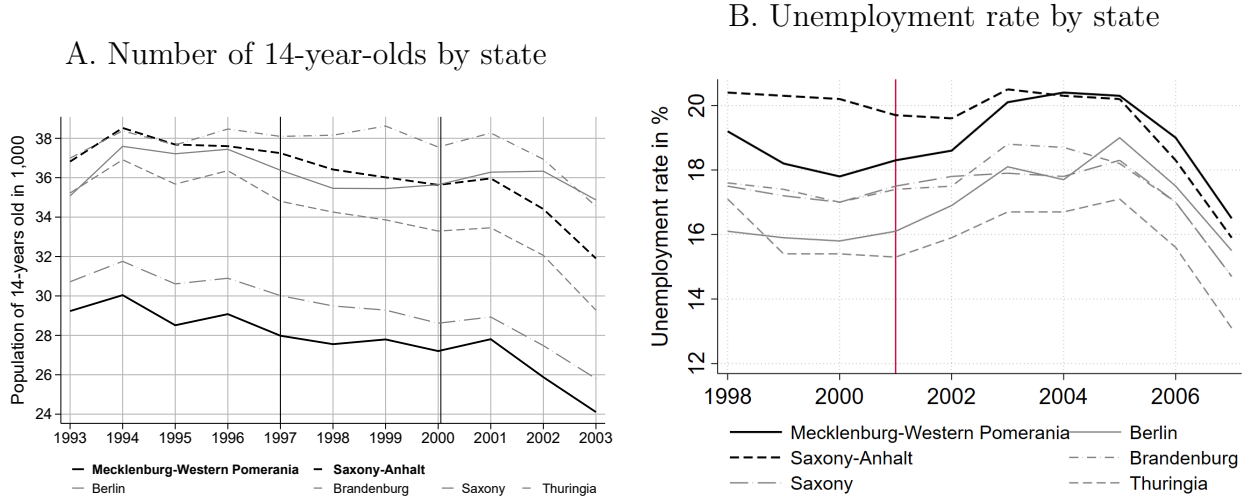
Figure A2: Firms' technical status of machinery and labor shortages at the regional level



Notes: Binscatter at the county level. Based on the IAB establishment panel survey, 1997–2006; years chosen for data availability reasons. N=1,961. Controlling for region and year fixed effects; labor shortages lagged by one year. The outcome is projected technical status of machinery, see Section 4. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

A.2 Other additional results

Figure A3: Demographic and economic trends across federal states



Notes: Panel A: Source: Federal Statistical Office, Genesis-Online (2022b). The number for Saxony is divided by two for better visibility. *Panel B:* Source: Federal Statistical Office (2023a).

Table A1: Correlations between technical status and other firm indicators

	Technical status		Organizational change
	(1)	(2)	(3)
Firm size – Reference category 1-19 employees			
20-99 employees	0.055***	0.096***	
	(0.015)	(0.015)	
100+ employees	0.068***	0.128***	
	(0.019)	(0.019)	
Firm AKM FE	0.070***	0.052***	
	(0.009)	(0.010)	
Technical status			0.037**
			(0.013)
Industry FE		✓	
Year FE + Firm FE			✓
N	46,614		18,991

Notes: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors clustered at the firm level.

Table A2: Projected technical status

	Regression coefficient	Relative importance of regression coefficient
Inv. per worker	0.0007	
	(0.0006)	
Inv. per worker $t-1$	0.0046***	0.50
	(0.0005)	
Inv. per worker $t-2$	0.0032***	0.35
	(0.0005)	
Inv. per worker $t-3$	0.0014**	0.15
	(0.0005)	
Inv. per worker $t-4$	0.0012*	
	(0.0005)	
Inv. per worker $t-5$	0.0014**	
	(0.0005)	
N	28,279	

Notes: Regression outcome: Technical status of machinery in t. Controlling for year and firm fixed effects. Standard errors clustered at the firm level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Relative importance of regression coefficient computed as the regression coefficient divided by the sum of the regression coefficients in $t-1, t-2$, and $t-3$ (0.0092).

Table A3: Correlation of the error term within states and within firms

	Overall SD	SD within states	% of SD within states
1997	7.256	7.163	98.7%
1998	17.051	12.524	73.4%
1999	8.504	7.907	93.0%
2000	7.854	7.267	92.5%
2001	8.561	7.965	93.0%
2002	8.427	8.295	98.4%
2003	7.048	6.822	96.8%
2004	10.919	9.197	84.2%
2005	7.104	7.051	99.3%
2006	7.941	7.633	84.2%
Average			91.4%

Notes: Standard deviations (SD) within years of the error term event study regression of investments per worker following equation (4).

Table A4: DiD Results – Wage and worker substitution effects based on the SIEED

	Shock	Absence of wage and entry responses					
	# highly educated trainees	Log wage highly educ. trainees	# low-educated trainees	% highly educ. commuting trainees	Log highly educ. VT employment	Log employment	Log wage
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treat × Post	-0.26*	0.01	-0.34	-0.05	0.05	-0.01	-0.00
	(0.11)	(0.03)	(0.89)	(0.03)	(0.08)	(0.07)	(0.03)
Mean dep. variable	0.90	3.21	2.63	0.06	0.73	2.10	4.45
N	4,289	2,156	4,289	4,289	2,588	4,100	4,086

Notes: DiD coefficients based on equation (5). Pre: 1997–2000. Post: 2001–2003. Standard errors clustered at the firm level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Mean dep. variable: Average outcome in 1997–2000. Column 4: a commuter is defined as a person living and working in two distinct federal states. This variable is available from 1999 onward. VT: completed vocational training. Log employment and log wages refer to all workers excluding trainees.

Table A5: DiD Results – Highly educated trainees by occupation group

	Production	Personal service	Business service	IT and scientific service	Other commercial service
	(1)	(2)	(3)	(4)	(5)
Treat × Post	-0.05*	-0.03	-0.17*	0.00	0.01
	(0.03)	(0.04)	(0.08)	(0.01)	(0.01)
Mean dep. variable	0.08	0.20	0.46	0.02	0.01
N	4,289	4,289	4,289	4,289	4,289

Notes: DiD coefficients based on equation (5). Pre: 1997–2000. Post: 2001–2003. Standard errors clustered at the firm level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Mean dep. variable: Average outcome in 1997–2000.

Table A6: Robustness of heterogeneity effects

	(1)	(2)	(3)	(4)
Treat × Post × Trainee retention	-3.98	-6.19	-5.71	-6.37
	(3.94)	(4.04)	(4.05)	(4.33)
Treat × Post × Exposure		✓	✓	✓
Treat × Post × Industry			✓	✓
Treat × Post × Firm size				✓
N	2,853	2,853	2,853	2,853

Notes: Standard errors clustered at the firm level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. For each triple interaction term, all corresponding one- and two-way interaction terms are included. Exposure: Number of highly educated trainees. Industry: ten categories. Firm size: three categories. All referring to the average of the pre-values 1997-1999.

B Potential reform confounders

Governing party – Social democrats. The education reform was an initiative of the Social Democratic Party, which entered the government in both treated states in 1994. In Mecklenburg-Western Pomerania, the Social Democrats were the junior governing party of a government led by the Christian Democratic Union. In Saxony-Anhalt, they were the senior governing party but shared power with the Greens. The Social Democrats also entered the government in one of the control states, Thuringia, in 1994 together with the Christian Democratic Union. Nonetheless, to exclude that the governance of the Social Democrats or other policy or socio-economic changes confound the effect of the education reform, I compare several state metrics including population size, education expenditure, unemployment rate, GDP, public debt and public investments between treated and control states before and after the reform, as well as between states governed by the Social Democrats and those not governed by the Social Democrats, see Table B1. Controlling for state and year fixed effects, I find no difference in any of these metrics between treated and control states post-reform compared to pre-reform that is statistically significantly different from zero, see Panel A. Turning to factors correlated with the governance of the Social Democrats, see Panel B, there is a significantly positive association between government of the Social Democrats and three indicators: education expenditure in % of the total state budget, unemployment rate, and log public investments. Higher education expenditure and public investments should, however, rather *in*crease instead of *de*crease firms technology investments. Regarding the unemployment rate, the relation to firm technology adoption is ambiguous. I conclude that major trends at the state level, potentially governed by the party composition of the government, are unlikely to cause the investment drop.

Investment subsidy programs. Two investment subsidy programs may have confounded the effects of the trainee supply shock. Below, I discuss each of them and how they may correlate with the trainee supply shock.

First, an investment tax credit policy was introduced in 1991 that aimed at supporting firms in former East Germany. In 1999, a policy change increased the tax credit rate for smaller manufacturing firms (with up to 250 employees) from 10% to 20%, and for larger firms from 5% to 10%, thereby reducing capital costs more significantly for smaller firms. Lerche (2025) exploits this reform and finds important increases in investments and employment in smaller compared to large firms in response to the reform. This reform is, however, unlikely, to confound the effect of the trainee supply shock given that my identification strategy relies on comparing firms across federal states but the tax credit reform had no regional variation.

Second, Germany's main regional policy, GRW, aimed at revitalizing underdeveloped regions, particularly East Germany, through investment subsidies for (mainly) manufacturing plants. The maximum subsidy rate varied based on counties' economic performance indicators, and was frequently reformed between 1997 and 2014. Siegloch et al. (2025) exploit these

Table B1: Correlation of state metrics with reform and Social Democratic Party

	(1) Log(Population)	(2) Log(Educ. expenditure)	(3) % education expenditure	(4) Unemployment rate	(5) Log (GDP)	(6) Log(Public Debt)	(7) Log(Public Investments)
			A. Education reform in 2001				
Treat × Post	-0.03	-0.01	1.37	-0.25	-0.00	0.15	0.15
	(0.11)	(0.19)	(2.46)	(0.84)	(0.17)	(0.32)	(0.15)
			B. Social democratic party in government				
Social Democrats	-0.01	0.05	0.94*	0.49*	-0.00	0.05	0.19**
	(0.01)	(0.04)	(0.50)	(0.28)	(0.02)	(0.04)	(0.07)
Mean dep. variable	14.83	21.52	27.68	18.08	10.74	9.21	6.09
N	84	66	66	84	84	84	84

Notes: Panel A: Treated: Mecklenburg-Western Pomerania and Saxony-Anhalt. Post: 2001 onward. *Panel B:* Social democrats among governing parties (1/0). Controlling for state and year fixed effects. Observations at the state-year level for East German states for 1992 until 2005, except for education expenditure (column 2 and 3) that is only observed from 1995 onward. Education expenditure: Total public expenditure on education. Share education expenditure: Public expenditure on education as a percentage of the total budget. Unemployment rate: Unemployment rate in % of the dependent civilian labor force. Debt: Debt of the overall public budget. Sources: (1) – Federal Statistical Office (2022) (2) & (3) – Federal Statistical Office (2023b) (4) – Federal Statistical Office (2023a) (5) – Federal Statistical Office (2023e) (6) – Federal Statistical Office (2023d) (7) – Federal Statistical Office (2023c) * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

reforms to study the effect of tax credit on investment and employment, finding important effects. This program had a regional component: while all counties were assigned the same maximum subsidy threshold in 1990, 27 counties that were previously assigned as high-subsidy counties were assigned as low-subsidy county in 1997. In 2000, nine further counties were assigned from high to low, while one was assigned from low to high.

Studying the regional correlation between this potentially confounding reform and the 2001 education reform suggests that the 1997 changes are unlikely to have caused the investment decline in Mecklenburg-Western Pomerania and Saxony-Anhalt following the 2001 trainee shortage: in Brandenburg, half of all counties were downgraded; Saxony saw 38% of its counties affected, and Thuringia experienced a change in 35% of its counties. In contrast, the treated states, Mecklenburg-Western Pomerania and Saxony-Anhalt, were the least affected, with only 25% and 21% of their counties downgraded, respectively. This should have led to increased investments in these states relative to the control states.

The 2000 reform primarily impacted Saxony, where 39% of counties were downgraded from high-subsidy status to low-subsidy status. Brandenburg and Western-Pomerania were not affected at all. Thuringia saw a downgrade in 9% of its counties, and Saxony-Anhalt had a net change of 7% with 14% of the counties changing from high status to low status, and 7% from low status to high status. Therefore, the decline in investments in Mecklenburg-Western Pomerania cannot be explained by these reforms, and it is very unlikely that the changes caused the decrease in Saxony-Anhalt.

C Supporting evidence based on additional datasets

BIBB Cost-Benefit Surveys. The Federal Institute for Vocational Education and Training (BIBB) surveys companies at intervals of several years on the benefits and costs of their vocational training, constituting a dataset called Cost-Benefit-Survey. A total of 2,518 companies took part in the 2000 survey. The data is representative of all German companies engaged in vocational training. The interviews are conducted with people who are primarily responsible for organizing and carrying out in-company training activities in these companies. For more information, see (Walden et al., 2009). For the analysis, I restrict the data to East German firms operating in industries with with above-median increases in log labor productivity over 1998–2003 as based on EU-KLEMS. Among other things, firms are asked "To what extent do the following statements apply to your company's own training?" on a five point scale from "fully applies" to "does not apply at all", with three of 17 statements being "ensures the constant inflow of new knowledge into our company", "significantly improves our ability to adapt to technical and market changes" and "improves the innovative capacity of our company" (question 13 in the 2000 benefit questionnaire).

IAB/BIBB/BAuA Qualification and Career Surveys. This analysis is based on the 1999–2012 waves of the IAB/BIBB/BAuA Qualification and Career Survey (QCS). The QCS are repeated cross sectional surveys conducted by BIBB, IAB, and BAuA.[39] The survey covers around 30,000 employees. I closely follow the data preparation by Arntz et al. (2025)—except that I don't exclude East Germany—and I use the harmonization and data restriction procedure by Rohrbach-Schmidt & Tiemann (2013). The surveys include questions regarding the main working tool used by each respondent. In the 1992 wave, these tools were categorized into (1) non-mechanical tools (e.g. handcart, pencil), (2) tools with some mechanization (e.g. telephone, hand drill machine), (3) tools with advanced mechanization (e.g. car, crane, copy machine), (4) semiautomatic tools (e.g. fax, milking installation, bottling machine) (5) and computer-based tools (e.g. computers, CNC machines). I adopt this categorization for all waves of the survey.

European Working Conditions Survey. I use the 2000 wave of the European Working Conditions Survey (EWCS), in which workers report how often their main job involves working with computers, laptops, or smartphones on a 7-point scale ranging from "all of the time" to "never". From this, I construct a binary outcome variable and regress it on a dummy variable indicating whether the worker is under 30 years old, and a range of detailed covariates, including occupation and sector, as above. I restrict the sample to the EU-15 countries (those countries that were members of the European Union before the 2004 enlargement) and exclude self-employed workers. The results are shown in Table C1: Being young significantly increases

[39] BIBB: Federal Institute for Vocational Education and Training; IAB: Institute for Employment Research; BAuA: Federal Institute for Occupational Safety and Health.

computer use at work in 2000— by between 2.4pp (6%) and 3.5pp (14%), depending on the specification of the outcome variable.

Table C1: Computer use by age group across EU-15

	(1)	(2)	(3)
$<$ 30 years	0.024*	0.035**	0.025*
	(0.012)	(0.011)	(0.011)
Mean dep. variable	0.43	0.25	0.20
N	17,056	17,056	17,056

Notes: Based on the European Working Conditions Survey in 2000, including Belgium, Denmark, Germany, Greece, Spain, France, Ireland, Italy, Luxembourg, Netherlands, Austria, Portugal, Finland, Sweden, and United Kingdom. Based on the survey question "Does your main job involve working with computers, laptops, or smartphones? – (1) All of the time, (2) Almost all of the time, (3) Around 3/4 of the time, (4) Around half of the time, (5) Around 1/4 of the time, (6) Almost never, (7) Never. Column 1: Computer use equal to one for (1)–(5); and zero otherwise. Column 2: Computer use equal to one for (1)–(3); and zero otherwise. Column 3: Computer use equal to one for (1)–(2); and zero otherwise. Controlling for country fixed effects, broad industry categories (4), gender, public sector employment, country-specific occupation fixed effects, and tenure at firm in three-year bins. Heteroscedasticity-robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

D Instrumental variable regression

In this Appendix, I employ a complementary identification strategy that serves three main purposes. First, it aims at closely linking the investment decline to the trainee shortage by studying whether firms that are more exposed to the trainee supply reduction decrease investments more. Second, it allows me to identify the treatment effect independent of the realized distribution of trainees across firms. Third, it identifies a different causal parameter: While the event study identifies the causal effect of facing a statewide reduction in trainee supply, this complementary analysis identifies the causal effect of having access to one additional trainee.

I estimate a two-stage-least-squares (2SLS) model of firm investments Inv on firm employment of highly educated trainees N^{Trainee} controlling for firm fixed effects π_j and year fixed effects ψ_t, see equation (D1). I instrument trainee employment as given in equation (D2):

$$\text{Inv}_{jt} = \alpha N_{jt}^{\text{Trainee}} + \psi_t + \pi_j + \epsilon_{jt} \tag{D1}$$

$$\begin{aligned} N_{jt}^{\text{Trainee}} = &\sum_t \gamma_t (N_{j,1996/97}^{\text{Trainee}} \times \text{Treat}_j \times \text{Year}_t) \\ &+ \sum_t \zeta_t (N_{j,1996/97}^{\text{Trainee}} \times \text{Year}_t) + \psi_t + \pi_j + \epsilon_{jt} \end{aligned} \tag{D2}$$

with j firms, and t calendar years. Treat takes the value one if a firm is located in a state undergoing the education reform and zero otherwise. I predict contemporaneous trainee employment by firm exposure to the shock, corresponding to the *shares* in a shift-share instrument, times the reform-induced *shifts* in the supply of trainees across states and years, Treat $\times$ Year.[40] Firm exposure to the shock, $N_{1996/97}^{\text{Trainee}}$, is measured by firms' mean employment of highly educated trainees in October 1996 and 1997 (based on information from June 1997 and June 1998). I control for year effects in firm trainee employment which are allowed to vary by firm exposure, $N_{j,1996/97}^{\text{Trainee}} \times \text{Year}_t$. Hence, the instrument exploits variation between two equally exposed firms located in a treated and a control state across time. The exogeneity of the instrument stems from the random assignment of the trainee supply shock, i.e. the education reform, to states and years. Since employment of highly educated trainees in 1996/97 is expected to directly impact investments around this time, which would violate the exclusion restriction, I run the regression for the years 1999 onward.[41]

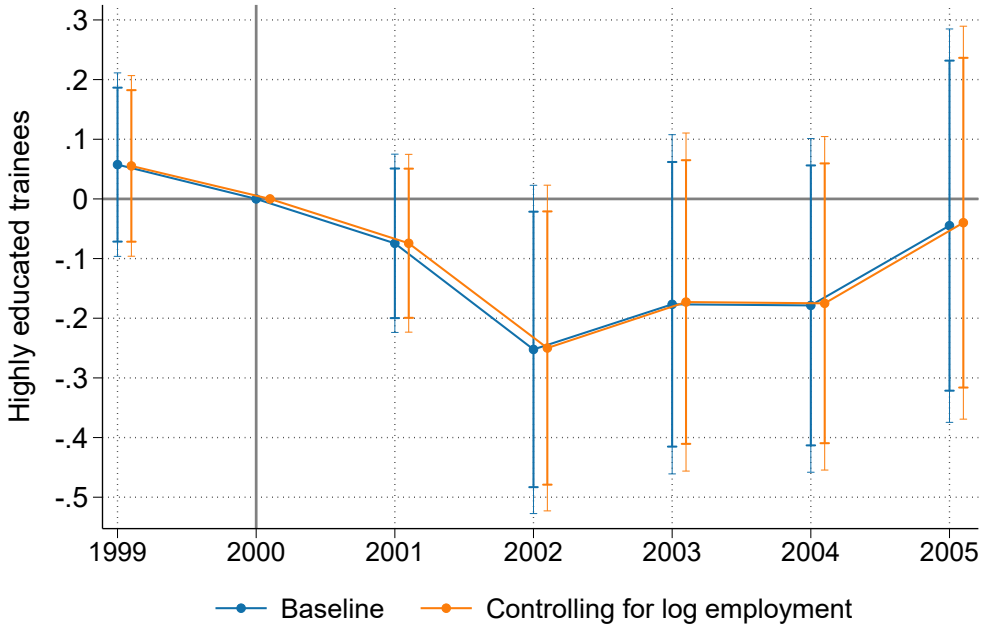
For this analysis, I use both training and non-training firms to exploit the variation in initial trainee employment between them. I use the sample of matched firms, repeating the same matching procedure as above for training and non-training firms separately, ensuring that

[40] While the investment information refers to the calendar year, trainee employment is recorded as of June 30 each year. To account for the fact that trainee employment in a given calendar year includes trainee employment as of June 30 for the months January to August (8 months), and employment as of June 30 of the next year for the months September to December (4 months), I construct contemporaneous trainee employment as a weighted average of this and next year's trainee employment, weighting this year's trainee employment by 8/12 and next year's trainee employment by 4/12. This approach forces me to stop the analysis in 2005. Results are robust to alternative definitions, such as using only this year's or next year's trainee employment.

[41] Results are robust to further restricting to the years 2000 onward.

treated and control firms are comparable in terms of sector, size, and employment of highly educated trainees.

Figure D1: IV results – First stage



Notes: Coefficients plus 90% and 95% confidence intervals of the term ($N^{\text{Trainee}}_{j,1996/97} \times \text{Treat}_j \times \text{Year}_t$) in equation (D2). Outcome: Number of highly educated trainees. Years refer to full calendar years, combining trainee employment of the trainee entry cohort from September last year and employment of the trainee entrant cohort from September of the current year. Standard errors clustered at the firm level.

Figure D1 shows the coefficients of interest of the first stage, γ_t. As expected, coefficients are negative in the years of the shock. One additional highly educated trainee prior to the reform is associated with approximately 0.25 fewer highly educated trainees in 2002 to 2004. Consistent with the timing of the shock, the coefficient in 2005 is zero. Please note that the year now refers to the calendar year, explaining why the trainee employment drop is more nuanced in 2001 than in the event study. The first stage and an F-statistic of 17.31, see Table D1, column 1, thus confirm the relevance of the instrument.

Table D1 shows the results of the second stage, looking at investments per worker as the outcome. One fewer trainee reduces investments by €830 per worker, equivalent to 5.8% of average investments (column 1). To ensure that the relationship between trainees and investments is not driven by the role trainees play in firm employment (growth), I control for time-variant log employment in column 2. Convincingly, the result remains very similar.

In summary, the complementary identification strategy confirms the negative impact of reduced trainee supply on firm investments. The effects are smaller than the ones implied by the ratio between missing trainees and missing investments in the event study regression. This discrepancy might be due to spill-over effects within treated states, i.e. firms decreasing investments beyond the first-order decrease related to foregone trainee employment. It might

Table D1: IV results – Second stage

	(1)	(2)
$N^{Trainee}$	0.829*	0.798*
	(0.354)	(0.349)
Firm log employment		X
F-Stat	17.31	17.01
N	5,419	5,419

Notes: Outcome: investments in €1,000 divided by total employment in 1997. F-Stat gives the robust Kleibergen-Paap Wald rk F statistic. Based on the matched sample, 1999-2005. Standard errors clustered at the firm level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

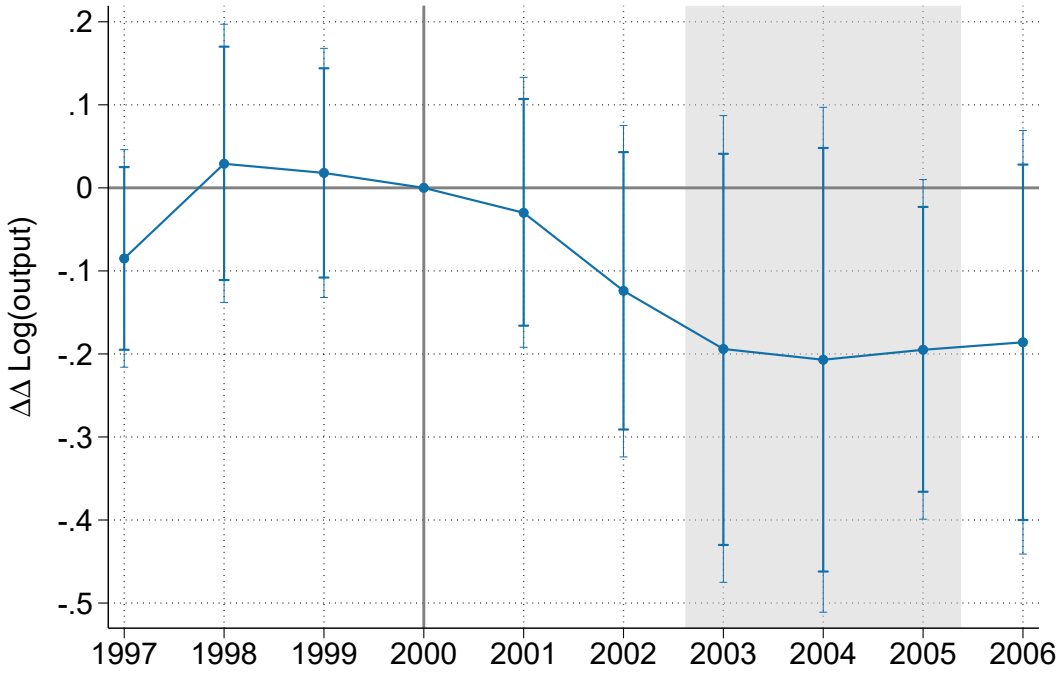
also hint at firm selection into trainee employment: If firms that would have invested in the absence of the supply shock employ fewer trainees than firms who would not have invested anyway, the parameter identified in the event study approach is large, while the parameter identified in the IV approach is unaffected by this selection. A third reason may be trainee quality: firms employing trainees during the shortage also reduce investments because the trainees they employ are of worse quality.

E Downstream effect on firm output

Foregone technology adoption should affect firm output. In the survey, firms can choose between reporting one of the following types of output measures: sales, balance sheet, premiums, budget. While the majority reports sales, I keep all values, only excluding firms that switch the measure during the time period of interest. This variable is less often filled than the investment variable, making me lose 12% of the observations. Firm output only reacts very indirectly to the education reform, is potentially more exposed to confounding demand shocks than investments, and suffers from inconsistent measurement as just described. I therefore run the triple differences analysis which cleans the effect from potential time-varying confounders at the industry-state level. However, even this specification should be treated with caution. The resulting estimates differ from the effect on investments in two ways: First, they occur with a time lag and are longer lasting given that output responds to changes in technology adoption only after the associated investments have materialized. Second, they are less precise, due to a lower fill rate of the variable, and consistent with these being indirect effects.

Figure E1 shows the effect, comparing the difference between training and non-training firms in treated states to the difference between training and non-training firms in control states. Output is reduced by approximately 20% from 2003 onward, with the coefficient being statistically significant at the 10% for 2005.

Figure E1: Triple differences specification – Effect on firm output

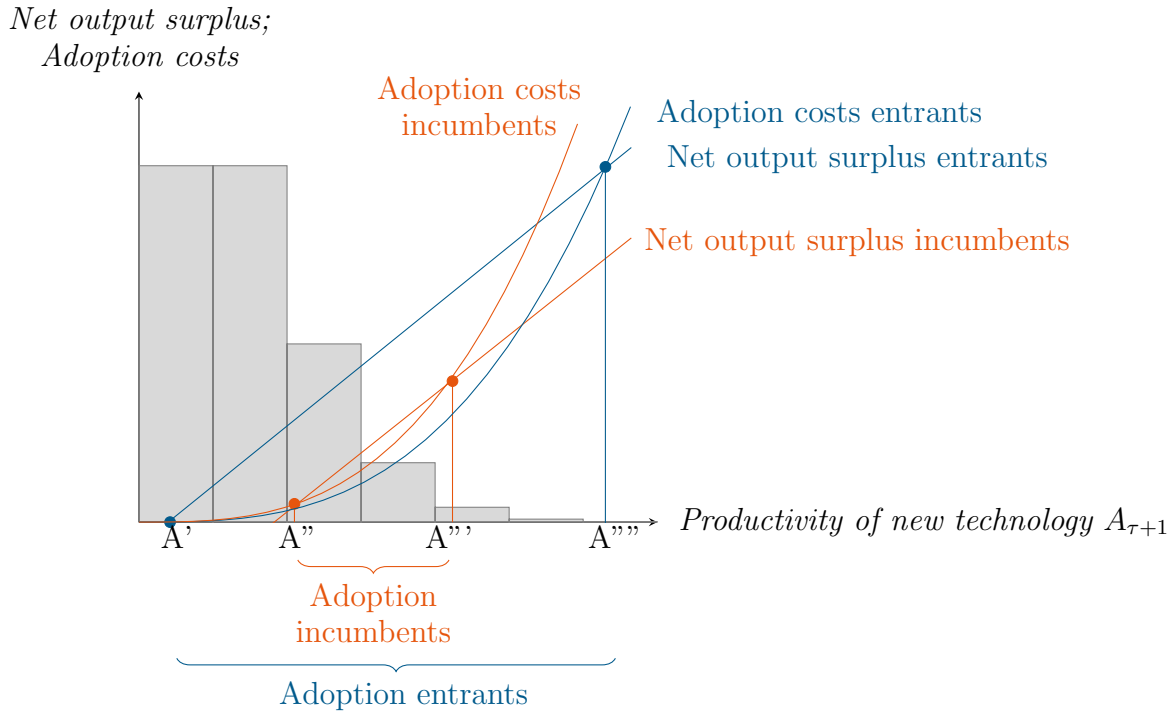


Notes: Event study coefficients of the interaction terms Train × Treat × Year plus 90% and 95% confidence bands. Based on equation (6). N=12,863. Standard errors clustered at the firm level.

F Extensions of the economic framework

Functional form of adoption costs. The main insight of the framework is that technology adoption is determined by the relationship between the productivity of a new technology and its adoption costs. It is therefore intuitive that the functional form of the adoption costs is decisive for the adoption decision: In the case of constant or linearly increasing adoption costs, as assumed until here, a missing entrant cohort will reduce investments for semi-productive technologies, while investments in unproductive and highly productive technologies are unchanged. We now turn to the case of adjustment costs being increasing and convex in investment size, as standard neoclassical investment models assume. Convex capital adjustment costs capture the plausible idea that rather unproductive technologies can be easily incorporated into the firm's structure without much training, while highly productive technologies tend to cause more pronounced disruptions requiring more intense training. The costs and benefits of technology adoption with increasing and convex adoption costs are shown in Figure F1. As before, new technologies below a productivity threshold A' are not adopted; new technologies above A' but below A'' are adopted by training labor market entrants only; and new technologies above a threshold A'' are adopted by retraining incumbent workers as well. As a new feature, highly productive technologies above a certain productivity threshold A''' are only adopted by entrants because retraining incumbents would be prohibitively costly, and even more productive technologies above a certain productivity threshold A'''' are not adopted for any worker group because training costs exceed productivity gains for all workers. Consequently, a lack of entrants not only hinders the adoption of rather unproductive new technologies in the range between A' and A'', but also of very productive technologies with productivities in the range between A''' and A''''.

Figure F1: Firms' costs and benefits of technology adoption with convex adjustment costs



Notes: Profitability of assigning new technology-using task to entrants versus incumbent workers when capital adjustment costs are increasing and convex in technology productivity. The histogram shows the distribution of the productivity increase related to the new technology vintage.