Industrial Policy at Work: Evidence from Romania's Income Tax Break for Workers in IT*

Isabela Manelici† and Smaranda Pantea‡

February 16, 2021

Abstract

We study the firm and sector-level effects of an industrial policy designed to support the development of the IT sector in Romania. In 2001, Romania introduced an unexpected personal income tax break to programmers with eligible bachelor's degrees and who work on software development for firms in eligible IT sector codes. In 2013, policy-makers suddenly expanded the scope of the original tax break to cover more bachelor's degrees and sector codes in IT. We first use firm-level data and difference-in-difference designs around each policy episode to show that treated firms experience strong and long-lasting growth. We then employ sector-level data and a synthetic control design to show that after the introduction of this policy in 2001, the IT sector grew faster in Romania than in otherwise similar countries. Finally, downstream sectors relying more on IT services also grew faster in Romania after 2001. Our results suggest that this policy has been effective in promoting the development of the IT sector, a sector typically seen as key to the transition to a knowledge economy.

Keywords: Industrial policy, Firm growth, Economic development, Information technology, Labor income taxation, Central and Eastern Europe, Downstream effects

JEL Codes: O25, O14, O38, L86, H24, D22, L25, O52, D57

*We thank Treb Allen, Alan Auerbach, Marc Bogdanowicz, Sebastian Buhai, Teresa Fort, Eva Hagsten, Douglas Irwin, Nathan Lane, Ethan Lewis, Dan Matei, Brian McCaig, Krisztina Orban, Nina Pavcnik, Andrés Rodríguez-Clare, Jose P. Vasquez, Emmanuel Saez, Benjamin Schoefer, Mine Senses, Danny Yagan, and Paolo Zacchia, two anonymous referees, and the editor for helpful comments and suggestions. We also appreciate the feedback received during seminar and conference presentations. We thank the National Center for Financial Information of Romania for granting us access to firm-level data, and Dan Matei and Ştefan Nanu (from the Ministry of Public Finance of Romania) for their decisive support. We thank Andrei Caprau, Mihai Codreanu, Micole De Vera, and Madeline Duhon for excellent research assistance. Last, we thank Ana-Maria Cazacu, Mihai Copaciu, and Irina Mihai (from the National Bank of Romania) for facilitating access to historical sector-level statistics. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The views expressed herein are those of the authors and do not necessarily reflect the views of either of the institutions to which the authors are affiliated.

†Department of Economics (Princeton University). Email: manelici@princeton.edu

‡Directorate of Policy, Analysis, and Research in Public Economics (Ministry of Public Finance of Romania) and the Departments of Entrepreneurship and of International Economic Relations (Prague University of Economics and Business). Email: smaranda.pantea@vse.cz

1 Introduction

"Nations have and will continue to shape their economies through industrial policy. Nevertheless, the empirical literature on these interventions is thin, dwarfed by the attention industrial policies receive from policymakers across the world" [according to a contemporaneous review of this literature by Lane, 2020]. In this article, we contribute with plausibly causal evidence on the effectiveness of an industrial policy introduced by Romania in 2001. This policy targeted the development of its information technology (IT) sector, a sector seen as key to the transition to a knowledge economy.

Since 2001, this policy provides a full personal income tax break to employees with an eligible bachelor's degree who work directly on software development and generate revenues from this activity for a firm in the eligible "Software consultancy and supply" sector. In 2013, an amendment to the 2001 law greatly expanded the pool of eligible firms and workers by adding several newly eligible sector codes for the firms and bachelor's degree specializations for the workers. This intricate set of rules – on the worker, the firm, and the activity performed by the worker in the firm – implies that the tax break rewards specific matches between workers and firms. The expectation was that lowering the tax burden on these matches would lead to an increase in their prevalence. Moreover, while the statutory incidence of the tax break is on workers' personal income, workers can only benefit from this break through their employer. Also, their eligibility is conditional on the characteristics of their employer and their role within the firm. These requirements of an explicit "buy-in" from the employer motivate the expectation of a shared economic incidence of this tax incentive. Finally, the restrictive set of conditions of the law ensures that the tax break actually subsidizes software development (as opposed to other misreported activities).

In the first part of the paper, we use firm-level data and difference-in-differences (DiD) designs to study the effects of the introduction of the tax break in 2001 and its 2013 reform. For the 2001 analysis, we use firm-level data from Amadeus. In the 2001 DiD strategy, a firm is considered treated if it belongs to the sector that became targeted by the law ("Software consultancy and supply" sector, NACE Rev 1 code 722). We compare the outcomes of firms in the eligible IT sector to the outcomes of firms in ineligible high-tech knowledge-intensive (HTKI) service sectors. For identification, we rely on the unexpected nature of the passing of the law and the lack of preexisting differential trends in outcomes between treated and comparison firms. Our estimates show that after 2001 firms in the eligible IT sector embark on a differential upward trend relative to firms in ineligible HTKI sectors. By 2005, firms in the eligible sector have, on average, 24% higher operating revenues than in 2000, 13% more workers, and 12% more assets than firms in comparison sectors. These DiD estimates measure the "intent-to-treat" effects of the policy on firms in the eligible IT sector. We use two alternative sets of comparison sectors to show that results are not an artifact of the baseline choice of comparison sectors.

Next, we examine the 2013 amendment to the income tax break law, which greatly expanded the set of eligible firms and workers. We base the 2013 analysis on more comprehensive administrative data that includes information on the number of income tax exempt employees. This allows us to estimate the effects of this reform on firms whose programmers actually benefit from the now more widely available income tax break. Therefore, we refine the definition of treatment from one based on the sector of the firm to one based on the extent of workforce exemption. We classify a firm as treated if it jumps from under 5% of workforce exemption before 2013 to over 20% of workforce exemption after 2013. The reference group contains other firms in HTKI service sectors (eligible and ineligible) that remain with under 5% of workforce exemption throughout the entire sample period.

The identification relies on the unexpected and generous expansion in firm and worker eligibility that occurred in 2013 – the most plausible driver of the sudden jump in the firm-level share of tax-exempt workers – and the lack of preexisting differential trends between treated and comparison firms. We find that firms treated by the 2013 reform experience large and long-lasting increases in size. In 2015, these firms have 20% more revenues than in 2012, 10% more employees, and 17% more assets (all relative to the comparison firms). These baseline estimates are robust to (i) using three alternative comparison group of firms in other ICT service sectors, in non-ICT HTKI service sectors, and only in the eligible IT sectors; (ii) running the baseline regression on a dataset from Amadeus that allows for a longer period over which to check for pre-trends; (iii) varying the threshold choice from 20 to 10, 15, 25 and 30%; and (iv) defining treatment based only on the sector of the firm, as in the 2001 DiD exercise.

These firm-level results corroborate the hypothesis of a shared economic incidence of the tax break between workers and firms. Lacking worker-level data, we do not attempt to estimate how the tax incentive is split between workers and firms.[1] To interpret the magnitudes of the estimates from the two analyses, we first assume a 25% take-up between 2001 and 2005 for firms in the eligible sector. We can then translate the 13% DiD ("intent-to-treat") 2005 estimate on employment into a 52% treatment-on-the-treated estimate. As the median firm in the eligible sector has three workers in 2000, a 52% increase in employment between 2001 and 2005 is sensible. This 52% estimate is larger than the corresponding 11% estimate from the 2013 analysis. This is most likely due to the fact that the "early adopters" (i.e., firms whose workers become exempt from the income tax just after 2001) are positively selected relative to the "late adopters." That said, the very purpose of the 2013 reform was to expand the scope of the tax break to include firms and workers unable to benefit from the tax break beforehand.

In the second part of the paper, we switch to a sector-level cross-country study of the impacts of the 2001 introduction of the tax break. Using data from Eurostat and the World Bank and the synthetic control method (SCM), we evaluate the effects of this policy on the relative growth of the IT sector. This analysis is complementary to the firm-level analysis in two ways. First, it captures not only the within-firm intensive margin of growth of the IT sector but also its extensive margin (i.e., the entry of new firms into the IT sector). Second, because we benchmark the relative growth of the IT sector in Romania to that in a synthetic Romania based on comparable Central and Eastern European (CEE) countries, we control for potentially confounding sector-specific productivity or demand shocks.

The SCM estimates indicate that, in 2015, the gross revenues (employment) in the IT sector of Romania are 6.52 (1.83) times larger than the gross revenues (employment) in 2000. This value reflects the exceptional growth of the IT sector in Romania – plausibly owed to the 2001 policy – as it is relative to the growth of gross revenues (employment) in all other sectors in Romania and relative to the same difference in growth rates in synthetic Romania. Given this extra double-differencing – which controls for broader trends in the rest of the Romanian economy and similar economies – the SCM estimates of the growth of the IT sector in Romania are smaller than its actual growth between 2000 and 2015 (14-fold in gross revenues and six-fold in employment). Placebo tests suggest that a similar growth cannot

[1] Using matched employer-employee data from Sweden, Saez et al. [2019] finds that a payroll tax reduction for young workers has had a full incidence on firms, who saw increases in their sales, number of workers, and capital (as in this paper).

be replicated in countries that did not implement this policy.

Next, we provide evidence on the inter-sector effects of the tax break, again using SCM on Eurostat, Comtrade, and World Bank data. The improvements in the prices, quality, and variety of IT services – which are likely to have occurred alongside the expansion of the IT sector – are expected to benefit more those downstream sectors which have a stronger reliance on IT services as inputs. We, therefore, ask whether sectors that relied more on IT services before the tax break expanded more than sectors with less of such a reliance (in Romania, relative to synthetic Romania). We group sectors into high- and low-intensity of use of IT services, based on the share of the IT sector in their total input expenditures (according to the input-output table of Romania for the year 2000).

These SCM results suggest that, after 2001, the high-intensity sectors of Romania grew more than its low-intensity sectors and more than in synthetic Romania (e.g., 0.75 times more in terms of gross revenues and 0.61 times more in terms of employment). Moreover, high-intensity sectors also improved their export performance more, which suggests a shift in Romania's comparative advantage. These results serve two purposes. First, they represent an additional indirect check on the effectiveness of the policy of interest, particularly due to the delayed onset of the downstream SCM effects and their smaller magnitude (both relative to the direct SCM effects). Second, while not a definitive test, the faster growth of downstream sectors relying more on IT services as inputs is a necessary condition for the IT sector to have generated inter-sector externalities (one of the theoretical conditions that justify industrial policies favoring a certain sector). The SCM results survive a battery of robustness checks.

Finally, we propose a back-of-the-envelope cost-benefit calculation for the 2013 reform. We estimate that for each extra euro in foregone tax revenues in 2015 due to the 2013 reform, treated firms generated 7 euros extra in revenues and accumulated 3.7 euros extra in assets. These numbers suggest that the direct benefits of the reform for treated firms outweigh its costs to the government.

Most directly, this article contributes to a long-standing academic debate on industrial policy. Academics are typically skeptical about the effectiveness of industrial policies due to possible capture by "sunset sectors" or lobbying firms, and difficulties to design and implement successful industrial policies (particularly so in countries with weak institutions). Despite this skepticism, the set of papers that provide well-identified reduced-form evidence on the effectiveness of industrial policy is rather small and recent [see the review of Lane, 2020]. Görg et al. [2008], Aghion et al. [2015], Juhász [2018], Criscuolo et al. [2019], Cai and Harrison [2019] are among these recent papers providing credible evidence, through only a minority pertain to emerging economies [a notable shortcoming of this literature Rodrik, 2008]. We bring to this debate evidence on the effectiveness of a Romanian industrial policy with a unique design[2] and targeting a sector widely regarded as systemically important.

The natural follow-up question is whether an effective industrial policy is also efficient. While answering this question lies outside the scope of this paper, this tax break appears to meet the theoretical criteria for welfare-improving industrial policy put forward by the literature. First, it encourages software development, a "new" activity for the domestic economy in 2001 [Rodrik, 2004], and one that is knowledge-intensive [Aghion et al., 2011, European Commission, 2017, Cherif and Hasanov, 2019]. Second, given its skill endowment in STEM inherited from communism, Romania most likely had a

[2] The most commonly studied industrial policies involve tax breaks or credits for capital and R&D [Fowkes et al., 2015, Boeing, 2016, Cai and Harrison, 2019], grants [Görg et al., 2008, Criscuolo et al., 2019], or trade tariffs [Aghion et al., 2015]

latent comparative advantage in this activity and only lacked a policy signal to tilt resources towards it [Rodrik, 1996, Harrison and Rodríguez-Clare, 2010, Itskhoki and Moll, 2018]. Third, we find that the growth of the IT sector has supported the growth of IT-using sectors – a necessary condition for the sector to generate inter-sector externalities.[3] Last, this policy benefits a sector that is competitive, and in that sector, all firms and workers meeting the eligibility criteria [Aghion et al., 2015].

By studying a policy that targets the IT sector, we naturally relate to research on this specific sector. One strand of this literature establishes the wide-ranging effects of IT, of which the effects on productivity have garnered the most attention.[4] In particular, Van Ark et al. [2008] makes the case that the later emergence and smaller size of the IT sector in the European Union (compared to the United States) explain its slower productivity growth. This makes the policy we study especially relevant to countries that grapple with the drawbacks of an underdeveloped IT sector. Another strand of this literature studies the drivers of growth of the IT sector. To our knowledge, this is the first paper to provide plausibly causal estimates of the effectiveness of tax incentives to boost firm size in the IT sector.[5]

Finally, given the specifics of the policy, we also relate to research on policies aimed at reducing non-wage labor costs. Most of this research evaluates the effects of reductions in non-wage labor costs for hard-to-employ workers (such as unemployed individuals, youth, parents returning to work, or people with disabilities) or wage subsidies that aim to support job creation in general. The typical finding is one of positive effects on firm size.[6] For instance, Kangasharju [2007] studies the effects of wage subsidies for hard-to-place job seekers (mostly long-term unemployed) in Finland and finds a pooled DiD estimate of a 9% increase in employment. Our larger estimates for IT firms (e.g, of 11% in the 2013 analysis) are most likely explained by the higher value of the incentive for the IT sector; the labor costs associated with programmers are likely to be more significant to IT firms than those associated with hard-to-employ workers (typically hired in low-skilled support positions, such as cleaners, by firms in all sectors).

By studying a personal income tax break for a high-skill/high-wage occupation, we also relate to a smaller set of papers that study reductions in taxes on the wages of R&D workers. As most papers study the effects over a short-term horizon – during which the supply of researchers is most likely inelastic – their most common finding is an increase in researcher wages [Goolsbee, 1998, Lokshin and Mohnen, 2013].[7] Our positive effects on employment most likely stem from a more elastic supply of programmers (relative to researchers), the fact that this policy reduced the incentives of Romanian programmers to emigrate (thus increasing the number of programmers available to firms in the country), and the fact that

[3] An important theoretical motive to deviate from policy neutrality requires that the targeted sector would later generate externalities [Succar, 1987, Greenwald and Stiglitz, 2006, Harrison and Rodríguez-Clare, 2010]. One such externality occurs between sectors, through the supply of specialized inputs used by (many or high-technology) downstream sectors. By studying not only the targeted sector, but also those affected through I-O linkages, we relate to both seminal [Hirschman, 1958, Pack and Westphal, 1986] and recent work on the I-O implications of industrial policy [Forslid and Midelfart, 2005, Du et al., 2014, Blonigen, 2016, Huremović and Vega-Redondo, 2016, Lane, 2017, Liu, 2019, Joya and Rougier, 2019].

[4] IT has been found to improve productivity in IT-using firms [Jorgenson et al., 2008, Syverson, 2011, Bloom et al., 2012, de Ridder, 2020], increase wages in high-skill locations [Forman et al., 2012], reduce information frictions [Steinwender, 2018], foster service exports [Kneller and Timmis, 2016], improve educational attainment [Beaudry et al., 2010], etc.

[5] Several studies provide descriptive evidence on the same industrial policy studied here [e.g., Grigoraş et al., 2017].

[6] For examples on positive effects on employment, see Kangasharju [2007], Girma et al. [2008], Lombardi et al. [2018], Saez et al. [2019], Banai et al. [2020]; on positive effects on profits, revenues, and investment, see Månsson and Quoreshi [2015], Saez et al. [2019], Banai et al. [2020]. For a review, see Eurofound [2017].

[7] For a review, see European Commission [2014]. In addition to wage effects, a subset of these studies also find positive effects on the employment of researchers [Guceri, 2018] or other high-skilled workers [Kaiser and Kuhn, 2016].

the 2013 amendment acted like a shock to the supply of programmers eligible for the income tax break.

The remainder of the article is organized as follows. Section 2 describes the two policy episodes of interest: the 2001 introduction of the income tax break to workers in IT and its 2013 reform. Section 3 presents our firm-level empirical strategy and findings. In Section 4, we bring sector-level cross-country evidence on both the direct and downstream effects of the 2001 introduction of the tax break. Section 5 provides a back-of-the-envelope cost-benefit calculation and Section 6 concludes.

2 Romania's Income Tax Break for Workers in IT

In 2001, high labor taxes were seen as a major constraint to the development of the IT sector in Romania. At the time, personal income taxes were progressive, with rates between 18 and 40%. Payroll taxes (social security contributions and insurance) were paid by both employers (up to 40%) and employees (17%) on gross salary. While the tax rates were the same across sectors, the progressive taxation combined with deductions based on household characteristics and gross salaries, led to varying ratios between net and gross salaries. In particular, the IT sector had a lower ratio (0.68) than that of the entire economy (0.75) or comparable sectors (over 0.70), or put differently, a higher burden of taxation. In addition, programmers are relatively better positioned to have their qualifications recognized abroad, due to lower language barriers and more standardized ways to appraise IT skills. Together, these reasons were seen as important drivers of the high emigration rates of programmers.[8] High labor taxes also led to relatively high labor costs for firms, limiting their growth.

Since 2001, Romania's IT sector has experienced dramatic shifts. The sector has greatly expanded, both in absolute terms and as a share of GDP. It has also become more integrated into the global economy through flows of foreign direct investment (FDI) into the sector and the growing importance of foreign revenues. IT-related bachelor's degrees remain among the most popular degrees to this day. A tax break introduced in 2001 and expanded in 2013, is widely perceived as having triggered these shifts.

In 2001, Romania introduced a personal income tax break for programmers at the initiative of Mr. Varujan Pambuccian, a member of the Chamber of Deputies. He expected that this tax break would address the concerns of both workers and firms in the IT sector. Mr. Pambuccian was, and still is, the representative of the Armenian minority in Romania in the Chamber of Deputies. For this reason, his proposal was not part of the program of a major party and widely discussed during the electoral campaign. It also did not benefit from automatic support in the Chamber of Deputies. Eventually, however, Mr. Pambuccian managed to convince the Prime Minister and the government about the potential benefits of this exemption. His first proposal had been to reduce the top marginal income tax rate on the salaries of programmers from 40 to 8%. The Ministry of Public Finance assessed this proposal difficult to implement. Instead, the final version entailed a full exemption from the income tax for programmers. The tax exemption was adopted by Government Emergency Ordinance 94/2001 and only later approved as law by the Parliament. Given the implementation difficulties and criticism related to the first version, in addition to the initial adoption by Government Emergency Ordinance, it is plausible that the

[8] Based on data generously shared by Giesing and Laurentsyeva [2021], we estimate that while the emigration rate of Romanian programmers was 16.7% in 2000, it fell to 4.8% by 2006 (the year before Romania acceded to the EU).

introduction of the tax exemption was an unexpected and positive shock for the IT sector.

To benefit from the tax break, workers had to simultaneously (i) have an eligible bachelor's degree (in either automation, computers, informatics, cybernetics, mathematics, or electronics); (ii) work for a firm in the "Software consultancy and supply" sector (NACE Rev 1 code 7220); (iii) work for the unit in charge of software development; (iv) have an eligible occupation title (e.g., "programmer" or "computer systems designer"); and finally (v) work for a firm that kept separate balance sheets recording revenues from software development and generated gross revenues of at least 10,000 U.S. dollars from this activity in the previous fiscal year (per exempted employee). An important feature of the policy is that, while the exemption applies to the income tax owed by workers, the firm is responsible for preparing the justifying documents, applying for the tax break, and archiving the documents for potential audits.

Four features of this tax break deserve emphasis. First, this break was particularly generous at its introduction. In 2001, the programmers' wages were among the highest in the country. Before the full income tax break, programmers faced a progressive personal income tax with rates between 18 and 40%. In 2005, all workers in the country saw a switch from progressive taxation to a flat rate of 16%. Between 2004 and 2005, the tax wedge for workers earning the average economy-wide salary decreased by two percentage points and for high wage earners by four.[9] Despite this change, the full exemption from the income tax of 16% and the trend of growing salaries in the IT sector meant that this break remained a sizable incentive for the IT sector. Second, the rules to benefit from the tax break were meant to ensure that the economic incidence of the tax break was shared between workers and firms. The policy rewarded specific types of matches between workers and firms engaged in software development activities. We expect that lowering the tax burden on these activities led to an increase in their prevalence. We document the equilibrium effects of this policy using firm- and sector-level data. Third, its strict accountability rules (explained above) ensured that exempted workers were actually developing software. Hence, the effects we estimate are plausibly real responses to the incentive and not a mere relabeling of activities.[10] Finally, this policy was designed to benefit all eligible workers and firms in the IT sector. Industrial policy works better when benefits are less concentrated [Aghion et al., 2011].

The first major amendment to the tax break law occurred in 2013. It was initially prompted by the need for Romania to transition from the NACE Rev 1.1 classification of sector codes to the NACE Rev 2 classification.[11] Similar to the initial introduction of the policy in 2001, the passing of this amendment and its final eligibility criteria were made uncertain by the fact that the negotiations spanned three different governments in a period of high political instability. The final form was first approved as Order 539/225/1479/2013, to be only later adopted as law. The amendment greatly expanded the scope of the

[9] The first statistic characterizes single workers without children who earned 100% of the economy-wide average salary and whose tax wedge decreased from 28.1% to 26.1%, whereas the second characterizes single workers without children who earned 167% of the average salary and whose tax wedge decreased from 32.5% to 28.5%. The source is INSSE. The tax wedge is defined as the ratio between the amount of taxes paid and the corresponding total labor cost for the employer.

[10] Chen et al. [2018] studies a Chinese policy that awards substantial corporate tax cuts to firms that increase R&D investment and finds that 30% of the increase in R&D comes from the relabeling of administrative expenses.

[11] In 2004, the law was revised to reflect the transition from the NACE Rev 1 to the NACE Rev 1.1 classification. This revision had no economic effects because the newly-beneficiary NACE Rev 1.1 sectors 7221 and 7222 were the same as the formerly-beneficiary NACE Rev 1 sector 7220.

income tax break by both increasing the number of eligible sector codes (for the firms)[12] and the number of eligible majors for the bachelor's degree (for the workers).[13]

3 Firm-Level Analysis

3.1 2001 Income Tax Break

We begin by studying the impact of the 2001 income tax break on firms already active in 2001. The main advantage of this analysis is that it studies the IT sector over a period when, still in its infancy, it received the unexpected positive news of the introduction of a dedicated tax break to programmers.

Data. For the 2001 analysis, we rely on Amadeus to construct a panel of Romanian firms with three outcome variables: log operating revenues, log number of workers and log total assets. The panel starts in 1997 and ends in 2005. These variables are both likely to react to the new incentive and are among the few variables whose values are less frequently missing. To remove outliers, the sample is winsorized at the 1^{st} and 99^{th} percentiles of the distribution of each of the outcome variables and the change in operating revenues per worker. Moreover, we only keep firms which are active in at least 2000, 2001 and 2002. Last, in the baseline sample we only keep firm-year observations with non-missing values for all three outcome variables (log revenue, log number of workers, and log total assets).[14]

Empirical Strategy. We estimate the firm-level effects of the introduction of the tax break on firms via a difference-in-differences (DiD) design. The first difference is taken between firm outcomes in a given year between 1997 and 2005 and the same firm outcomes in the year 2000 (the reference year). The second difference is taken between the contemporaneous outcomes of firms in the sector with NACE Rev 1 code 722 (the eligible sector) and the outcomes of firms in comparable sectors. Formally,

$$y_{ist} = \alpha_i + \sum_{k=1997}^{2005} \delta_k \mathbb{1}[t=k] + \sum_{k=1997}^{2005} \beta_{DiD,k} \mathbb{1}[t=k] Eligible\_sector_{is} + \varepsilon_{ist}, \tag{1}$$

where i stands for firm, s for the sector of firm i, t for the calendar year. α_i is the firm fixed effect. $\mathbb{1}[t=k]$ is a dummy taking value 1 if an observation pertains to calendar year k and is meant to capture common shocks to all firms that year. We use as outcome variables, y_{ist}, the firm i, year t, log operating revenue, log number of workers and log total assets. We cluster standard errors at the firm level.

The treated firms are those in the eligible IT sector (NACE Rev 1 sector 722, software consultancy and supply); for those firms, $Eligible\_sector_{is} = 1$. The comparison group for our baseline results

[12] In 2013, Romania had to transition from the NACE Rev 1 classification of sector codes to the NACE Rev 2 classification. Because the crosswalks between classifications are not bijective, some sector codes became newly eligible in the transition from one classification to another. See Table A2 in Online Appendix A.1 for details.

[13] While two additional amendments were introduced in the second half of 2015 and 2016, we do not study them due to data availability constraints and their more limited scope. See Online Appendix A.1 for details. Online Appendix A.2 summarizes other policies relevant to the Romanian IT sector and our approach to isolating the effects of the income tax break.

[14] Of the 27 European countries in Amadeus, Romania is in the top tier of countries with the highest coverage [Kalemli-Özcan et al., 2015]. For Romania, the data provider to Amadeus is the Chamber of Commerce and Industry, which reports on an annual basis on the account of more than 500,000 companies (joint stock companies, partnerships limited by shares, limited liability companies, state owned companies, co-operative companies). The variables used in our analysis (operating revenue, employment and total assets) are those with the most consistent coverage in Amadeus [Kalemli-Özcan et al., 2015].

($Eligible\_sector_{is} = 0$) includes firms in (ineligible) high-tech knowledge intensive (HTKI) sectors (as classified by Eurostat).[15] These comparison sectors resemble the eligible IT sector in their focus on high value-added services and reliance on high-skilled workers and technology. Moreover, Table B1 (Online Appendix B.1.1) shows that both the median and average size of firms in treated and comparison sectors are similar. For instance, both the median treated and comparison firms had revenues of 27,000 euros in 2000 and employed three workers.

To deliver credible estimates of the treatment effect of this policy, this strategy relies on two assumptions. First, it relies on the introduction of the tax break being unexpected. As discussed in Section 2, this policy was introduced by Emergency Ordinance at the initiative of one independent policy-maker alone. The lack of wide debate over the policy and the decision to grant a full personal income tax break to programmers (as opposed to the initially-proposed partial break) turned the policy into an unanticipated shock to the IT sector. Second, identification hinges on the assumption that firms in the comparison group form a suitable counterfactual for firms in the IT sector, after accounting for time-invariant differences between firms and common year-specific shocks. The lack of differential pre-trends between treated and comparison firms is an important test for the validity of both assumptions.

As Amadeus does not include information on the extent to which the workers of a given firm have actually benefited from the income tax break, these DiD estimates do not measure the impact of the introduction of the exemption on firms that start having exempted workers, but instead measure the average effects on firms that are part of a sector with a newly available exemption for their workers involved in software development (i.e., “intent-to-treat” effects).

Baseline Results. Figure 1 plots the DiD estimates from the model in Equation (1). These estimates pertain to our baseline comparison group (i.e., firms in HTKI sectors). Reassuringly, across all outcome variables, we observe a lack of preexisting differential trends between treated and control firms, and between 1997 and 2000. After 2001, however, firms in the eligible sector experience significant improvements in all three measures of firm size, such that by 2005, firms in the IT sector have a 24% higher operating revenue, employ 13% more workers, and have 12% more assets than firms in comparable sectors (relative to 2000). Table 1 provides more details.

Robustness Checks. The strong and lasting effect of the policy on firm size survives different comparison groups. We first use an alternative comparison group which contains firms in the subset of HTKI sectors that are the closest to the treated IT sector 722, namely R&D sectors (as focused on knowledge creation as sector 722) and IT sectors that became newly eligible for the income tax break after the 2013 reform (due to their perceived similarity to the already tax-exempt sector 722).[16] The second alternative comparison

[15] Specifically, the baseline comparison group contains firms in the following sectors: 2214 (Publishing of sound recordings), 721 (Hardware consultancy), 723 (Data processing), 724 (Database activities), 725 (Maintenance and repair of office, accounting, and computing machinery), 726 (Other computer related activities), 3002 (Manufacture of computers and other information processing equipment), 731 (Research and experimental development on natural sciences and engineering), 732 (Research and experimental development on social sciences and humanities), 7487 (Miscellaneous business activities n.e.c.), 921 (Motion picture and video activities), 922 (Radio and television activities), and 924 (News agency activities). We exclude sector 642 (Telecommunications) as it underwent a massive liberalization in 2002 in Romania.

[16] Sectors 721 (Hardware consultancy), 724 (Database activities), 726 (Other computer related activities), and 3002 (Manufacture of computers and other information processing equipment) became eligible for the income tax break in 2013. In addition to these sectors, the first comparison group also contains sectors 731 (Research and experimental development on natural sciences and engineering) and 732 (Research and experimental development on social sciences and humanities).

group is comprised of HTKI sectors that *are not* ICT service sectors.[17] This second group allows us to rule out concerns of spillovers from the eligible software development sector to other ineligible yet related sectors (such as the hardware consultancy sector).

Columns (4)-(6) and (7)-(9) in Table 1 report the DiD estimates for the two alternative comparison groups. Contrasting these estimates with those using the baseline comparison group suggests that our results are driven by the policy itself and not by the choice of the comparison group. First, we note that the sign and significance of the DiD coefficients remains unchanged. Second, given that the 95% confidence intervals overlap between specifications for the same variable, we do not find compelling evidence of contamination between the treated IT sector and similar sectors.

3.2 2013 Reform to the Income Tax Break Law

We now move on to study the impact of the 2013 reform to the conditions of eligibility for the income tax break for workers in IT. The main advantage of this analysis is that it is performed on administrative data recording the actual firm-level exemption rate from the income tax of its workers. This allows us to estimate the effect of the actual exemption as opposed to the "intent-to-treat" effect on firms in the eligible sector, irrespective of the actual exemption status of their workers.

Data. The firm-level analysis of the impact of the 2013 reform is based on administrative datasets collected by the National Agency of Fiscal Administration of Romania. The first dataset contains firms' balance sheets, which give us information on yearly revenue and total assets. We add firm-level information coming from two compulsory fiscal forms that record the income taxes paid by workers. In particular, these forms track the firm-level number of workers exempted from paying any personal income tax. The resulting dataset starts in 2011 and ends in 2015.[18]

We first remove from the sample firms that benefited from major State Aid programs (797/2012 and 332/2014) during the period studied. Second, we limit the sample to firm-observations with a non-missing number of income tax-exempt workers. Third, we only keep firms which are active at least in 2012, 2013 and 2014. Last, in the baseline sample, we only keep firm-year observations with non-missing values for all three outcome variables: log revenue, log number of workers and log total assets.

Empirical Strategy. In order to estimate the firm-level effect of the 2013 extension of the tax break to new firm activities and bachelor's degree majors, we estimate the following DiD specification:

$$y_{ist} = \alpha_i + \lambda_{st} + \sum_{k=2011}^{2015} \delta_k \mathbb{1}[t=k] + \sum_{k=2011}^{2015} \beta_{DiD,k} \mathbb{1}[t=k] Exempted_{isk} + \varepsilon_{ist}, \quad (2)$$

where, with the exception of λ_{st}, $Exempted_{isk}$ and the new analysis period, the specification in Equation (2) is identical to the specification in Equation (1). λ_{st} is the sector-by-year fixed effect.

[17] The second alternative comparison group contains NACE Rev 1.1 sector codes 2214, 731, 732, 7487, 921, 922, and 924.

[18] Firms fill in and submit the D112 and D205 forms, retain the owed income taxes from the wages of their workers, and transfer these taxes to the tax authority (all on behalf of their workers). The 2013 analysis relies heavily on these forms. As these forms were first introduced in 2011, this analysis can only start that year. The original dataset includes 2016 as well, but we are excluding 2016 from the analysis because the outcomes in that year are likely to be affected by the 2015 amendment to the income tax break law (see Section 2 and Online Appendix A.1 for details).

$Exempted_{isk}$ is a dummy equal to 1 whenever in year k firm i in sector s has more than 20% of its workforce exempted from the income tax. We keep in the baseline sample only firms who in 2011 and 2012, had less than 5% of workers exempted from the income tax. To improve the interpretation of the estimates, we exclude firms who never reach the 20% threshold after 2013, while at the same time surpassing the 5% threshold at least once between 2011 and 2015. Hence, all firms in the baseline sample have $Exempted_{isk} = 0$ for $k < 2013$ and only treated firms have $Exempted_{isk} = 1$ for $k \geq 2013$.

We choose 5% – instead of 0% – because there are other categories of workers who are eligible for this tax exemption (in particular those with disabilities).[19] We keep only firms with at most 5% workforce exemption before 2013 to mimic the tax conditions of IT firms before the initial introduction of the policy in 2001. The 20% threshold after 2013 is chosen to avoid inadvertently measuring the effect of other income tax exemptions that are unrelated to the 2013 reform, and to ensure that the exemption from the income tax has a non-trivial effect on firm labor costs. This choice is meant to capture both firms whose sector code became suddenly eligible in 2013 and firms (with an eligible NACE Rev 1.1 sector code before 2013) for whom a significantly larger share of workers became eligible in 2013.[20] 20% is also a typical percentage of exemption in the eligible sector.

In addition to these sample restrictions, in the baseline exercise, we only keep firms in high-tech knowledge-intensive (HTKI) service sectors (as classified by Eurostat).[21] Online Appendix B.1.2 contains descriptive statistics on our baseline sample. Table B3 compares the firm size and demographic characteristics in 2012 for the three groups of firms: firms in ineligible sectors, firms in eligible sectors with less than 5% of employees exempted from the income tax throughout the entire sample period, and firms with less than 5% of exempted employees in 2011 and 2012, and which jumped to over 20% of exempted employees after 2013. Firms in ineligible sectors are, on average, the largest; firms in eligible sectors for which a large share of workers became exempted after 2013 have the second-highest average size; while firms in eligible sectors that remain under 5% exemption rate have the smallest average size. In all our specifications, we control for time-invariant differences in size using firm fixed effects.

Table B4 shows that the percentage of firms in eligible sectors with at least one worker exempted from the income tax has increased from 36% in 2011 to 45% in 2015. If we require firms to have more than 20% of their workers exempted from the income tax, the share of such firms increases from 28% in 2011 to 36% in 2015.[22] Table B5 documents the predictors of the firm-level share of workers who are exempt from the income tax. Firms that are foreign-owned, larger, and/or located in counties with a higher overall share of workforce exemption are more likely to have a higher share of workforce

[19] See Online Appendix A.2 for details on these other categories and the reasons why they do not jeopardize our empirical strategy.

[20] We cannot separate these two scenarios for two reasons. First, we do not have worker-level information. Second, we only observe firms' sector codes in 2016, already in the NACE Rev 2 classification. The NACE Rev 2 codes that became eligible in 2013 contain both NACE Rev 1.1 codes that were eligible before 2013 and NACE Rev 1.1 codes that were non-eligible before 2013 (see Table A2, Online Appendix A.1).

[21] Specifically, the following NACE Rev 2 sectors are classified as HTKI service sectors: 582, 59, 60, 61, 62, 63, and 72.

[22] This relatively low level of take-up of the tax exemption is likely due to: (i) the restrictive set of conditions that must be jointly met by firms and programmers in order for programmers to qualify for the tax break; (ii) difficulties in hiring eligible programmers in a tight labor market; (iii) a lack of knowledge of the administrative procedures to apply for the break; or (iv) the high perceived cost of preparing the required documentation. While we cannot distinguish between these scenarios, the first two are the most plausible. The restrictive nature of the criteria to qualify for the break and the need to implement the NACE Rev 2 classification are the main motivations of the 2013 amendment.

exemption from the income tax. These correlations may reflect the ability of these firms to meet the restrictive conditions of the exemption (for instance, by having access to a larger local pool of eligible programmers) or to assemble the documentation necessary to solicit the tax break.

Table 2 studies the "first-stage" effects of the 2013 reform on the firm-level share of workforce exemption from the income tax. We find that in 2013 and after, firms in eligible sectors (according to the new definition of eligibility post-2013) experience a marked increase in their share of workforce exemption relative to 2012. In 2015, firms in eligible sectors have a share of exempted workers that is 4% higher than in 2012. Also, in 2015 there are 7% more firms in eligible sectors whose share of workforce exemption is larger than 20% (again relative to 2012). We find no evidence of either growth or decline in these measures of workforce exemption between 2011 and 2012. These findings suggest that the 2013 reform was effective in its goal to broaden the access of firms and workers to the tax break.

Before proceeding to the results, we want to highlight two key differences between the DiD specification in Equation (2) and the one in Equation (1). First, while in the previous specification firms were deemed treated after 2001 when they were part of the eligible sector, firms are now deemed treated if they are part of the eligible sectors *and* if they start with less than 5% of workers exempted from the income tax pre-2013 and suddenly exceed the 20% threshold of workforce exemption after 2013. Second, we now provide estimates that characterize a subset of the firms in the eligible IT sectors, containing firms that were most likely ineligible for the tax exemption before the 2013 amendment (either due to their sector code or due to their workers' bachelor's degrees).

This new definition of treatment has several advantages. While in the Amadeus data used for the 2001 exercise we did not observe how many of a firm's employees actually benefited from the tax break (if any), the administrative data we use in this exercise tracks this number and allows us to focus on firms whose workforce became treated to a sizable degree. In addition, defining treatment as firm-specific allows us to control for sector-by-year FEs, in addition to firm FEs. This new set of FEs control for potential sector-specific demand and/or technology shocks contemporaneous to the 2013 reform.

This new definition also raises concerns over the extent to which a jump after 2013 in the share of tax-exempt employees is endogenous to firm characteristics. Lacking worker-level data and data on the NACE Rev 1.1 sector code of firms pre-2013, we cannot unequivocally address these concerns. Notwithstanding, we rely on the unexpected timing and generous nature of the 2013 reform (see Section 2). We also build on background information about the state of the IT sector around 2013. At that time, the sector faced a notably scarce labor supply relative to demand, which forced IT firms to hire programmers who were ineligible for the tax break. It is therefore plausible that if a firm suddenly jumped from under 5 to over 20% of workforce exemption after 2013, this jump was caused by the expansion of the lists of eligible sector codes and bachelor's degree majors. Conversely, firms who stayed under the 5% threshold throughout 2011 to 2015 either had a sector code that did not become eligible for the tax break or employed programmers whose bachelor's degree remained ineligible.

Moreover, all our specifications include firm fixed effects to control for time-invariant firm-specific unobservables. What we cannot rule out directly are firm-specific shocks contemporaneous with the 2013 reform (such as a change in management), which may be the actual driver of both the jump in the firm-level workforce exemption share and the estimated firm growth effects. That said, the jumps we study occurred after 2013, implying that they were most likely driven by the reform and not by large-scale changes in management. Last, the lack of differential trends between 2011 and 2012 (relative to firms in the control group) suggests that firms with a sudden increase after 2013 in their share of exempted workers were unlikely to be undergoing productive or organizational changes just before that increase.

Last, we will show that our baseline estimates are robust to (i) using three alternative comparison group of firms in other ICT service sectors, in non-ICT HTKI service sectors, and only in the eligible IT sectors; (ii) running the baseline regression on a dataset from Amadeus that starts in 2008 and allows us to observe a longer pre-reform period; (iii) varying the threshold choice from 20 to 10, 15, 25 and 30%; and (iv) defining treatment based only on the sector of the firm.

Baseline Results. Figure 2 (sub-figures on the left) plots the DiD estimates from the model in Equation (2). We find that across all outcome variables the estimate of $\beta_{DiD,2011}$ is not statistically different from zero. Hence, firms whose workforce became significantly exempted from the income tax after 2013 were not embarked on differential trends between 2011 and 2012 relative to firms whose exemption rate was left unaffected by the reform. This finding goes against anticipatory effects and is in line with the unexpected nature of the expansion. After 2013 however, the treated firms experience a gradual growth along all three measures of firm size, such that in 2015, they have a 20% higher revenues, 10% more workers, and 17% more assets relative to 2012. Columns (1)-(3) in Table 3 provide more details.

Heterogeneity Analysis. The DiD estimates so far refer to the average effect on treated firms. However, an important policy concern is that the tax break expansion benefited only specific types of firms, for instance large foreign firms. We estimate the baseline specification on four groups of firms, defined based on their size and age in 2011: (i) employing strictly less than ten workers (micro firms); (ii) employing at least ten workers (small, medium, or large firms); (iii) strictly less than five years old (young firms); (iv) at least five years old (old firms). As we do not observe foreign ownership in 2011, we cannot directly alleviate the concern that foreign firms benefited more from this policy. That said, most foreign-owned firms in the sector are likely to be large and most micro firms are likely to be domestically-owned.

Table 5 reports the results of the heterogeneity analysis. We find that the tax exemption had a positive effect on all types of firms. The point estimate for revenues is smaller for micro firms than for larger firms (though these estimates are not statistically different). The point estimates for workers and assets are very similar across size categories (though the smaller sample for larger firms leads to noisier estimates). The tax exemption has had stronger effects on younger firms than on older firms; this is likely to reflect the fact that the older firms treated by the 2013 reform may be negatively selected. Most older firms already had more than 5% exemption rate in 2011 and, hence, are not part of the analysis.

Robustness Checks. All robustness checks for our 2013 baseline results are in Tables 3 and 4, and Online Appendix B.2. In our first check, we show that our baseline estimates are robust to three alternative control groups. In columns (7)-(9) of Table 3, the alternative control group 1 contains firms in ICT service sectors (as classified by the OECD).[23] In columns (1)-(3) of Table 4, the alternative control group 2 contains firms in HTKI service sectors *excluding* ICT service sectors. In columns (7)-(9) of Table 4,

[23]This control group includes the following NACE Rev 2 codes: 582 (Software publishing), 61 (Telecommunications), 62 (Computer programming, consultancy and related activities), 631 (Data processing, hosting and related activities, and web portals), and 951 (Repair of computers and communication equipment).

the alternative control group 3 only contains firms in eligible sectors. The similarity of the results across control groups suggests that the results are driven by the tax exemption and not by unrelated trends in the control groups. Note that results are not sensitive to whether we include firms in ICT service sectors, alleviating concerns over potential spillovers from treated firms to firms in related sectors.

Second, we repeat our 2013 analysis on a panel from Amadeus that starts in 2008. We first merge the raw Amadeus data with each of the four samples used in the regressions on the administrative data (i.e., the baseline sample and the three samples with alternative control groups). The merge with the administrative data allows us to classify firms in Amadeus into treatment groups. We then clean the Amadeus data using the same rules applied to the 2001 Amadeus sample.

Tables 3 and 4 (in addition to the sub-figures on the right of Figure 2) present the results from the Amadeus data alongside those from the administrative data. The main advantage of the Amadeus results is that they allow us to observe firms for up to five years before the reform (relative to two years in the baseline results). Across alternative control groups, treated and comparison firms continue to exhibit parallel trends before 2013. Regarding the magnitude of the effects, the 95% confidence intervals overlap between specifications for the same variable, implying similar conclusions.[24]

Estimates of the DiD coefficients are also robust to the choice of the share of exempted workers above which we consider a firm to become treated. Table B6 shows the results of estimations where treatment arises when firms jump after 2013 to at least 10, 15, 25, or 30% of workers exempted from income tax. For each of these thresholds, we report the estimates for both the yearly and pooled DiD coefficients. We present the results for both the administrative (upper panel of Table B6) and Amadeus data (lower panel of Table B6). The point estimates for the pooled DiD coefficients have overlapping confidence intervals, which cautions against over-interpreting the pattern of the point estimates. That said, the overall impression is that the higher the threshold, the lower the estimate of the effects. One plausible explanation is that the firms which experience a higher jump in the rate of workforce exemption tend to be smaller. As we have learned from the heterogeneity analysis, smaller firms experience slightly lower effects. The 20% threshold appears as a conservative choice (compared to the lower thresholds).

In a last robustness check, we revert to the definition of treatment based on a firm's sector (used to study the initial introduction of the policy in 2001, see Equation (1)). This definition mitigates concerns over the potentially-endogenous firm-level jump in the share of exempted employees. Table B8 reports the results from this exercise implemented on two samples: the baseline sample from columns (1)-(3) of Table 3 and an enlarged sample. To ensure comparability with the baseline sample, the enlarged sample only keeps firms with under 5% exemption rate before 2013 – irrespective of their post-2013 exemption rate. Moreover, in Table B9 we also drop all threshold restrictions (both pre- and post-2013)

[24]We undertook various checks on the quality of the Amadeus data. First, Table B2 (Online Appendix B.1.2) reports correlations for 2012 for revenues, number of workers, and total assets. The correlations are between values for the same variable and firm in 2012, where the values come from two different sources: the administrative data and the Amadeus data. We find correlations of at least 0.97. Second, Table B3 (Online Appendix B.1.2) presents summary statistics for 2012 from both the administrative and Amadeus data for the Amadeus sample of 4,837 firms. As anticipated by Table B2, between data sources, the summary statistics are very similar. Third, in columns (1)-(3) and (7)-(9) of Table B7 (Online Appendix B.2) we show that the baseline results (using the administrative data on the 6,146 firms from the baseline administrative sample) are comparable to those from the administrative data but using the baseline Amadeus sample of 4,837 firms. Hence, the reassuring findings from the Amadeus robustness check (which uses the Amadeus data and Amadeus sample of 4,837 firms) are not an artifact of the sample change.

and use Amadeus for a longer panel pre-2013. From these tables, we learn that firms in eligible sectors are not engaged in differential trends before 2013 (relative to firms in ineligible HTKI service sectors). However, after 2013, firms in eligible sectors grow relatively more. Their post-2013 growth is weaker than that of the treated firms in the baseline specification for an intuitive reason. In 2015, the average firm in eligible sectors has 4% more exempted workers than in 2012 (see Table 2, column (8)), whereas treated firms in the baseline specification jump from under 5% exempted workers pre-2013 to over 20% post-2013. We prefer the baseline results for two reasons: (i) they focus on the firms actually treated by the 2013 reform; and (ii) they control for sector-by-year shocks unrelated to this reform.

3.3 Discussion of the Firm-Level Findings

Overall, we find that both policy episodes have led to strong and lasting growth for firms in eligible IT sectors. A first natural question is whether this finding is consistent with the tax break's statutory incidence on workers' personal income. We provide three arguments in favor of an affirmative answer.

First, the restrictive requirements of the tax break ensure that it only applies to particular worker-firm matches, which is likely to grant bargaining power to firms over a shared economic incidence of the tax break. Second, the tax break law stipulates that it is the firm that is responsible for preparing the necessary paperwork and applying for its workers' tax break. This requirement of an explicit "buy-in" from the firm is another reason to expect a shared economic incidence of this tax incentive. Third and last, both policy episodes have led to increases in labor productivity (measured as revenue per worker).[25] There are three plausible drivers of these improvements in labor productivity. First, this policy allows firms to pay workers higher net wages and, by doing so, to improve their motivation and efficiency. Second, the increase in employment in IT is likely to have led to sector-level economies of scale.[26] Finally, it is also likely that both policy episodes have attracted higher-ability workers to software development.

While we find evidence of a shared economic incidence of the tax break, estimating the split of the tax between firms and workers is outside the scope of this paper. As we lack worker-level data, any estimate of this split would need to be model-dependent. Moreover, the split is likely to change with time, depending on the entry of new firms into the IT sector and the long-run elasticity of the supply of programmers. Nonetheless, the firm-level effects are in line with the incentive structure of this policy.

Before comparing the magnitude of the estimates from the two firm-level exercises, let us first highlight the distinguishing features of each. When this policy was introduced in 2001, the Romanian IT sector was in its infancy. Before 2001, Romanian programmers were emigrating in large numbers, lacking confidence that the Romanian IT sector was poised for growth. The introduction of the 2001 tax break shifted that perception, as it signaled that the development of the sector had become a priority for policy-makers. Whereas firms and workers who became eligible in 2001 for the tax break are still eligible, in the 2001 exercise, we only focus on the 1997 to 2005 period. This aims to isolate the effects of this tax break from those of other policies that may interact with it (such as the nationwide switch in 2005 from progressive income taxation to a flat income tax) or global shocks to the IT sector.

[25] In 2015, the labor productivity of Romania's IT and information services (value added per hour worked, PPP) was the third-highest in the EU28, only behind that of Ireland and Luxembourg [van der Marel et al., 2020].

[26] Bartelme et al. [2018] finds large sector-level economies of scale in the "Computers and electronics" manufacturing sector.

Moreover, the Amadeus data used in the 2001 exercise does not contain the firm-level exemption rate of workers but only the sector of the firm. Hence, this exercise estimates the effects on firms whose sector becomes eligible, i.e., "intent-to-treat" effects. The typical firm in the eligible sector employed 13% more workers in 2005 relative to 2000 and relative to the typical firm in ineligible HTKI sectors. Assuming a 25% take-up among firms in the eligible sector, we convert the ITT estimates to "treatment-on-the-treated" (TOT) estimates of 52%. Given that in 2000, the median firm in the eligible sector employed three workers, a 52% increase in employment is sensible. Furthermore, the firms who took advantage of this new tax break during those initial years were likely to be positively selected from all firms in the IT sector, lending further credibility to these magnitudes.

The 2013 exercise estimates the effects on firms who actually experience a sizable decrease in labor costs. Despite this merit, this exercise also has the disadvantage of focusing on a sample of firms that in 2012 – 11 years after the initial introduction of the policy – have under 5% exempted employees. Hence, this sample of firms is likely to be negatively selected (e.g., be younger and less experienced, or smaller). This is likely to explain the smaller 10% estimate for the increase in workforce after the 2013 reform, relative to the 52% TOT estimate for 2001. Nevertheless, the purpose of the 2013 reform was precisely to improve the reach of the tax break to relevant but not yet eligible firms and workers. In addition, the size of the tax incentive (as a share of the wage) is larger between 2001 and 2004 than after 2004 (hence after 2013). This is because – while programmers are fully exempt from the income tax in all years after 2001 – between 2001 and 2004 Romania had a progressive income tax (with a 40% marginal top tax rate), whereas, in 2005, it switched to a flat income tax rate of 16%. This is another plausible reason why the treatment effects are smaller in the 2013 exercise than in the 2001 exercise.

A last feature of our findings that merits comment is the swift onset of the positive effects after both policy episodes. For these swift expansions of firms to be plausible, one requires an immediate ability to hire eligible workers. After the 2001 introduction of the tax break, we believe that this immediate ability to hire was made possible by a slowed-down emigration of Romanian programmers. The shift in the supply of eligible workers contemporaneous to the 2013 reform is built in the reform, as the reform has suddenly doubled the list of workers' university degrees eligible for the tax break. The newly-hired workers helped firms expand their revenues while requiring complementary capital (such as computers or licensed software) to do so. This explains that revenues and assets increase alongside employment.

4 Sector-Level Cross-Country Analysis: 2001 Income Tax Break

4.1 Direct Effects on the Expansion of the IT Sector

This sector-level cross-country analysis of the growth of the IT sector relative to the rest of the economy complements the firm-level analysis from Section 3 in two ways. First, this analysis captures not only the intensive margin of growth (as the firm-level analysis does) but also the extensive margin (the number of firms in the IT sector). Online Appendix C.1 presents descriptive evidence on the extensive margin, suggesting that after 2001 the number of firms in the IT sector has increased faster in Romania than in comparable CEE countries. This increase has occurred through relatively higher entry rates and stable exit rates in Romania. Part of the entry is attributable to entirely new firms, with the remainder accounted for by incumbent firms that switched to the IT sector. While we lack the data to separately estimate the causal effects of the tax break on the extensive margin, the stronger effects that we estimate here reflect the fact that a sector-level analysis encompasses this margin as well.

Second, this synthetic control method (SCM) analysis allows us to alleviate concerns about potential confounding factors that may affect the IT sector globally in 2001 and that may be the true cause of the effects measured with firm-level data. One such confounding factor could be the dot-com crash of 2001. In its aftermath, U.S.-based companies may have chosen to mitigate some of the losses incurred during the crash by offshoring part of their operations in CEE countries. Global productivity or demand shocks to the IT sector could be another confounding factor for its growth in Romania. We therefore benchmark the growth of the IT sector in Romania to that in similar neighboring countries, that are likely to have been similarly affected by such confounding factors.

Empirical Strategy. We use SCM to measure the effect of the income tax break on the growth of the IT sector in Romania. SCM is a data-driven approach to small-sample studies proposed by Abadie and Gardeazabal [2003] and used to estimate treatment effects. Its intuition is that a weighted combination of countries provides a better comparison for Romania than any individual country. SCM makes explicit the relative contribution of each control country to the counterfactual, and the degree of similarity between Romania and synthetic Romania, in terms of pre-intervention outcomes and other predictors of post-intervention outcomes [Abadie et al., 2010]. The choice of weights is such that the resulting unit closely matches the treated unit over the pre-treatment period. Outcomes for the synthetic unit are then projected into the post-treatment period based on these weights. Inference is conducted using placebo tests. The same model is estimated on each untreated country, assuming that it was treated in 2001. The result is a distribution of placebo effects. If this procedure does not yield effects for untreated countries as large as the effects for Romania, then it is unlikely that the estimated effect for Romania is a result of chance.

Formally, SCM entails the following. Let J be the number of control countries ("the donor pool"), where J equals 12 in our case. Let $\boldsymbol{W} = (w_2, ..., w_{J+1})'$ be a $J \times 1$ vector of weights w_j, such that $w_j \geq 0$ and $\sum_{j=2}^{J+1} w_j = 1$. w_j is the weight of country j in synthetic Romania. SCM chooses W such that synthetic Romania most closely matches the real Romania before 2001 (i.e., in 1999 and 2000). Let $\boldsymbol{X_1}$ be a $(K \times 1)$ vector of pre-2001 values of K predictors for the relative growth of the Romanian IT sector. Similar to $\boldsymbol{X_1}$, we define $\boldsymbol{X_0}$ as the $(K \times J)$ matrix containing the values for the same predictors for the J candidate control countries. Let V be a diagonal matrix with non-negative components, whose diagonal elements represent the relative importance of the predictors in the construction of synthetic Romania. The vector of weights $\boldsymbol{W}$ is chosen to minimize the objective function $(\boldsymbol{X_1} - \boldsymbol{X_0}\boldsymbol{W})'\boldsymbol{V}(\boldsymbol{X_1} - \boldsymbol{X_0}\boldsymbol{W})$, such that $w_j \geq 0$. As in Abadie and Gardeazabal [2003], Abadie et al. [2010], we allow for the choice of the weighting matrix, $\boldsymbol{V}$, to be data-driven. $\boldsymbol{V}$ allows for the pre-2001 outcomes of Romania to be closest to the outcomes of the synthetic control obtained from $\boldsymbol{W}^*(\boldsymbol{V})$.[27]

Treatment effects, α_{1t}, for Romania in post-treatment years t (2002 to 2015) are estimated as the

[27]Throughout the paper, we implement SCM with the help of the **synth** and **synth_runner** packages in Stata [Quistorff and Galiani, 2017]. We depart from the default option of these packages by selecting the *nested* option. Hence **synth** embarks on a fully nested optimization procedure that searches among all (diagonal) positive semidefinite $\boldsymbol{V}$ matrices and sets of $\boldsymbol{W}^*$-weights for the best fitting convex combination of the control units. The fully nested optimization contains the regression-based $\boldsymbol{V}$ as a starting point, but produces convex combinations that achieve even lower mean squared prediction error.

difference between the year t outcomes for Romania and those for synthetic Romania:

$$\hat{\alpha}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt}, \tag{3}$$

where Y_{1t} is the year t outcome for Romania and Y_{jt} is the same outcome in year t for control country j.

Concretely, we first contrast the within-country growth of the IT sector to the growth of all other sectors in the economy. Through SCM, we then compare this relative growth of the IT sector in Romania to the relative growth of the IT sector in synthetic Romania. An advantage of SCM is that it delivers the optimal set of weights to construct synthetic Romania and limits the researchers' degrees of freedom in the choice of the comparison group. Last, we ask whether the relative growth of the IT sector in Romania is exceptional compared to the relative growth experienced by this sector in untreated countries.

Data. The data source for the dependent variables is Eurostat, Structural Business Statistics. We require from these variables to be available in 1999 and 2000 for both Romania and all other countries in the donor pool. Also, we require these variables to appear consistently in the following years and meet minimal data quality standards. Finally, we want these variables to be relevant for this study. These conditions are met by the following three variables: number of employees, gross revenues (turnover or gross premiums written, in million euros) and production value (in million euros). Data for the predictors comes from the World Bank, World Development Indicators. We use as predictors the GDP per capita (constant LCU), the share of manufacturing value-added coming from the medium and high-tech sector, and the share of GDP coming from services.

Due to frequent missing data at the three-digit level, the IT sector is defined more broadly (than the set of eligible three-digit codes) as K72 ("Computer and related activities"). We use as comparison sectors all other sectors in the economy. To obtain normalized values for the outcome variables in each year and country, the yearly absolute value of the variable in the treated sector is divided by its value in 2000, the year prior to the introduction of the policy in Romania. From these resulting yearly ratios we subtract the corresponding ratios for the comparison sectors.[28]

The donor pool of countries for the synthetic control contains Bulgaria, the Czech Republic, Estonia, Hungary, Ireland, Latvia, Lithuania, Poland, Portugal, the Slovak Republic, and Slovenia. These countries were chosen based on their geographic proximity, similarity in development, performance in the IT sector pre-2001, and data availability. See Online Appendix D for details on data construction.

Baseline Results. Figures 3 and 4 present the output of our SCM analysis for two dependent variables: gross revenues and employment (both normalized). The upper left panels show the evolution of these outcomes in Romania and synthetic Romania. By construction, before 2001, the growth of the IT sector of synthetic Romania closely mimics that of Romania. From 2001 onward, however, both the gross revenues and employment of Romania's IT sector experience a marked relative growth. This pattern is no longer a mechanical output of SCM but points to the effect of the 2001 policy on Romania's IT sector. The upper right panels show the difference between the outcomes of Romania and those of synthetic Romania. Fourteen years after, the gross revenues (employment) in the IT sector in Romania

[28] A value of two in year t means that the multiplication factor of the value of the dependent variable in the treated sector in t, relative to 2000, is larger by two units than the counterfactual multiplication factor in comparison sectors.

had expanded 6.52 (1.83) times more relative to the gross revenues (employment) in all other sectors, relative to the year 2000, and relative to the corresponding relative growth in synthetic Romania.

The lower left panels plot the raw paths of these normalized outcomes of the IT sector in Romania and the 11 other donor countries. We notice how exceptional the growth was in Romania, compared to that in donor countries. Last, we implement a battery of placebo tests that considers all other donor countries as potentially treated and proposes synthetic controls for each. Reassuringly, the lower right panels show that the relative growth for Romania is starker than the relative growth for all other donor countries. At least until 2008, the actual treatment differences of gross revenues and employment growth for Romania lie outside the range of placebo differences. Formally, these results are confirmed by the almost-zero p-values until 2008 (see Table C1, Online Appendix C.2.1).

One might be concerned that synthetic Romania is an unreasonable proposition of SCM. Synthetic Romania is a combination of Bulgaria, the Czech Republic, and Slovakia (with weights varying with the outcome variable). Table C2 shows that synthetic Romania is reasonably similar in terms of the share of services in GDP and the share of high-tech manufacturing in total manufacturing value added. While synthetic Romania is different to Romania in terms of its GDP, our SCM analysis is relative to each country's level in 2000. Hence, this proposal of synthetic Romania seems appropriate.

Robustness Checks. The robustness checks figures and tables can be found in Online Appendix C.2.1. To begin with, the SCM findings presented above are robust to different choices of the weighting matrix $\boldsymbol{V}$ and matrix of weights $\boldsymbol{W}$. This is a particularly important check in this exercise where the SCM chooses the weights behind synthetic Romania based on data from only one pre-treatment year (1999).

We then show that these findings are not unique to gross revenues and employment. Figure C6 (and its associated Tables C1 and C2) shows similar patterns of outstanding growth in the IT sector of Romania, this time in terms of production value. Again, we employ permutation methods to assess the statistical likelihood of our results. In the first seven years after the introduction of the policy, almost-zero p-values allow us to rule out a treatment effect of zero. Visually, the lower right panel of Figure C6 shows that the actual difference for Romania is consistently above the upper limit of placebo differences.

One concern with deriving results from the entire 1999 to 2015 time series of Eurostat data may come from the need to rely on a crosswalk between NACE Rev 1.1 and Rev 2 sector codes. Because pre-2007 data is reported for NACE Rev 1.1 sectors and post-2007 data for NACE Rev 2 sectors, one needs a crosswalk to stitch together the time series. As the relationship between classifications is not bijective, there is no widely-used crosswalk. Online Appendix D provides details on own crosswalk construction.

To test whether our results are driven by the use of the full time series, we truncate the time series in 2007. We can thus study the effect of the 2001 policy using data that is consistently reported in one classification. Fortunately, we find that results for pre-2007 years are not affected by the addition of post-2007 years (see Figure C7). However, this finding does not imply that post-2007 results are not affected by the stitching of the sector-level time series. We cannot distinguish whether our weaker post-2007 results are driven by the differential effects of the financial crisis on the IT sector (compared to the rest of the economy), by a later introduction by donor countries of other policies that also favor the IT sector, or by an imprecise stitching of the sector-level time series.

A last robustness check is one in which we would contrast the growth of the IT sector to that of the same three-digit comparison sectors used in the firm-level analysis. A first constraint comes from the fact that Eurostat data is at the two-digit level. Second, the two-digit sectors 64 and 92 cannot be used, as their data is frequently missing across years and countries. Third, sector 73 is small and with noisy data, and sector 74 contains several three-digit sectors other than those we use as control to sector 722, sectors we believe are dissimilar to 722. Despite these caveats, sector 72 still exhibits a faster growth than that of comparison sectors 73 and 74. Results are available upon request.

4.2 Downstream Effects of the Expansion of the IT Sector

Given that, in 2017, the IT sector accounted for less than 2% of Romania's total employment and that only specific workers in this sector are eligible for the income tax break, one might question the wider effects of the policy. The IT sector is a sector whose inputs are broadly used by households and firms in all sectors. While the development of the IT sector most likely led to level effects as well (given its broad use), we propose a research design that allows us to credibly estimate the differential effects of the policy on sectors that relied more heavily on IT services relative to sectors that relied less.

The development of the IT sector after 2001 is likely to have boosted the development of sectors relying more heavily on IT services in two ways.[29] First, after the tax break, labor productivity (measured as revenues per worker) increased. As labor is the main input in the production of IT services, this is also likely to have improved the quality of IT services. Hence, the tax break is likely to have lowered the quality-adjusted price of IT services. Under a plausible market structure and well-behaved cost and demand conditions, this should lead to output increases in sectors purchasing more inputs from the targeted (IT) sector [Lane, 2017]. To the extent that the increase in the quality of IT inputs was not fully priced, then IT-using sectors not only experienced increases in output but also in productivity.

Second, the IT sector has also expanded through the entry of new firms (see Section 4.1 and Online Appendix C.1). Figure C5 (Online Appendix C.1.3) suggests that part of this entry occurred through foreign direct investment in the IT sector of Romania, which has intensified since the early 2000s. This implies that the IT sector has considerably expanded the set of varieties proposed to downstream sectors. Whenever downstream sectors have a love of variety for intermediate inputs, this expansion in varieties leads to productivity gains.[30] In addition, varieties proposed by foreign-owned firms are likely to have been of higher quality than those proposed by domestic firms.

All in all, improvements in the prices, quality, and variety of IT inputs are likely to have provided a boost to downstream sectors relying more on IT. We study not only the evolution of the size of these sectors, but also their export performance. To the extent that a stronger IT sector generates productivity gains for IT-using downstream sectors, this is likely to shift trade patterns.

Empirical Strategy. To study the effect of the 2001 tax break on the expansion of downstream sectors, we employ a similar SCM to the one described in Section 4.1. The only difference between these

[29] One scenario in which this policy might have *hurt* the development of sectors relying more heavily on IT services, is one in which firms in these sectors used to produce programming services in-house. As programmers only benefit from the tax break if working for a firm in the eligible IT sector, this may have made them less likely to join ineligible sectors. This scenario is unlikely, as it was uncommon for firms in non-IT sectors to develop software in-house. Note that employees who work in maintenance are not eligible for this tax break, irrespective of their employer's sector.

[30] See Rodríguez-Clare [1996], Goldberg et al. [2010], Carluccio and Fally [2013], Kee [2015]

exercises is in the definition of treated and comparison sectors. Hereafter, we define treated sectors as those downstream sectors for which the IT sector is most important as an input supplier. Conversely, comparison sectors are those relying relatively less on the IT sector as an input supplier.

Data. We start from the input-output table (I-O table, henceforth) of Romania for 2000. We use the harmonized I-O table provided by the OECD, which tracks the flows of goods and services between all two-digit NACE Rev 1 sectors. Given that 2000 is the year before the unexpected introduction of the policy, inter-sector linkages are not yet affected by this policy. We then compute the share of the total input expenditures of a given sector purchased from the IT sector (NACE Rev 1 sector 72, "Computer and related activities"). Based on these shares we identify the sectors for which IT services are the most important inputs in 2000. Based on their position in this sector-level distribution of shares, we assign sectors to either a high- or low-intensity category of use of IT services. The treated high-intensity category contains sectors that are among the top 25% users of IT services. All other sectors lying in the bottom 75% constitute the control category.[31] See Table C3 in Online Appendix C.2.2 for details.

Similar to the SCM in Section 4.1, we rely on Eurostat data to construct the same normalized dependent variables and World Bank data for the same predictor variables. In addition, we use UN Comtrade data to study the export performance of sectors relying more on IT services, relative to those relying less. Given data availability, we use SCM to study the exports of goods alone. One notable advantage of UN Comtrade data is that it starts in 1996, offering three more years than the Eurostat data of pre-treatment years. See Online Appendix D for details on data construction.

Baseline Results. We first ask whether sectors with stronger upstream linkages to the IT sector experienced a more pronounced growth than sectors with weaker linkages. Figures 5 and 6 provide a visual answer for gross revenues and employment. Fourteen years after the introduction of the policy, gross revenues (employment) in sectors with high-intensity use of IT services have grown 0.75 (0.61) times more than gross revenues (employment) in low-intensity sectors (compared to year 2000 and compared to the equivalent difference in synthetic Romania). When implementing the permutation method suggested by Abadie et al. [2010], we find that our SCM estimates lie at the upper limit of the distribution of placebo estimates. The low p-values in Table C4 (Online Appendix C.2.2) rule out null effects.

Next, we study the export performance of sectors relying more on IT services, compared to those relying less. Because high-intensity sectors (defined until now as those over the third quartile of the IT-usage intensity distribution) are all service sectors, we now define high-intensity sectors as those manufacturing sectors between the second and third quartile. Figure 7 depicts a striking relative growth in the export trade values of high-intensity sectors in Romania (relative to those from low-intensity sectors and to synthetic Romania).[32] Placebo tests show that this relative growth in Romania is exceptional

[31] We also calculate the share of the total sales of the IT sector purchased by each sector. Again, we assign sectors to quarters based on these new shares. While these two classifications are conceptually different, given the I-O table of Romania for 2000, we find that there is no practical difference in the final split of sectors between the top and bottom three quartiles.

[32] Most of this growth is explained by the SITC Rev 1 commodity codes 54 (Medicinal and pharmaceutical products), 62 (Rubber manufactures), 73 (Road vehicles, other than motor vehicles), and 86 (Watches and clocks).

compared to that predicted for all other countries in the donor pool.[33] Table C6 (Online Appendix C.2.2) makes the same argument formally. This suggests that the development of the IT sector in Romania not only increased the output of IT-using sectors but also improved their comparative advantage.[34]

Robustness Checks All robustness checks figures and tables are in Online Appendix C.2.2. First, we show that our baseline findings on the relative growth of downstream sectors with heavier usage of IT services are not specific to gross revenues, employment, or export value. For instance, we find that the production value of IT-using sectors has also grown significantly more in Romania (see Figure C8).

Second, the SCM findings presented above are robust to different choices of the weighting matrix $\boldsymbol{V}$ and matrix of weights $\boldsymbol{W}$. Another potential concern with SCM relates to its sensitivity to the number of pre-intervention periods used in the computation of the weights. We cannot rule out this concern with the Eurostat data, as the panel only starts in 1999. However, a benefit of Comtrade data is that it allows us to observe export patterns since 1996. We run the Comtrade SCM exercise varying the number of pre-treatment years used in the estimation. Reassuringly, treatment effects remain unaltered.

Third, one might be concerned that our results are driven by specific high-intensity IT using sectors. From the beginning, in the analysis using Eurostat data, we exclude NACE Rev 1 sector 72, as this two-digit sector contains the three-digit sector 722 eligible for the income tax break. This avoids the risk of a mechanical result. In addition, as a robustness check, we also exclude NACE Rev 1 sectors 73 and 74; these sectors belong to the list of high-intensity IT users, but also contain three-digit NACE Rev 1 sectors that we use as comparison sectors in the firm-level analysis. If sectors 73 and 74 – or their subset of three-digit sectors comparable to sector 722 – were experiencing correlated shocks with those of sector 722, our findings could be affected by such shocks. Figure C9 shows that when we exclude sectors 73 and 74 the treatment effect is actually larger than the one found in the baseline Figure 6.[35]

Fourth, we also show that results have qualitatively similar patterns when we change the threshold of the grouping of sectors into the high- and low-intensity categories. Figure C10 presents results from the grouping of sectors under the median of usage of inputs from the IT sector into the low-intensity category (as opposed to under the third quartile). As expected, while the difference in the development of the high and low intensity categories becomes less stark, the general pattern is maintained.

Fifth, one might also worry that relying on Romania's I-O table from 2000 to construct the high- and low-intensity treatment categories is a concern in itself. As an alternative, we use the classification of sectors proposed by van Ark et al. [2003]. Sectors are assigned one of the following six categories based on U.S. measures of pre-2000 ICT (information and communication technology) intensity from Stiroh [2002]: ICT-producing manufacturing, ICT-producing services, ICT-using manufacturing, ICT-using services, non-ICT manufacturing, and non-ICT services. We exclude sectors in ICT-producing

[33] Figure C12 (Online Appendix C.2.2) shows that after 2000 the exports of service sectors relying more on IT services also grew noticeably faster in Romania than in comparable countries. Among these sectors, those under NACE Rev 1 sector 74 (e.g., call centers, advertising, business and management consultancy, secretarial and translation activities etc.) experienced the most impressive growth. Romania's trend is compared to that of the five countries that constitute the typical synthetic Romania in all SCM exercises thus far, i.e. Bulgaria, the Czech Republic, Hungary, Lithuania, and Slovakia.

[34] The behavior of FDI flows to Romania also supports this claim. Figure C5 (Online Appendix C.1.3) shows that FDI in high-intensity IT using sectors grew faster than FDI in low-intensity sectors. While not the only driver behind this relative growth in FDI, Romania's IT sector is frequently mentioned among those that are most significant. For instance, this article describes how Romania's IT sector played an important role in Renault's 2007 decision to build its Technocentre in Romania.

[35] The main SCM exercise using Comtrade data excludes, by construction, all service sectors, hence sectors 72, 73, and 74.

manufacturing and ICT-producing services, as they might be directly affected by the tax break. We group sectors in ICT-using manufacturing and ICT-using services into the high-intensity category, and sectors in non-ICT manufacturing and non-ICT services into the low-intensity category. While the patterns obtained with this grouping are noisier than those obtained with our preferred grouping, we still find a stronger relative growth in Romania in ICT-using sectors compared to non-ICT using, and compared to synthetic Romania. We assess our initial grouping to be superior, as it is more narrowly defined around the treated sector (NACE Rev 1 sector 722) than the one proposed by Stiroh [2002].[36]

Finally, we check whether our results are robust to the exclusion of the second half of the Eurostat time series. As explained in Section 4.1, the lengthening of sector-level time series to include 2007 to 2015 relies on an inherently-imprecise crosswalk between the NACE Rev 2 and NACE Rev 1 classifications. Figure C11 shows that results for the years under the NACE Rev 1 classification (1999 to 2006) are identical to those for the same years obtained using the full time series (1999 to 2015). This concern does not affect the Comtrade SCM results, as the Comtrade data is available in a unique classification.

In addition to these robustness checks, the timing of the relative growth of IT-using sectors speaks against concerns of reverse causality, i.e., it is the development of downstream sectors using IT intensively that actually boosted the development of the IT sector.[37]

4.3 Discussion of the Sector-Level Cross-Country Findings

In Section 4.1, we show that, since 2001, the IT sector in Romania has grown significantly faster compared to the rest of the sectors in Romania and compared to the same relative growth in similar countries. This finding gives us confidence that the effects we measure are plausibly caused by the tax break to programmers introduced in Romania, and not by other global supply- or demand-side shocks to the IT sector. This (relative) growth of the IT sector has occurred both on the intensive (through the growth of incumbent firms) and extensive margin (through the entry of new firms in the sector). In Section 4.2, we find that since 2001, IT-using sectors in Romania have grown significantly faster compared to non-IT using sectors in Romania and the same relative growth in similar countries. This pattern provides support to the conjecture that improvements in the quality-adjusted prices and variety of IT services benefit more sectors relying more heavily on these services.

Several pieces of evidence lend credibility to the magnitude of our SCM estimates. First, while the estimated magnitude of the relative growth rate of the IT sector in Romania might seem impressive, this magnitude is lower than the actual growth of the IT sector in Romania. Our SCM estimate for 2015 for the number of workers is 1.83, i.e., the number of workers in IT grew 1.83 times faster between 2000 and 2015 compared to the number of workers in the rest of the Romanian economy and compared to the same relative growth in synthetic Romania. In the raw data, the employment in IT in Romania grew six times, from 13,691 workers in 2000 to 81,780 workers in 2015. Similarly, gross revenues grew 14-fold over the same period. Hence, our estimates attenuate the actual growth of the IT sector in Romania after 2001, as they control for broader trends in the rest of the Romanian economy and similar economies. Second,

[36] ICT contains several other (significantly larger) sectors than 722, unrelated to the policy we study.

[37] This does not exclude the possibility of a feedback loop between the development of IT-using sectors and the IT sector itself. In a 2016 *Reuters* article, Florin Talpeş (a pioneer in Romania's IT sector) advised new entrants in the IT sector to focus on developing technology for the now-mature automotive sector (e.g., driver-less technology or car connectivity).

consistent with the firm-level results, we find stronger increases in revenues in the IT sector than in employment. Third, the relative growth of downstream sectors with stronger links to the IT sector is not as large as the relative growth of the IT sector itself. As one would expect, the sector directly receiving the tax incentive grows faster than downstream sectors benefiting from the incentive indirectly.[38]

It is important to emphasize that there are features of either the available data or of the empirical strategies that do not recommend a direct comparison of the sector-level and the firm-level evidence. First, due to data constraints, the sector-level and firm-level growth rates in IT are measured relative to different comparison groups.[39] The choice of the reference group can affect the exact magnitude of the estimates. Second, the sector-level evidence also allows for growth through firm entry, whereas the firm-level evidence characterizes incumbent firms alone.

Finally, the time frame of analysis is also important for the magnitudes and their interpretation. We conduct each firm-level exercise in a relatively narrow time window around the two policy of interest: 1997 to 2005 for the 2001 introduction of the income tax break, and 2011 to 2015 for the 2013 amendment to the tax break. To improve identification, the firm-level analysis, therefore, uses short-term variation in the tax conditions of firms in the IT sector in Romania. In the long term, there are general equilibrium effects (such as those on the supply of programmers) or unrelated shocks (such as the global financial crisis) that would have hindered the interpretation of long-term firm-level growth estimates.

By virtue of their long time frame (1999 to 2015), the sector-level cross-country findings are likely to capture not only the direct effects of the initial introduction of the policy in 2001, but also those of other developments in Romania and abroad that differentially affect the IT sector of Romania since 2001. For instance, to the extent that in the early 2000s the development of the IT sector in Central and Eastern Europe was at the cusp of multiple potential equilibria, this policy is likely to have acted as a timely signal to both local and foreign firms. In a world with first-mover advantage and path dependence, this policy is likely to have tilted the balance towards an equilibrium favorable to the IT sector of Romania. In this light, it is likely that a sizable part of the subsequent growth of the IT sector is due to a snowballing effect of the signal of the policy, as opposed to the size of the actual incentive it provides.

Relatedly, the magnitude of the effects on IT-using sectors captures more than the indirect incentive granted by this tax break. It also reflects the idea that in an economy with coordination failures – due to economies of scale and imperfect tradability of services associated to skill-intensive manufacturing – government policy can move the economy towards the "high-wage, high-tech equilibrium" [Rodrik, 1996]. Moreover, it suggests the possibility of strong complementarities between this industrial policy and the FDI attraction and trade opportunities that followed Romania's joining of the EU in 2007 [as in Topalova and Khandelwal, 2011]. While Romania's comparison countries also joined the EU in 2004 or 2007, only Romania saw such a distinctive growth in the exports of IT-intensive downstream sectors.

To conclude, while it is outside the scope of this paper to disentangle the direct effects of this income tax break and those circumstances that may have amplified or dampened its effects in the aggregate and in the long-run, it is reassuring that the sector-level cross-country evidence and the firm-level evi-

[38] Figure C5 (Online Appendix C.1.3) confirms the same intuition: the IT sector itself is the one becoming more attractive for FDI, followed by downstream sectors using IT services intensively.

[39] The sector-level cross-country data is at the two-digit level, with frequent missing values in the two-digits containing the three-digit codes used for comparison in the firm-level analysis. Hence, our baseline sector-level estimates are with respect to the rest of the economy, whereas the firm-level estimates are with respect to firms in certain three-digit sector codes.

dence paint an overall consistent picture by which the income tax break has been effective in its objective to boost the development of the IT sector in Romania.

5 Back-of-the-Envelope Cost-Benefit Analysis

In Table 6, we propose a back-of-the-envelope cost-benefit analysis for the 2013 reform to the tax break law [similar to the one in Munch and Schaur, 2018].[40] We focus on this latter policy episode (as opposed to the initial introduction of the policy in 2001) as it is only around this episode that we observe firms' shares of tax-exempt workers. We estimate the total and average costs and benefits for 2015 for the 462 treated firms from the baseline 2013 analysis, i.e., those 462 firms that had less than 5% tax-exempt workers pre-2013 but that experienced a jump to more than 20% tax-exempt workers post-2013 (plausibly due to the sudden and sizable expansion in the scope of the tax break).

We calculate the benefits by multiplying the values for revenues, employment, and assets in 2012 (the base year before the tax break reform) for the treated firms with their respective estimates of the average DiD effects for 2015 (the end year in our analysis). The average values for 2012 are the same as those reported in Table B3 (Online Appendix B.1.2). The baseline average DiD effects are those from columns (1) to (3) in Table 3. We find that in 2015, firms treated by the 2013 reform generated 29.7 million euros more revenues, employed 300.7 more workers, and had 15.9 million euros more assets than in 2012.

We take the total cost to the government of the 2013 reform to be the foregone taxes from the full personal income tax break of newly-exempted workers in treated firms. This number of newly-exempted workers in treated firms is equal to the number of tax-exempt workers in these firms in 2015 (1,868) minus the same number in 2012 (16). We assume that all these 1,852 workers owe their tax break to the 2013 reform (e.g., because the sector code of their firm became eligible for the tax break thanks to the reform). Note that part of these workers are not necessarily new hires but newly-exempted incumbent workers. We also assume that the costs of implementing the tax break or of any potential distortions are negligible. In sum, the total cost is equal to the value of the flat personal income tax of 16% on the 2015 taxable income of all 1,852 newly-exempted workers in treated firms. This amounts to 4,3 million euros.

Therefore, for each extra euro in foregone tax revenues in 2015 due to the 2013 reform, treated firms generated 7 euros extra in revenues and accumulated 3.7 euros extra in assets. This cost-benefit comparison suggests that for firms treated by the 2013 reform, the benefits of the reform outweigh its costs to the government. One might wonder why the tax break has been so beneficial to firms relative to the fiscal incentive granted. We conjecture that the expansion of the tax break has not only provided a fiscal incentive to productive matches between firms in IT and programmers but has also served as a signal to the sector of the renewed commitment of the government to support its development. This signal is likely to have reduced uncertainty and helped address coordination failures [Harrison and Rodríguez-Clare, 2010], which may have otherwise hindered the formation of those productive matches.

[40] This analysis is not a welfare analysis, as the estimation of general equilibrium effects would require us to lean heavily on the structure of a model. Moreover, the lack of matched employer-employee data precludes us from having credible estimates of how gains and costs are distributed across the population. This analysis limits itself to the estimation of direct benefits and costs pertaining to the firms treated by the 2013 reform.

6 Conclusion

This paper examines the effects of a unique industrial policy introduced by Romania in 2001: a personal income tax break for workers with specific IT-relevant bachelor's degrees and who work directly on software development for a firm with an eligible IT sector code. In 2013, the tax break law was amended to allow for a significantly larger list of eligible sector codes for firms and eligible bachelor's degrees for workers. We use both policy episodes to bring plausibly-causal estimates of their firm and sector-level effects. Across empirical strategies and measures of firm size, we find that the policy and its reform led to strong and lasting growth for IT firms. This is in line with a shared economic incidence of the tax incentive between firms and workers. Sector-level cross-country evidence makes the additional point that the growth after 2001 of the IT sector in Romania is unlikely to be driven by confounding factors, as similar neighboring countries fail to show comparable growth. Moreover, we find evidence of wider benefits from this policy for the Romanian economy. Namely, we find that sectors relying relatively more on IT inputs grew faster than those relying less.

Our results suggest that this policy has been effective in supporting the development of the IT sector – a sector seen as key to the transition to a knowledge economy. Moreover, the policy reallocated resources towards a "good jobs" (high-skill/high-wage) sector, a policy priority in both developed and developing countries. Beyond the special nature of this sector, these are notable achievements, as many industrial policies only allow specific groups to extract rents, without actually affecting resource allocations. While this policy helped the expansion of its IT sector, Romania still lags behind the EU average in other dimensions of the digital economy and society (e.g., in the use of ICT by households, businesses, and public services). Future policies should consider addressing these remaining challenges.

Establishing whether this policy was also efficient is beyond the scope of this project, but a fruitful area for future research. A number of special features of this policy motivate such additional research. First, this policy was not designed to rescue a "sunset sector" (as has been the case in several East Asian industrial policies). On the contrary, it targeted a sector under-developed in Romania at the time, but generally seen as crucial for growth. Second, the effects on IT-using downstream sectors suggest that this policy may have also mitigated (inter-sector) coordination failures coming from scale economies and imperfectly-tradable services (here, IT services) useful for skill-intensive downstream sectors [as in Rodrik, 2004]. Last, this policy involved reductions in labor taxes, as opposed to the vastly more common reductions in corporate taxes or state aid packages. With worker-level data, one could study the distributional implications of this policy design relative to the common designs.

References

Alberto Abadie and Javier Gardeazabal. The Economic Costs of Conflict: A Case Study of the Basque Country. American Economic Review, 93(1):113–132, March 2003.

Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. Journal of the American Statistical Association, 105(490):493–505, 2010.

Philippe Aghion, Jing Cai, Mathias Dewatripont, Luosha Du, Ann Harrison, and Patrick Legros. Industrial Policy and Competition. American Economic Journal: Macroeconomics, 7(4):1–32, October 2015.

Phillipe Aghion, Julian Boulanger, and Elie Cohen. Rethinking Industrial Policy. Bruegel Policy Brief, 4, 2011.

Adam Banai, Peter Lang, Gabor Nagy, and Martin Stancsics. Waste of Money or Growth Opportunity: The Causal Effect of EU Subsidies on Hungarian SMEs. Economic Systems, 44, 2020.

Dominick Bartelme, Arnaud Costinot, Dave Donaldson, and Andres Rodriguez-Clare. External Economies of Scale and Industrial Policy: A View from Trade. Working Paper, 2018.

Paul Beaudry, Mark Doms, and Ethan Lewis. Should the Personal Computer Be Considered a Technological Revolution? Evidence from U.S. Metropolitan Areas. Journal of Political Economy, 118(5):988–1036, 2010.

Bruce A Blonigen. Industrial Policy and Downstream Export Performance. Economic Journal, 126(595):1635–1659, 2016.

Nicholas Bloom, Raffaella Sadun, and John Van Reenen. Americans Do IT Better: US Multinationals and the Productivity Miracle. American Economic Review, 102(1):167–201, 2012.

Philipp Boeing. The Allocation and Effectiveness of China's R&D Subsidies – Evidence From Listed Firms. Research Policy, 45(9):1774–1789, 2016.

Jing Cai and Ann Harrison. Industrial Policy in China: Some Unintended Consequences? Industrial and Labor Relations Review, 2019.

Juan Carluccio and Thibault Fally. Foreign Entry and Spillovers with Technological Incompatibilities in the Supply Chain. Journal of International Economics, 90(1):123–135, 2013.

Zhao Chen, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Xu. Notching R&D Investment with Corporate Income Tax Cuts in China. Working Paper, November 2018.

Reda Cherif and Fuad Hasanov. The Return of the Policy That Shall Not Be Named: Principles of Industrial Policy. Technical report, International Monetary Fund, 2019.

Chiara Criscuolo, Ralf Martin, Henry G Overman, and John Van Reenen. Some Causal Effects of an Industrial Policy. American Economic Review, 109(1):48–85, 2019.

Maarten de Ridder. Market Power and Innovation in the Intangible Economy. Working Paper, 2020.

Luosha Du, Ann Harrison, and Gary Jefferson. FDI Spillovers and Industrial Policy: The Role of Tariffs and Tax Holidays. World Development, 64:366–383, 2014.

Eurofound. Employment Effects of Reduced Non-Wage Labour Costs. Luxembourg, 2017.

European Commission. A Study on R&D Tax Incentives. Taxation Papers, Working Paper 52, 2014.

European Commission. Industrial Policy Strategy: Investing in a Smart, Innovative and Sustainable Industry. European Commission Policy Report, 2017.

Chris Forman, Avi Goldfarb, and Shane Greenstein. The Internet and Local Wages: A Puzzle. American Economic Review, 102(1):556–75, 2012.

Rikard Forslid and Karen Helene Midelfart. Internationalisation, Industrial Policy and Clusters. Journal of International Economics, 66(1):197–213, 2005.

Rigmor Kringelholt Fowkes, João Sousa, and Neil Duncan. Evaluation of Research and Development Tax Credit. HM Revenue and Customs HMRC Working Paper, 17, 2015.

Yvonne Giesing and Nadzeya Laurentsyeva. Firms Left Behind: Emigration and Firm Productivity. Working Paper, 2021.

Sourafel Girma, Holger Görg, Eric Strobl, and Frank Walsh. Creating Jobs through Public Subsidies: An Empirical Analysis. Labour Economics, 15:1179–1199, 2008.

Pinelopi Koujianou Goldberg, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova. Imported Intermediate Inputs and Domestic Product Growth: Evidence from India. Quarterly Journal of Economics, 125(4):1727–1767, 2010.

Austan Goolsbee. Does Government R&D Policy Mainly Benefit Scientists and Engineers? American Economic Review, 88:298–302, 1998.

Holger Görg, Michael Henry, and Eric Strobl. Grant Support and Exporting Activity. Review of Economics and Statistics, 90(1):168–174, 2008.

Bruce Greenwald and Joseph E Stiglitz. Helping Infant Economies Grow: Foundations of Trade Policies for Developing Countries. American Economic Review, 96(2):141–146, 2006.

Veaceslav Grigoraş, Andrei Tănase, and Alexandru Leonte. Studiu al Evoluțiilor Sectorului IT&C în România. Technical report, National Bank of Romania, 2017.

Irem Guceri. Will the Real R&D Employees Please Stand Up? Effects of Tax breaks on Firm-Level Outcomes. International Tax and Public Finance, 25:1–63, 2018.

Ann Harrison and Andrés Rodríguez-Clare. Trade, Foreign Investment, and Industrial Policy. Handbook of Development Economics, 5, 2010.

Albert Hirschman. The Strategy of Economic Development. Yale University Press, 1958.

Kenan Huremović and Fernando Vega-Redondo. Production Networks. Working Paper, 2016.

Oleg Itskhoki and Benjamin Moll. Optimal Development Policies with Financial Frictions. Econometrica, 2018.

Dale W Jorgenson, Mun S Ho, and Kevin J Stiroh. A Retrospective Look at the US Productivity Growth Resurgence. Journal of Economic Perspectives, 22(1):3–24, 2008.

Omar Joya and Eric Rougier. Do (All) Sectoral Shocks Lead to Aggregate Volatility? Empirics from a Production Network Perspective. European Economic Review, 113:77–107, 2019.

Réka Juhász. Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade. American Economic Review, 2018.

Ulrich Kaiser and Johan M. Kuhn. Worker-Level and Firm-Level Effects of a Wage Subsidy Program for Highly Educated Labor: Evidence from Denmark. Research Policy, 45:1939–1943, 2016.

Şebnem Kalemli-Özcan, Bent Sørensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yeşiltaş. How to Construct Nationally Representative Firm Level Data from the ORBIS Global Database. Note, 2015.

Aki Kangasharju. Do Wage Subsidies Increase Employment in Subsidized Firms? Economica, 74:51–67, 2007.

Hiau Looi Kee. Local Intermediate Inputs and the Shared Supplier Spillovers of Foreign Direct Investment. Journal of Development Economics, 112:56 – 71, 2015. ISSN 0304-3878.

Richard Kneller and Jonathan Timmis. ICT and Exporting: The Effects of Broadband on the Extensive Margin of Business Service Exports. Review of International Economics, 24(4):757–796, 2016.

Nathaniel Lane. Manufacturing Revolutions: Industrial Policy and Networks in South Korea. Working Paper, 2017.

Nathaniel Lane. The New Empirics of Industrial Policy. Review Article Requested by Dani Rodrik for a Special Issue on Industrial Policy. Journal of Industry, Competition and Trade, 20(3), 2020.

Ernest Liu. Industrial Policies in Production Networks. Quarterly Journal of Economics, 134(4), 2019.

Boris Lokshin and Pierre Mohnen. Do R&D Tax Incentives Lead to Higher Wages for R&D Workers? Evidence from The Netherlands. Research Policy, 42(3):823 – 830, 2013.

Stefano Lombardi, Oskar Nordström Skans, and Johan Vikström. Targeted Wage Subsidies and Firm Performance. Labour Economics, 53:33–45, 2018.

Jonas Månsson and A.M.M. Shahiduzzaman Quoreshi. Evaluating Regional Cuts in the Payroll Tax from a Firm Perspective. Annals of Regional Science, 54:323–347, 2015.

Jakob Munch and Georg Schaur. The Effect of Export Promotion on Firm-Level Performance. American Economic Journal: Economic Policy, 10(1):357–87, 2018.

Howard Pack and Larry E Westphal. Industrial Strategy and Technological Change: Theory Versus Reality. Journal of development economics, 22(1):87–128, 1986.

Brian Quistorff and Sebastian Galiani. The synth_runner Package: Utilities to Automate Synthetic Control Estimation Using **synth**, 2017. URL https://github.com/bquistorff/synth_runner. Version 1.6.0.

Andrés Rodríguez-Clare. Multinationals, Linkages, and Economic Development. American Economic Review, 86(4):852–873, 1996.

Dani Rodrik. Coordination Failures and Government policy: A Model with Applications to East Asia and Eastern Europe. Journal of International Economics, 40(1):1 – 22, 1996.

Dani Rodrik. Industrial Policy for the Twenty-first Century. Technical report, UNIDO Working Paper, 2004.

Dani Rodrik. Normalizing Industrial Policy. Commission on Growth and Development Working Paper No. 3, 2008.

Emmanuel Saez, Benjamin Schoefer, and David Seim. Payroll Taxes, Firm Behavior, and Rent Sharing: Evidence from a Young Workers' Tax Cut in Sweden. American Economic Review, 109(5):1717–1763, 2019.

Claudia Steinwender. Real Effects of Information Frictions: When the States and the Kingdom Became United. American Economic Review, 108(3):657–96, 2018.

Kevin J. Stiroh. Information Technology and the U.S. Productivity Revival: What Do the Industry Data Say? American Economic Review, 92(5):1559–1576, 2002.

Patricia Succar. The Need for Industrial Policy in LDC's-A Re-statement of the Infant Industry Argument. International Economic Review, pages 521–534, 1987.

Chad Syverson. What Determines Productivity? Journal of Economic Literature, 49(2):326–365, 2011.

Petia Topalova and Amit Khandelwal. Trade Liberalization and Firm Productivity: The Case of India. Review of Economics and Statistics, 93(3):995–1009, 2011.

Bart van Ark, Robert Inklaar, and Robert H. McGuckin. ICT and Productivity in Europe and the United States: Where Do the Differences Come From? Working Paper, 2003.

Bart Van Ark, Mary O'Mahony, and Marcel P. Timmer. The Productivity Gap between Europe and the United States: Trends and Causes. Journal of Economic Perspectives, 22(1):25–44, 2008.

Erik van der Marel, Fredrik Erixon, Oscar Guinea, and Philipp Lamprecht. Are Services Sick? How Going Digital Can Cure Services Performance. Global Economic Dynamics, Bertelsmann Stiftung Report, 2020.

Figures

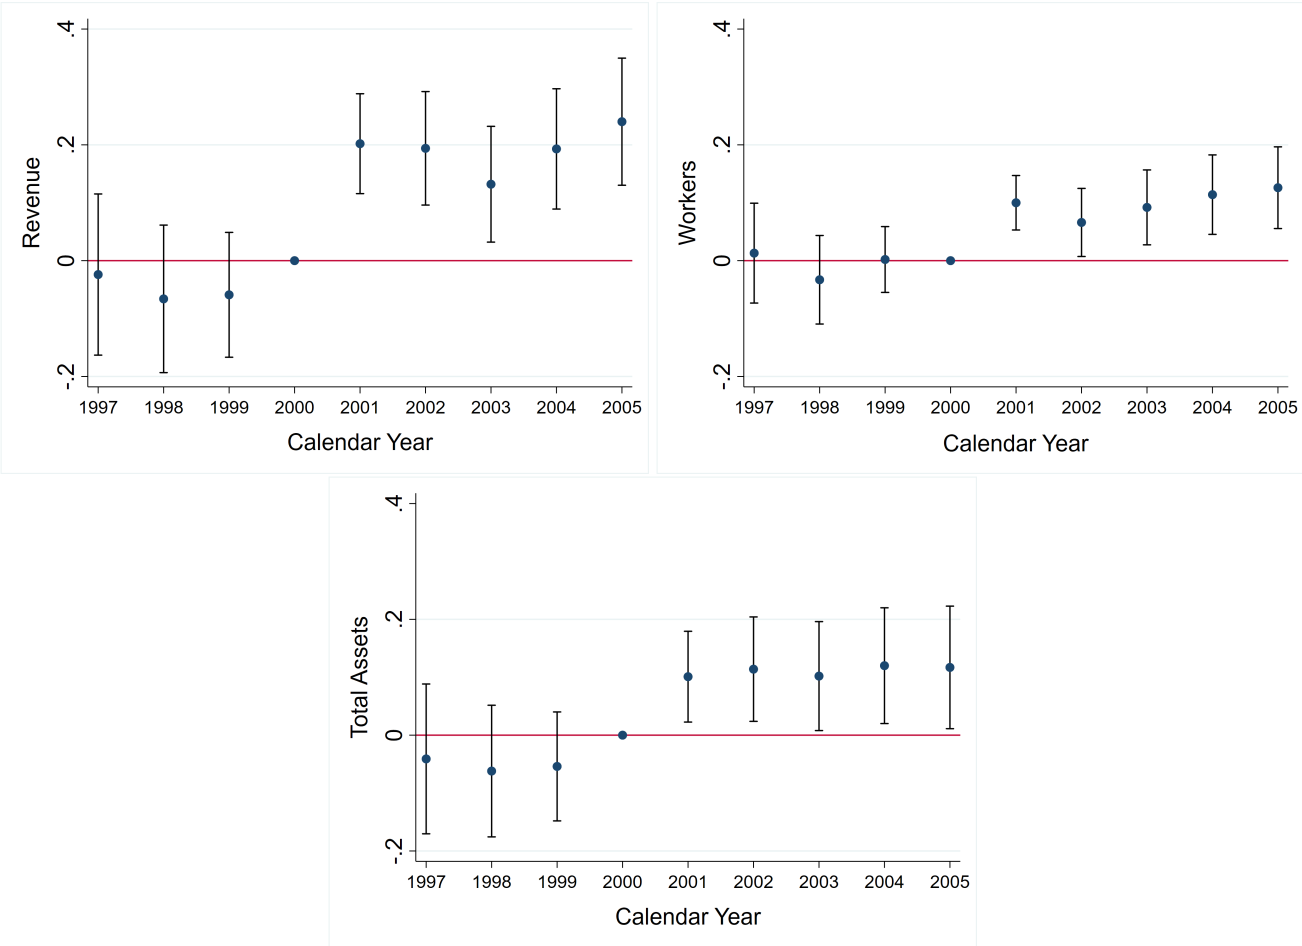

Figure 1: DiD Estimates of the Effects of the 2001 Income Tax Break. Amadeus Data

Notes: In this exercise, we study the firm-level impact of the introduction of the 2001 income tax break to workers in IT. Figure 1 plots the baseline estimates of the yearly DiD coefficients from Equation (1), $\beta_{DiD,t}$, together with their 95% confidence intervals. We consider three firm-level outcome variables: log operating revenues, log number of workers, and log total assets. The coefficients for the year 2000, the year prior to the introduction of the tax break, are normalized to zero. Treated firms are those in the NACE Rev 1 sector 722 (software consultancy and supply). Firms in the baseline control group (used in these figures) are in high-tech knowledge-intensive service sectors (as classified by the Eurostat). All regressions in this figure use Amadeus data. The regression model includes firm and calendar year fixed effects. See columns (1)-(3) from Table 1 for more details.

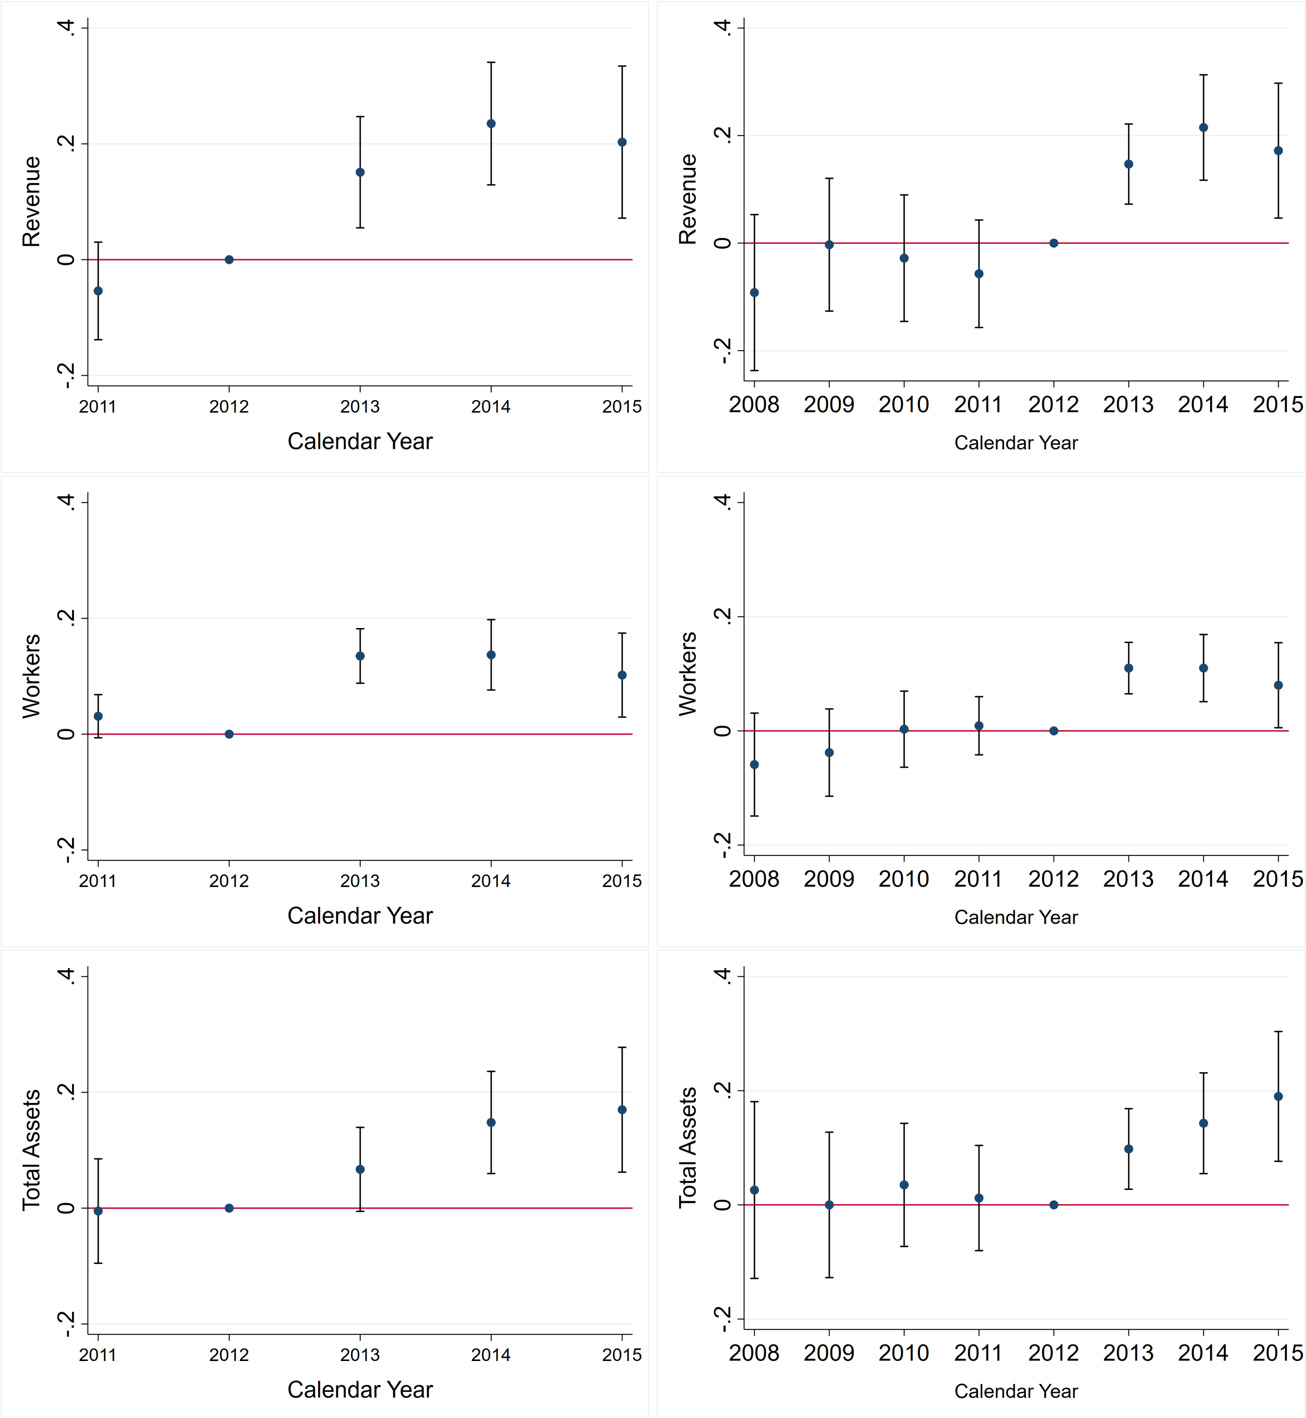

Figure 2: DiD Estimates of the Effects of the 2013 Reform. Administrative (Left) and Amadeus Data (Right)

Notes: In this exercise, we study the firm-level impact of the introduction of the 2013 reform to the 2001 tax break. Figure 2 plots the baseline estimates of the yearly DiD coefficients from Equation (2), $\beta_{DiD,t}$, together with their 95% confidence intervals for two samples: the one based on administrative data from the Ministry of Public Finance (sub-figures on the left) and the one based on Amadeus data from Bureau van Dijk (sub-figures on the right). We consider three firm-level outcome variables: log revenues, log number of workers, and log total assets. The coefficients for the year 2012, the year prior to the 2013 reform, are normalized to zero. Treated firms are those whose share of income tax-exempt workers jumps from under 5% to over 20% after 2013. Firms in the baseline control group (used in these figures) are in high-tech knowledge-intensive service sectors (as classified by the Eurostat) and have an under 5% share of income tax-exempt workers throughout the entire 2011–2015 period. The regression model includes firm and sector-by-year fixed effects. See columns (1)-(3) and (4)-(6) from Table 3 for more details.

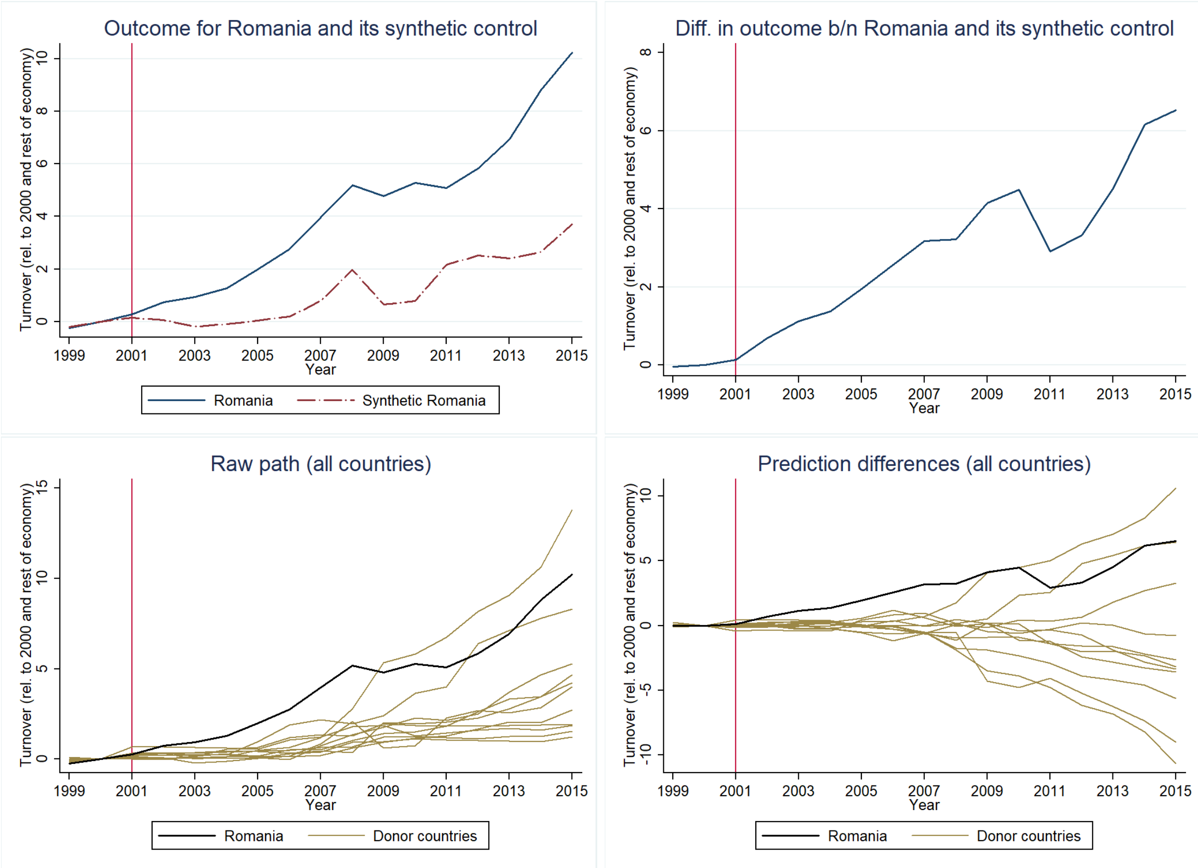

Figure 3: IT Sector Vs. Rest of the Economy. SCM with Outcome Variable: Gross Revenues ("Turnover or Gross Premiums Written") - Million Euros (Normalized)

Notes: In this exercise we use the synthetic control method introduced in Section 4.1 to study the sector-level direct effects of the introduction of the 2001 law granting an income tax break to workers in IT. All figures have as dependent variable the country-level (normalized) gross revenues ("Turnover or gross premiums written - million euros"). The yearly absolute value of the dependent variable in the treated sector is divided by its value in 2000, the year prior to the introduction of the income tax break in Romania. From these resulting yearly ratios we subtract the corresponding ratios for the comparison sectors. The treated sector is K72 ("Computer and related activities," including "Software consultancy and supply " and "Publishing of software"). We use as comparison sectors all other sectors in the economy (all except K72). The data source for the dependent variable is Eurostat, Structural Business Statistics, Annual detailed enterprise statistics on services (NACE Rev 1.1). Data for the predictors comes from the World Bank, World Development Indicators. We use as predictors the "GDP per capita (constant LCU)," "Medium and high-tech sector (% manufacturing value added)" and "Services, etc., value added (% of GDP)." All figures are an output of the **synth_runner** package for Stata [Quistorff and Galiani, 2017], with the *nested* option specified.

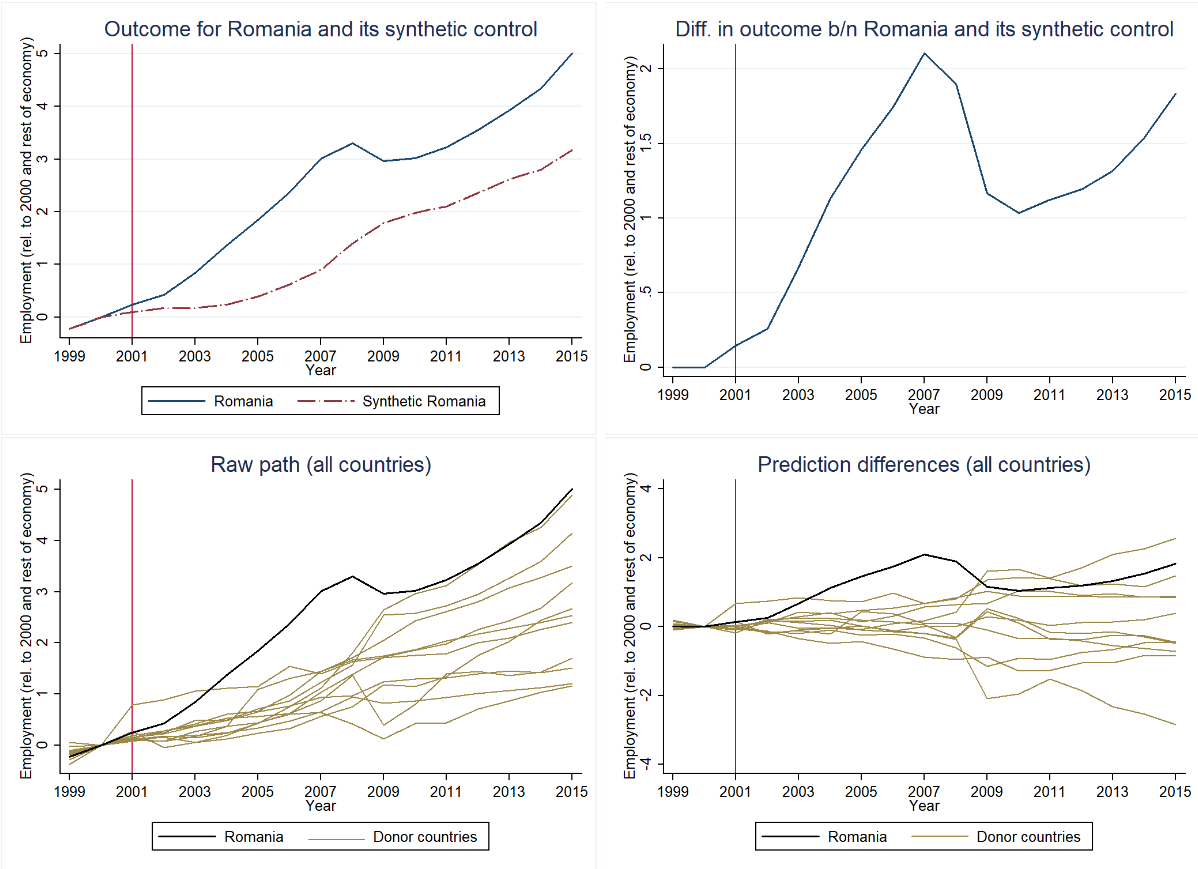

Figure 4: IT Sector Vs. Rest of the Economy. SCM with Outcome Variable: "Employees - Number" (Normalized)

Notes: In this exercise we use the synthetic control method introduced in Section 4.1 to study the sector-level direct effects of the introduction of the 2001 law granting an income tax break to workers in IT. All figures have as dependent variable the country-level (normalized) "Employees - number" The yearly absolute value of the dependent variable in the treated sector is divided by its value in 2000, the year prior to the introduction of the income tax break in Romania. From these resulting yearly ratios we subtract the corresponding ratios for the comparison sectors. The treated sector is K72 ("Computer and related activities," including "Software consultancy and supply " and "Publishing of software"). We use as comparison sectors all other sectors in the economy (all except K72). The data source for the dependent variable is Eurostat, Structural Business Statistics, Annual detailed enterprise statistics on services (NACE Rev 1.1). Data for the predictors comes from the World Bank, World Development Indicators. We use as predictors the "GDP per capita (constant LCU)," "Medium and high-tech sector (% manufacturing value added)" and "Services, etc., value added (% of GDP)." All figures are an output of the **synth_runner** package for Stata [Quistorff and Galiani, 2017], with the *nested* option specified.

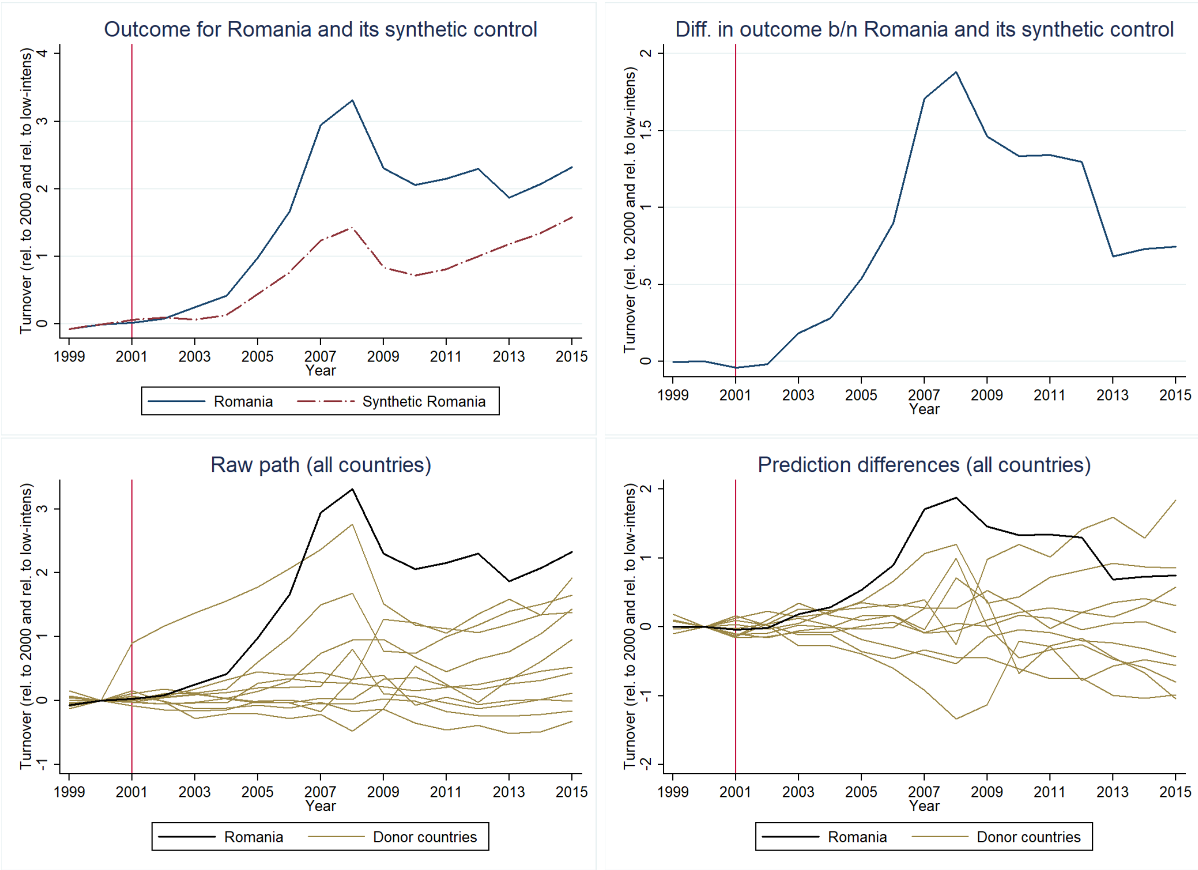

Figure 5: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: Gross Revenues ("Turnover or Gross Premiums Written") - Million Euros (Normalized)

Notes: In this exercise we use the synthetic control method introduced in Section 4.2 to study the sector-level downstream effects of the introduction of the 2001 law granting an income tax break to workers in IT. All figures have as dependent variable the country-level (normalized) gross revenues ("Turnover or gross premiums written - million euros"). The yearly absolute value of the dependent variable in the treated sector is divided by its value in 2000, the year prior to the introduction of the income tax break in Romania. From these resulting yearly ratios we subtract the corresponding ratios for the comparison sectors. The treated sectors are those that use K72 ("Computer and related activities," including "Software consultancy and supply" and "Publishing of software") services at high-intensity. We exclude K72 itself from this category. Sectors that have a low-intensity of use of K72 services serve as comparison sectors. The data source for the dependent variable is Eurostat, Structural Business Statistics, Annual detailed enterprise statistics on services (NACE Rev 1.1). Data for the predictors comes from the World Bank, World Development Indicators. We use as predictors the "GDP per capita (constant LCU)," "Medium and high-tech sector (% manufacturing value added)" and "Services, etc., value added (% of GDP)." All figures are an output of the **synth_runner** package for Stata [Quistorff and Galiani, 2017], with the *nested* option specified.

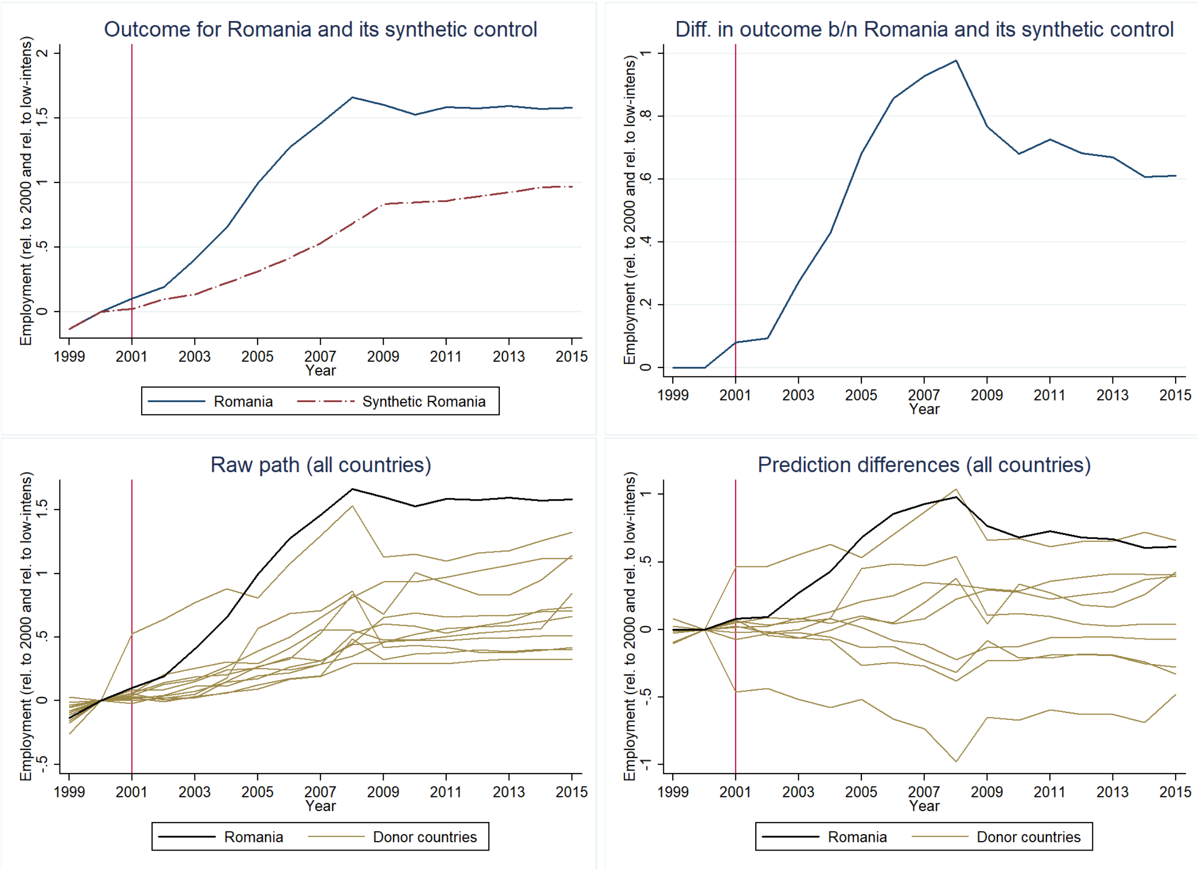

Figure 6: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: "Employees - Number" (Normalized)

Notes: In this exercise we use the synthetic control method introduced in Section 4.2 to study the sector-level downstream effects of the introduction of the 2001 law granting an income tax break to workers in IT. All figures have as dependent variable the country-level (normalized) "Employees - number." The yearly absolute value of the dependent variable in the treated sector is divided by its value in 2000, the year prior to the introduction of the income tax break in Romania. From these resulting yearly ratios we subtract the corresponding ratios for the comparison sectors. The treated sectors are those that use K72 ("Computer and related activities," including "Software consultancy and supply" and "Publishing of software") services at high-intensity. We exclude K72 itself from this category. Sectors that have a low-intensity of use of K72 services serve as comparison sectors. The data source for the dependent variable is Eurostat, Structural Business Statistics, Annual detailed enterprise statistics on services (NACE Rev 1.1). Data for the predictors comes from the World Bank, World Development Indicators. We use as predictors the "GDP per capita (constant LCU)," "Medium and high-tech sector (% manufacturing value added)" and "Services, etc., value added (% of GDP)." All figures are an output of the **synth_runner** package for Stata [Quistorff and Galiani, 2017], with the *nested* option specified.

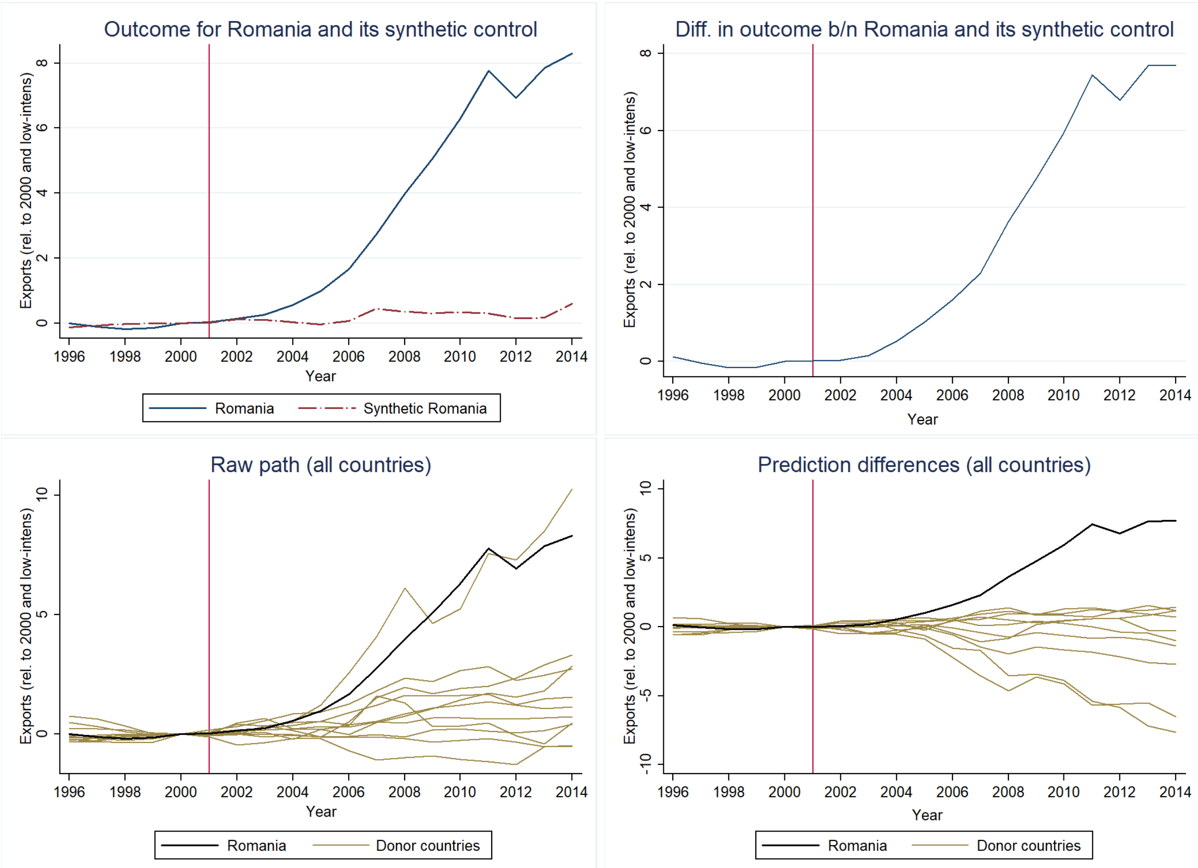

Figure 7: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: "Goods Export Value" (Normalized)

Notes: In this exercise we use the synthetic control method introduced in Section 4.2 to study the sector-level downstream effects of the introduction of the 2001 law granting an income tax break to workers in IT. All figures have as dependent variable the country-level (normalized) "(Goods Export) Trade Value (US$)." The yearly absolute value of the dependent variable in the treated sector is divided by its value in 2000, the year prior to the introduction of the income tax break in Romania. From these resulting yearly ratios we subtract the corresponding ratios for the comparison sectors. The treated sectors are those that use K72 ("Computer and related activities," including "Software consultancy and supply" and "Publishing of software") services at high-intensity. K72 itself is excluded from this category. Sectors that have a low-intensity of use of K72 services serve as comparison sectors. The data source for the dependent variable is UN Comtrade, Goods Exports. Data for the predictors comes from the World Bank, World Development Indicators. We use as predictors the "GDP per capita (constant LCU)," "Medium and high-tech sector (% manufacturing value added)" and "Services, etc., value added (% of GDP)." All figures are an output of the **synth_runner** package for Stata [Quistorff and Galiani, 2017], with the *nested* option specified.

Tables

Table 1: Difference-in-Difference Around 2001 Income Tax Break: Main Results for the Baseline and Alternative Control Groups

	Baseline sample			Alternative control group 1			Alternative control group 2		
Dep. var.	Revenues (1)	Workers (2)	Assets (3)	Revenues (4)	Workers (5)	Assets (6)	Revenues (7)	Workers (8)	Assets (9)
$\beta_{DiD,1997}$	-0.024	0.013	-0.041	-0.026	-0.005	-0.005	-0.043	0.014	0.005
	(0.071)	(0.044)	(0.066)	(0.101)	(0.063)	(0.091)	(0.083)	(0.050)	(0.077)
$\beta_{DiD,1998}$	-0.066	-0.033	-0.062	-0.124	-0.064	-0.038	-0.075	-0.023	-0.026
	(0.065)	(0.039)	(0.058)	(0.089)	(0.053)	(0.081)	(0.074)	(0.044)	(0.068)
$\beta_{DiD,1999}$	-0.059	0.002	-0.054	-0.026	0.010	0.053	-0.066	0.014	-0.029
	(0.055)	(0.029)	(0.048)	(0.074)	(0.040)	(0.066)	(0.062)	(0.033)	(0.057)
$\beta_{DiD,2001}$	0.202***	0.100***	0.101**	0.198***	0.109***	0.114**	0.250***	0.122***	0.166***
	(0.044)	(0.024)	(0.040)	(0.059)	(0.032)	(0.053)	(0.051)	(0.028)	(0.048)
$\beta_{DiD,2002}$	0.194***	0.066**	0.114**	0.218***	0.088**	0.197***	0.252***	0.098***	0.183***
	(0.050)	(0.030)	(0.046)	(0.066)	(0.039)	(0.061)	(0.058)	(0.035)	(0.055)
$\beta_{DiD,2003}$	0.132***	0.092***	0.102**	0.208***	0.099**	0.230***	0.143**	0.144***	0.142**
	(0.051)	(0.033)	(0.048)	(0.069)	(0.042)	(0.064)	(0.059)	(0.038)	(0.057)
$\beta_{DiD,2004}$	0.193***	0.114***	0.120**	0.298***	0.136***	0.266***	0.191***	0.163***	0.147**
	(0.053)	(0.035)	(0.051)	(0.074)	(0.046)	(0.069)	(0.062)	(0.040)	(0.060)
$\beta_{DiD,2005}$	0.240***	0.126***	0.117**	0.301***	0.140***	0.259***	0.229***	0.183***	0.135**
	(0.056)	(0.036)	(0.054)	(0.077)	(0.048)	(0.072)	(0.066)	(0.043)	(0.063)
Adjusted R^2	0.819	0.822	0.833	0.832	0.849	0.849	0.813	0.836	0.830
# Observations	25,518	25,518	25,518	12,203	12,203	12,203	15,931	15,931	15,931
# Firms	4,336	4,336	4,336	2,120	2,120	2,120	2,696	2,696	2,696

Notes: In this exercise we study the firm-level impact of the introduction of the 2001 law granting an income tax break to workers in IT. All regressions in this table use Amadeus data. In this table we report the point estimates of the difference-in-difference coefficients of interest from Equation (1), i.e., $\beta_{DiD,t}$. The outcome variables used are log operating revenues, log number of workers, and log total assets. The year prior to the introduction of the law (2000) is the reference year. Treated firms are those in the NACE Rev 1 sector 722 (software consultancy and supply). Firms in the baseline control group – columns (1)-(3) – are in high-tech knowledge-intensive service sectors (as classified by the Eurostat). Firms in the alternative control group 1 – columns (4)-(6) – are either part of untreated IT service sectors or R&D sectors. Firms in the alternative control group 2 – columns (7)-(9) – are part of high-tech knowledge-intensive service sectors (as classified by the Eurostat) *excluding* ICT service sectors. All specifications include firm and calendar year fixed effects. Robust standard errors, clustered at the firm-level, in parenthesis. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 2: Difference-in-Differences Design Around 2013 Reform: "First Stage" Effects on the Share of Income Tax-Exempt Employees

Dep. var.	1 if more than 20% exempted employees				Share of exempted employees			
Sample	Baseline (1)	Alternative 1 (2)	Alternative 2 (3)	Eligible (4)	Baseline (5)	Alternative 1 (6)	Alternative 2 (7)	Eligible (8)
Elig. sect. × d_{2011}	-0.002	-0.002	-0.002		-0.003	-0.002	-0.003	
	(0.006)	(0.006)	(0.006)		(0.003)	(0.003)	(0.003)	
Elig. sect. × d_{2013}	0.023***	0.024***	0.023***		0.012***	0.012***	0.012***	
	(0.005)	(0.005)	(0.005)		(0.003)	(0.003)	(0.003)	
Elig. sect. × d_{2014}	0.050***	0.050***	0.049***		0.028***	0.028***	0.027***	
	(0.006)	(0.006)	(0.006)		(0.004)	(0.004)	(0.004)	
Elig. sect. × d_{2015}	0.071***	0.072***	0.071***		0.038***	0.039***	0.038***	
	(0.007)	(0.007)	(0.007)		(0.005)	(0.005)	(0.005)	
d_{2011}				-0.002				-0.002
				(0.006)				(0.003)
d_{2013}				0.024***				0.012***
				(0.005)				(0.003)
d_{2014}				0.049***				0.028***
				(0.006)				(0.004)
d_{2015}				0.071***				0.038***
				(0.007)				(0.005)
Adjusted R^2	0.174	0.166	0.102	0.004	0.152	0.145	0.087	0.003
# Observations	39,052	37,349	28,269	21,601	39,052	37,349	28,269	21,601
# Firms	8,990	8,551	6,547	4,986	8,990	8,551	6,547	4,986

Notes: In this exercise, we study the "first stage" effects of the 2013 income tax reform on the firm-level share of workforce exemption from the income tax. All regressions in this table use administrative data. Columns (1)-(4) use as the dependent variable a dummy variable that takes value 1 for firm i in year t if firm i has more than 20% of its workers exempted from the income tax in year t. Columns (5)-(8) use as the dependent variable the share of workers of firm i who are exempted from the income tax in year t. The year prior to the amendment of 2013 (2012) is the reference year. Firms in the baseline control group – columns (1) and (5) – are in high-tech knowledge-intensive service sectors (as classified by the Eurostat). Firms in the alternative control group 1 – columns (2) and (6) – are in ICT service sectors (as classified by the OECD). Firms in the alternative control group 2 – columns (3) and (7) – are part of high-tech knowledge-intensive service sectors (as classified by the Eurostat) *excluding* ineligible ICT service sectors (as classified by the OECD). Firms in alternative control 3 are only part of eligible sectors. The samples used in this table are different from the main samples in Table 3 and Table 4 in that the former keep *all* firms in HTKI (high-tech knowledge-intensive) service sectors, in ICT service sectors, in HTKI service sectors *excluding* the ICT service sectors, and in eligible sectors, whereas the latter bring additional restrictions on the firm-level share of workforce exemption before and after 2013. In addition to interaction terms between a dummy taking value one if the sector of firm i is eligible for the income tax break of its eligible workers (Eligible sector) and a year dummy (d_t), columns (1), (2), (3), (5), (6) and (7) include firm fixed effects and year fixed effects (d_t, whose estimates are omitted from the table). In addition to the year fixed effects (d_t) whose estimates are reported in the table, columns (4) and (8) also include firm fixed effects. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 3: Difference-in-Differences Design Around 2013 Reform: Main Results for the Baseline and Alternative Control Groups (1/2)

	Administrative, baseline sample			Amadeus, baseline sample			Administrative, alternative control 1			Amadeus, alternative control 1		
	Revenues	Workers	Assets	Revenues	Workers	Assets	Revenues	Workers	Assets	Revenues	Workers	Assets
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\beta_{DiD,2008}$				-0.092	-0.059	0.026				-0.098	-0.061	0.023
				(0.074)	(0.046)	(0.079)				(0.074)	(0.046)	(0.080)
$\beta_{DiD,2009}$				-0.003	-0.038	-0.000				-0.006	-0.039	-0.004
				(0.063)	(0.039)	(0.065)				(0.063)	(0.039)	(0.065)
$\beta_{DiD,2010}$				-0.028	0.003	0.035				-0.029	0.004	0.033
				(0.060)	(0.034)	(0.055)				(0.060)	(0.034)	(0.055)
$\beta_{DiD,2011}$	-0.054	0.031	-0.005	-0.057	0.009	0.012	-0.065	0.031	-0.008	-0.070	0.009	0.008
	(0.043)	(0.019)	(0.046)	(0.051)	(0.026)	(0.047)	(0.043)	(0.019)	(0.046)	(0.052)	(0.026)	(0.047)
$\beta_{DiD,2013}$	0.151***	0.135***	0.067*	0.147***	0.110***	0.098***	0.147***	0.136***	0.064*	0.143***	0.111***	0.094***
	(0.049)	(0.024)	(0.037)	(0.038)	(0.023)	(0.036)	(0.049)	(0.024)	(0.037)	(0.039)	(0.023)	(0.036)
$\beta_{DiD,2014}$	0.235***	0.137***	0.148***	0.215***	0.110***	0.143***	0.232***	0.138***	0.143***	0.212***	0.111***	0.137***
	(0.054)	(0.031)	(0.045)	(0.050)	(0.030)	(0.045)	(0.054)	(0.031)	(0.045)	(0.050)	(0.030)	(0.045)
$\beta_{DiD,2015}$	0.203***	0.102***	0.170***	0.172***	0.080**	0.190***	0.203***	0.105***	0.168***	0.173***	0.082**	0.188***
	(0.067)	(0.037)	(0.055)	(0.064)	(0.038)	(0.058)	(0.067)	(0.037)	(0.055)	(0.064)	(0.038)	(0.058)
Adjusted R^2	0.848	0.910	0.905	0.792	0.855	0.849	0.851	0.899	0.896	0.798	0.847	0.843
# Observations	27,491	27,491	27,491	32,715	32,715	32,715	25,916	25,916	25,916	31,651	31,651	31,651
# Firms	6,146	6,146	6,146	4,837	4,837	4,837	5,769	5,769	5,769	4,676	4,676	4,676

Notes: In this exercise we study the firm-level impact of the 2013 expansion to the income tax break law of 2001. We present both the baseline results on the administrative in columns (1)-(3) and their robustness to alternative data sources and control groups. The outcome variables used are log revenues, log number of workers, and log total assets. The difference between columns (1)-(3) and columns (4)-(6) (between columns (7)-(9) and (10)-(12)) stems from the source of the data over which we estimate the model in Equation (2). Columns (1)-(3) and (7)-(9) use the administrative data described in Section 3.2. Columns (4)-(6) and (10)-(12) use data from Amadeus, a Bureau van Dijk product. Firms in the baseline control group – columns (1)-(6) – are part of high-tech knowledge-intensive service sectors (as classified by the Eurostat). Firms in the alternative control group 1 – columns (7)-(12) – are in ICT service sectors (as classified by the OECD). Robust standard errors, clustered at the firm-level, in parenthesis. All specifications include firm and sector-by-year fixed effects. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4: Difference-in-Differences Design Around 2013 Reform: Robustness of the Results to Alternative Control Groups (2/2)

	Administrative, alternative control 2			Amadeus, alternative control 2			Administrative, alternative control 3			Amadeus, alternative control 3		
	Revenues	Workers	Assets	Revenues	Workers	Assets	Revenues	Workers	Assets	Revenues	Workers	Assets
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\beta_{DiD,2008}$				-0.084	-0.045	0.028				-0.087	-0.046	0.025
				(0.077)	(0.046)	(0.083)				(0.077)	(0.046)	(0.083)
$\beta_{DiD,2009}$				-0.004	-0.030	-0.009				-0.006	-0.031	-0.013
				(0.065)	(0.039)	(0.068)				(0.066)	(0.039)	(0.068)
$\beta_{DiD,2010}$				-0.034	0.011	0.008				-0.032	0.013	0.007
				(0.062)	(0.034)	(0.056)				(0.062)	(0.034)	(0.056)
$\beta_{DiD,2011}$	-0.051	0.025	-0.019	-0.074	0.003	-0.009	-0.055	0.026	-0.023	-0.079	0.003	-0.014
	(0.044)	(0.020)	(0.048)	(0.053)	(0.026)	(0.049)	(0.044)	(0.020)	(0.048)	(0.053)	(0.026)	(0.049)
$\beta_{DiD,2013}$	0.173***	0.137***	0.069*	0.151***	0.115***	0.097***	0.170***	0.139***	0.065*	0.149***	0.118***	0.093**
	(0.047)	(0.024)	(0.037)	(0.039)	(0.024)	(0.037)	(0.048)	(0.024)	(0.037)	(0.040)	(0.024)	(0.037)
$\beta_{DiD,2014}$	0.242***	0.134***	0.147***	0.213***	0.114***	0.141***	0.239***	0.136***	0.143***	0.212***	0.116***	0.136***
	(0.056)	(0.031)	(0.046)	(0.051)	(0.030)	(0.046)	(0.056)	(0.032)	(0.046)	(0.052)	(0.031)	(0.046)
$\beta_{DiD,2015}$	0.203***	0.104***	0.166***	0.170***	0.090**	0.192***	0.204***	0.106***	0.164***	0.171***	0.092**	0.189***
	(0.068)	(0.037)	(0.056)	(0.066)	(0.039)	(0.059)	(0.068)	(0.037)	(0.056)	(0.066)	(0.039)	(0.059)
Adjusted R^2	0.835	0.919	0.904	0.787	0.868	0.847	0.822	0.902	0.875	0.779	0.859	0.817
# Observations	17,238	17,238	17,238	20,550	20,550	20,550	10,934	10,934	10,934	13,227	13,227	13,227
# Firms	3,882	3,882	3,882	3,050	3,050	3,050	2,473	2,473	2,473	1,994	1,994	1,994

Notes: In this exercise we check the robustness of the baseline results from Table 3 on the firm-level impact of the 2013 amendment to the income tax break law of 2001. The outcome variables used are log revenues, log number of workers, and log total assets. The difference between columns (1)-(3) and columns (4)-(6) (between columns (7)-(9) and (10)-(12)) stems from the source of the data over which we estimate the model in Equation (2). Columns (1)-(3) and (7)-(9) use the administrative data described in Section 3.2. Columns (4)-(6) and (10)-(12) use data from Amadeus, a Bureau van Dijk product. Firms in the alternative control group 2 – columns (1)-(6) – are part of high-tech knowledge-intensive service sectors (as classified by the Eurostat) *excluding* ineligible ICT service sectors. Firms in the alternative control group 3 – columns (7)-(12) – belong only to eligible sectors. Thus the comparison is between firms under the 5% threshold of exempted employees with firms that jump over the 20% threshold after 2013 (all in eligible sectors). Robust standard errors, clustered at the firm-level, in parenthesis. All specifications include firm and sector-by-year fixed effects. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 5: Difference-in-Differences Design Around 2013 Reform: Heterogeneity of the Baseline Results

	Revenues (1)	Workers (2)	Assets (3)	Revenues (4)	Workers (5)	Assets (6)
Panel A: Size	*Micro: < 10 workers*			*Small, medium, large: $\geq$ 10 workers*		
β_{DiD}	0.213***	0.113***	0.128***	0.251***	0.110	0.129
	(0.055)	(0.028)	(0.046)	(0.085)	(0.082)	(0.082)
Adjusted R^2	0.753	0.772	0.850	0.893	0.878	0.949
# Observations	23,903	23,903	23,903	3,588	3,588	3,588
# Firms	5,384	5,384	5,384	762	762	762
Panel B: Age	*Young: < Five years old*			*Old: $\geq$ Five years old*		
β_{DiD}	0.271***	0.149***	0.163**	0.141***	0.072**	0.065
	(0.088)	(0.045)	(0.068)	(0.053)	(0.032)	(0.050)
Adjusted R^2	0.758	0.836	0.828	0.890	0.934	0.937
# Observations	10,244	10,244	10,244	17,247	17,247	17,247
# Firms	2,358	2,358	2,358	3,788	3,788	3,788

Notes: In this exercise we explore the heterogeneity of the baseline effects of the 2013 expansion to the income tax break law of 2001. The outcome variables used are log revenues, log number of workers, and log total assets. For brevity, we implement the pooled version of the DiD Equation (2) on the baseline sample of 6,146 firms (see Table 3, columns (1)-(3)). Firms in the baseline control group are in high-tech knowledge-intensive service sectors (as classified by the Eurostat). The sample is split in two parts based on the number of workers or age of the firm (both in 2011). The DiD regressions are run separately on each part of the baseline sample. All specifications include firm and sector-by-year fixed effects. Robust standard errors, clustered at the firm-level, in parenthesis. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Back-of-the-Envelope Cost-Benefit Calculation for the 2013 Reform

Outcome	Base year 2012 (1)	End year 2015 (2)
Benefits		
Revenues		
Average revenues (ths euros)	317.03	
Total revenues (ths euros)	146,467.86	
Estimated average DiD effect on revenues		0.203
Extra revenues of the average treated firm (ths euros)		64.36
Total extra revenues of treated firms (ths euros)		29,732.98
Employment		
Average number of workers	6.38	
Total number of workers	2,947.56	
Estimated average DiD effect on employment		0.102
Extra workers in the average treated firm		0.65
Total extra workers in treated firms		300.65
Assets		
Average assets (ths euros)	203.07	
Total assets (ths euros)	93,818.34	
Estimated average DiD effect on assets		0.170
Extra assets of the average treated firm (ths euros)		34.52
Total extra assets of treated firms (ths euros)		15,949.12
Costs		
Average annual labor income (ths euros)		17.20
Annual exemption per exempted worker (ths euros)		2.30
Number of exempted workers in treated firms	16	1,868
Total extra cost (=extra annual foregone income taxes) (ths euros)		4,256.72

Notes: This table reports a back-of-the-envelope calculation of the costs and benefits of the 2013 reform for the year 2015. All values in this table characterize only the 462 firms treated by the 2013 reform – according to our baseline definition of treatment (i.e., firms that had less than 5% of tax-exempt workers pre-2013 and that experienced a jump to more than 20% tax-exempt workers post-2013, plausibly due to the sudden and sizable expansion in the scope of the tax break). The 2012 average revenues, number of workers, and assets for treated firms are the same averages reported in the upper panel of Table B3 (Online Appendix B.1.2). The estimated difference-in-difference (DiD) effects on revenues, employment, and assets for 2015 are the baseline DiD effects reported in columns (1)-(3) of Table 3. We obtain the total extra revenues (workers or assets) of treated firms in 2015 by multiplying the total revenues (workers or assets) of treated firms in 2012 by the estimated average DiD effect for 2015 (relative to 2012). To compute the total extra cost to the government of the 2013 reform, we multiply the 2015 value of the annual personal income tax exemption (=2.3 thousand euros per exempted worker) by the extra number of exempted workers in treated firms (=1,868-16). The 2015 value of the annual personal income tax exemption is equal to 16% (the flat personal income tax rate) of the taxable part of the 2015 average annual labor income of exempted workers in treated firms (=(1-0.165)$\times$17.2 thousand euros).