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

Isabela Manelici† and Smaranda Pantea‡

June 13, 2019

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 are grateful to Treb Allen, Alan Auerbach, Marc Bogdanowicz, Sebastian Buhai, Teresa Fort, Eva Hagsten, Douglas Irwin, Nathan Lane, Ethan Lewis, Dan Matei, Krisztina Orban, Nina Pavcnik, Andrés Rodríguez-Clare, Jose P. Vasquez, Emmanuel Saez, Benjamin Schoefer, Mine Senses, Danny Yagan, and Paolo Zacchia for feedback. We also appreciate the helpful suggestions of numerous seminar participants and colleagues. 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 Finance of Romania) for their decisive support. We thank Andrei Caprau, 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 (University of California, Berkeley). Email: imanelici@berkeley.edu

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

1 Introduction

> *"When listing the world's most promising places for tech and start-ups, you could be forgiven for overlooking Romania. During the almost three decades since a revolution lifted the nation out of communism, the country has maintained a low profile internationally. Nonetheless, a quieter type of revolution has been percolating behind the scenes." – 2016 article in Quartz, a New York-based business news portal.*

According to both business and technology journalists, Romania has emerged over the past two decades as the "unlikely Silicon Valley of Europe."[1] The comparison with Silicon Valley is noteworthy given that few topics have been researched as extensively as the rise of Silicon Valley and the types of policies that countries pursue in the hope of developing their own high-technology sectors.[2] In this article, we provide firm and sector-level evidence that an industrial policy introduced by Romania in 2001 has been key for the impressive development of its information technology (IT) sector.

Since 2001, this policy provides a full personal income tax break to employees with an eligible bachelor's degree specialization who work directly on software development and generate revenues from this activity for a firm in the IT sector (specifically, in the "Software consultancy and supply" sector, with the NACE Rev 1 code 722). 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.[3] 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 very 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, this restrictive set of conditions ensures that the tax break actually subsidizes software development (as opposed to other misreported activities). Finally, for workers to benefit from this tax break, their employer has to prepare the necessary paperwork and apply for their income tax break. This requirement of an explicit "buy-in" from the employer is another reason to expect a shared economic incidence of this tax incentive.

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. The unexpected nature of both policy events enables us to credibly estimate their impact. The analysis of the 2001 event allows us to study the behavior of firms in a set of comparable sectors (including the IT sector) just before and after the introduction of a differential tax treatment of the IT sector. These DiD estimates measure the intent-to-treat effects of the policy on firms in the eligible IT sector. The analysis of the 2013 reform is based 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.

[1] For example, see 2019 HeadHuntingIT blog post, December 2018 *France 24* article, September 2018 *Accace* article, March 2018 *Outsourcing Portal* article, September 2017 *Financial Times* article, May 2017 *IDG Connect* article, 2017 *Teamfound* blog post, May 2016 *Forbes* article, or April 2016 *TechCrunch* article.

[2] A Google Scholar search on the keywords "Silicon Valley" returns 898,000 papers.

[3] We will use the terms "IT" and "software development" to refer to the sector initially targeted by the 2001 tax break law ("Software consultancy and supply" or NACE Rev 1 722) and to the activity typically performed in that sector. We will also use the term "IT" to refer to the enlarged set of eligible sectors, as per the 2013 amendment. We therefore use the term IT loosely. In a stricter sense, not all IT sectors are eligible for the income tax break of their workers (particularly in 2001).

To study the impact of the 2001 introduction of an income tax break for programmers, 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 (i.e., if in the "Software consultancy and supply" sector, with the NACE Rev 1 code 722).[4] The outcomes of firms in eligible sectors are compared to the outcomes of firms in comparable non-eligible sectors, such as NACE Rev 1 sector 731 (Research and experimental development on natural sciences and engineering). 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. However, after 2001, firms in the eligible IT sector embark on a differential upward trend relative to firms in comparable non-eligible sectors. By 2005, firms in the eligible sector have, on average, 31% higher operating revenues than in 2000, hire 17% more workers, and have 17% more assets than firms in comparable non-eligible sectors. We also find a relative improvement of 12 points in the solvency ratio of firms in the eligible sector. We use six other combinations 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. Given that the administrative data available for this exercise includes the number of employees that are actually exempt from income tax each year, 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 ICT service sectors (eligible and non-eligible) that remain at 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 differential preexisting trends between treated and comparison firms.

Our DiD estimates suggest that firms treated by the 2013 reform experience large and long-lasting increases in size. In 2015, these firms have 37% more revenues than in 2012, 38% more sales, 21% more employees, and 26% more assets (all relative to the comparison firms). These baseline estimates are robust to (i) running the baseline regression on a dataset from Amadeus that starts in 2008 and allows us to observe a longer pre-reform period; (ii) using an alternative comparison group of firms in other high-tech knowledge-intensive service sectors; (iii) restricting the sample only to firms in the eligible IT sector; (iv) using an alternative event-study research design; (v) varying the threshold choice (e.g., from 20 to 15%); and (vi) 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.[5] To interpret the magnitudes of the estimates from the two analyses, we first assume a 25% take-up between 2001 and 2015 for firms in the eligible sector. We can then translate the 17% 2001 pooled DiD (intent-to-treat) estimate on employment into a 68% treatment-on-the-treated estimate. As the median firm in the eligible sector has four workers in 2000, a

[4] In Amadeus one cannot observe whether a firm has workers who benefit from the income tax break or not.

[5] It is not surprising to find at least a partial incidence on firms. Using linked employer-employee data from Sweden, Saez et al. [2017] find that a payroll tax reduction for young workers (from 31.4% to 15.5%) has had a full incidence on firms. The authors find no effect on net-of-tax wages of young treated workers relative to older untreated workers. Firms employing many young workers receive a larger tax windfall and increase their employment, capital, and sales after the reform.

68% increase in employment between 2001 and 2005 is sensible. This 68% estimate is larger than the corresponding 21% 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. The research design is based on the synthetic control method (SCM). We first use data from Eurostat and the World Bank to 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 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 revenue (employment) in the IT sector of Romania is 6.52 (1.83) times larger than the gross revenue (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 revenue (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 in similar economies – the SCM estimates of the expansion of the IT sector in Romania are smaller than its actual expansion.[6] Placebo tests suggest that a similar growth cannot be replicated in countries that did not implement this policy.

Next, we provide evidence on the inter-industry 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 service 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 (for instance, 0.75 times more in terms of gross revenue 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 service inputs is a necessary condition for the IT sector to have generated inter-industry externalities (one of the theoretical conditions that justify industrial policies favoring a certain sector). All the SCM results survive a battery of robustness checks.

[6]The actual gross revenue in IT in Romania grew 14 times (from 282 million euros in 2000 to 4,031 million euros in 2015) and the actual number of workers grew six times (from 13,691 workers in 2000 to 81,780 workers in 2015).

Finally, we propose a set of back-of-the-envelope cost estimates of this policy to the government. The corresponding amount of foregone fiscal revenues is equivalent to 4.7 to 6.4% of the total gross wage bill in the IT sector. In terms of value added of the sector, these foregone tax revenues range between 2.8 and 3.8%. These numbers suggest that the income tax break did not only signal a commitment of the government to the development of the IT sector, but was also a sizable incentive.

Most directly, this article contributes to a long-standing academic debate on industrial policy (IP). To begin with, academics are skeptical that IPs can be effective, particularly in less developed countries, where public resources are often captured by "sunset industries" or lobbying firms. Moreover, even when governments identify the high-potential sectors and genuinely intend to support them, they may still lack the capacity to design and implement IPs that meet that goal. However, despite this skepticism, the set of papers that provide well-identified reduced-form evidence on the effectiveness of IP is rather small and recent.[7] We contribute to this debate by bringing evidence on the effectiveness of a Romanian IP with a unique design[8] and targeting a sector of general interest. The natural follow-up question is whether an effective IP is also efficient. While answering this question lies outside the scope of this paper, the tax break we study appears to meet the theoretical criteria for welfare-improving IP 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, Romania most likely had a latent comparative advantage in this activity [Rodrik, 1996, Harrison and Rodríguez-Clare, 2010] and only lacked a policy signal to tilt resources towards it.[9] 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-industry

[7] For such examples, see Görg et al. [2008], Aghion et al. [2015], Juhász [2018], Criscuolo et al. [2018], Cai and Harrison [2019]. Rodrik [2008] justifies the difficulty to conduct statistical inference on the effects of IPs and notes the scarcity of credible evidence, particularly from less developed countries. Until recently, the empirical evidence on IP came in two forms: (i) detailed country studies [Amsden, 1989, Wade, 1990, Evans, 1995], or (ii) cross-industry (country) econometric studies, which regress a measure of economic performance on indicators of government support [Krueger and Tuncer, 1982, Lee, 1996, Beason and Weinstein, 1996]. The main drawback of the former type of evidence is that causality is typically defended using a narrative approach, whereas the main shortcoming of the latter is one of misspecification [Rodrik, 2012].

[8] The most commonly studied IPs 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., 2018], or trade tariffs [Aghion et al., 2015]

[9] Although in practice, it is challenging to identify latent comparative advantage sectors, in theory, social planners can identify these sectors and design policies which tilt resource allocation towards them [Itskhoki and Moll, 2018]. This IP was meant to boost a sector that, while in its infancy, was perceived to have an unrealized potential. IT specialists and policy-makers of the time believed that Romania had the appropriate conditions to develop a strong IT sector, consistent with having a latent comparative advantage in this sector. Romania's rigorous education in STEM, inherited from the country's communist past, was a particularly valuable asset. However, short of a clear policy push to this sector, a large share of Romanian engineers were either emigrating or working in lower value-added branches of engineering. This insight is formally developed by Rodrik [1996]: Eastern Europe countries – while relatively well-endowed with skills and human capital – can be stuck in a low-income, low-tech equilibrium (even when a high-tech equilibrium is viable) due to coordination failures.

externalities.[10] Last, this policy benefits a sector that is competitive, and within 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 research establishes the wide-ranging effects of the IT sector, of which the effects on productivity have garnered the most attention.[11] 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) explains its slower productivity growth. This makes the policy under study especially relevant to countries that grapple with the drawbacks of an underdeveloped IT sector. Another strand of this literature studies the determinants of firm growth in the IT sector.[12] We add to this literature by estimating the plausibly-causal effects of tax incentives on the growth of IT firms. While we are not the first to point to tax incentives as options for countries to spur the growth of their IT sector,[13] we are unaware of other studies that bring econometric evidence on their effectiveness.

Finally, given the specifics of the policy, we also relate to research on policies aimed at reducing non-wage labor costs. Most papers in this literature evaluate 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) and find positive effects on firm size.[14] 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 (the pooled DiD estimate from the 2013 exercise is 15%) 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 or secretaries, by firms in all sectors). By studying a personal income tax break for programmers, 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

[10] 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. It is common for policy-makers to label upstream sectors that provide such inputs as "strategic" and grant them special policy treatment. The premise behind such special treatment is that, in its absence, "strategic" sectors would under-develop, at the detriment of their downstream sectors. This type of argument was frequently used in favor of protecting the semiconductor industry in the U.S., as discussed in research on this industry [Borrus et al., 1986, Irwin and Klenow, 1994, 1996]. 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, Rodríguez-Clare, 1996b] and recent work on the I-O implications of IP [Forslid and Midelfart, 2005, Du et al., 2014, Blonigen, 2016, Huremović and Vega-Redondo, 2016, Lane, 2017, Liu, 2018, Joya and Rougier, 2019].

[11] IT – or ICT, more broadly – has been found to improve productivity in IT-using firms [Wilson, 2009, Jorgenson et al., 2008, Syverson, 2011, Bloom et al., 2012], increase wages in high-skill locations [Forman et al., 2012], lead to the fragmentation of production [Fort, 2017], reduce information frictions [Steinwender, 2018], foster service exports [Kneller and Timmis, 2016], connect rural markets [Couture et al., 2018], and improve educational attainment [Beaudry et al., 2010] etc.

[12] The following factors have a strong explanatory power for firm growth in the IT sector: the founders' human capital [Colombo and Grilli, 2005, Ganotakis, 2012], firm internationalization [Ganesan and Samii, 2014, Falk and Hagsten, 2018], the initial size, age, and productivity of the firm [Johansson, 2004, Falk and Hagsten, 2018].

[13] See, for instance, Tigre and Botelho [2001], Biswas [2004], Tan and Leewongcharoen [2005], Chen et al. [2018a].

[14] For examples on positive effects on employment, see Crépon and Desplatz [2002], Rotger and Arendt [2010], Kangasharju [2007], Moczall [2013], Cahuc et al. [2014], Saez et al. [2017]; on positive effects on profits, revenue, and long-term investment, see Månsson and Quoreshi [2015], Saez et al. [2017]. For a review, see Eurofound [2017].

common finding is an increase in researcher wages.[15] Our findings of 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 (emigration which decreased the number of programmers available to work for firms in the country), and the fact that 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 back-of-the-envelope cost estimates of this tax break. Section 6 concludes.

2 Romania's Income Tax Break for Workers in IT

Before 2001, high labor taxes were seen as a major constraint to the development of the IT sector in Romania. These taxes lowered net wages for programmers and were seen as a root cause of the high emigration rates of programmers. These same taxes also led to relatively high labor costs for firms, limiting their growth.[16] For these reasons, Romania's IT sector did not stand out in Europe at the time.

Since 2001 however, 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 revenue. IT-related bachelor's degrees remain among the most popular degrees to this day. A tax break, effective since 2001, is widely perceived as having triggered these shifts.

In 2001, Romania introduced a personal income tax break for programmers at the proposal of Mr. Varujan Pambuccian, a member of the Chamber of Deputies (one of the two legislative bodies of the bicameral legislature of Romania). Mr. Pambuccian expected that this tax break would address both the concerns of workers and firms in the IT sector. He first proposed to reduce the top marginal income tax rate on wages of IT professionals from 40 to 8%.[17] Mr. Pambuccian was and still is an independent legislator, meaning that his proposal was not automatically backed up by a majority in the Chamber of Deputies.[18] In fact, his first proposal was deemed unrealistic and was rejected. Eventually, however, Mr. Pambuccian managed to rally support for his initiative. Surprisingly, in the final version of the law, the tax break became a full tax break (i.e., no income tax to be paid overall). Given the rejection of the first version of the law and the uncertainty around the approval of the final version, its adoption and its more generous nature came as unexpected and positive shocks to the IT sector.

To benefit from the tax break, workers had to fulfill all of the following criteria: (i) they had to have

[15] See Hægeland and Møen [2007], Dumont [2013], Lokshin and Mohnen [2013], Lelarge et al. [2015].

[16] This article (in Romanian) from January 2001 summarizes how the news on the introduction of this tax break was received by firms and workers in the IT sector. For instance: "All stakeholders in the IT industry are looking forward to [...] this measure. Why? Because of the generally high levels of migration of the workforce from company to company or from Romania abroad [due to low wages] and the substantial burden of labor taxation on firms."

[17] In 2001, personal income taxes were progressive, with rates between 18 and 40%. Payroll taxes (social security contributions and insurance) were paid both by employers (up to 40%) and employees (17%) on gross salary.

[18] Mr. Pambuccian is part of the Armenian minority in Romania, which he represents in the Chamber of Deputies.

an eligible bachelor's degree (in either automation, computers, informatics, cybernetics, mathematics, or electronics), (ii) they had to work for a firm whose sector code is "Software consultancy and supply" (7220 in NACE Rev 1), (iii) they had to work for the unit in charge of software development, (iv) they had to have an eligible occupation title (such as "programmer" or "computer systems designer"), and finally (v) they had to work for a firm that kept separate balance sheets recording revenues from software development and that generated a gross revenue 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 future audits.

Four features of this tax break deserve emphasis. First, this break was particularly generous at the time of its introduction. In 2001 the wages of programmers were among the highest in the country. Before the full income tax break, programmers faced a top marginal tax rate of 40%. In 2004, all workers in the country saw a generalized reduction in their income tax rate due to the switch from progressive income taxation to a flat rate of 16%. Despite this change, the trend of growing wages in the IT sector meant that this tax break remained a sizable incentive. Second, the rules for benefiting 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 this activity led to an increase in its 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.[19] Finally, this policy was designed to benefit all eligible workers and firms in the IT sector. Sectoral policy works better when benefits are less concentrated [Aghion et al., 2011].

The first major amendment to the tax break law occurred in 2013.[20] As for the initial introduction of the policy back in 2001, the passing of this amendment and its final eligibility criteria were also unexpected and uncertain. Negotiations were initiated in 2012 and were completed in 2013, years during which Romania had three different prime-ministers.[21] The amendment expanded the lists of eligible sector codes (for the firms) and of eligible majors for the bachelor's degree (for the workers). The new list of eligible sector codes consisted of the following NACE Rev 2 codes: 5821, 5829, 6201, 6202, and 6209. This turned previously ineligible activities in sectors such as "Database activities" or "Other computer related activities" into newly eligible activities.[22] The number of eligible bachelor's degree majors increased from six to 14 (adding majors such as "Cybernetics and economics"). Combined, these

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

[20] In 2004, the law had a minor revision without any economic effects.

[21] Mr. Pambuccian had initiated the negotiations toward the introduction of this amendment with Prime Minister Emil Boc in 2012, continued them with Prime Minister Mihai Ungureanu in 2012, and finally completed them with Prime Minister Victor Ponta in 2013. Hence, these disruptions made the success of these negotiations unpredictable.

[22] Romania had to officially transition from the NACE Rev 1 classification of sector codes to the NACE Rev 2 classification. Because the crosswalks between classifications are not bijective, and because policy-makers also intended to expand the scope of the tax break, some sectoral codes became newly eligible due to the transition from the NACE Rev 1 classification to the NACE Rev 2 classification. For instance, NACE Rev 1 code 726 (Other computer related activities) was not eligible pre-2013 but became eligible post-2013 as NACE Rev 2 code 6209 (which contains, in addition to NACE Rev 1 code 726, NACE Rev 1 code 7222, Other software consultancy and supply, a code eligible since 2001). See Table E3 in Appendix E.1.

changes significantly increased the number of eligible firms and workers.

We study the effects of the initial introduction in 2001 of the income tax break for workers in IT and of its 2013 reform. Two additional amendments were introduced in the second half of 2015 and 2016 respectively, which we do not study due to data availability constraints and their more limited scope. Appendix E.1 presents more details on both the 2001 tax break law and its subsequent amendments. Appendix E.2 summarizes other policies relevant to the Romanian IT sector and argues why these other policies do not jeopardize our ability to isolate the effects of the income tax break.

3 Firm-Level Analysis

Our analysis of the effects of the income tax break to workers in IT proceeds in two steps: a firm-level analysis and a sector-level cross-country analysis. In this section, we conduct a firm-level analysis centered around each of the two main events in the history of this policy: its initial introduction in 2001 and its only major reform in 2013. At the end of the section, we provide a discussion of the findings.

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 this analysis, we use Amadeus data on firms in Romania.[23] We construct a panel of firms between 1999 and 2005, with four outcome variables: operating revenues, number of workers, total assets, and solvency ratio. These variables are both likely to react to the new incentive and are among the few variables whose values are less frequently missing. As reporting to Bureau Van Dijk is not consistent across years – with many firms not complying with continuous annual reporting – we focus on firms who report these four outcomes at least in 2000, 2001, and 2002. In order to trim outliers, the sample is winsorized at the 1st and 99th percentiles of the distribution of operating revenues per worker.

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 1999 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, we use the following model:

$$y_{ist} = \alpha_i + \sum_{k=1999}^{2005} \delta_k \mathbb{1}[t=k] + \sum_{k=1999}^{2005} \beta_{DiD,k} \mathbb{1}[t=k] Target\_sector_{is} + \varepsilon_{ist}, \quad (1)$$

[23] Amadeus – a commercial dataset provided by Bureau Van Dijk – contains balance sheet information on firms in Europe. Amadeus data comes from official business registers, annual reports, newswires, and webpages. The Amadeus coverage of firms in Romania is good relative to the census of firms in Romania, particularly for larger firms.

where i stands for firm, s for the sector of firm i, t for the calendar year. α_i is the firm fixed effect.[24] $\mathbb{1}[t = k]$ is an indicator function that takes value 1 whenever an observation is in calendar year k. It is meant to capture common shocks across all firms in a given calendar year. We set 2000 as the reference year for the DiD coefficient estimates. We use as outcome variables, y_{ist}, the firm i, year t, log(operating revenue), log(number of workers), log(total assets), and the solvency ratio.

The treated firms are those whose NACE Rev 1 sector is 722 (Software consultancy and supply); hence, for firms whose sector s is 722, $Target\_sector_{is} = 1$. The comparison group for our baseline results has $Target\_sector_{is} = 0$ and includes firms from NACE Rev 1 sectors 721 (Hardware consultancy), 723 (Data processing), 724 (Database activities), 725 (Maintenance and repair of office, accounting and computing machinery), 726 (Other computer related activities), 731 (Research and experimental development on natural sciences and engineering), and 732 (Research and experimental development on social sciences and humanities). These comparison sectors share several features in common with 722, including their focus on high value-added services and their reliance on high-skilled workers and technology. We later show that our results are robust to six alternative comparison groups.

For this strategy to deliver credible estimates of the treatment effect of this policy, we first rely on its unexpected introduction. As discussed in Section 2, this policy was introduced at the initiative of one independent policy-maker alone. The success of his initiative and the ultimate generosity of the tax break came as a surprise to the IT industry. 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 fixed differences between firms and common year-specific shocks.[25] 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 impact on firms that are part of a sector with a newly available exemption for their workers involved in software development.

Baseline Results. Figure 1 plots the DiD estimates from the model in Equation (1). These estimates pertain to our baseline choice of the comparison group (see definition above). Reassuringly, across all outcome variables, we observe a lack of preexisting differential trends between treated and control firms, and between 1999 and 2000. After 2001, however, firms in the eligible sector experience significant improvements in all three measures of firm size (operating revenue, number of workers, and total assets) and in their financial health (solvency ratio). These improvements are mostly gradual, with part of the improvement already taking place in 2001. By 2005, firms in the IT sector have a 31% higher operating revenue, hire 17% more workers, have 17% more assets, and a 12-points higher solvency ratio than firms

[24] Given that in this dataset firms cannot be observed switching sectors (as we only observe their 2005 sector), adding or not adding sector fixed effects is inconsequential.

[25] Table C1 (Appendix C.1) shows that the median firm in the eligible sector is comparable to the median firm in the comparison sectors, while the average firm is significantly different (smaller or larger, depending on the variable). Given our use of firm fixed effects, time-invariant differences are not a threat to the identification of the effects of interest. We also control in some specifications for the initial size of the firm in order to account for potentially-heterogeneous effects based on initial firm size.

in comparable sectors (relative to 2000). The upper panel of Table 1 provides more details.[26]

Robustness Checks. This finding of strong and lasting boosts in firm size and financial performance is not driven by the choice of the baseline comparison group. We propose six other comparison groups, which are combinations of sectors akin to the IT sector. The new sectors in the pool of candidates (in addition to those forming the baseline set of comparison sectors) were selected on similar criteria as the baseline set: their focus on high value-added services and the type of inputs employed (workers in particular).[27] For instance, one of our proposed sets of comparison sectors is exactly the set of sectors whose firms became newly eligible for the income tax break of their programmers after the 2013 reform.[28] Policy-makers considered that these sectors were similar to the IT sector and were therefore deserving of the same tax treatment. Another set of comparison sectors excludes all IT-relevant sectors, in order to avoid concerns of spillovers from the eligible software development sector to other non-eligible sectors such as the hardware consultancy sector. Table A1 (Appendix A.1.1) shows that the main takeaways from our baseline estimates survive across these six other comparison groups. While the magnitude of the coefficients varies, the sign and significance of the DiD coefficients remains largely unchanged.

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 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 (*Agenţia Naţională de Administrare Fiscală*). The first dataset contains company balance sheets, which give us information on yearly revenue, sales, and total assets. We add firm-level information coming from two compulsory fiscal forms that record the income taxes paid by workers.[29] In particular, these forms track the firm-level number of

[26] The lower panel of Table 1 shows the estimate of a pooled DiD coefficient that measures the average increase in an outcome from the 1999-2000 period to the 2001-2005 period. These estimates are 31% for operating revenue, 11% for the number of workers, 22% for total assets, and nine points for the solvency ratio.

[27] In order to construct these six sets of comparison sectors we draw from the following list of NACE Rev 1 codes: 721 (Hardware consultancy), 722 (Software consultancy and supply), 723 (Data processing), 724 (Database activities), 725 (Maintenance and repair of office, accounting and computing machinery), 726 (Other computer related activities), 731 (Research and experimental development on natural sciences and engineering), 732 (Research and experimental development on social sciences and humanities), 741 (Legal, accounting, book-keeping and auditing activities, tax consultancy, market research and public opinion polling, business and management consultancy, and holdings), 742 (Architectural and engineering activities and related technical consultancy), 743 (Technical testing and analysis), 744 (Advertising), 748 (Miscellaneous business activities), and 3002 (Manufacture of computers and other information processing equipment).

[28] We refer to sectors 721 (Hardware consultancy), 724 (Database activities), 726 (Other computer related activities), and 3002 (Manufacture of computers and other information processing equipment). See Table E3 in Appendix E.1 for details. We do not prefer this set of sectors over the baseline set because it contains 3002: this is not only a manufacturing sector, but also one that provides inputs to the treated sector (therefore indirectly treated itself).

[29] Firms fill in and submit these 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).

workers exempted from paying any income tax. The resulting dataset starts in 2011 and ends in 2015.[30]

We remove from the analysis sample any firms with negative or missing values for the main variables of interest: revenue, sales, number of workers, and assets. Observations in the top and bottom 1% of the distribution of labor productivity (sales per worker) are also excluded. Moreover, we remove firms that benefited from major State Aid programs (797/2012 and 332/2014) during the period studied. We also remove from the analysis sample firms that we do not observe at least in 2012, 2013, and 2014.

Next, we keep in the analysis sample only firms who, in 2011 and 2012, had less than 5% of workers exempted from the income tax. In our baseline specification below, we will turn our attention to firms who experience a jump in their share of worker exemption from under 5 to over 20% after 2013. In order 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.

Appendix C.2 contains descriptive statistics on our main analysis samples. Table C2 compares the firm size and productivity in 2011 for the three groups of firms in our baseline sample: firms in non-eligible 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 non-eligible 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. The three types of firms do not differ in their average relative productivity. In our baseline specification, we control for time-invariant differences in size using firm fixed effects. In a variant of our baseline specification, we also control for initial differences in the size category, relative productivity, and age.

Table C3 shows that the percentage of firms in the eligible sectors with at least one worker exempted from the income tax has increased from 36% in 2011 to 41% 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 29% in 2011 to 35% in 2015.[31] Table C4 documents the predictors of firm-level share of workers who are exempt from the income tax. The reference category contains firms that operate in non-eligible sectors, are domestically-owned, have a micro size, and have been in business for more than five years. The reference year is 2011. Firms that are foreign-owned and/or larger are more likely to have a higher share of workforce exemptions from the income tax. This may reflect either their ability to meet the restrictive conditions of the exemption (for instance, by attracting eligible workers) or to assemble the documentation necessary to solicit the tax break.

Table C4 also shows that in 2013 and after, firm-level shares of workforce exemption experience a

[30] The baseline version of this analysis relies heavily on the forms recording the number of workers exempted from the income tax. As these forms were first introduced in 2011, this dataset 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 Appendix E.1 for details).

[31] This relatively low level of take-up of the tax exemption can be explained by the restrictive set of conditions that must be jointly met by firms and programmers in order for programmers to qualify for the tax break; by difficulties in hiring eligible programmers in a tight labor market; by a lack of knowledge of the administrative procedures to apply for the break; or by 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.

significant jump with respect to their 2011 level, whereas no such jump occurs in 2012. Table A2 further emphasizes these first-stage effects of the 2013 reform on the firm-level share of workforce exemption. The main takeaway is that in 2013 and after, firms in eligible sectors (according to the new definition of eligibility post-2013) experience an increase in their share of workers exempted from the income tax relative to their 2012 share. In 2015, the average firm in eligible sectors has a share of exempted workers that is 3% higher than in 2012. Also, there are 6% more firms in eligible sectors whose share of workforce exemption is larger than 20% in 2015 relative to 2011. Again, we find no evidence of trends 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.

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:

$$y_{ist} = \alpha_i + \boldsymbol{X'}_{ist}\beta_c + \lambda_{st} + \sum_{k=2011}^{2015} \delta_k \mathbb{1}[t=k] + \sum_{k=2011}^{2015} \beta_{DiD,k} \mathbb{1}[t=k] Exempted_{isk} + \varepsilon_{ist}, \quad (2)$$

where i stands for firm i, s for the sector of firm i, and t for the calendar year. α_i is the firm fixed effect (FE), and λ_{st} is the sector-by-year FE. $\mathbb{1}[t = k]$ is an indicator function that takes value 1 whenever an observation is in calendar year k. It is meant to capture common shocks across all firms in a given calendar year. We set 2012 as the reference year for the DiD coefficient estimates. We use as the outcome variables, y_{ist}, the firm i, year t, log(revenue), log(sales), log(number of workers), and log(total assets).

$Exempted_{isk}$ is a dummy variable that takes value 1 whenever in year k the firm i in sector s has more than 20% of its workforce exempted from the income tax. As mentioned above, we only keep firms with under 5% of workers exempted from the income tax in 2011 and 2012. Hence, by construction, all firms in the analysis sample have $Exempted_{isk} = 0$ for $k < 2013$. We choose 5% – instead of 0% – because there are other categories of workers who can benefit from this tax exemption, particularly those with disabilities.[32] We keep only firms with at most 5% workforce exemption before 2013 – as opposed to higher percentages – to mimic the tax conditions of IT firms before the initial introduction of the policy in 2001. That said, this threshold is ad-hoc and we run robustness checks that vary this threshold. Conversely, the 20% threshold after 2013 is chosen to avoid inadvertently measuring the effect of other income tax exemptions that are unrelated to the policy under study, 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 sectoral code became suddenly eligible in 2013 and firms (with an eligible NACE Rev 1.1 sectoral code before 2013) for whom a significantly larger share of workers became eligible in 2013.[33]

[32] For firms with more than 50 workers, at least 4% of their workers must be workers with a disability. In our data, we can only observe the number of employees exempted from the income tax and not the basis of the exemption. That said, when a firm has more than 5% of its workers exempted from the income tax, it is very likely that at least part of these exempted workers owe their exemption specifically to the income tax break for workers in IT. Indeed, in October 2017, when information on the justification for the tax exemption became available, 96% of the exempted employees in eligible sectors were exempted due to the tax break for programmers. See Appendix E.2 for details.

[33] Our data does not allow us to separate between these two possibilities for two reasons. First, we do not have worker-level information. Second, we only observe firms' sector codes in 2016, already translated to the NACE Rev 2 classification. Put differently, we do not observe the sector codes that firms had before Romania transitioned to the NACE Rev 2 classification. As shown in Table E3 (Appendix E.1), the NACE Rev 2 codes that became eligible in 2013 contain NACE Rev 1.1 codes that were both eligible and ineligible before 2013.

20% is a typical percentage of exemption in the eligible sector. This threshold is also ad-hoc and we run robustness checks around this threshold as well.

In addition to the sample restrictions introduced in the data section above, in the baseline exercise, we only keep firms in ICT service sectors.[34] These firms have similar economic activities to the IT sector and face similar technology and demand shocks. We compare the average change in outcomes for firms for which more than 20% of workers become exempted from the income tax for the first time after 2013 (after having less than 5% exempted workers pre-2013) to the average change in outcomes for firms in ICT service sectors with less than 5% of workers exempted throughout 2011 to 2015.

There are two important differences between this DiD specification 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 (NACE Rev 1 code 722), firms are now deemed treated if they are part of the eligible sectors (NACE Rev 2 codes 5821, 5829, 6201, 6202, and 6209) *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 IT sector, which 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 exempted employees is endogenous to firm characteristics. Lacking worker-level data and data on the NACE Rev 1 code of the sector 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 general 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 not eligible for the tax break. It is therefore very plausible that if a firm suddenly jumped from under 5 to over 20% of workforce exemptions 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 2001 to 2015 either had a sector code that did not become eligible for the tax break or hired programmers whose bachelor's degree majors remained ineligible after 2013.

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 policy shock, which may be the actual driver of both the jump in the firm-level workforce exemption share and the estimated firm growth effects. A change in management is a shock that may potentially

[34] We use the OECD definition of ICT service sectors, which 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). Of these ICT sectors, only sectors 582 and 62 benefited from an income tax exemption for employees working on software development.

explain both behaviors. That said, the jumps we study occurred after 2013, implying that they were most likely driven by the policy shock and not by large-scale coordinated changes in management. Last, the lack of preexisting differential trends with respect to firms in the comparison group also suggests that firms that experienced a sudden increase in the share of exempted workers were unlikely to be undergoing notable productive or organizational changes.

Following the literature on firm growth, in part of our specifications we also control for lagged size, age, and relative productivity (found in the firm-specific vector of time-varying characteristics, $\boldsymbol{X}_{ist}$) [Doms et al., 1995, Lotti et al., 2003, Coad, 2009, Barba Navaretti et al., 2014]. Categories for firm size follow the Eurostat definition: micro (1-9 workers), small (10-49 workers), medium (50-249 workers), and large (250 or more workers). Young is an age dummy variable that takes the value 1 if the firm is five years or younger [as in Lotti et al., 2003, Falk and Hagsten, 2018]. Relative productivity is defined as a firm's labor productivity relative to the most productive firm in the same sector [as in Falk and Hagsten, 2018].

Last, we will show that our baseline estimates are robust to (i) running the regression in Equation (2) on a dataset from Amadeus that starts in 2008 and allows us to observe a longer pre-reform period, (ii) using an alternative comparison group of firms in other high-tech knowledge-intensive service sectors, (iii) restricting the sample only to firms in the eligible IT sector, (iv) using an alternative event-study research design, (v) varying the threshold choice (e.g., from 20 to 15%), and (vi) defining treatment only based on the sector of the firm, as in the 2001 DiD specification.

Baseline Results. Figure 2 plots the DiD estimates from the model in Equation (2). These estimates pertain to our baseline choice of the comparison group (see definition above). We find that across all outcome variables the estimate of $\beta_{DiD,2011}$ is not statistically different from zero. This lack of anticipation effects is in line with the unexpected nature of the expansion. Moreover, firms whose workforce became significantly exempted from the income tax after 2013 did not embark on differential growth trends relative to firms whose exemption rate was left unaffected by the reform. After 2013 however, the treated firms experience a gradual growth along all four measures of firm size, such that in 2015, they have a 37% higher revenue, 38% higher sales, 21% more workers, and 26% more assets relative to 2012.[35]

Table 2 provides more details on these estimations. Columns (1)-(4) show the results of the estimations without firm controls (other than firm FEs) and plotted in Figure 2, whereas columns (5)-(8) show the results from adding firm-specific controls (initial size, relative productivity, and age). As expected, controlling for year $t-1$ productivity, size, and age decreases the magnitude of the coefficients, but they all remain statistically and economically significant. As in Doms et al. [1995], Falk and Hagsten [2018], firms that were initially larger or more productive grew faster. Age is found to have an insignificant effect on firm growth.

Heterogeneity Analysis. The DiD estimates so far refer to the average effect on treated firms. However,

[35]The lower panel of Table 2 shows the estimate of a pooled DiD coefficient that measures the average increase in an outcome from the 2011-2012 period to the 2013-2015 period. These estimates are 25% for revenue, 27% for sales, 15% for number of workers, and 18% for total assets. Controlling for other firm characteristics reduces the estimated effects to 13% for revenue, 16% for sales, 11% for number of workers, and 8% for total assets.

an important policy concern is that the tax break expansion benefited only specific groups 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) hiring strictly less than ten workers (micro firms), (ii) hiring 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). Because 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 3 reports the results of this heterogeneity analysis. The main takeaway is that the tax exemption had a positive effect on all four types of firms that are considered. The effects are as large for micro-sized firms (typically start-ups) as they are for larger firms. Younger firms experience higher growth than older firms. This last finding 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% exempted employees in 2011 and, hence, were not included in this analysis. To conclude, the 2013 expansion of the tax break has been beneficial for firms of all sizes and for young firms in particular.

Robustness Checks. All robustness checks for our 2013 baseline results are in Appendix A.1.2. In our first exercise, we run our baseline regression from Equation (2) on an Amadeus dataset that starts in 2008. The firms in the Amadeus sample are those firms from the baseline sample based on administrative data (5,177 firms) that were found in Amadeus (3,889 of these 5,177 firms). Table A4 presents the results from this Amadeus sample using data from two sources: the Amadeus data and the administrative data. The similarity between the estimates in columns (5)-(8) and those in columns (1)-(4) from Table 2 suggest that the sample matched with Amadeus is representative for the baseline sample. The similarity between the estimates in columns (1)-(4) and those in columns (5)-(8) suggest that the Amadeus data is of comparable quality to the administrative data. The main advantage of the results from columns (1)-(4) (relative to those in columns (1)-(4) from Table 2) is that they allow us to observe firms for up to five years before the reform (relative to 2 years). Treated and comparison firms continue to exhibit parallel trends before 2013, which is reassuring in regard to the suitability of the comparison group.

Next, we show that our baseline estimates are robust to reasonable alternative comparison groups. In Table A5, the sample of analysis uses firms in high-tech knowledge-intensive service sectors (as classified by Eurostat). In addition to other similar sectors, this category includes the eligible sectors. These new results corroborate with the baseline results. To the extent that one might be concerned that results are driven by the inclusion of firms in non-eligible sectors, Table A6 repeats the analysis using only firms in eligible sectors. Despite losing half of the baseline sample, results are strikingly similar.

Another concern is that our baseline results are driven by the contrast to firms who always remain under the 5% threshold of worker exemption. To overcome this concern, we use an event-study design where we exploit the staggered timing of the moment when firms jump over the 20% threshold. Even in the restricted sample, which only keeps the firms experiencing the jump, the event-study estimates from Columns (5)-(8) in Table A7 continue to display a lack of pre-trends and a clear pattern of growth after the expansion in the tax exemption. These findings suggest that our baseline results are not an artifact of the choice of the comparison group. Appendix A.1.2 provides more details on this event-study analysis.

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 A8 shows the results of estimations where treatment arises when firms jump after 2013 to at least 15% of workers exempted from income tax. These results are only slightly smaller in magnitude to those obtained for the 20% threshold, which is consistent with the fact that we are including firms with a lower intensity of treatment in the sample. Results from other values of thresholds are available upon request.

In a last robustness check, we revert to the definition of treatment based on a firm's sector, that we use 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. Because this definition does not depend on a firm's share of exempted employees, we estimate this model on two samples: the full sample of firms in ICT service sectors, unrestricted based on firms' exemption shares (columns (1)-(4) in Table A9), and the baseline sample, for direct comparability (columns (5)-(8)). Before 2013, we continue to find no evidence of differential trends between firms in eligible sectors and firms in non-eligible ICT service sectors. After 2013 however, we find evidence of growth for firms in eligible sectors. Although statistically significant, these effects are smaller than our baseline effects in Table 2. This is intuitive: in 2015, the average firm in eligible sectors has a share of exempted workers that is 3% higher than that in 2012 (see Table A2). This jump is to be compared with the jump from under 5% pre-2013 to over 20% post-2013 that characterizes the treated firms in our baseline specification. The smaller effect is therefore in line with the smaller jump. Notwithstanding, the results from this exercise also lend support to the suitability of the baseline DiD design used for the 2013 analysis.

3.3 Discussion of the Firm-Level Findings

Overall, we find that both policy episodes – the introduction of the tax break in 2001 and its 2013 reform – have led to strong and long-lasting growth for firms in the eligible IT sectors. A natural first question is whether these firm-level effects are consistent with the statutory incidence of the tax break on the personal income of workers. We provide three arguments in favor of an affirmative answer.

First, the tax break only applies to the income of workers with an eligible bachelor's degree specialization who are matched with firms operating in an eligible sector. To benefit from the tax break, these workers are then supposed to develop software for these firms that generates at least 10,000 U.S. dollars in business revenue per year and per exempted worker. These restrictive requirements imposed on worker-firm matches are likely to turn the qualifying firms into desirable employers and to improve their bargaining power over a shared economic incidence of the tax break.

Second, the tax break law stipulates that it is the responsibility of the firm to prepare the necessary paperwork and apply for the income tax break of its workers. This requirement of an explicit "buy-in" from the employer is another reason to expect a shared economic incidence of this tax incentive. To the extent that access to financing is scarce in Romania (as in other emerging economies) and firms are cash-constrained, firms can use their savings in labor costs from the tax break towards expansion.

Third and last, both policy episodes have led to improvements in the labor productivity of treated firms (measured as revenues or sales per worker). For instance, Table A3 (Appendix A.1.2) shows that in 2015, firms treated by the 2013 reform became 16% more productive than comparison firms. 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 level of motivation and efficiency. Second, the increase in the number of workers in the IT sector (see Figure 4) is likely to have led to sector-level economies of scale.[36] Finally, it is also likely that both policy episodes have attracted higher-ability workers to software development.

While all these factors suggest a shared economic incidence of the tax break, the estimation of how the tax incentive is split between firms and workers is outside the scope of this paper. Most important, we lack worker-level data, meaning that any estimate of this split would lean too heavily on the structure of a model. Moreover, the split is likely to change with time, depending on factors such as the entry of new firms into the IT sector and the long-run elasticity of the supply of programmers. That said, we conclude that firm-level effects are in line with the incentive structure of this policy.

Before comparing the magnitude of the estimates derived from the two firm-level exercises, it is important to first highlight the distinguishing features of each exercise. When this policy was introduced in 2001, the IT sector of Romania was in its infancy. Before 2001, Romanian programmers were emigrating at high rates, lacking confidence that the Romanian IT sector was poised for growth. According to media articles from 2001, both programmers and incumbent IT firms perceived the policy as a signal 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 to this day, in the 2001 exercise, we only focus on the 1999 to 2005 period. This aims to isolate the effects of this tax break from either those of other policies that may interact with it (such as the switch in 2004 from progressive income taxation to a flat income tax) or other global shocks that may differentially affect the IT sector.

Moreover, in this 2001 exercise, we use Amadeus data that does not contain the actual exemption rate of the workers of a firm but only the sector of the firm. Hence, our 2001 exercise estimates the effects of the introduction of the policy on firms whose sector is eligible (NACE Rev 1 722), which are "intent-to-treat" (ITT) effects. Table 1 indicates that the typical firm in the eligible sector hired 17% more workers in 2005 relative to 2000 and relative to the typical firm in non-eligible comparable sectors. Under the assumption of 25% take-up among firms in the eligible sector, we can convert our ITT estimates to "treatment-on-the-treated" (TOT) estimates. This back-of-the-envelope calculation suggests that firms whose workers actually benefited from the income tax break grew their workforce by 68%. Given that in 2000, the median firm in the eligible sector hired four workers, such an increase in the number of workers is not unreasonable. 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 key advantage of the exercise studying the 2013 reform is that it estimates the effects on firms who actually benefit from the tax break expansion by increasing their share of workforce exemption from under 5% pre-2013 to over 20% post-2013. The benefit of this definition is that it allows us to isolate firms experiencing a sizable (relative) decrease in labor costs. That said, this definition 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 (most likely either due to their sector code not being eligible or their not being able to hire eligible workers, both according to pre-2013 eligibility rules). This

[36] While measuring external economies of scale for manufacturing sectors alone, Bartelme et al. [2018] find large sector-level economies of scale, particularly in high-technology sectors such as "Computers and Electronics."

suggests that this sample of firms is likely to be negatively selected relative to the sample of firms with higher exemption rates (e.g., be younger and less experienced, or smaller). This is likely to explain the smaller 21% estimate for the increase in workforce after the 2013 reform (see Table 2), relative to the 68% back-of-the-envelope TOT estimate for 2001. Nevertheless, the purpose of the 2013 reform was exactly that of improving 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 2004, 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.

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

The firm-level analysis in the previous section focuses on the behavior of incumbent firms in the IT sector in Romania, either at the time of the introduction of the 2001 income tax break or of its expansion in 2013. We now present two sector-level cross-country analyses: one that estimates the overall growth after 2001 of the eligible IT sector relative to the rest of the economy, the other that estimates the growth after 2001 of sectors using IT services intensively relative to those using IT services less.

4.1 Direct Effects on the Expansion of the IT Sector

This sector-level study of the overall 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 (through the entry of new firms into the IT sector). While the available data does not allow us to conduct a complete analysis of the patterns of entry (exit) into (out of) the IT sector, Appendix B presents suggestive evidence that entry rates into the IT sector have increased after 2001 (both the entry of entirely new firms and the entry of firms previously in other sectors), whereas the exit rates have not.

Second, this synthetic control method (SCM) analysis allows us to alleviate concerns about potential confounding factors that may affect the IT sector globally 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.[37] Sector-specific global technological or demand shocks could be another confounding factor for the growth of the Romanian IT sector. 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 third factors.

Empirical Strategy. We use SCM to measure the effect of the income tax break on the growth of the

[37] Given that the typical firm in the IT sector is high skill-intensive and low capital-intensive, researchers have argued that IT sectors in CEE had suitable conditions for development [Radosevic, 2006, Grigoraş et al., 2017]: CEE countries tend to have a highly-skilled labor force in relevant technical and scientific fields, while lacking modern physical capital. This explains why it is important to benchmark the growth of the Romanian IT sector to that in other CEE countries.

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. The intuition of SCM is that a weighted combination of countries provides a better comparison for Romania than any single country alone. SCM makes explicit the relative contribution of each control country to the counterfactual of interest, and the similarities (or lack thereof) 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 control 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 on 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 available control countries ("the donor pool"), where J equals 13 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 (more specifically, 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. We include in $\boldsymbol{X_1}$ and $\boldsymbol{X_0}$ the three predictor variables just mentioned.

Let V be a diagonal matrix with non-negative components, whose diagonal elements represent the relative importance of these three predictor variables 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 $w_j \geq 0$. We follow Abadie and Gardeazabal [2003] and Abadie et al. [2010] and allow for the choice of the weighting matrix, $\boldsymbol{V}$, to be data-driven. $\boldsymbol{V}$ is the matrix that allows for the pre-2001 outcome of Romania to be closest to the outcome for the synthetic control obtained from $\boldsymbol{W}^*(\boldsymbol{V})$.[38]

Treatment effects, α_{1t}, for Romania in post-treatment years t (2002 to 2015) are estimated as the difference between the year t outcomes for Romania and those for synthetic Romania:

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

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

[38] 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.

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 the 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. Finally, we want these variables to have at least a minimum relevance for this study. These conditions are met by the following five variables: number of employees, gross revenue (turnover or gross premiums written, in million euro), production value (in million euro), gross investment in tangible goods (in million euro), and the number of enterprises. Data for the predictors (pre-treatment characteristics used to generate the weights) 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 industry, and the share of GDP coming from services.

For data availability reasons, the IT sector is defined broadly as K72 (Computer and related activities).[39] We use as comparison sectors all other sectors in the economy. To obtain normalized values for the outcome variables 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.[40]

The donor pool of countries from which to create the synthetic control for Romania contains Bulgaria, the Czech Republic, Estonia, Hungary, Ireland, Latvia, Lithuania, Poland, Portugal, Slovenia, and the Slovak Republic. These countries were chosen based on their geographic proximity, similarity in development, performance in the IT sector pre-2001, and data availability. See 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 revenue and employment (both normalized). The upper left panels show the evolution of these outcomes in Romania and in synthetic Romania. Before 2001, the growth of the IT sector of synthetic Romania closely mimics that of Romania. From 2001 onward, both the gross revenue and employment of the Romanian IT sector experience a marked relative growth. The upper right panels show the difference between the outcomes of Romania and those of synthetic Romania. Fourteen years after, the gross revenue (employment) in the IT industry in Romania had expanded 6.52 (1.83) times more relative to the gross revenue (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 donor countries. We notice how exceptional the growth was in Romania, compared to that in all donor countries. Last, we implement a battery of placebo tests that considers all other donor

[39] K72 includes K721 (Hardware consultancy), K722 (Software consultancy and supply, including K7221 – Publishing of software – and K722 –Other software consultancy and supply), K723 (Data processing), K724 (Database activities), K725 (Maintenance and repair of office, accounting and computing machinery), and K726 (Other computer related activities). We use the two-digit aggregation of the data due to frequent missing values at the three-digit level.

[40] For example, a value of 2 in year t for a given country means that the multiplication factor of the value of the dependent variable in the treated sector in year t, relative to year 2000, is larger by 2 units than the counterfactual multiplication factor in comparison sectors.

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. For at least until 2008, the actual treatment differences of gross revenue 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 A10, Appendix A.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 A11 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. All robustness checks figures and tables can be found in Appendix A.2.1. We first show that these findings are not unique to gross revenue and employment. Figure A1 (and its associated Tables A10 and A11) shows similar patterns of outstanding growth in the IT sector of Romania, this time in terms of production value.[41] 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 A1 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. While we envisaged several options, we find that results are not driven by the crosswalk choice.[42]

We test whether our results are driven by the use of the full time series, by truncating the time series in 2007. This way we can study the effect of the 2001 policy using data that is consistently reported within one classification. Fortunately, we find that results for pre-2007 years are not affected by the addition of post-2007 years (see Figure A2). 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 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.

Last, an alternative exercise 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 (see Sections 3.1 and 3.2). 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

[41] While the underlying data is of lower quality, we obtain qualitatively similar results for the number of firms and gross investment in tangible goods (both normalized). Results are available upon request.

[42] See Appendix D for details on our crosswalk construction and choice. SCM results using different crosswalk choices are available upon request.

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, across sectors and households. While the development of the IT sector most likely led to level effects – given its large base of users – 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.[43] First, after the tax break, labor productivity (measured as sales per worker) increased (see Section 3.3). 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 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 Appendix B). Figure B12 (Appendix B.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.[44] 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 main difference between these 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 the supplier of inputs. 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

[43] One scenario in which the introduction of 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 hired by a firm in the eligible IT sector, this may have made them less likely to join ineligible sectors. We believe 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.

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

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 income tax break for workers in IT, inter-industry linkages are not yet affected by this break. 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.[45]

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.[46] See 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 revenue and employment. Fourteen years after the introduction of the policy, gross revenue (employment) in sectors with high-intensity use of IT services has grown 0.75 (0.61) times more than gross revenue (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 A13 (Appendix A.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 (see Table A12 in Appendix A.2.2). Figure 7 depicts a striking relative growth in the export trade value of goods from high-intensity sectors in Romania (relative to those from low-intensity sectors and relative to synthetic Romania).[47] Placebo tests show that this relative growth in Romania is exceptional compared to that predicted for all other

[45] 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 a full overlap in the sets of sectors that belong to the top 25% of the two distributions. Hence, there is no practical difference in the final split of sectors between the top and bottom three quartiles.

[46] We cannot study service exports, as the EBOPS data on service exports only starts in 2000. Studying the exports of the treated IT sector (72 NACE Rev 1 code) is not possible either, as the output of 72 is mostly in the form of services. Moreover, Romania only starts reporting exports for 72 in 2005.

[47] 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), 86 (Watches and clocks), and a combination of Chemical Industry codes such as 59, 55, 53.

countries in the donor pool.[48] Table A15 (Appendix A.2.2) makes the same argument, formally. This evidence 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.[49]

Robustness Checks All robustness checks figures and tables can be found in Appendix A.2.2. We first show that our baseline findings on the relative growth of downstream sectors with heavier usage of IT services are not specific to gross revenue, 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 A3).[50]

Second, it is important to note that the SCM findings presented above are robust to different choices of the weighting matrix $\boldsymbol{V}$ and matrix of weights $\boldsymbol{W}$ matrices (defined in Section 4.1).[51] Another potential concern with SCM relates to its sensitivity to the number of pre-intervention periods used in the computation of the weights. This concern cannot be ruled out 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 consistently 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 sectors from the treated category. 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 by 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 two sectors belong to the treated category of sectors (see Table A12), but also contain three-digit NACE Rev 1 sectors that we use as comparison sectors in the firm-level analysis of Section 3.1. To the extent that 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 A4 shows that when we exclude sectors 73 and 74 the treatment effect is actually larger than the one found in the baseline Figure 6.[52]

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 A5 presents results from the grouping of sectors *under the median* of usage of inputs from the IT sector into the low-intensity

[48] Figure A7 (Appendix A.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.

[49] The behavior of FDI flows to Romania also supports this claim. Figure B12 (Appendix B.3) shows that FDI in high-intensity 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, Renault decided in 2007 to build its Technocentre in Romania, which took over a large share of the R&D previously done in France. In various articles motivating this decision, Romania's IT sector seems to have played an important role. As an example, this article describes how Romania's accession to the EU benefits France, e.g., through its IT skills and their importance to firms such as Renault.

[50] We find similar, yet noisier, patterns for gross investment and the number of firms. Results are available upon request.

[51] We implement these different choices by selecting different options of the **synth** and **synth_runner** packages in Stata. Our main results are obtained using the *nested* option, which maximizes the fit during pre-intervention periods.

[52] We decide to keep 73 and 74 in the main specification for two reasons. First, results obtained when including these two two-digit sectors are more conservative. Second, both sectors contain many other three-digit sectors that are not comparable to 722, and, hence, are less likely to experience correlated shocks with 722. The main SCM exercise using Comtrade data excludes, by construction, all service sectors, hence sectors 72, 73, and 74 are not a concern.

category (as opposed to *under the third quartile*) and sectors *over the median* into the high-intensity category. As expected, while the difference in the development of these two new 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 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.[53] 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].[54]

Finally, as for the SCM exercise in Section 4.1, we verify whether our results are sensitive 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 industry classification and NACE Rev 1. Figure A6 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 apply to the Comtrade SCM exercise as the data is reported within the same classification for all years.

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.[55]

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 that benefit the IT sector in Romania (and in all other similar countries). This (relative) growth of the IT sector has occurred both on the intensive (through the growth of incumbent firms, see Section 3) and on the extensive margin (through the entry of new firms in the sector, both domestic and foreign). In Section 4.2 we find that since 2001, IT-using sectors in Romania have grown significantly faster compared to

[53] Results are available upon request.

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

[55] 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 industry to focus on developing technology for the now-mature automotive industry (e.g., driver-less technology or car connectivity).

non-IT using sectors in Romania and compared to 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.

As to the magnitude of these relative growth rates, several pieces of evidence lend them credibility. 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. For example, 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.[56] 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 in similar economies. Second, as in the case of the firm-level results, we find stronger increases in revenue in the IT sector than in employment. This points to a consistency between our sector-level and firm-level evidence. 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 receiving the tax incentive directly grows faster than downstream sectors benefiting from the incentive indirectly, through cheaper or more diverse inputs.[57]

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.[58] The choice of the reference group can obviously 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: 1999 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

[56] The gross revenue grew 14 times (from 282 million euros in 2000 to 4,031 million euros in 2015), the number of firms grew four times (from 3,174 firms in 2000 to 12,432 firms in 2015).

[57] Figure B12 (Appendix B.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.

[58] 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.

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.

In a similar vein, the exact 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 (such as IT) 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 IP and the FDI attraction and trade opportunities that followed Romania's joining of the EU in 2007.[59] While Romania's comparison countries in the SCM 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 evidence 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 Estimates of the Policy

One important policy concern is the cost of this policy to the government, given that the IT workers who benefit from this income tax exemption are among the highest-paid workers in Romania. Another concern might be that our findings of large and persistent firm- and sector-level growth after the introduction and extension of the income tax break, may be disproportionate compared to the incentive offered by this tax break. To alleviate these concerns, we estimate the cost of this policy to the government, or put differently, the implicit incentive to the IT sector. Under the assumption that the administrative costs of this policy are negligible (both to firms and to the government), we define this cost as the amount foregone in tax revenues due to the tax exemption. We estimate this amount for 2015, the last year of our firm-level analysis. In 2015, these foregone tax receipts refer to the flat personal income tax of 16% owed on the taxable income of all exempted employees.

Assumptions on the wages and employment of the exempted employees are necessary. If the tax exemption were to be unexpectedly removed at the beginning of 2015, we assume that exempted employees would retain full-time employment. Then, in our first scenario we assume that in the short term, gross wages would remain unchanged. Therefore, our first estimate is based on the actual average wages of the exempted employees. In the absence of the policy, the currently exempted employees may also move to other ICT services sectors or other high-tech knowledge-intensive sectors. Our second and third estimates are based on the average wages in those two types of sectors, respectively. Table 4 shows our estimates for these three scenarios.

[59] Topalova and Khandelwal [2011] also provides evidence "that there may be strong complementarity among different industrial policies." In their case, "the impact of trade reforms appears to be magnified as FDI was allowed or restrictive licensing requirements were removed."

Overall, the estimates show that the tax receipts foregone in 2015 due to the policy vary between 62 and 80 million euros. To put these numbers into perspective, they represent between 4.7 and 6.4% of the total gross wage bill in the IT sector, which suggests that the income tax break is likely to represent an important incentive to the sector. Note that these shares are computed out from the total wage bill of the sector (which includes the wage bill of firms with little to no exemption of their employees). For firms with a large share of employees exempt from the income tax, the share between the "forgiven" tax bill and their total wage bill is significantly larger. In terms of the value added of the sector, the foregone tax revenues are in the range of 2.8 and 3.8% of the value added. These estimates suggest that the income tax exemption does not only signal a commitment of the Romanian government to the development of the IT sector but is also a sizable incentive. This incentive is also likely to act as an indirect incentive to other sectors, particularly those sectors that rely heavily on IT services.

6 Conclusion

This paper examines the effects of a unique industrial policy (IP) that was 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 law of the tax break was amended to allow for a significantly larger list of eligible sector codes for firms and eligible bachelor's degrees for workers. We exploit both policy episodes to bring plausibly-causal estimates of the firm and sector-level effects of this policy. Across various empirical strategies and measures of firm size, we bring evidence of a strong and lasting growth of IT firms in Romania. 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 factors unrelated to this IP, as similar neighboring countries fail to show comparable growth. Moreover, in Romania, sectors relying relatively more on IT inputs also grew faster.

Our results suggest that this policy has been effective in its goal to support the development of the IT sector and to reallocate resources (such as high-skill labor) towards this sector. This is an important achievement, because many IPs only allow specific groups to extract rents, without actually affecting resource allocations. Moreover, this shift was made towards a high-skill/high-wage sector, a policy priority in both developed and developing countries.

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 industry" (as has been the case in several East Asian IPs). On the contrary, it targeted an industry 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-industry) 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.

Alice H. Amsden. Asia's Next Giant: South Korea and Late Industrialization. Oxford University Press, 1989.

Giorgio Barba Navaretti, Davide Castellani, and Fabio Pieri. Age and Firm Growth: Evidence from Three European Countries. Small Business Economics, 43(4):823–837, 2014.

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

Richard Beason and David E Weinstein. Growth, Economies of Scale, and Targeting in Japan (1955-1990). Review of Economics and Statistics, 78(2):286–295, 1996.

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.

Radha Roy Biswas. Making a Technopolis in Hyderabad, India: The Role of Government IT Policy. Technological Forecasting and Social Change, 71(8):823–835, 2004.

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 chinas r&d subsidies – evidence from listed firms. Research Policy, 45(9):1774–1789, 2016.

Michael Borrus, Laura D'Andrea Tyson, and John Zysman. Creating Advantage: How Government Policies Shape International Trade in the Semiconductor Industry. Strategic Trade Policy and the New International Economics, pages 99–114, 1986.

Pierre Cahuc, Stéphane Carcillo, and Thomas Le Barbanchon. Do Hiring Credits Work in Recessions? Evidence from France. IZA Discussion Papers 8330, 2014.

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.

Jin Chen, Cheng Suang Heng, Bernard CY Tan, and Zhijie Lin. The Distinct Signaling Effects of R&D Subsidy and Non-R&D Subsidy on IPO Performance of IT Entrepreneurial Firms in China. Research Policy, 47(1): 108–120, 2018a.

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 2018b.

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.

Alex Coad. The Growth of Firms: A Survey of Theories and Empirical Evidence. In The Growth of Firms: A Survey of Theories and Empirical Evidence. Edward Elgar Pub, 2009.

Massimo G. Colombo and Luca Grilli. Founders' Human Capital and the Growth of New Technology-based Firms: A Competence-based View. Research Policy, 34(6):795–816, 2005.

Victor Couture, Benjamin Faber, Yizhen Gu, and Lizhi Liu. E-Commerce Integration and Economic Development: Evidence from China. NBER Working Paper No. 24384, 2018.

Bruno Crépon and Rozenn Desplatz. Evaluation of the Effects of Payroll Tax Subsidies for Low Wage Workers. CREST Working Paper, 2002.

Chiara Criscuolo, Ralf Martin, Henry Overman, and John Van Reenen. Some Causal Effects of an Industrial Policy. American Economic Review, 2018.

Mark Doms, Timothy Dunne, and Mark J. Roberts. The Role of Technology Use in the Survival and Growth of Manufacturing Plants. International Journal of Industrial Organization, 13(4):523 – 542, 1995.

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.

Michel Dumont. The Impact of Subsidies and Fiscal Incentives on Corporate R&D Expenditures in Belgium (2001-2009). Reflets et Perspectives de la Vie Economique, LII:69–91, 2013.

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

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

Peter Evans. Embedded Autonomy: States and Industrial Transformation. Princeton University Press, 1995.

Martin Falk and Eva Hagsten. An Exploration of Growth in Computer Software Micro Firms. Journal of Small Business and Entrepreneurship, pages 1–20, 2018.

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.

Teresa C. Fort. Technology and Production Fragmentation: Domestic versus Foreign Sourcing. Review of Economic Studies, 84(2):650–687, 2017.

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

Vedavinayagam Ganesan and Massood Samii. Internationalization and Longevity of Small- and Medium-Sized Software Firms. Journal of Small Business and Entrepreneurship, 27:537–555, 2014.

Panagiotis Ganotakis. Founders' Human Capital and the Performance of UK New Technology Based Firms. Small Business Economics, 39:495–515, 2012.

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.

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.

Torbjørn Hægeland and Jarle Møen. Input Additionality in the Norwegian R&D Tax Credit Scheme. Reports 2007/47, 2007.

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.

Douglas A Irwin and Peter J Klenow. Learning-by-Doing Spillovers in the Semiconductor Industry. Journal of Political Economy, 102(6):1200–1227, 1994.

Douglas A Irwin and Peter J Klenow. High-tech R&D Subsidies: Estimating the Effects of Sematech. Journal of International Economics, 40(3-4):323–344, 1996.

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

Dan Johansson. Is Small Beautiful? The Case of the Swedish IT Industry. Entrepreneurship & Regional Development, 16:271–287, 2004.

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.

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.

Anne O Krueger and Baran Tuncer. An Empirical Test of the Infant Industry Argument. American Economic Review, 72(5):1142–1152, 1982.

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

Jong-Wha Lee. Government Interventions and Productivity Growth. Journal of Economic Growth, 1(3):391–414, 1996.

Claire Lelarge, Erwan Gautier, and François-Charles Wolff. Alleviating the Burden of Entrepreneurial Risk? Evidence from Two French Targeted Programs. Working Paper, 2015.

Ernest Liu. Industrial Policies in Production Networks. Working Paper, 2018.

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.

Francesca Lotti, Enrico Santarelli, and Marco Vivarelli. Does Gibrat's Law Hold Among Young, Small Firms? Journal of Evolutionary Economics, 13(3):213–235, 2003.

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.

Andreas Moczall. Subsidies for Substitutes? New Evidence on Deadweight Loss and Substitution Effects of a Wage Subsidy for Hard-to-place Job-seekers. IAB Discussion Paper 5/2013, Institute for Employment Research, Nuremberg, Germany, 2013.

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

Madalina Pruna and John Soleanicov. Rethinking Romanian Industrial Policy. Technical report, MPA/ID Thesis, John F. Kennedy School of Government, Harvard University, 2012.

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.

Slavo Radosevic. Growth, Integration and Spillovers in the Central and East European Software Industry. Economics Working Papers 69, 2006.

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

Andrés Rodríguez-Clare. The Division of Labor and Economic Development. Journal of Development Economics, 49(1):3 – 32, 1996b.

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.

Dani Rodrik. Why We Learn Nothing From Regressing Economic Growth on Policies. Seoul Journal of Economics, 25(2), 2012.

Gabriel Pons Rotger and Jacob Nielsen Arendt. The Effect of a Wage Subsidy on Subsidized Firm's Ordinary Employment. AKF Working Paper No. 2010, Danish Institute of Governmental Research, 2010.

Emmanuel Saez, Benjamin Schoefer, and David Seim. Payroll Taxes, Firm Behavior, and Rent Sharing: Evidence from a Young Workers' Tax Cut in Sweden. NBER Working Paper, 2017.

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.

Felix B Tan and Kallaya Leewongcharoen. Factors Contributing to IT Industry Success in Developing Countries: The Case of Thailand. Information Technology for Development, 11(2):161–194, 2005.

Paulo Bastos Tigre and Antonio JosÉjunqueira Botelho. Brazil Meets the Global Challenge: IT Policy in a Postliberalization Environment. Information Society, 17(2):91–103, 2001.

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.

Robert Wade. Governing the Market: Economic Theory and the Role of Government in East Asia's Industrialization. Princeton University Press, 1990.

Daniel J. Wilson. IT and Beyond: The Contribution of Heterogeneous Capital to Productivity. Journal of Business & Economic Statistics, 27:52–70, 2009.

Figures

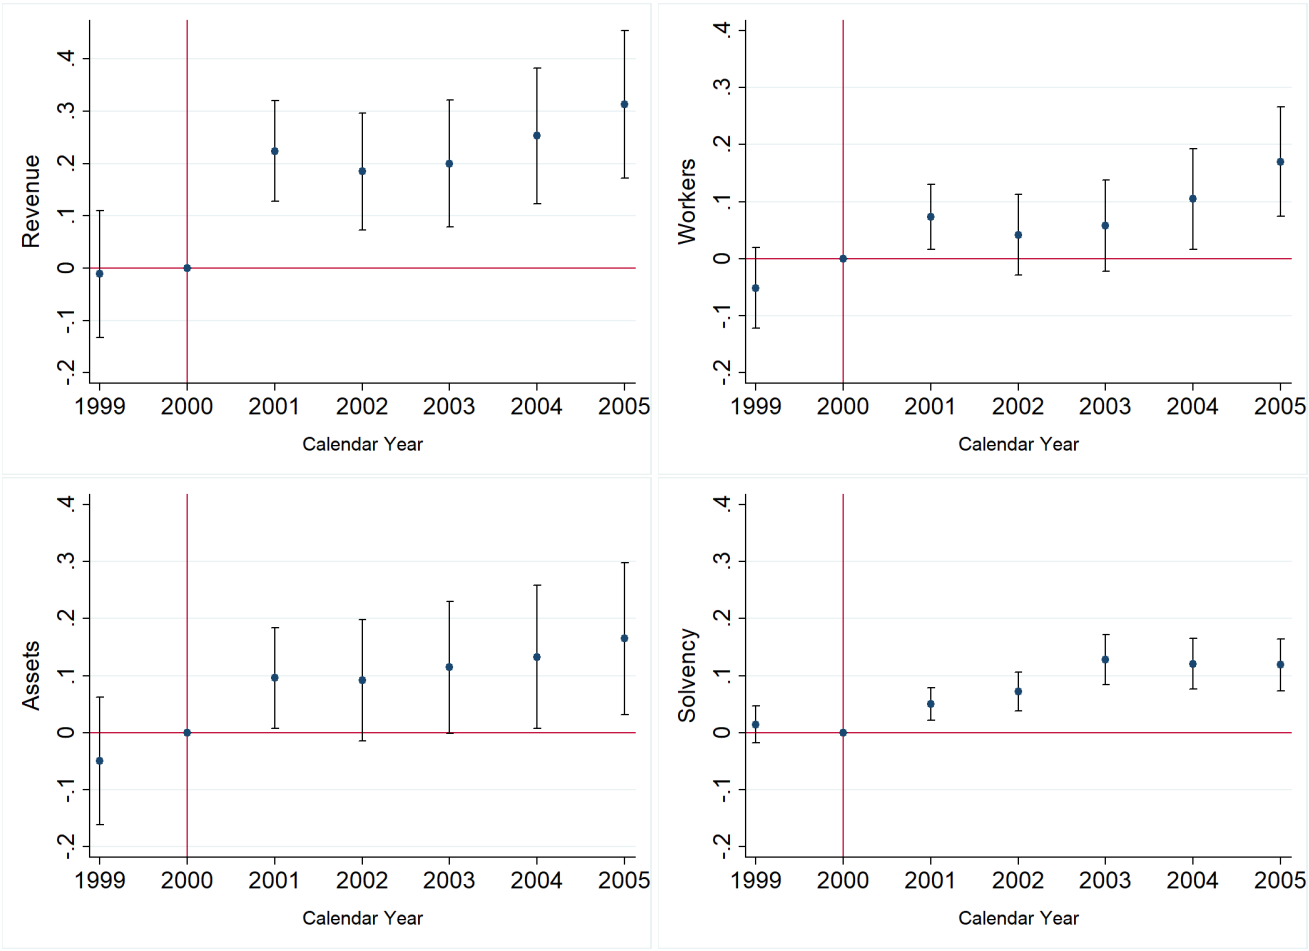

Figure 1: DiD Estimates of the Effects of the 2001 Income Tax Break

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 four firm-level outcome variables: log(operating revenue), log(number of workers), log(total assets), and the solvency ratio. 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 join the baseline comparison group if their NACE Rev 1 sector code is either 721 (Hardware consultancy), 723 (Data processing), 724 (Database activities), 725 (Maintenance and repair of office, accounting and computing machinery), 726 (Other computer related activities), 731 (Research and experimental development on natural sciences and engineering), or 732 (Research and experimental development on social sciences and humanities). The data source is Amadeus, a commercial database from Bureau Van Dijk. The regression model includes firm and calendar year fixed effects. See Table 1 for more details.

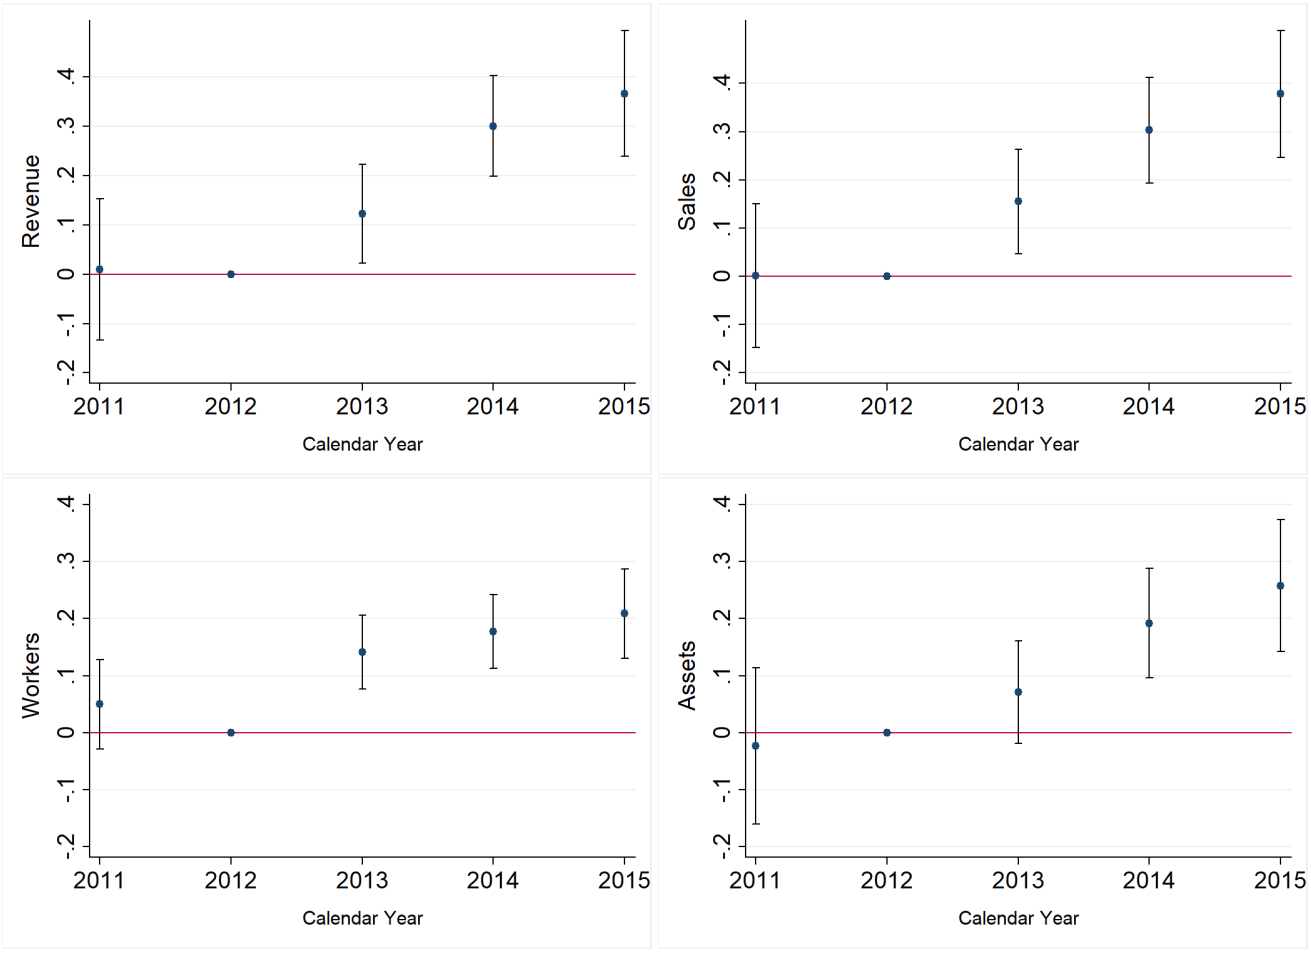

Figure 2: DiD Estimates of the Effects of the 2013 Reform

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. We consider four firm-level outcome variables: log(revenue), log(sales), 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 exempted workers jumps from under 5% to over 20% after 2013. This exercise builds on administrative tax data collected by the Ministry of Finance. The regression model includes firm and sector-by-year fixed effects. See columns (1)-(4) from Table 2 for more details.

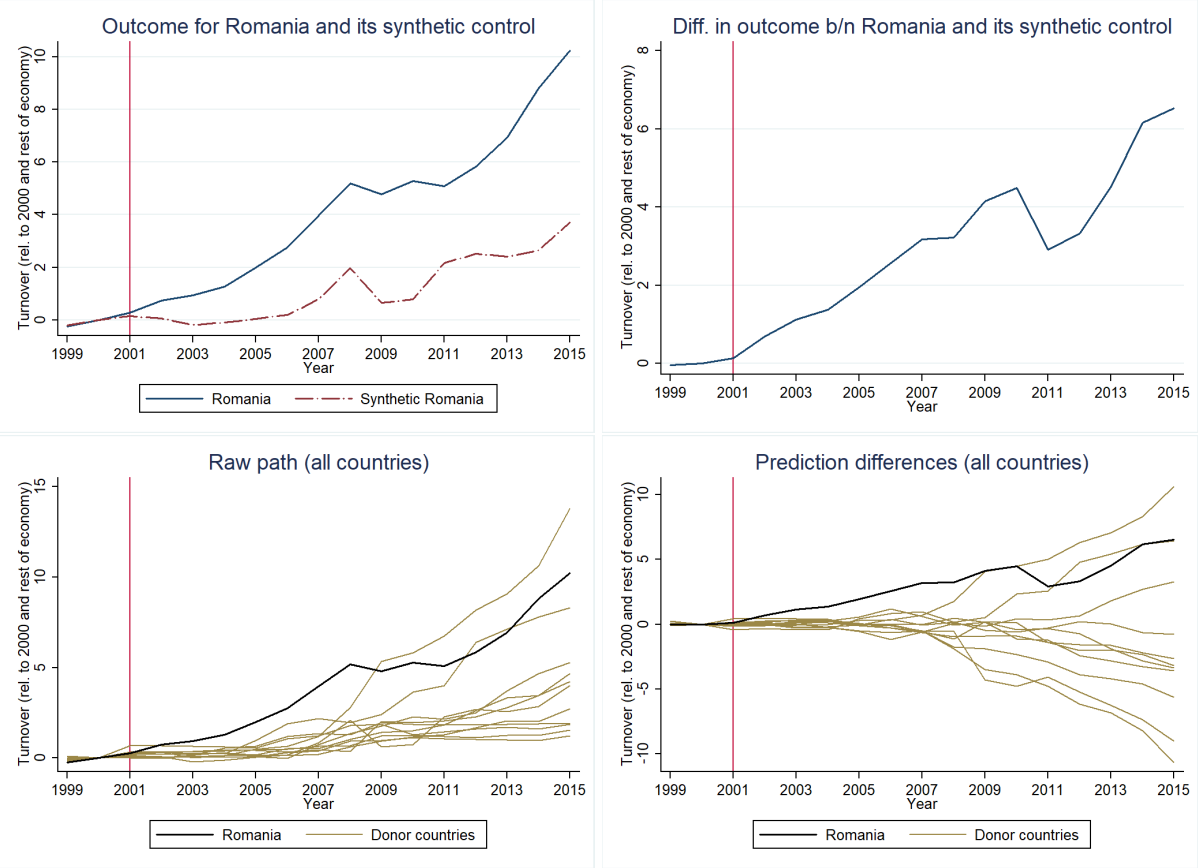

Figure 3: IT Sector Vs. Rest of the Economy. SCM with Outcome Variable: Gross Revenue ("Turnover or Gross Premiums Written") - Million Euro (Normalized)

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 revenue ("Turnover or gross premiums written - million euro"). 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 industry (% 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.

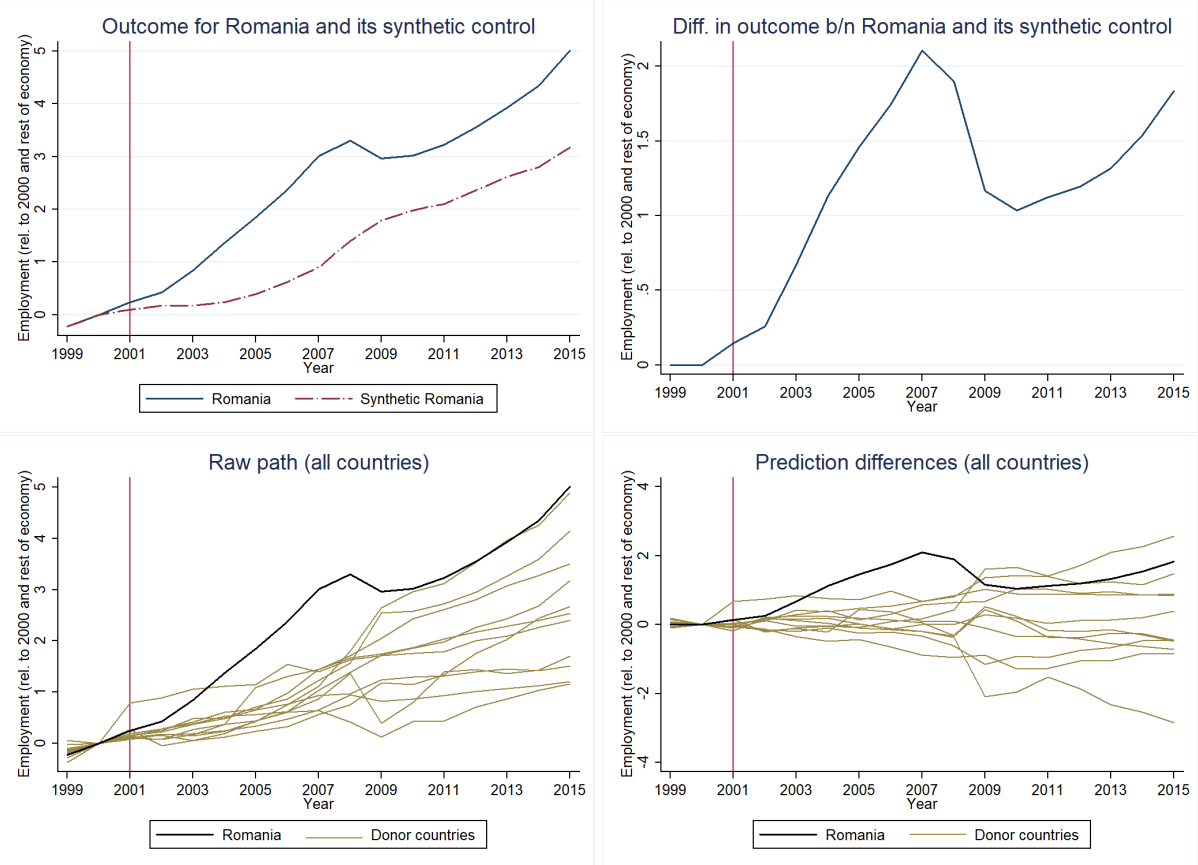

Figure 4: IT Sector Vs. Rest of the Economy. SCM with Outcome Variable: "Employees - Number" (Normalized)

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 industry (% 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.

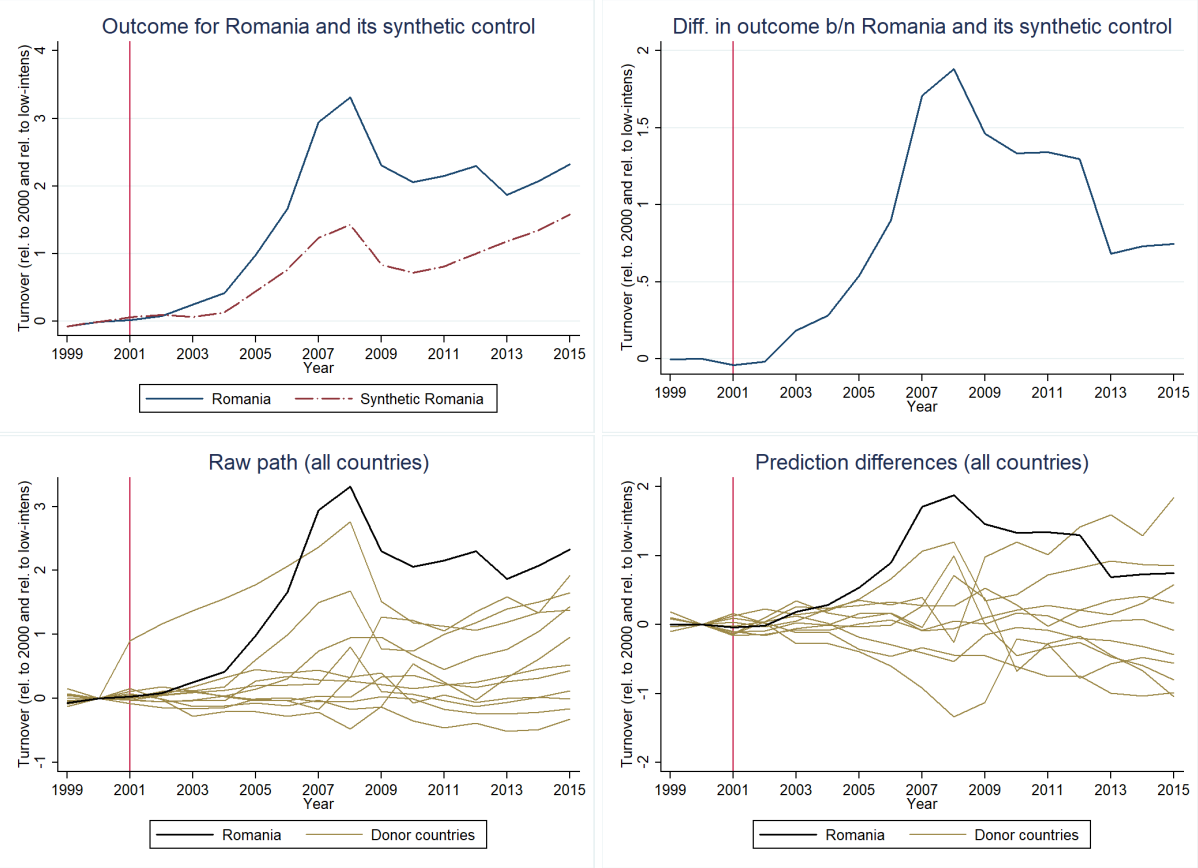

Figure 5: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: Gross Revenue ("Turnover or Gross Premiums Written") - Million Euro (Normalized)

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 revenue ("Turnover or gross premiums written - million euro"). 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 industry (% 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.

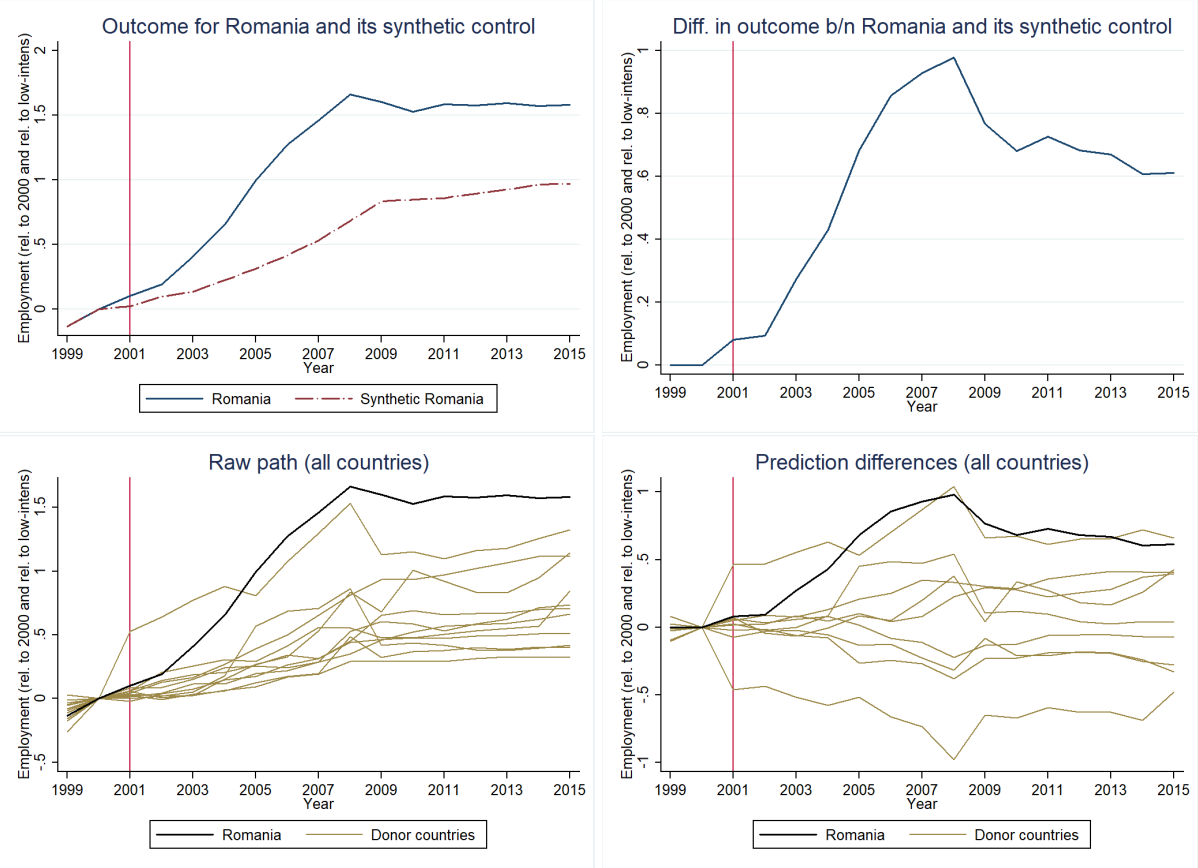

Figure 6: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: "Employees - Number" (Normalized)

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 industry (% 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.

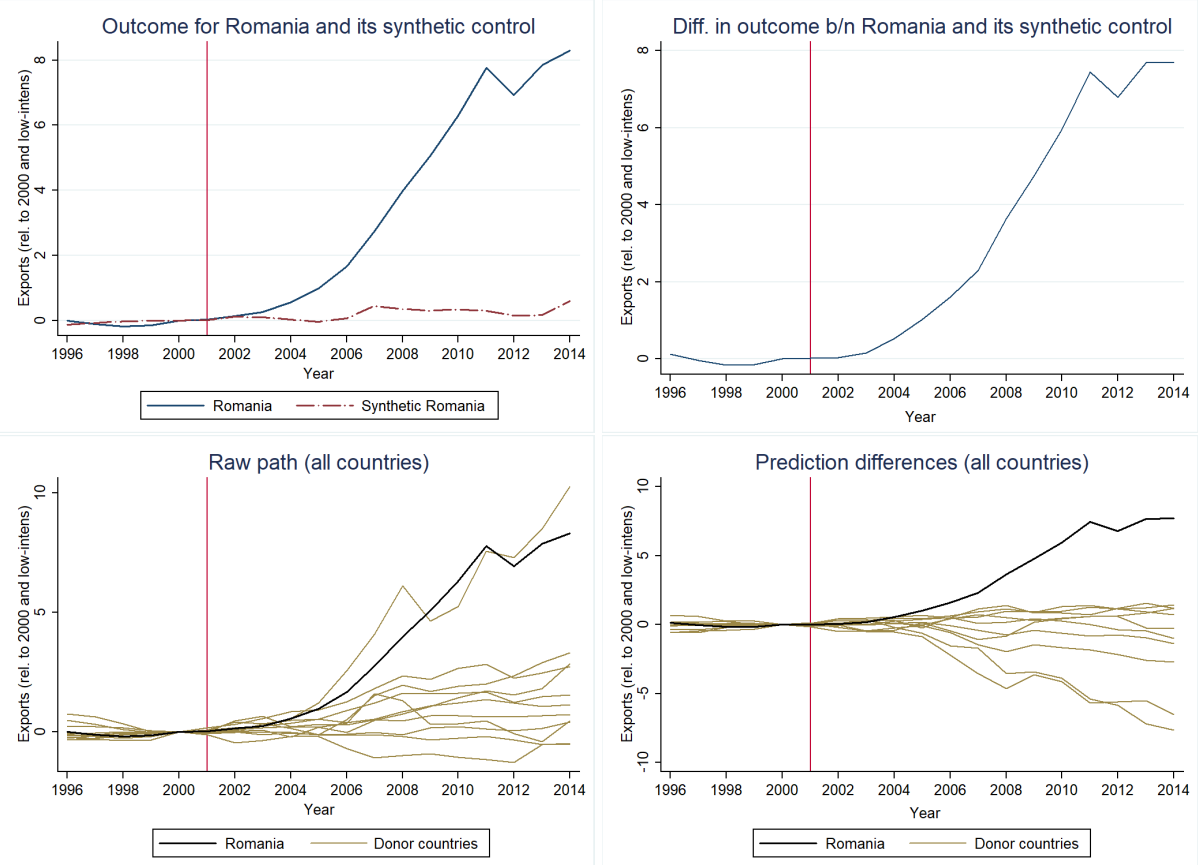

Figure 7: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: "Goods Export Value" (Normalized)

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 industry (% 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: Baseline Results

	Revenue (1)	Workers (2)	Assets (3)	Solvency (4)
Yearly Regression				
$\beta_{DiD,1999}$	-0.011	-0.051	-0.049	0.014
	(0.062)	(0.036)	(0.057)	(0.017)
$\beta_{DiD,2001}$	0.224***	0.073**	0.096**	0.050***
	(0.049)	(0.029)	(0.045)	(0.015)
$\beta_{DiD,2002}$	0.185***	0.042	0.092*	0.072***
	(0.057)	(0.036)	(0.054)	(0.018)
$\beta_{DiD,2003}$	0.200***	0.058	0.115*	0.128***
	(0.062)	(0.041)	(0.059)	(0.022)
$\beta_{DiD,2004}$	0.253***	0.105**	0.133**	0.121***
	(0.066)	(0.045)	(0.064)	(0.023)
$\beta_{DiD,2005}$	0.313***	0.170***	0.165**	0.119***
	(0.072)	(0.049)	(0.068)	(0.023)
R^2	0.868	0.878	0.882	0.582
Pooled Regression				
β_{DiD}	0.312***	0.106***	0.221***	0.093***
	(0.054)	(0.034)	(0.052)	(0.016)
R^2	0.792	0.877	0.796	0.561
# Observations	10,534	10,576	10,401	10,101
# Firms	1,622	1,622	1,614	1,604

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. The data source is Amadeus, a commercial database from Bureau Van Dijk. In this table we report the baseline point estimates of the difference-in-difference coefficients of interest from the Equation (1) (upper panel) and from the pooled version of Equation (1) (lower panel), i.e., $\beta_{DiD,t}$ and β_{DiD} respectively. The outcome variables used are log(operating revenue), log(number of workers), log(total assets), and the solvency ratio. 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. Firms in the baseline comparison group are those whose NACE Rev 1 sector code is either 721 (Hardware consultancy), 723 (Data processing), 724 (Database activities), 725 (Maintenance and repair of office, accounting and computing machinery), 726 (Other computer related activities), 731 (Research and experimental development on natural sciences and engineering), or 732 (Research and experimental development on social sciences and humanities). All specifications include firm and calendar year fixed effects. The number of observations and firms is the same for the yearly and pooled versions of the same regression. 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: Baseline Results

	Revenue (1)	Sales (2)	Workers (3)	Assets (4)	Revenue (5)	Sales (6)	Workers (7)	Assets (8)
Yearly Regression								
$\beta_{DiD,2011}$	0.010	0.001	0.050	-0.023	0.000	0.000	0.000	0.000
	(0.073)	(0.076)	(0.040)	(0.070)	(.)	(.)	(.)	(.)
$\beta_{DiD,2013}$	0.122**	0.155***	0.141***	0.071	0.021	0.063	0.084**	0.016
	(0.051)	(0.055)	(0.033)	(0.046)	(0.052)	(0.057)	(0.037)	(0.050)
$\beta_{DiD,2014}$	0.300***	0.303***	0.177***	0.192***	0.164***	0.183***	0.107***	0.093*
	(0.052)	(0.056)	(0.033)	(0.049)	(0.050)	(0.055)	(0.036)	(0.051)
$\beta_{DiD,2015}$	0.366***	0.378***	0.209***	0.258***	0.221***	0.231***	0.127***	0.133**
	(0.065)	(0.067)	(0.040)	(0.059)	(0.058)	(0.062)	(0.041)	(0.057)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Rel. prod.					1.593***	1.278***	0.218**	1.765***
					(0.225)	(0.197)	(0.085)	(0.132)
Small					0.346***	0.344***	0.373***	0.223***
					(0.048)	(0.047)	(0.037)	(0.038)
Medium					0.631***	0.662***	0.752***	0.634***
					(0.118)	(0.106)	(0.116)	(0.117)
Large					0.969**	0.993**	1.089**	1.034**
					(0.394)	(0.386)	(0.466)	(0.406)
Young					-0.037	-0.037	-0.001	-0.041**
					(0.023)	(0.023)	(0.014)	(0.020)
Adjusted R^2	0.869	0.874	0.907	0.901	0.906	0.909	0.924	0.931
Pooled Regression								
β_{DiD}	0.253***	0.273***	0.153***	0.179***	0.131***	0.155***	0.105***	0.081*
	(0.044)	(0.046)	(0.025)	(0.040)	(0.046)	(0.051)	(0.034)	(0.047)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Adjusted R^2	0.868	0.874	0.907	0.900	0.906	0.909	0.924	0.930
# Observations	22,592	22,212	22,598	22,587	16,820	16,536	16,824	16,813
# Firms	5,177	5,128	5,177	5,177	4,864	4,804	4,865	4,863

Notes: In this exercise we study the firm-level impact of the 2013 expansion to the income tax break law of 2001. Here, we report the baseline point estimates of the difference-in-difference coefficients of interest from the Equation (2) (upper panel) and from the pooled version of Equation (2) (lower panel), i.e., $\beta_{DiD,t}$ and β_{DiD} respectively. The outcome variables used are log(revenue), log(sales), log(number of workers), and log(total assets). The year prior to the amendment of 2013 (2012) is the reference year. The only difference between columns (1)-(4) and columns (5)-(8) is that the latter include firm-specific time-variant controls in addition to firm fixed effects, whereas the former include firm fixed effects alone. All specifications include firm and sector-by-year fixed effects. Heteroskedasticity robust errors in parentheses. The number of observations and firms is the same for the yearly and pooled versions of the same regression. The reference category for the firm size is "Micro" and for "Young" firms the reference category are firms "Older than five years." ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 3: Difference-in-Differences Design Around 2013 Reform: Heterogeneity of the Baseline Results

	Revenue (1)	Sales (2)	Workers (3)	Assets (4)	Revenue (5)	Sales (6)	Workers (7)	Assets (8)
Panel A: Size		*Micro: < 10 workers*				*Small, medium, large: ≥ 10 workers*		
β_{DiD}	0.248***	0.276***	0.161***	0.173***	0.283***	0.264***	0.144***	0.235***
	(0.051)	(0.053)	(0.028)	(0.046)	(0.078)	(0.080)	(0.055)	(0.083)
Adjusted R^2	0.790	0.795	0.790	0.853	0.944	0.942	0.945	0.960
# Observations	19,664	19,310	19,670	19,659	2,928	2,902	2,928	2,928
# Firms	4,451	4,410	4,451	4,451	726	718	726	726
Panel B: Age		*Young: < 5 years old*				*Old: ≥ 5 years old*		
β_{DiD}	0.341***	0.389***	0.229***	0.198***	0.119**	0.113**	0.055*	0.104**
	(0.072)	(0.074)	(0.038)	(0.064)	(0.047)	(0.050)	(0.031)	(0.045)
Adjusted R^2	0.807	0.824	0.862	0.843	0.911	0.910	0.934	0.939
# Observations	10,148	9,972	10,152	10,144	12,444	12,240	12,446	12,443
# Firms	2,410	2,391	2,410	2,410	2,767	2,737	2,767	2,767

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(revenue), log(sales), 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 5,177 firms (see Table 2). 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. Heteroskedasticity robust errors in parentheses. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4: Back-of-the-Envelope Calculation of Foregone Tax Revenues in 2015

Scenario	(1)	(2)	(3)
Equivalent Subsidy (Mil. RON)	355.55	262.50	273.62
Equivalent Subsidy (Mil. Euro)	79.98	59.05	61.55
% Value Added	3.84%	2.84%	2.96%
% Revenue	1.92%	1.64%	1.48%
% Production	2.02%	1.82%	1.55%
% Wage Bill	6.39%	4.72%	4.92%

Notes: Data sources: Administrative data from the National Agency of Fiscal Administration of Romania (*Agenţia Naţională de Administrare Fiscală*) and Structural Business Statistics from Eurostat. In all three scenarios the employees exempted in 2015 are assumed to maintain full employment were the exemption to be removed at the beginning of that year. Scenario (1) for the back-of-the-envelope calculation of foregone tax revenues from the income tax break is one in which the wages of exempted employees remain unchanged. Scenario (2) is one in which exempted employees are paid the average wage in non-eligible ICT services sectors. Scenario (3) is one in which exempted employees are paid the average wage in non-eligible high-tech knowledge-intensive sectors.

Appendix A Additional Evidence

Appendix A.1 Firm-Level Evidence

Appendix A.1.1 2001 Income Tax Break: Amadeus Data

Table A1: Robustness of the Baseline Results to the Choice of the Comparison Group

	Revenue *C1*	Workers *C1*	Assets *C1*	Solvency *C1*	Revenue *C2*	Workers *C2*	Assets *C2*	Solvency *C2*
β_{DiD}	0.304***	0.040	0.202***	0.105***	0.255***	0.062*	0.145***	0.073***
	(0.069)	(0.043)	(0.064)	(0.019)	(0.056)	(0.035)	(0.053)	(0.016)
R^2	0.804	0.856	0.800	0.555	0.777	0.845	0.769	0.555
# Observations	6,947	6,977	6,866	6,659	9,480	9,503	9,350	9,090
	Revenue *C3*	Workers *C3*	Assets *C3*	Solvency *C3*	Revenue *C4*	Workers *C4*	Assets *C4*	Solvency *C4*
β_{DiD}	0.569***	0.319***	0.563***	0.180***	0.131***	0.065**	0.063	0.050***
	(0.091)	(0.057)	(0.094)	(0.030)	(0.046)	(0.029)	(0.043)	(0.013)
R^2	0.800	0.890	0.811	0.550	0.759	0.846	0.763	0.519
# Observations	5,365	5,394	5,326	5,169	36,457	36,649	35,778	34,486
	Revenue *C5*	Workers *C5*	Assets *C5*	Solvency *C5*	Revenue *C6*	Workers *C6*	Assets *C6*	Solvency *C6*
β_{DiD}	0.145***	0.073**	0.080*	0.054***	0.149***	0.062**	0.075*	0.053***
	(0.046)	(0.029)	(0.043)	(0.013)	(0.045)	(0.029)	(0.042)	(0.013)
R^2	0.766	0.857	0.771	0.521	0.761	0.844	0.762	0.524
# Observations	37,493	37,691	36,815	35,482	41,619	41,829	40,850	39,415

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. The data source is Amadeus, a commercial database from Bureau Van Dijk. The outcome variables used are log(operating revenue), log(number of workers), log(total assets), and the solvency ratio. Treated firms are those in the NACE Rev 1 sector 722. The difference between the six sets of results comes from the set of firms allowed in the comparison group. Comparison firm i joins a sample based on its NACE Rev 1 sector. For the regressions labelled *C1*, the comparison group contains firms whose NACE Rev 1 sector is either 3002, 721, 724, or 726. For the regressions labelled *C2*, the comparison group contains firms whose NACE Rev 1 sector is either 721, 723, 724, 725, or 726. For the regressions labelled *C3*, the comparison group contains firms whose NACE Rev 1 sector is either 731 or 732. For the regressions labelled *C4*, the comparison group contains firms whose NACE Rev 1 sector is either 741, 742, 743, 744, or 748. For the regressions labelled *C5*, the comparison group contains firms whose NACE Rev 1 sector is either 731, 732, 741, 742, 743, 744, or 748. For the regressions labelled *C6*, the comparison group contains firms whose NACE Rev 1 sector is either 721, 723, 724, 725, 726, 741, 742, 743, 744, or 748. The NACE Rev 1 sector codes stand for the following sectors (in parenthesis): 721 (Hardware consultancy), 722 (Software consultancy and supply), 723 (Data processing), 724 (Database activities), 725 (Maintenance and repair of office, accounting and computing machinery), 726 (Other computer related activities), 731 (Research and experimental development on natural sciences and engineering), 732 (Research and experimental development on social sciences and humanities), 741 (Legal, accounting, book-keeping and auditing activities; tax consultancy; market research and public opinion polling; business and management consultancy; holdings), 742 (Architectural and engineering activities and related technical consultancy), 743 (Technical testing and analysis), 744 (Advertising), and 748 (Miscellaneous business activities). All specifications include firm fixed effects. Robust standard errors, clustered at the firm-level, in parenthesis. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Appendix A.1.2 2013 Reform to the Income Tax Break

Table A2: 2013 Reform: First Stage Effects on the Share of Income Tax Exempted Employees

Dependent variable	1 if more than 20% exempted employees			Share of exempted employees		
Sectors of Firms	ICT services	HKTI services	Eligible	ICT services	HKTI services	Eligible
	(1)	(2)	(3)	(4)	(5)	(6)
Eligible sector $\times$ d_{2011}	0.005	0.005		0.004	0.004	
	(0.007)	(0.007)		(0.004)	(0.004)	
Eligible sector $\times$ d_{2013}	0.020***	0.020***		0.011***	0.011***	
	(0.006)	(0.006)		(0.004)	(0.004)	
Eligible sector $\times$ d_{2014}	0.035***	0.035***		0.023***	0.023***	
	(0.006)	(0.006)		(0.004)	(0.004)	
Eligible sector $\times$ d_{2015}	0.055***	0.056***		0.026***	0.026***	
	(0.007)	(0.007)		(0.004)	(0.004)	
d_{2011}			0.004			0.004
			(0.007)			(0.004)
d_{2013}			0.020***			0.011***
			(0.006)			(0.004)
d_{2014}			0.034***			0.023***
			(0.006)			(0.004)
d_{2015}			0.055***			0.026***
			(0.007)			(0.004)
Adjusted R^2	0.727	0.729	0.682	0.772	0.774	0.740
# Observations	26,507	27,407	16,207	26,507	27,407	16,207
# Firms	5,721	5,932	3,499	5,721	5,932	3,499

Notes: In this exercise, we focus on the 2013 tax reform and study its first stage effects on the firm-level share of workforce exemption from the income tax. Columns (1)-(3) 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 (4)-(6) use as the dependent variable the share of workers of firm i who are exempted from the income tax in year t. The samples used in this table are different from the baseline sample in Table 2 and the samples in Tables A5 and A6 in that the former keep *all* firms in ICT service sectors, in HTKI (high-tech knowledge-intensive) service sectors, and in eligible sectors, whereas the latter bring additional restrictions on the firm-level share of workforce exemption before and after 2013. The dependent variables and samples used in this table are the same as those used in Table C4, which explores the predictors of firm-level workforce exemption. In addition to interaction terms between a dummy taking value 1 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), (4), (5) 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 (3) and (6) also include firm fixed effects. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table A3: Difference-in-Differences Design Around 2013 Reform: Sales per Worker

	Sales per Worker	Sales per Worker
$\beta_{DiD,2011}$	-0.056	0.000
	(0.068)	(.)
$\beta_{DiD,2013}$	0.001	-0.026
	(0.053)	(0.058)
$\beta_{DiD,2014}$	0.142***	0.090
	(0.052)	(0.055)
$\beta_{DiD,2015}$	0.161***	0.119*
	(0.062)	(0.063)
Controls	No	Yes
Rel. prod.		1.199***
		(0.228)
Small		-0.012
		(0.043)
Medium		-0.070
		(0.089)
Large		-0.070
		(0.133)
Young		-0.024
		(0.026)
Adjusted R^2	0.732	0.781
# Observations	22,356	16,652
# Firms	5,142	4,830

Notes: In this exercise, we focus on the 2013 tax reform and study the evolution of labor productivity (defined as sales per worker) for the baseline sample defined in Section 3.2. We report point estimates of the DiD coefficient from Equation (2). The only difference between Column (1) and Column (2) is that the latter includes firm-specific time-variant controls in addition to firm fixed effects, whereas the former includes firm fixed effects alone. Heteroskedasticity robust errors in parentheses. All specifications include firm and sector-by-year fixed effects. The reference category for the firm size is "Micro" and for "Young" firms the reference category are firms "Older than five years." ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table A4: Robustness of the Baseline Results to Adding More Pre-Reform Years Using Amadeus Data

	Amadeus Data				**Administrative Data**			
	Revenue (1)	Sales (2)	Workers (3)	Assets (4)	Revenue (5)	Sales (6)	Workers (7)	Assets (8)
$\beta_{DiD,2008}$	-0.070	-0.077	0.027	-0.020				
	(0.087)	(0.086)	(0.059)	(0.086)				
$\beta_{DiD,2009}$	0.029	0.005	0.048	0.005				
	(0.072)	(0.071)	(0.043)	(0.074)				
$\beta_{DiD,2010}$	0.026	0.018	0.037	0.063				
	(0.072)	(0.072)	(0.040)	(0.065)				
$\beta_{DiD,2011}$	0.012	0.006	0.055	-0.009	0.018	0.015	0.068	-0.015
	(0.076)	(0.076)	(0.040)	(0.072)	(0.077)	(0.081)	(0.042)	(0.074)
$\beta_{DiD,2013}$	0.107*	0.101	0.122***	0.113*	0.078	0.107*	0.143***	0.062
	(0.062)	(0.062)	(0.038)	(0.058)	(0.055)	(0.061)	(0.035)	(0.049)
$\beta_{DiD,2014}$	0.264***	0.258***	0.165***	0.192***	0.240***	0.257***	0.181***	0.142***
	(0.063)	(0.063)	(0.039)	(0.062)	(0.056)	(0.061)	(0.035)	(0.052)
$\beta_{DiD,2015}$	0.324***	0.328***	0.210***	0.306***	0.302***	0.315***	0.209***	0.233***
	(0.077)	(0.076)	(0.046)	(0.069)	(0.070)	(0.074)	(0.043)	(0.062)
Adjusted R^2	0.828	0.831	0.878	0.866	0.870	0.860	0.912	0.900
# Observations	25,693	25,684	25,775	25,452	17,471	17,302	17,471	17,471
# Firms	3,889	3,889	3,889	3,874	3,888	3,868	3,888	3,888

Notes: In this exercise we check the robustness of the baseline results from Table 2 on the firm-level impact of the 2013 amendment to the income tax break law of 2001. The outcome variables used are log(revenue), log(sales), log(number of workers), and log(total assets). The only difference between columns (1)-(4), columns (5)-(8), and columns (1)-(4) in Table 2 stems from the source of the data over which we estimate the model in Equation (2). Columns (1)-(4) use data from Amadeus, a Bureau van Dijk product. Columns (5)-(8) use the administrative data described in Section 3.2. The only difference between the sample used in Columns (5)-(8) and the baseline sample (used in columns (1)-(4) from Table 2) is that the former is the subset of the latter that we were able to match with Amadeus. The similarity between the estimates in columns (5)-(8) and those in columns (1)-(4) from Table 2 suggest that the sample matched with Amadeus is representative for the baseline sample. The similarity between the estimates in columns (1)-(4) and those in columns (5)-(8) suggest that the Amadeus data is of comparable quality to the administrative data. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table A5: Robustness of the Baseline Results to the Comparison Group Choice. Here Comparison Group in High-Tech, Knowledge-Intensive Sectors

	Revenue (1)	Sales (2)	Workers (3)	Assets (4)	Revenue (5)	Sales (6)	Workers (7)	Assets (8)
$\beta_{DiD,2011}$	0.008	0.001	0.052	-0.024	0.000	0.000	0.000	0.000
	(0.073)	(0.076)	(0.040)	(0.069)	(.)	(.)	(.)	(.)
$\beta_{DiD,2013}$	0.123**	0.156***	0.141***	0.071	0.024	0.066	0.083**	0.017
	(0.051)	(0.055)	(0.033)	(0.046)	(0.052)	(0.057)	(0.037)	(0.050)
$\beta_{DiD,2014}$	0.300***	0.304***	0.176***	0.191***	0.165***	0.185***	0.106***	0.094*
	(0.052)	(0.056)	(0.033)	(0.049)	(0.050)	(0.055)	(0.036)	(0.051)
$\beta_{DiD,2015}$	0.365***	0.379***	0.207***	0.258***	0.225***	0.236***	0.125***	0.136**
	(0.065)	(0.066)	(0.040)	(0.058)	(0.058)	(0.062)	(0.041)	(0.057)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Rel. prod.					1.393***	1.064***	0.204**	1.688***
					(0.244)	(0.210)	(0.086)	(0.133)
Small					0.331***	0.331***	0.393***	0.219***
					(0.045)	(0.044)	(0.036)	(0.036)
Medium					0.630***	0.635***	0.703***	0.512***
					(0.105)	(0.098)	(0.104)	(0.093)
Large					0.870***	0.856***	0.960***	0.753***
					(0.241)	(0.237)	(0.279)	(0.251)
Young					-0.011	-0.024	-0.011	-0.022
					(0.024)	(0.024)	(0.014)	(0.020)
Adjusted R^2	0.869	0.872	0.918	0.911	0.902	0.904	0.932	0.937
# Observations	23,517	23,164	23,526	23,516	17,528	17,263	17,535	17,525
# Firms	5,388	5,339	5,388	5,388	5,082	5,024	5,084	5,082

, *Notes:* In this exercise we check the robustness of the baseline results from Table 2 on the firm-level impact of the 2013 amendment to the income tax break law of 2001. The outcome variables used are log(revenue), log(sales), log(number of workers), and log(total assets). The only difference between this Table and Table 2 stems from the choice of the comparison group. In Table 2 we allowed firms whose sector was classified as part of ICT services, according to the OECD definition: Sectors 582 (Software publishing), 61 (Telecommunications), 62 (Computer programming, consultancy and related activities), 631 (Data processing, hosting and related activities; web portals), 951 (Repair of computers and communication equipment. In this Table we include firms which are part of sectors defined as high-tech, knowledge-intensive service sectors, according to the Eurostat definition: Motion picture, video and television programme production, sound recording and music publishing activities (59), Programming and broadcasting activities (60); Telecommunications (61), Computer programming, consultancy and related activities (62), Information service activities (63), and Scientific research and development (72). All other specification features are the same as those in Equation (2). The only difference between columns (1)-(4) and columns (5)-(8) is that the latter include firm-specific time-variant controls in addition to firm fixed effects, whereas the former include firm fixed effects alone. Heteroskedasticity robust errors in parentheses. All specifications include firm and sector-by-year fixed effects. The reference category for the firm size is "Micro" and for "Young" firms the reference category are firms "Older than five years." ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table A6: Robustness of the Baseline Results to the Comparison Group Choice. Here Comparison Group Only in Eligible Sectors

	Revenue (1)	Sales (2)	Workers (3)	Assets (4)	Revenue (5)	Sales (6)	Workers (7)	Assets (8)
$\beta_{DiD,2011}$	-0.001	-0.009	0.048	-0.041	0.000	0.000	0.000	0.000
	(0.075)	(0.078)	(0.040)	(0.071)	(.)	(.)	(.)	(.)
$\beta_{DiD,2013}$	0.123**	0.155***	0.144***	0.069	0.019	0.060	0.085**	0.012
	(0.052)	(0.056)	(0.033)	(0.047)	(0.052)	(0.058)	(0.038)	(0.051)
$\beta_{DiD,2014}$	0.302***	0.304***	0.179***	0.192***	0.162***	0.180***	0.108***	0.092*
	(0.053)	(0.057)	(0.034)	(0.050)	(0.051)	(0.056)	(0.036)	(0.052)
$\beta_{DiD,2015}$	0.364***	0.376***	0.211***	0.256***	0.217***	0.228***	0.126***	0.131**
	(0.066)	(0.068)	(0.040)	(0.060)	(0.059)	(0.063)	(0.041)	(0.059)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Rel. prod.					1.618***	1.032***	0.253**	1.787***
					(0.342)	(0.281)	(0.103)	(0.179)
Small					0.303***	0.319***	0.364***	0.168***
					(0.059)	(0.059)	(0.048)	(0.054)
Medium					0.690***	0.654***	0.729***	0.559***
					(0.150)	(0.123)	(0.100)	(0.179)
Large					0.000	0.000	0.000	0.000
					(.)	(.)	(.)	(.)
Young					-0.065*	-0.057*	-0.020	-0.040
					(0.034)	(0.033)	(0.018)	(0.029)
Adjusted R^2	0.840	0.847	0.904	0.879	0.883	0.891	0.925	0.916
# Observations	11,251	11,093	11,255	11,248	8,343	8,222	8,346	8,340
# Firms	2,600	2,579	2,600	2,600	2,429	2,404	2,430	2,429

Notes: In this exercise we check the robustness of the baseline results from Table 2 on the firm-level impact of the 2013 amendment to the income tax break law of 2001. The outcome variables used are log(revenue), log(sales), log(number of workers), and log(total assets). The only difference between this Table and Table 2 stems from the choice of the comparison group. In Table 2 we allowed firms whose sector was classified as part of ICT services, according to the OECD definition: Sectors 582 (Software publishing), 61 (Telecommunications), 62 (Computer programming, consultancy and related activities), 631 (Data processing, hosting and related activities; web portals), 951 (Repair of computers and communication equipment. In this Table we only include firms that belong 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). All other specification features are the same as those in Equation (2). The only difference between columns (1)-(4) and columns (5)-(8) is that the latter include firm-specific time-variant controls in addition to firm fixed effects, whereas the former include firm fixed effects alone. Heteroskedasticity robust errors in parentheses. All specifications include firm and sector-by-year fixed effects. The reference category for the firm size is "Micro" and for "Young" firms the reference category are firms "Older than five years." ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Event-Study Design. The 2013 reform to the income tax break law substantially expanded the scope of the income tax exemption. However, the treated firms in the main DiD exercise increased the share of exempted employees at different times after the amendment in 2013. We apply a more general version of the DiD design, i.e., an event-study design that allows us to exploit the differential timing of the first large-scale exemption across firms. Specifically,

$$y_{ist} = \alpha_i + \lambda_{st} + \sum_{k=-4}^{+2} \theta_k D_{ist}^k + \varepsilon_{ist}, \tag{A3}$$

where D_{ist}^k is an event year dummy variable. We map calendar years t to their event year analog. τ_i is the event year of firm i, i.e., the year when i jumped from less than 5 to more than 20% of its workers exempted from the income tax. Then $D_{ist}^k := \mathbb{1}[t = \tau_i + k]\ \forall k$, where $\mathbb{1}[.]$ is an indicator function for the expression in the brackets being true.[60] All other variables are defined as in Equation (2).

A key advantage of this design is that it allows us to lengthen the period over which we observe the firms-to-be-treated. While the main DiD design can only exploit 2011 as a pre-period year, in the event-study design firms that become treated in 2015 are observed for three years before their event.[61]

The sequence of θ_k is our object of interest. The interpretation of the θ_k depends on the sample of estimation.[62] We use two samples that enable complementary findings. In the "baseline sample" we allow firms that are in ICT service sectors but do not experience an event. This sample has the benefit of a significantly larger size. This is also the same sample as the one used for our baseline results in Table 2. The "restricted sample" only uses firms that eventually experience the event. With the second sample, θ_k for $k \geq 0$ captures the average increase in the outcome of interest in event year k across eventually treated firms, compared to the average level of that outcome for the same firms in the year before their event. The baseline sample adds to the reference outcomes those of firms that have less than 5% of workers exempted from the income tax, while still in ICT service sectors.[63]

One concern relative to this design is related to the extent to which a delay in the onset of a large-scale exemption can be plausibly exogenous to the firm. We claim that the most plausible scenarios do not pose threats to the identification of the causal effects of the event. It is very likely for the 2013 amendment to not have received as much media coverage as the initial introduction of the policy. Hence, firms are likely to have learned about the 2013 amendment at different points in time, which may explain the differential timing of the onset of treatment. As long as the reason for the delay of the event is unrelated to future trends of expansion after the event, our findings have a causal interpretation.

Table A7 reports the event-study estimates from the model in Equation (A3). Columns (1)-(4) report the results for the full baseline sample, whereas columns (5)-(8) report the results for the restricted sample. We first notice that overall, prior to their event, firms are not embarked on a trend of growth. In columns (1)-(4), in event year *-3* we find evidence of some statistically-significant differences between eventually exempted firms and firms never exempted. For all other years however, we fail to find distinc-

[60] As we use the same baseline sample as for the DiD design, by construction $\tau_i \geq 2013\ \forall i$.

[61] Given the length of our panel in calendar years and the fact that events are constrained to occur between 2013 to 2015, we observe a maximum of two years after the event (for firms whose events occur in 2013) and a maximum of four years prior to the event (for firms whose events occur in 2015).

[62] The sample on which we estimate Equation (A3) implicitly assumes a counterfactual path for the eventually-exempted firms.

[63] As not all parameters are identified as written, we normalize θ_{-1} to zero to simplify the test for an effect on impact.

tive growth trends. Moreover, when we restrict in columns (5)-(8) to the restricted sample, we find that firms were not embarked on growth trends prior to their event.

Second, all models show that firms experiencing an event expand along the four measures of size: revenue, sales, employment, and total assets. This expansion already occurs in the first year of treatment and intensifies in the next two years. Reassuringly, the magnitudes from the event-study design are comparable to those from the historical DiD design applied on Amadeus data and from the DiD design applied to the same administrative data from 2011 to 2015.

Table A7: Robustness of the Baseline Results to an Event-Study Design

	Baseline Sample				**Restricted Sample**			
	Revenue (1)	Sales (2)	Workers (3)	Assets (4)	Revenue (5)	Sales (6)	Workers (7)	Assets (8)
θ_{-4}	0.025	-0.007	0.021	0.011				
	(0.121)	(0.112)	(0.067)	(0.118)				
θ_{-3}	-0.198**	-0.175**	-0.068	-0.081	-0.205**	-0.159	-0.066	-0.084
	(0.079)	(0.085)	(0.044)	(0.073)	(0.100)	(0.100)	(0.056)	(0.099)
θ_{-2}	-0.087	-0.043	-0.015	-0.050	-0.082	-0.027	-0.018	-0.045
	(0.055)	(0.058)	(0.033)	(0.051)	(0.062)	(0.063)	(0.036)	(0.059)
θ_0	0.109**	0.152***	0.100***	0.099**	0.113	0.151**	0.086*	0.111
	(0.051)	(0.053)	(0.031)	(0.043)	(0.074)	(0.077)	(0.045)	(0.071)
θ_{+1}	0.286***	0.308***	0.161***	0.255***	0.287**	0.314**	0.171**	0.278**
	(0.057)	(0.060)	(0.036)	(0.052)	(0.126)	(0.126)	(0.072)	(0.118)
θ_{+2}	0.422***	0.458***	0.208***	0.312***	0.470**	0.510***	0.248**	0.367**
	(0.085)	(0.086)	(0.056)	(0.088)	(0.200)	(0.197)	(0.114)	(0.186)
Adjusted R^2	0.869	0.874	0.907	0.901	0.868	0.871	0.908	0.872
# Observations	22,592	22,212	22,598	22,587	1,393	1,383	1,393	1,393
# Firms	5,177	5,128	5,177	5,177	317	316	317	317

Notes: We use an event-study design to study the impact of the 2013 amendment to the income tax break law of 2001. A firm experiences an "event" when it jumps after 2013 from under 5 to over 20% of workers exempted from the income tax. This table reports the point estimates of the event-study coefficients from Equation (A3). The outcome variables used are log(revenue), log(sales), log(number of workers), and log(total assets). We propose two samples: the baseline sample includes firms in ICT services not experiencing the event, in addition to those experiencing the event (columns (1)-(4)), while the restricted sample is constrained to only include the firms experiencing the event (columns (5)-(8)). The reference year for each firm is the one prior to the year of its event. All specifications include firm and sector-by-year fixed effects. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Last, reporting alongside the point estimates of θ_k from the two samples allows us to check whether results are driven by the contrast with the "never-exempted" (i.e., firms with under 5% of their workforce exempted in all years, or firms in non-eligible sectors) or by the staggered timing of the event across treated firms. As point estimates do not seem to vary much with the sample, we conclude that our results are driven by the differential firm-level timing of this large-scale exemption.[64]

[64] Only the standard errors are affected by the choice of the sample, as the restricted sample contains only 6% of the firms in the baseline sample.

Table A8: Robustness of the Baseline Results to the Threshold Choice

	Revenue (1)	Sales (2)	Workers (3)	Assets (4)	Revenue (5)	Sales (6)	Workers (7)	Assets (8)
$\beta_{DiD,2011}$	-0.004	-0.012	0.043	-0.026	0.000	0.000	0.000	0.000
	(0.070)	(0.073)	(0.039)	(0.066)	(.)	(.)	(.)	(.)
$\beta_{DiD,2013}$	0.115**	0.131**	0.136***	0.046	0.035	0.055	0.081**	-0.002
	(0.048)	(0.053)	(0.032)	(0.044)	(0.048)	(0.054)	(0.036)	(0.047)
$\beta_{DiD,2014}$	0.281***	0.302***	0.184***	0.161***	0.161***	0.186***	0.115***	0.063
	(0.050)	(0.054)	(0.032)	(0.046)	(0.047)	(0.053)	(0.035)	(0.048)
$\beta_{DiD,2015}$	0.336***	0.359***	0.218***	0.212***	0.209***	0.240***	0.139***	0.091*
	(0.061)	(0.065)	(0.038)	(0.056)	(0.054)	(0.060)	(0.039)	(0.054)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
Rel. prod.					1.595***	1.414***	0.218***	1.792***
					(0.222)	(0.229)	(0.085)	(0.133)
Small					0.358***	0.357***	0.373***	0.232***
					(0.046)	(0.046)	(0.037)	(0.037)
Medium					0.639***	0.672***	0.752***	0.633***
					(0.114)	(0.103)	(0.112)	(0.114)
Large					0.925***	0.937***	1.042**	0.898**
					(0.344)	(0.341)	(0.407)	(0.363)
Young					-0.035	-0.022	-0.001	-0.044**
					(0.022)	(0.026)	(0.014)	(0.021)
Adjusted R^2	0.868	0.859	0.908	0.901	0.907	0.898	0.924	0.931
# Observations	22,678	22,434	22,678	22,677	16,888	16,715	16,888	16,887
# Firms	5,194	5,159	5,194	5,194	4,881	4,846	4,881	4,881

Notes: In this exercise we check the robustness of the baseline results from Table 2 on the firm-level impact of the 2013 amendment to the income tax break law of 2001. The outcome variables used are log(revenue), log(sales), log(number of workers), and log(total assets). The only difference between this Table and Table 2 stems from the choice of the threshold above which we consider a firm to be treated by the 2013 amendment, i.e., instead of allowing $Exempted_{isk}$ to become 1 only when the firm passes above the 20% threshold of exempted employees, we lower this threshold to 15%. All other specification features are the same as those in Equation (2). Similar results are obtained when we lower this threshold further to 10%. The only difference between columns (1)-(4) and columns (5)-(8) is that the latter include firm-specific time-variant controls in addition to firm fixed effects, whereas the former include firm fixed effects alone. Heteroskedasticity robust errors in parentheses. All specifications include firm and sector-by-year fixed effects. The reference category for the firm size is "Micro" and for "Young" firms the reference category are firms "Older than five years." ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Table A9: Robustness of the Baseline Results to the Definition of Treatment. Here Treatment Based on Firm Sector, Full and Baseline Sample

	Full Sample				**Baseline Sample**			
	Revenue (1)	Sales (2)	Workers (3)	Assets (4)	Revenue (5)	Sales (6)	Workers (7)	Assets (8)
$\beta_{DiD,2011}$	-0.004	-0.007	-0.006	-0.021	0.016	-0.008	0.012	-0.011
	(0.026)	(0.028)	(0.015)	(0.025)	(0.029)	(0.029)	(0.015)	(0.027)
$\beta_{DiD,2013}$	0.038*	0.038*	0.023*	-0.000	0.007	0.014	0.012	-0.021
	(0.020)	(0.022)	(0.012)	(0.018)	(0.023)	(0.023)	(0.013)	(0.020)
$\beta_{DiD,2014}$	0.084***	0.091***	0.073***	0.027	0.050**	0.054**	0.050***	0.000
	(0.022)	(0.024)	(0.013)	(0.020)	(0.024)	(0.024)	(0.014)	(0.022)
$\beta_{DiD,2015}$	0.094***	0.109***	0.100***	0.067***	0.028	0.065**	0.064***	0.020
	(0.027)	(0.029)	(0.016)	(0.025)	(0.029)	(0.029)	(0.016)	(0.027)
Adjusted R^2	0.890	0.883	0.926	0.911	0.868	0.873	0.907	0.901
# Observations	26,507	26,281	26,507	26,506	22,592	22,212	22,598	22,587
# Firms	5,721	5,696	5,721	5,721	5,177	5,128	5,177	5,177

Notes: In this exercise we study the effects of the 2013 expansion to the income tax break law of 2001 using a DiD design similar to the DiD design used to study the initial introduction of the law in 2001. Specifically, treatment is assigned to firms based on whether their sector was targeted by the income tax break law or not, interacted with a dummy activated when $t \geq 2013$ (as in Equation (1), for year 2001), as opposed to their workforce exemption status. Note that this design is oblivious to a firm's workforce exemption status both before and after the 2013 reform. The outcome variables used are log(revenue), log(sales), log(number of workers), and log(total assets). Columns (5)-(8) use the exact same baseline sample used in Table 2. Columns (1)-(4) use an expanded version of the baseline sample. Because the definition of treatment used here does not impose any restrictions on the firm-level share of exempted workers, neither before nor after 2013, the full sample used in columns (1)-(4) contains *all* firms in the eligible sectors (NACE codes 5821, 5829, 6201, 6202, or 6209) and *all* firms in the control group (i.e. all firms in other ICT service sectors, based on the OECD definition). This explains why the sample sizes in columns (1)-(4) are larger than the sample sizes in columns (5)-(8). Heteroskedasticity robust errors in parentheses. All specifications include firm, year, and sector fixed effects. ***,**,* denotes statistical significance at the 1%, 5%, and 10% levels, respectively.

Appendix A.2 Sector-Level Cross-Country Effects of 2001 Tax Break

Appendix A.2.1 IT Sector Vs. Rest of the Economy

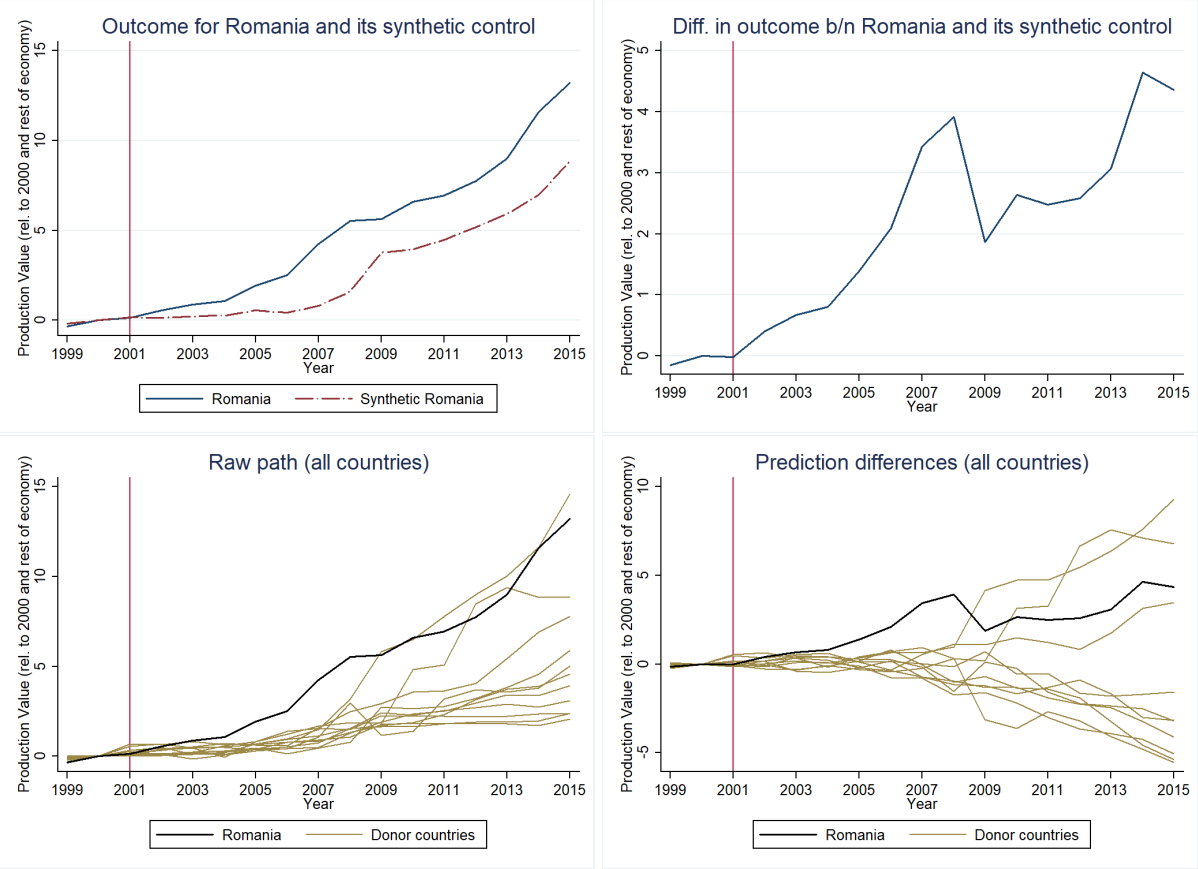

Figure A1: IT Sector Vs. Rest of Economy. SCM with Outcome Variable: "Production Value" - Million Euro (Normalized)

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) " Production value - million euro" 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 industry (% 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.

Table A10: Post-treatment Results: Effects, p-values, Standardized p-values. Outcomes: Gross Revenue, Employment, and Production Value

	Gross Revenue			**Employment**			**Production Value**		
Post-treatment year	$\hat{\alpha}_{1t}$	p-values	Standardized p-values	$\hat{\alpha}_{1t}$	p-values	Standardized p-values	$\hat{\alpha}_{1t}$	p-values	Standardized
2001	0.13	0.27	0.55	0.14	0.27	0.00	-0.01	1.00	1.00
2002	0.69	0.00	0.27	0.26	0.09	0.00	0.41	0.10	0.50
2003	1.13	0.00	0.27	0.67	0.09	0.00	0.67	0.00	0.80
2004	1.37	0.00	0.27	1.13	0.00	0.00	0.80	0.00	0.30
2005	1.95	0.00	0.36	1.46	0.00	0.00	1.39	0.00	0.40
2006	2.57	0.00	0.36	1.75	0.00	0.00	2.10	0.00	0.50
2007	3.17	0.00	0.36	2.10	0.00	0.00	3.43	0.00	0.30
2008	3.21	0.00	0.27	1.90	0.00	0.00	3.92	0.00	0.40
2009	4.14	0.09	0.27	1.17	0.27	0.00	1.87	0.20	0.70
2010	4.49	0.09	0.27	1.03	0.45	0.00	2.64	0.30	0.70
2011	2.92	0.36	0.27	1.12	0.36	0.00	2.48	0.40	0.70
2012	3.32	0.45	0.36	1.19	0.18	0.00	2.59	0.40	0.90
2013	4.52	0.36	0.36	1.32	0.18	0.00	3.06	0.50	0.90
2014	6.15	0.36	0.36	1.54	0.18	0.00	4.64	0.30	0.70
2015	6.52	0.27	0.36	1.83	0.18	0.00	4.36	0.50	0.80

Table A11: Predictor balance, averaged between 1999 and 2000. Synthetic Romania = 0.104 Czech Republic and 0.896 Slovakia (Gross Revenue); = 0.534 Bulgaria and 0.466 Czech Republic (Employment); = 0.339 Bulgaria and 0.661 Lithuania (Production Value)

	Romania	Synthetic Romania		
		Gross Revenue	Employment	Prod. Value
GDP per capita (constant LCU, $)	9,663	9,796	15,972	5,341
Services (% of GDP)	53.23	60.32	60.07	62.59
High-tech manufacturing (% mfg. value added)	23.85	35.69	34.05	23.85

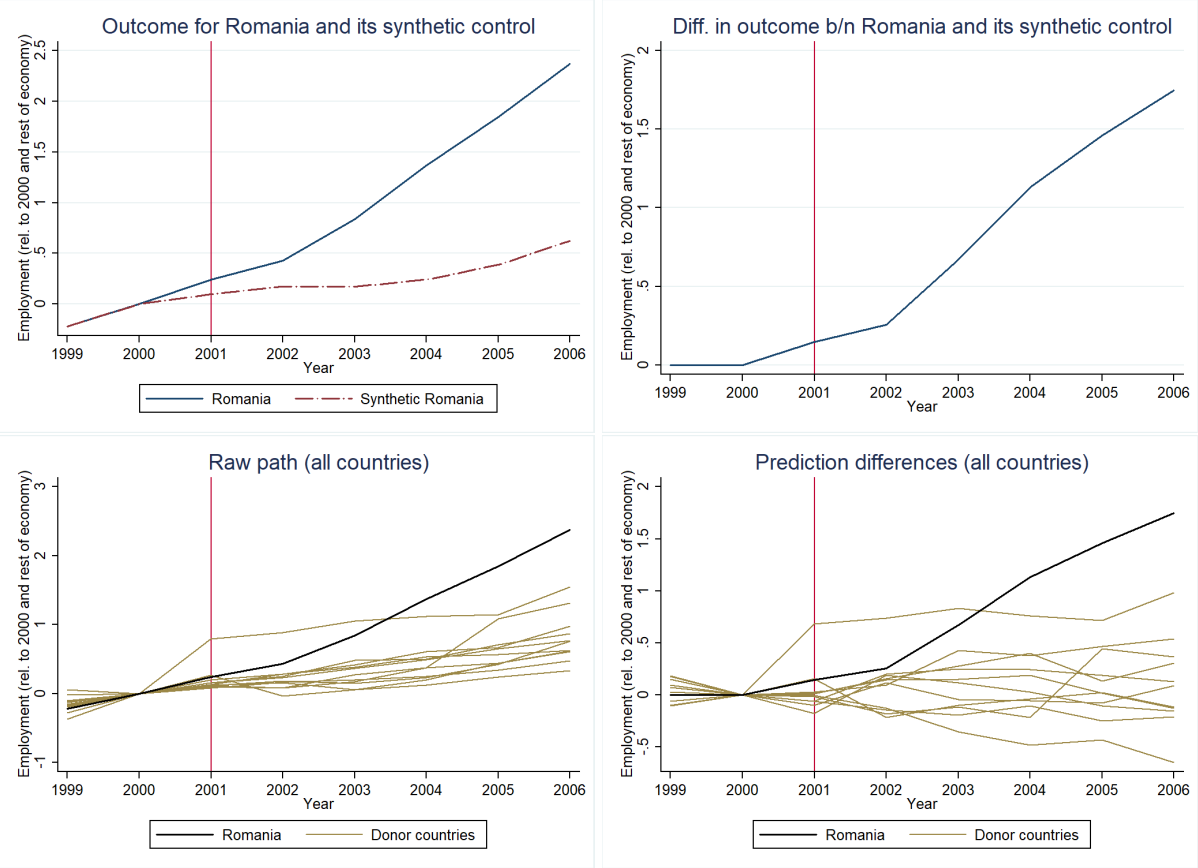

Figure A2: IT Sector Vs. Rest of Economy. SCM with Outcome Variable: "Employees - Number" (Normalized). Robustness to the Correspondence between NACE Industry Classifications

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. This exercise differs from the one presented in Figure 4 in the period of analysis. Instead of considering the entire 1999 to 2015 period, we cut the analysis in 2006 (the last year where the data is available in the NACE Rev 1 classification). This tests the robustness of our results to the correspondence we develop between the NACE Rev 1 and NACE Rev 2 classifications. 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 industry (% 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.

Appendix A.2.2 IT-Using Sectors Vs. Non-IT Using Sectors

Table A12: Sectors with Consistently-Available Data in Eurostat, by Quarter of IT-use Intensity

Under Quartile 1	**Between Quartiles 1-2**
Wood and products of wood and cork (20)	Manufacture of textiles and textile products (17)
Basic metals (27)	Manufacture of wearing apparel; dressing, dyeing of fur (18)
Fabricated metal products (28)	Manufacture of pulp, paper and paper products (21)
Electricity, gas and water supply (40-41)	Publishing, printing and reproduction of recorded media (22)
	Manufacture of machinery and equipment n.e.c. (29)
	Manufacture of furniture; manufacturing n.e.c. (36)
	Recycling (37)
	Construction (45)
Between Quartiles 2-3	**Over Quartile 3**
Manufacture of chemicals and chemical products (24)	Wholesale and retail trade (50-52)
Manufacture of rubber and plastic products (25)	Transport and storage (60-63)
Manufacture of other non-metallic mineral products (26)	Real estate activities (70)
Manufacture of office machinery and computers (30)	Renting of machinery and equipment (71)
Manufacture of electrical machinery and apparatus (31)	Computer and related activities (72)
Manufacture of radio, television and communication (32)	R&D and other business activities (73-74)
Manufacture of medical, precision, optical instruments (33)	
Manufacture of motor vehicles, trailers, semi-trailers (34)	
Manufacture of other transport equipment (35)	
Hotels and restaurants (55)	

Notes: This table groups sectors with consistently-available data in Eurostat (across variables, years, and countries - see Appendix D for details) based on their dependence on the IT sector for inputs. To establish this dependence, we start from the input-output table (I-O table, henceforth) of Romania for the year 2000. We use the harmonized I-O table provided by the *OECD*, that tracks the flows of goods and services between all two-digit NACE Rev 1 sectors in Romania. 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"). It is based on this distribution of shares that we compute the quartiles mentioned in this table. In the main specification of the inter-industry SCM analysis we exclude sector 72 altogether from the analysis, as this is the treated sector itself. We also assign sectors over quartile three into the high-intensity category and all other sectors into the low-intensity. Note that the sectors actually used in the analysis and presented in this table are a subset of all sectors in the economy, as not all sectors had consistently-available data. Nevertheless, the grouping of these sectors by quarters was based on the full I-O matrix, which includes all sectors in the economy.

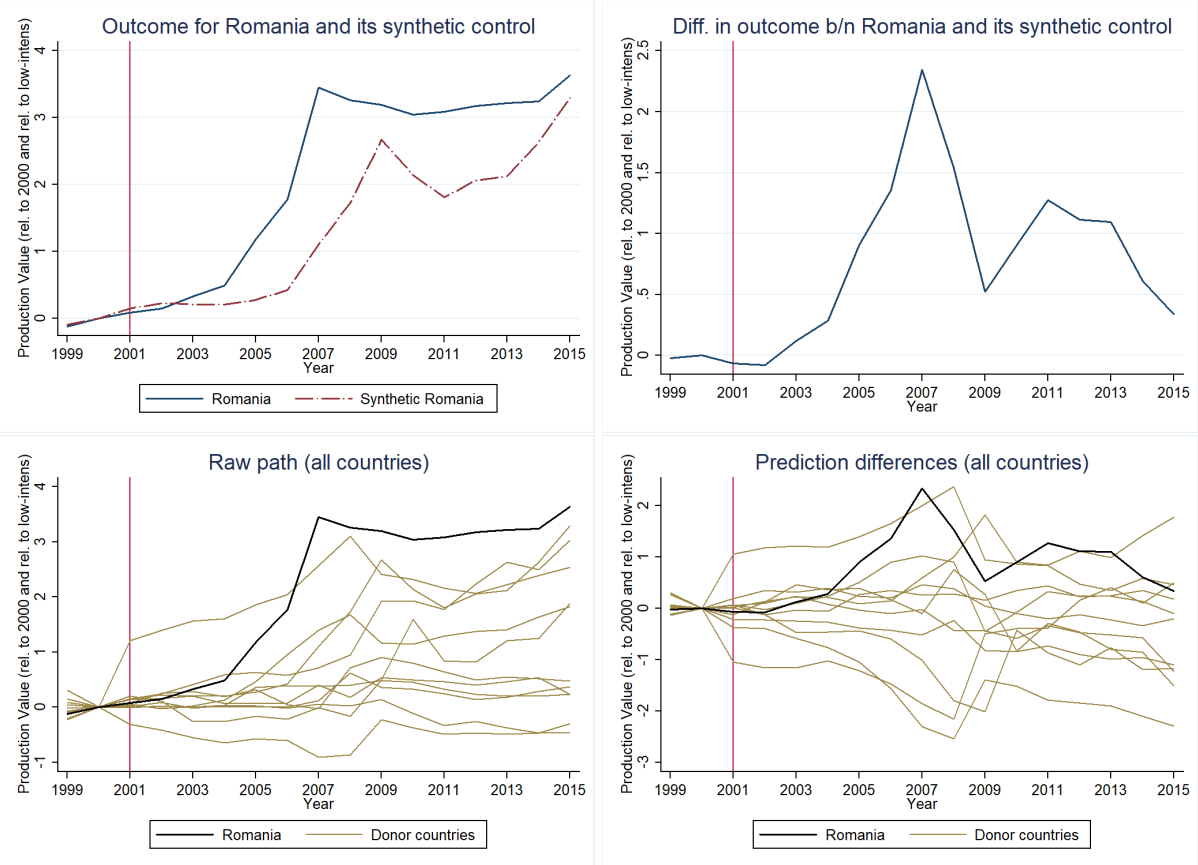

Figure A3: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: "Production Value" - Million Euro (Normalized)

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) " Production value - million euro." 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 industry (% 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.

Table A13: Post-treatment Results: Effects, p-values, Standardized p-values. Outcomes: Gross Revenue, Employment, and Production Value

	Gross Revenue			**Employment**			**Production Value**		
Post-treatment year	$\hat{\alpha}_{1t}$	p-values	Standardized p-values	$\hat{\alpha}_{1t}$	p-values	Standardized p-values	$\hat{\alpha}_{1t}$	p-values	Standardized
2001	-0.04	0.80	0.50	0.08	0.22	0.22	-0.07	0.64	0.64
2002	-0.02	1.00	0.50	0.09	0.22	0.11	-0.08	0.91	0.64
2003	0.18	0.30	0.50	0.27	0.22	0.11	0.12	0.82	0.55
2004	0.28	0.00	0.40	0.43	0.22	0.11	0.29	0.55	0.45
2005	0.54	0.00	0.40	0.68	0.00	0.11	0.90	0.27	0.27
2006	0.90	0.00	0.40	0.86	0.00	0.11	1.36	0.27	0.18
2007	1.71	0.00	0.40	0.93	0.00	0.11	2.34	0.00	0.18
2008	1.88	0.00	0.40	0.98	0.11	0.11	1.53	0.36	0.36
2009	1.46	0.00	0.40	0.77	0.00	0.11	0.52	0.45	0.27
2010	1.33	0.00	0.40	0.68	0.00	0.11	0.90	0.09	0.27
2011	1.34	0.00	0.40	0.73	0.00	0.11	1.27	0.09	0.27
2012	1.30	0.10	0.40	0.68	0.00	0.11	1.11	0.18	0.27
2013	0.68	0.30	0.40	0.67	0.00	0.11	1.10	0.09	0.27
2014	0.73	0.30	0.50	0.61	0.22	0.11	0.61	0.45	0.36
2015	0.75	0.50	0.50	0.61	0.11	0.11	0.34	0.73	0.55

Table A14: Predictor balance, averaged between 1999 and 2000. Synthetic Romania = 0.341 Bulgaria and 0.659 Lithuania (Gross Revenue); = 0.363 Bulgaria + 0.002 Hungary + 0.625 Lithuania (Employment) ; = Bulgaria (Production Value)

	Romania	Synthetic Romania		
		Gross Revenue	Employment	Prod. Value
GDP per capita (constant LCU, $)	9,663	5,342	9,616	5,771
Services (% of GDP)	53.23	62.59	62.53	60.85
High-tech manufacturing (% mfg. value added)	23.85	23.87	24.15	30.53

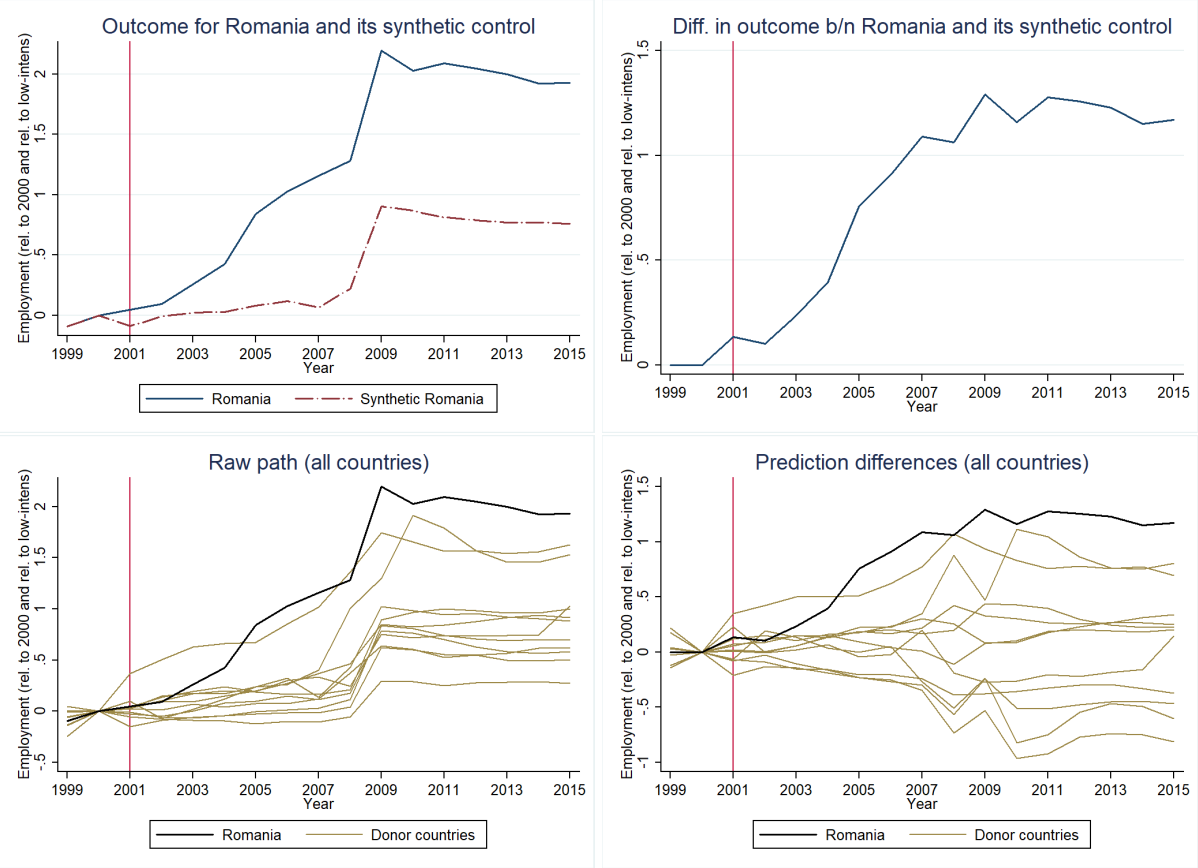

Figure A4: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: "Employees - Number" (Normalized). Robustness to Exclusion of Sectors Comparable to IT from High-Intensity Sectors Category

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. This exercise differs from the one presented in Figure 6 in its exclusion of sectors K73 and K74 from the analysis. Both sectors are otherwise part of the list of sectors that rely heavily on IT services (top 25% users of IT services). We used these sectors as comparison sectors in the firm-level analysis. 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 industry (% 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.

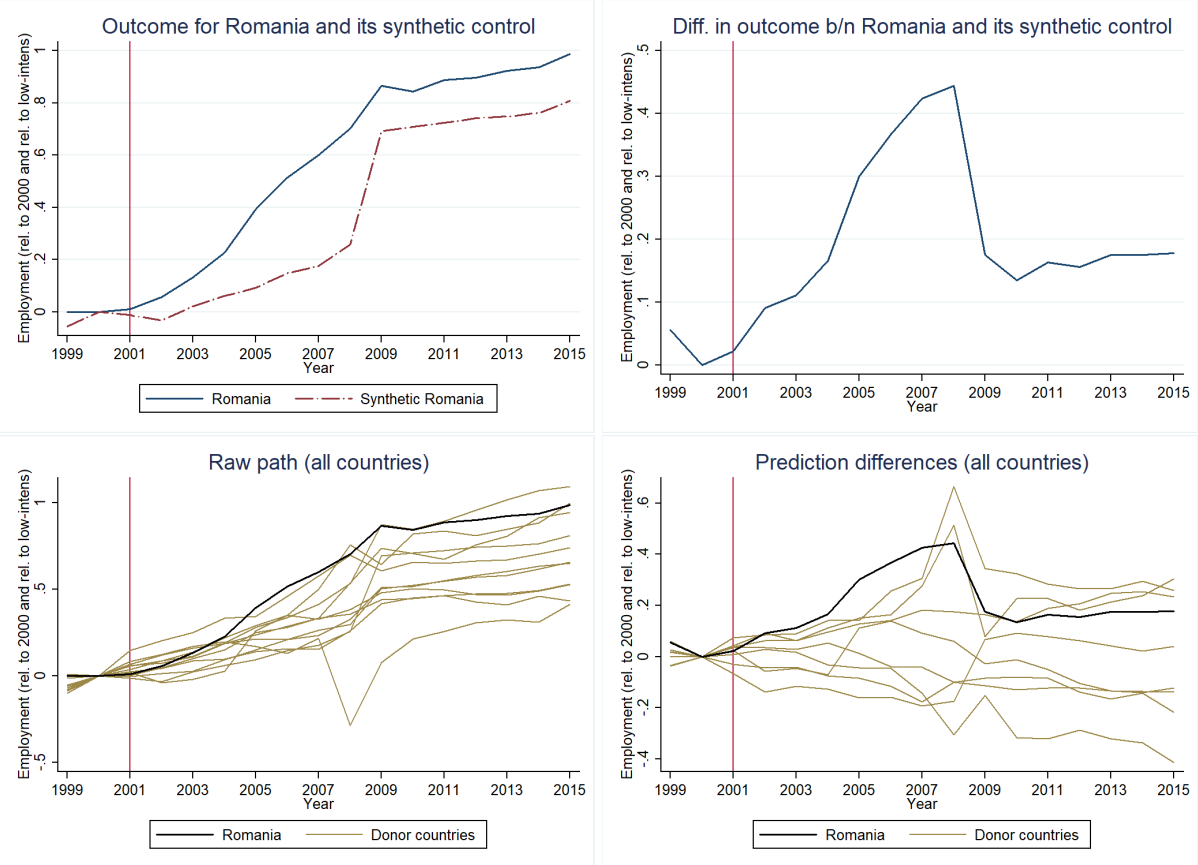

Figure A5: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: "Employees - Number" (Normalized). Robustness to Grouping of Quarters into High- and Low-Intensity

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. This exercise differs from the one presented in Figure 6 in the grouping of quarters into high- and low-intensity IT users. The main results compare the top 25% sectors in terms of IT services use (high-intensity category) with the bottom 75% (low-intensity category). In this exercise we compare the top 50% (high-intensity category) to the bottom 50% (low-intensity category). 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 industry (% 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.

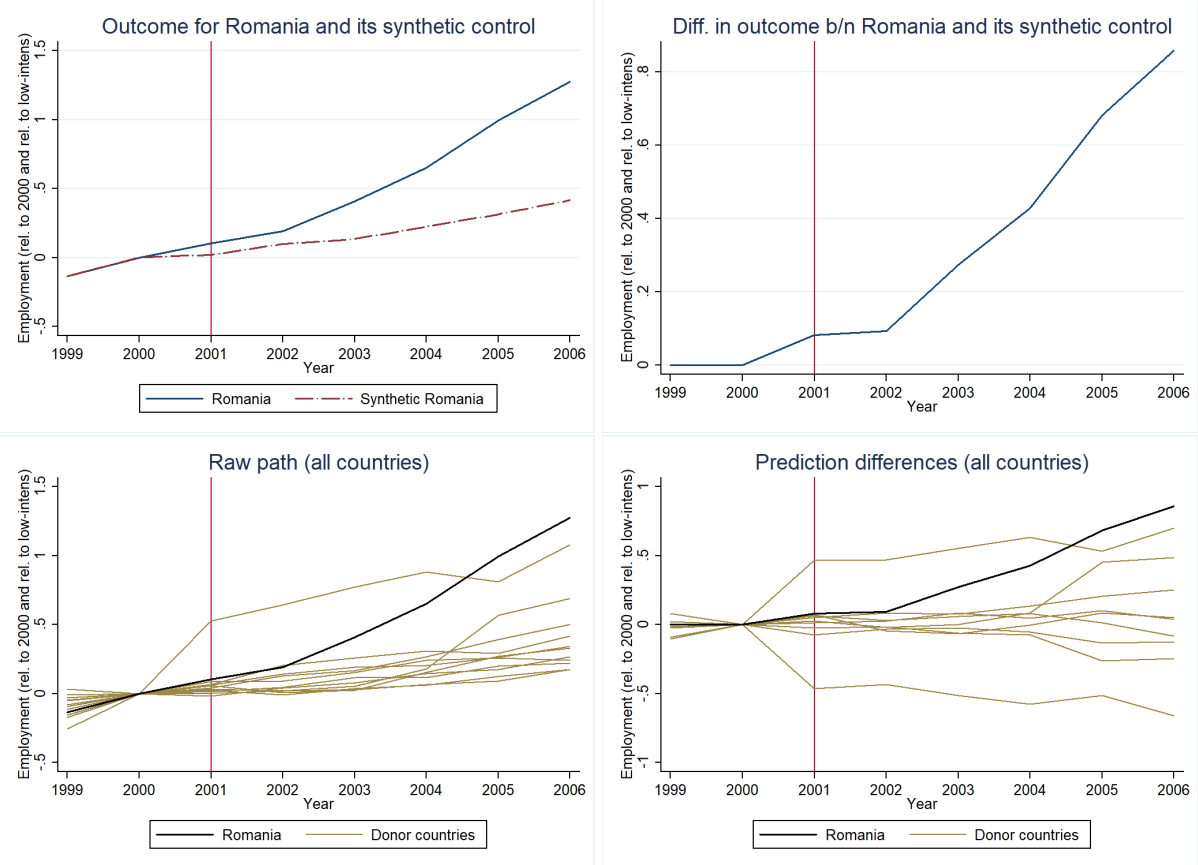

Figure A6: IT-Using Sectors Vs. Non-IT Using Sectors. SCM with Outcome Variable: "Employees - Number" (Normalized). Robustness to the Correspondence between NACE Industry Classifications

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. This exercise differs from the one presented in Figure 6 in the period of analysis. Instead of considering the entire 1999 to 2015 period, we cut the analysis in 2006 (the last year where the data is available in the NACE Rev 1 classification). This tests the robustness of our results to the correspondence we develop between the NACE Rev 1 and NACE Rev 2 classifications. 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 industry (% 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.

Table A15: Post-treatment Results: Effects, p-values, Standardized p-values. Outcome: (Goods Export) Trade Values

Post-treatment year	$\hat{\alpha}_{1t}$	p-values	Standardized p-values
2001	0.02	0.90	0.90
2002	0.03	0.80	0.70
2003	0.16	0.80	0.40
2004	0.53	0.10	0.10
2005	1.03	0.00	0.20
2006	1.60	0.10	0.10
2007	2.30	0.10	0.10
2008	3.64	0.10	0.10
2009	4.76	0.00	0.10
2010	5.96	0.00	0.10
2011	7.46	0.00	0.10
2012	6.79	0.00	0.10
2013	7.69	0.00	0.10
2014	7.70	0.00	0.10

Table A16: Predictor balance, averaged between 1996 and 2000. Synthetic Romania = 0.706 Czech Republic, 0.240 Lithuania, and 0.053 Slovakia

	Romania	
	Actual	Synthetic
GDP per capita (constant LCU, $)	9,771	20,689
Services (% of GDP)	47.89	58.41
High-tech manufacturing (% mfg. value added)	24.33	32.60

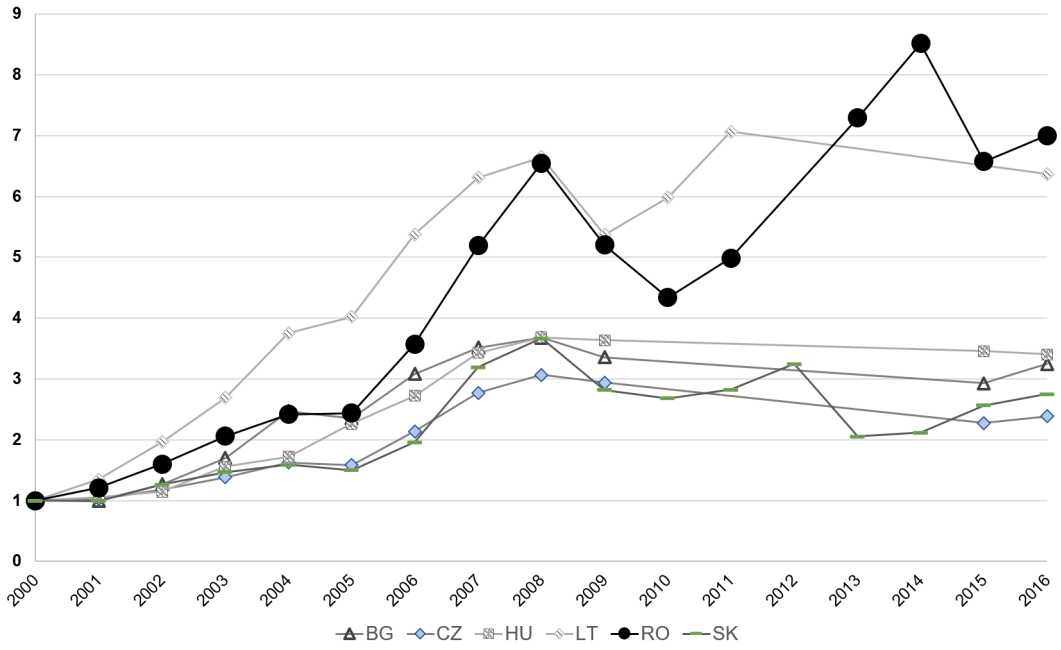

Figure A7: Trends in Export Value of IT-using Service Sectors

Notes: This figure plots the evolution of export trade values for service sectors relying more on IT services. These sectors are NACE Rev 1 sectors 51 (Wholesale trade and commission trade), 60 (Land transport), 62 (Air transport), 63 (Cargo handling and storage), 74 (Other business activities), and 75 (Public administration and defence). In Romania, most of the impressive growth experienced between 2001 and 2016 is explained by sectors under NACE Rev 1 sector 74 (e.g., call centers, advertising, business and management consultancy, secretarial and translation activities etc.). The figure includes Romania and five other countries that appear in SCM exercises as part of synthetic Romania, i.e. Bulgaria, the Czech Republic, Hungary, Lithuania, and Slovakia. This evolution is with respect to the export value in year 2000. This figure uses UN Comtrade data, EBOPS 2002. A similar graph can be made comparing the growth of IT-using service sectors to the growth of less IT-using service sectors. Romania continues to stand out in a relative comparison as well.

Appendix B Growth of the Sector on the Extensive Margin

Appendix B.1 Sectoral Switches of Existing Firms to the IT Sector

One channel through which the income tax break law could have boosted the growth of the Romanian IT sector is by changing incentives of firms (workers) to conduct their activity in the eligible sector. Using a sector-level transition matrix shared by the *National Bank of Romania*, we study transitions into the eligible sector (NACE Rev 1 code 722), around the introduction of the 2001 income tax break. While this evidence is descriptive, it highlights abnormal trends in transitions concomitant to the introduction of the income tax break.

Figure B8 plots the absolute number of firms alive in year $t-1$ (with an economic activity different from 722) that switched in year t to the NACE Code 722 (7221 and 7222, Software consultancy and supply). Figure B9 plots the year t share of firms in NACE Code 722 alive in year $t-1$ coming from firms with a different economic activity in $t-1$ that have switched their year t economic activity to 722. We notice that while switching one's sector towards 722 was a trend occurring both before and long after the passing of the 2001 and 2004 laws, immediately after the initial passing of the law, switching to NACE Rev 1 sector 722 became a visibly more popular practice. From 2004 onward, the momentum of this switching practice is lost, which suggests that most justifiable switches occurred immediately after

the passage of the law.

One caveat to this exercise is that we cannot pin down the exact reasons behind these switches in firms' main sectoral code. Hence, this type of growth of the sector may be to some extent artificial, if firms switching their main sectoral activity were already conducting most of their activity in the eligible sector, but had a misassigned main sector. As the law made the income tax break eligibility dependent on a firm's main sectoral code, the law may have incentivized corrections in firms' sectoral codes.

Most switches in firms' main sectoral code occurred from sectoral codes complementary to the eligible sector (e.g., 721 - "Consulting in the field of computing equipment / hardware" and 726 - "Other activities related to computer science"). We conjecture that these switching firms were multi-activity firms that decided to focus on software development once the income tax break for programmers became available. One might be concerned that these sectoral switches were meant to deceit tax authorities. We do not believe that this concern is justified, because the sector of the firm was only one of several strict requirements for a worker to qualify for the exemption. In particular, the firm had to provide thorough evidence that a revenue of at least 10,000 U.S. dollars per exempted employee came from software development and that employees involved in this software development had eligible bachelor degrees.

For these reasons, the evidence in Figures B8 and B9 suggests that switches were likely to be motivated by an actual focus of the firm on software development. Even if part of these switches may not have resulted in a stronger focus on software development (had some of these firms had a previously misassigned main sector), it is plausible to expect that the income tax break law has strengthened incentives to focus on software development. Moreover, the stark jump in switches around 2001 is a convincing piece of evidence on the salience of the law.

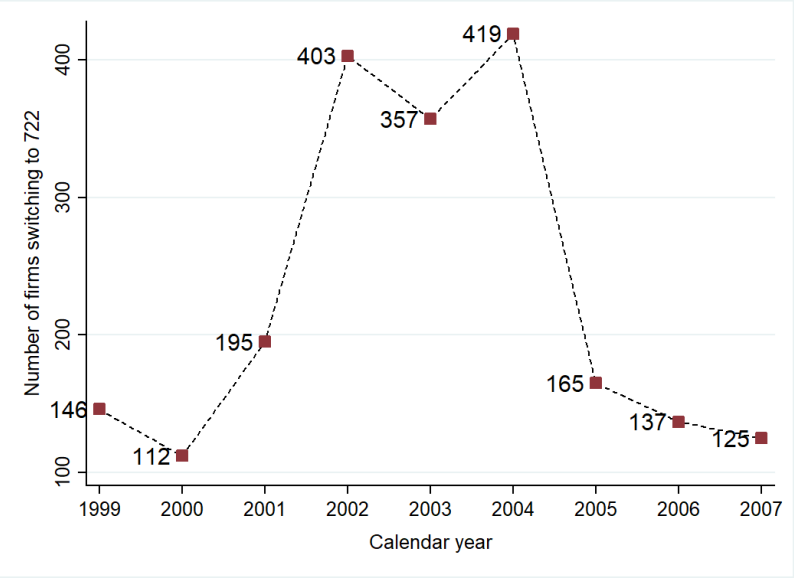

Figure B8: Romania: Number of firms alive in year $t-1$ and switching their economic activity in t towards the beneficiary sector, 722

Notes: Data source: *National Bank of Romania*. This graph plots the number of firms alive in year $t-1$ (with an economic activity different from 722) that switched in year t to the NACE Code 722 (7221 and 7222, Software consultancy and supply). We notice two peaks in this practice in 2002 and 2004, the years after the first and second laws of income tax exemption for workers in sector 722 were passed in Romania.

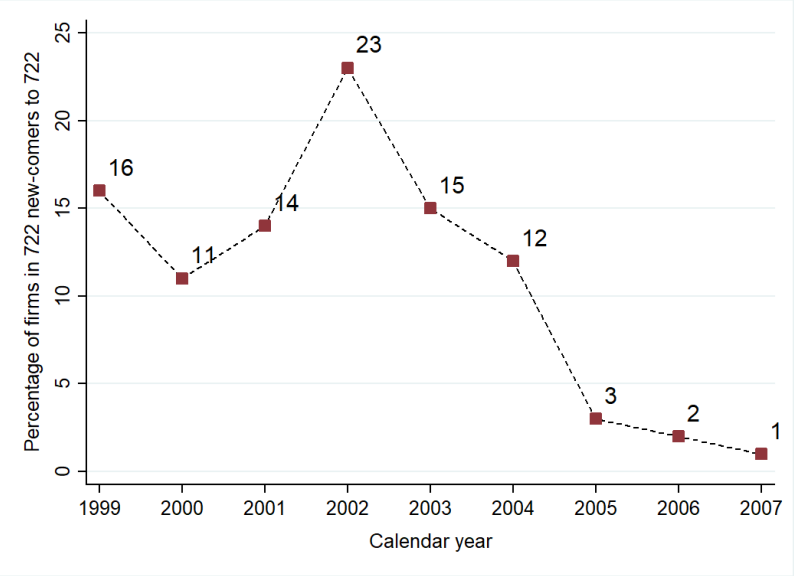

Figure B9: Romania: Among firms in NACE code 722 in year t and alive in year $t-1$ (with any economic activity), the percentage of new-comers in t to 722

Notes: Data source: *National Bank of Romania*. This graph plots the year t share of firms in NACE Code 722 (7221 and 7222, Software consultancy and supply) alive in year $t-1$ coming from firms with a different economic activity in $t-1$ that have switched their year t economic activity to 722. We notice a clear peak in this practice in 2002, the year after the first law of income tax exemption for workers in sector 722 was passed in Romania.

Appendix B.2 Firm Entries and Exits in the IT Sector

The income tax break may have also affected the growth of the software sector by increasing the birth rate of new firms in the sector and/or decreasing firm death rates. In this section, we provide descriptive evidence on the evolution of firm birth and death rates in the sector (722, in NACE Rev 1.1 classification) in the years before and after the introduction of the policy. The birth/entry (death/exit) rate at time t is defined as the number of firm births/entries (deaths/exits) relative to the population of active firms at the beginning of time t. Our analysis is based on sector-level data from the Eurostat Business Demography database. Figures B10 and B11 show the evolution of birth and death rates in the software sector between 2000 and 2007 for Romania and four other comparable Central and Eastern European (CEE) countries. These other countries were chosen based on data availability for the entire 2000 to 2007 period and on their similarity to Romania (i.e., they are transition economies, at similar stages of development, and with similar economic structures).

The main takeaway from Figure B10 is that the birth rate of firms in the IT sector in Romania experienced a notable increase in 2003 and 2004 relative to its pre-tax break levels.[65] No similar increase can be observed for any of the other four CEE countries. There is a two-year lag in the peak of the firm birth rate in Romania, which is consistent with a need for entrepreneurs to be reassured that the tax break was not to be reverted in the short-term. Figure B11 shows the evolution of firm death rates in the IT sector in Romania and in comparable CEE countries. While overall the death rate in Romania does not seem to be affected by the introduction of the tax break, throughout the period of analysis Romanian firm death rates remain less volatile than those in comparable countries. The combined trends in firm birth

[65] As shown in Appendix B.3, from 2003 to 2016, the stock of FDI in the IT sector of Romania has been multiplied by twenty. A sizable share of the firm entry into the IT sector is likely to be due to foreign firm entry.

and death rates lead to an increasing trend in the stock of firms in the IT sector in Romania.

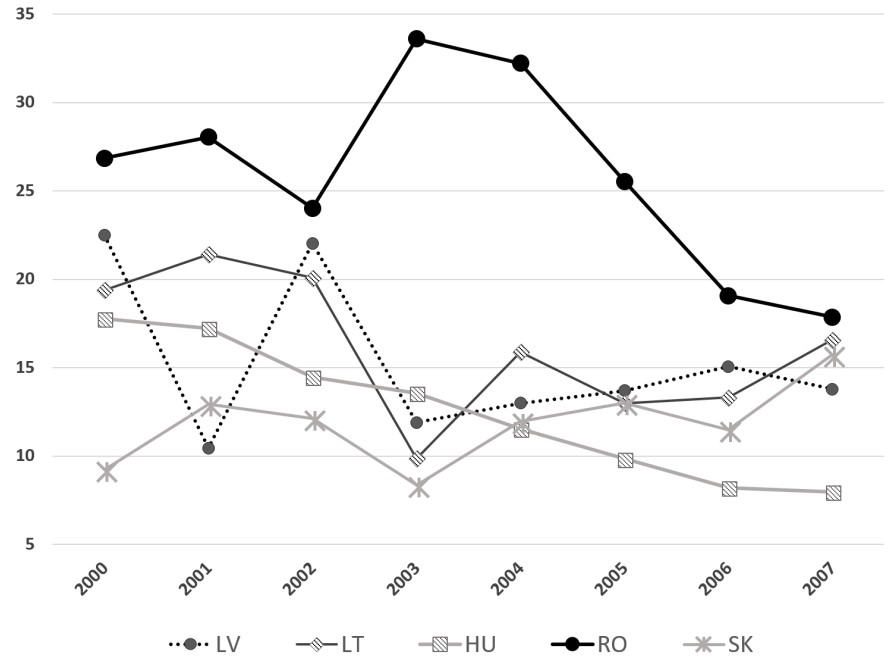

Figure B10: Entry Rates of Firms in the "Software Consultancy and Supply" Sector in Romania and Comparable Countries from Central and Eastern Europe

Notes: Data source: Sector-level data from the Eurostat Business Demography database. The entry rate at time t in a given country is defined as the number of firm entries (births) relative to the population of active firms in that country at the beginning of time t.

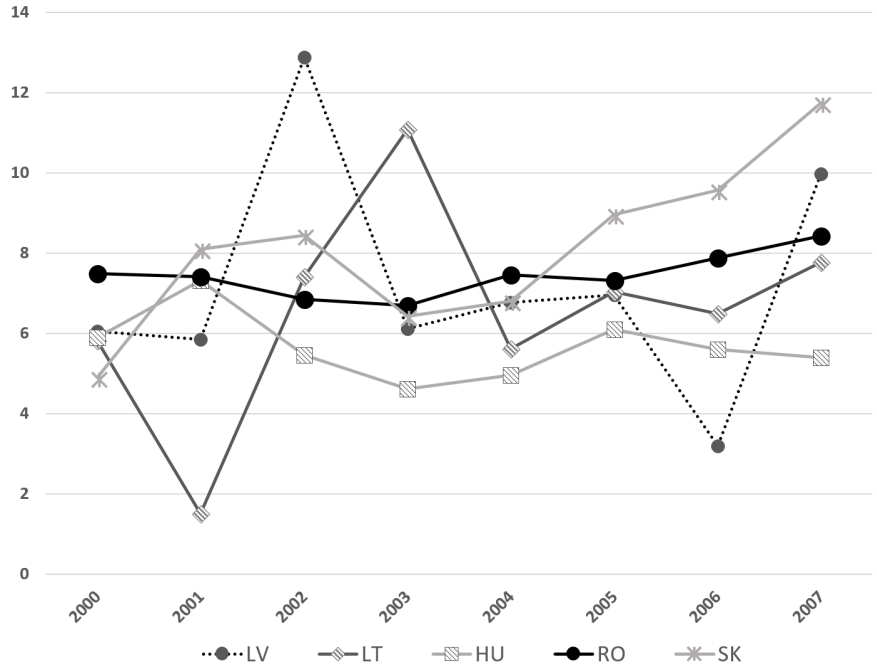

Figure B11: Death Rates of Firms in the "Software Consultancy and Supply" Sector in Romania and Comparable Countries from Central and Eastern Europe

Notes: Data source: Sector-level data from the Eurostat Business Demography database. The exit rate at time t in a given country is defined as the number of firm exits (deaths) relative to the population of active firms in that country at the beginning of time t.

Appendix B.3 Firm Entry through Foreign Direct Investment in IT

In terms of foreign ownership in the IT sector of Romania, the early 2000s were the turning point. Between 2000 and 2004, the first multinational firms decided to offshore part of their operations to Romania [Pruna and Soleanicov, 2012]. Moreover, during the same period, there were also smaller foreign companies that acquired Romanian firms, e.g., Adobe Systems Inc. who acquired InterAKT [Pruna and Soleanicov, 2012]. Ever since, Romania has continued attracting steady FDI inflows into the IT sector, to the point that in 2017 the biggest players in the IT sector of Romania were multinationals (Oracle, IBM, Ericsson, and Endava). Bitdefender is the only Romanian firm with comparable operating revenue from its operations in Romania.[66] In 2017 foreign-owned companies produced 73% of the gross revenues and hired 59% of total employees in the Software and IT Services of Romania.

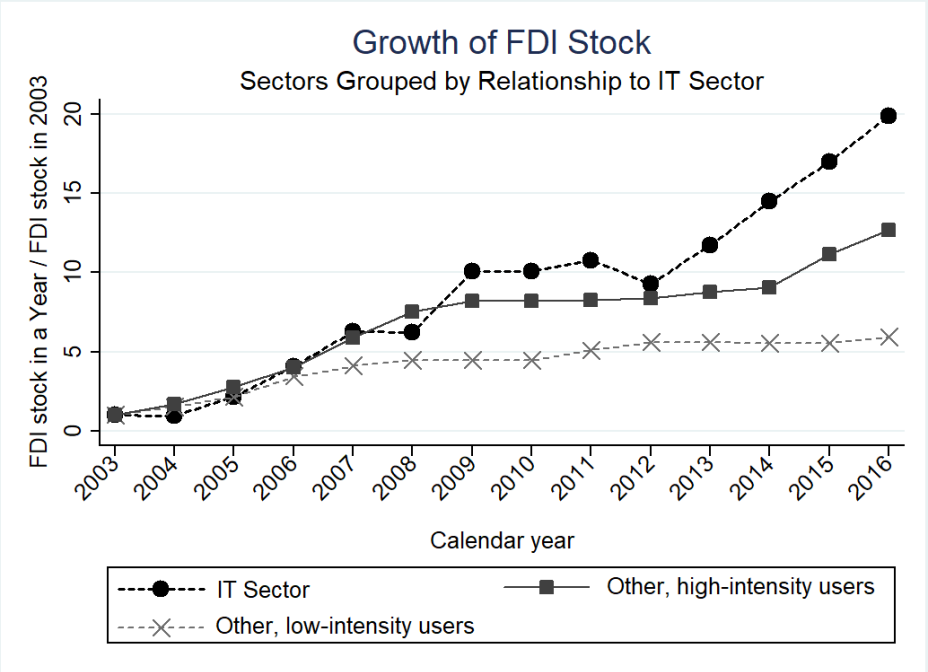

Figure B12: Growth of the Stock of FDI in Romania, by Relationship to the IT Sector

Notes: Data sources: *National Bank of Romania* (for historical sector-level FDI stocks) and *OECD* (for the Romanian input-output table for the year 2000, at the two-digit NACE Rev 1 level). This figure plots the growth of the FDI stock in Romania between 2003 and 2016 (defined as the yearly value of the FDI stock divided by the relevant FDI stock value in 2003) for three categories of sectors. The first category is the sector 72 ("Computer and related activities") itself. The second category called "Other, high-intensity users" contains all two-digit NACE Rev 1 sectors that are in the top quarter of the distribution of sector-level shares of input purchases from sector 72 in a sector's total input expenditure (except sector 72). These shares are computed based on Romania's input-output table for the year 2000. The last category of sectors called "Other, low-intensity users" contains all other sectors for which IT inputs are less important. The following are the NACE Rev 1 codes and names of the sectors in the "Other, high-intensity users" category: 50-52 (Wholesale and retail trade; repairs), 60-63 (Transport and storage), 70 (Real estate activities), 71 (Renting of machinery and equipment), 73-74 (R&D and other business activities), 75 (Public administration and defence; compulsory social security), 80 (Education), and 85 (Health and social work).

An ideal dataset to study the causal effects of the income tax break for workers in IT on Foreign Direct Investment (FDI, henceforth) into the IT sector of Romania would have had to start before 2001, include not only Romania but also comparable countries, and be disaggregated at the sector-level. Unfortunately, such a dataset was not found. For Romania, the longest sector-level FDI time series (from

[66] See December 2017 article from the *Romanian Journal*.

the National Bank of Romania) starts in 2003. With this dataset, we study the growth of the FDI stock for three categories of sectors: the IT sector itself, the group of sectors that are "high-intensity users" of IT services, and the "low-intensity users".[67] Figure B12 plots the growth of the FDI stock in Romania between 2003 and 2016 (defined as the yearly value of the FDI stock divided by the relevant FDI stock value in 2003) for these three categories of sectors. The figure shows that sectors with the highest share of IT services in their input expenditure are those experiencing the highest growth of their FDI stock since 2003. The IT sector itself is experiencing an even starker growth in its FDI stock.

The importance of foreign revenue has also been growing steadily over the years, to the point that in 2017 total foreign revenues were three times larger than domestic revenues.[68]

Appendix C Descriptive Statistics of Firm-Level Data

Appendix C.1 Amadeus Data: 1999 - 2005

Table C1: Descriptive Statistics of Firms in Treated and Comparison Sectors in Year 2000

Variable	# Obs.	Mean	Median	SD
Firms in eligible sector				
Operating Revenue	669	175.46	37.00	588.81
Number of Workers	669	34.38	4.00	646.07
Total Assets	669	92.44	16.00	300.20
Solvency Ratio	652	0.20	0.16	0.33
Firms in baseline set of comparison sectors				
Operating Revenue	953	321.92	39.00	1,631.15
Number of Workers	953	16.78	3.00	58.92
Total Assets	953	248.32	12.00	2,361.41
Solvency Ratio	926	0.20	0.13	0.33

Notes: Table C1 reports descriptive statistics for the baseline sample used to study the firm-level impact of the introduction in 2001 of the income tax break for workers in IT. The data source is Amadeus. These statistics pertain to year 2000, the year prior to the introduction of the policy of interest. The upper panel reports summary statistics for the firms in the eligible sector (722, NACE Rev 1), for the four main outcome variables of interest. The lower panel reports summary statistics for the firms that belong to the baseline set of comparison sectors: 721, 723, 724, 725, 726, 731, 732, all NACE Rev 1. The unit of measure for the operating revenue and assets variables is thousands of euros.

[67] See Section 4.2 for details on how we assign sectors into either the high- or low-usage groups.
[68] Source of last two sets of statistics: the Association of Employers in the Software and Services Industry or ANIS.

Appendix C.2 Administrative Data: 2011 - 2015

Table C2: Descriptive Statistics of Firms in Year 2011, by Firm Category

Sector eligible Exempted after 2013	Not eligible Not exempted		Eligible Not exempted		Eligible Exempted	
	Mean	SD	Mean	SD	Mean	SD
Employees	21.9	253.1	4.1	13.4	10.4	34.4
Revenue ($\times$ 1,000 RON)	8,718.9	143,729.8	890.0	4,464.1	2,100.6	11,125.0
Sales ($\times$ 1,000 RON)	8,179.2	137,021.7	653.0	3,006.2	2,015.6	11,091.8
Assets ($\times$ 1,000 RON)	11,230.0	177,337.7	1,001.0	8,221.4	1,365.6	8,433.1
Rel. Prod	0.60	0.11	0.58	0.11	0.58	0.11
Micro	0.84	0.36	0.91	0.29	0.81	0.39
Small	0.11	0.31	0.06	0.24	0.13	0.34
Medium	0.02	0.14	0.01	0.09	0.03	0.18
Large	0.03	0.17	0.02	0.15	0.02	0.15
Young	0.39	0.49	0.42	0.49	0.47	0.50

Notes: This table reports descriptive statistics for 2011 (two years before the expansion of the income tax break law) and for the three types of firms that we include in our baseline sample. The first sample includes firms in sectors comparable to the IT sector. These comparable sectors were, however, not targeted by the income tax break law. The second and third samples contain firms in the eligible sectors. The second sample contains firms with less than 5% of employees exempted from the income tax in each year of our 2011 to 2015 sample. The third sample contains firms that had less than 5% of exempted employees in 2011 and 2012, but jumped to over 20% of exempted employees after 2013. 1 RON $\approx$ 0.2 euros.

Table C2 presents descriptive statistics on firm size (measured in terms of employment, revenue, production, or assets) and relative productivity in 2011 for the three groups of firms in our baseline sample: firms in non-eligible sectors, firms in eligible sectors with less than 5% of employees exempted from the income tax throughout the entire sample period, and firms that had less than 5% of exempted employees in 2011 and 2012, but jumped to over 20% of exempted employees after 2013. Firms in non-eligible sectors are on average the largest. Firms in eligible sectors for which a large share of workers become exempted after 2013 have the second-highest average size. Last, firms in eligible sectors whose workers never became large-scale exempted have the smallest average size. In this last "eligible not exempted" category there is a larger share of micro firms (firms with less than ten employees): 91%, compared to 84% and 81%. The three types of firms do not differ in their average relative productivity. In the econometric analysis, we control for time-invariant differences in size using firm fixed effects and time-variant controls for size category, relative productivity, and age.

Table C3: Share of Firms with Exempted Employees in Different Samples

Share of firms with at least 1 exempted employee, raw dataset

Year	% of firms in ICT services	% of firms in HTKI services	% of firms in eligible sectors
2011	20%	18%	35%
2012	19%	18%	34%
2013	21%	20%	37%
2014	23%	22%	41%
2015	24%	23%	42%

Share of firms with at least 1 exempted employee, clean dataset used as main sample

Year	% of firms in ICT services	% of firms in HTKI services	% of firms in eligible sectors
2011	22%	21%	36%
2012	21%	20%	35%
2013	23%	22%	38%
2014	23%	22%	38%
2015	25%	25%	41%

Share of firms with 20% or more exempted employees, clean dataset used as main sample

Year	% of firms in ICT services	% of firms in HTKI services	% of firms in eligible sectors
2011	18%	17%	29%
2012	17%	17%	28%
2013	18%	18%	30%
2014	19%	19%	32%
2015	21%	21%	35%

Notes: Table C3 reports the share of firms in a given sample that have at least 1 exempted employee or that have that more than 20% of their employees exempted from the income tax for each year between 2011 and 2015. By construction, all firms in non-eligible sectors have less than 5% of their employees exempted from the income tax (employees who are likely to have a disability, another criterion on which this tax break can be granted). The sample of "firms in ICT services" used in the two lower panels is the baseline sample whose results are reported in Table 2. The sample of "firms in HTKI services" (HTKI stands for "high-tech knowledge-intensive") used in the two lower panels is the robustness check sample whose results are reported in Table A5, (Appendix A.1.2). The sample of "firms in eligible sectors" used in the two lower panels contains only firms in sectors whose IT workers can be eligible to the income tax break and is the robustness check sample whose results are reported in Table A6 (Appendix A.1.2).

Table C4: Predictors of Firm-level Exemption Performance

Dependent variable	1 if more than 20% exempted employees			Share of exempted employees		
Sectors of Firms	ICT services (1)	HTKI services (2)	Eligible (3)	ICT services (4)	HTKI services (5)	Eligible (6)
In an eligible sector	0.273***	0.285***		0.170***	0.178***	
	(0.007)	(0.007)		(0.005)	(0.005)	
Foreign-owned	0.153***	0.143***	0.182***	0.096***	0.089***	0.115***
	(0.022)	(0.020)	(0.027)	(0.014)	(0.013)	(0.018)
Small	0.049***	0.044***	0.077***	0.021***	0.019***	0.032***
	(0.010)	(0.009)	(0.016)	(0.005)	(0.005)	(0.009)
Medium	0.126***	0.107***	0.178***	0.054***	0.047***	0.074***
	(0.024)	(0.020)	(0.033)	(0.013)	(0.011)	(0.017)
Large	0.136***	0.113***	0.223***	0.081***	0.069***	0.127***
	(0.043)	(0.034)	(0.061)	(0.023)	(0.019)	(0.033)
Young	0.004	0.004	0.005	0.006	0.006	0.010
	(0.006)	(0.006)	(0.010)	(0.004)	(0.004)	(0.007)
Rel. prod.	0.028	0.022	-0.001	0.001	-0.001	-0.032
	(0.034)	(0.032)	(0.058)	(0.026)	(0.024)	(0.045)
County-level % exemption	0.270***	0.273***	0.444***	0.182***	0.183***	0.299***
	(0.032)	(0.032)	(0.053)	(0.022)	(0.022)	(0.037)
d_{2012}	0.001	0.003	0.003	0.001	0.002	0.003
	(0.004)	(0.004)	(0.007)	(0.002)	(0.002)	(0.004)
d_{2013}	0.011**	0.012***	0.018**	0.007**	0.007**	0.012**
	(0.004)	(0.004)	(0.008)	(0.003)	(0.003)	(0.005)
d_{2014}	0.016***	0.016***	0.026***	0.011***	0.011***	0.019***
	(0.005)	(0.005)	(0.008)	(0.003)	(0.003)	(0.005)
d_{2015}	0.022***	0.022***	0.035***	0.009***	0.009***	0.014**
	(0.005)	(0.005)	(0.009)	(0.003)	(0.003)	(0.006)
Adjusted R^2	0.186	0.189	0.067	0.158	0.162	0.050
# Observations	26,514	27,414	16,234	26,514	27,414	16,234
# Firms	5,728	5,939	3,526	5,728	5,939	3,526

Notes: Table C4 reports the results of OLS regressions run on three samples used in the analysis in Section 3.2 for 2011 to 2015. Columns (1)-(3) 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 (4)-(6) use as the dependent variable the share of workers of firm i who are exempted from the income tax in year t. The samples used in this table are different from the baseline sample in Table 2 and the samples in Tables A5 and A6 in that the former keep *all* firms in ICT service sectors, in HTKI (high-tech knowledge-intensive) service sectors, and in eligible sectors, whereas the latter bring additional restrictions on the firm-level share of workforce exemption before and after 2013. The dependent variables and samples used in this table are the same as those used in Table A2, which explores the first stage effects of the 2013 reform on the firm-level share of workforce exemption from the income tax. The explanatory variables of productivity, size, and age are contemporaneous. Data on foreign ownership in 2016 comes from Amadeus. We assume that if a firm is foreign-owned in 2016 it has been foreign-owned throughout 2011 to 2015. Columns (3) and (6) include only firms whose sector is targeted by the income tax break. The reference category contains firms that operate in non-eligible sectors, are domestically-owned, have a micro size, and are older than five years. The reference year in 2011. We find that firms that are foreign-owned, that are larger, and/or that operate in counties (*județe*) with a higher-level of workforce exemption are more likely to have a higher share of workforce exemption from the income tax. We also show that after 2013, firm-level shares of expansion have significantly increased with respect to their 2011 level.

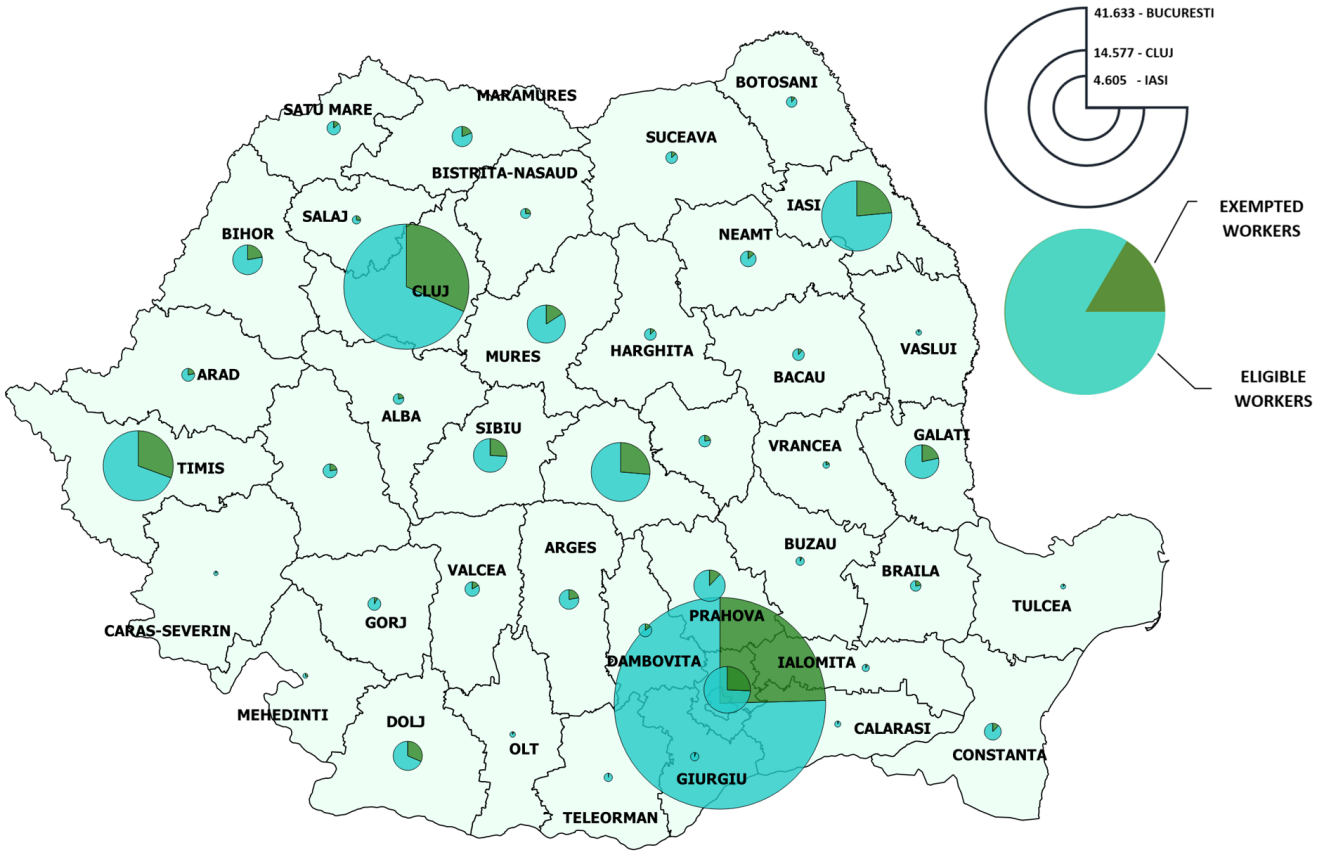

Figure C13: County-Level Number of Employees in the Eligible IT Sector in 2015 and the Share of these Employees Actually Exempted from the Income Tax

Notes: The legend reports the total number of employees in the eligible IT sectors in 2015 in the three counties with the largest number of such employees, i.e., București, Cluj, and Iași. This map uses the sample from the 2011 to 2015 administrative dataset kept for the econometric analysis, as described in Section 3.2.

Figure C13 is informative on the spatial distribution in 2015 of the number of workers in the sectors targeted by the income tax exemption and the share of these workers who are actually benefiting from this exemption. The unit of analysis is a county (*judeţ*). This map highlights the striking spatial concentration of the IT industry in Romania's leading counties, București (the capital), Cluj, Iași, and Timiș. While Bucureşti stands out as a giant in terms of its agglomeration of workers in IT, the second-placed county (Cluj) clusters only a third of the same type of employees. The third- and fourth-placed counties (Iaşi and Timiş) house a tenth of the number of workers in Bucureşti.

This map also draws attention to heterogeneity in the share of employees in the IT sector who actually benefit from the tax exemption. This share reaches a peak in Timiș (around 46%) and a low in Teleorman (under 5%). This heterogeneity is likely to reflect differences across workers and firms in their ability to meet the requirements of the tax break (firms in less developed counties are more likely to hire workers with a profile that does not meet the educational criteria, such firms are also less likely to earn more than 10,000 U.S. dollars per eligible employee etc.). Discrepancies are also likely to reflect differential abilities across firms to fill in the necessary paperwork for a worker to be granted the tax break.

From Table C4 we learn that IT firms located in counties with a higher share of exemption of the total workforce in IT are more likely to have a higher share of workforce exemption themselves. This

finding points to potentially localized knowledge spillovers on the necessary procedures to benefit from the tax break and the importance of a strong local pool of skilled programmers.

Another map can be drawn for the number of IT firms in each county and the share of these firms actually employing exempted employees. Such a map (available upon request) shows a slightly more equal distribution of the number of firms, which captures the tendency of the median firm in leading counties to hire significantly more workers than the median firm in lagging counties.

The absolute number of firms in the eligible sectors in a given county is to a certain extent capturing the size of the population of that county. For comparability across counties, we compute the county-level number of workers in the eligible sectors per 100,000 inhabitants. While Bucureşti remains the leader in the country with 2,257 workers in the eligible sectors per 100,000 inhabitants, Cluj continues to stand out in a comparable position with 2,078 such workers per 100,000 inhabitants. The following three performers are Timiş, Braşov, and Iaşi, each between 580 to 660 workers in IT per 100,000 inhabitants. The lowest-ranked 30 counties are home to less than 100 workers in the eligible sectors per 100,000 inhabitants.

Appendix D Data Construction for Synthetic Control

Appendix D.1 Eurostat Data

Downloading the Data. The first step in this process was to download data from Eurostat, using the Structural Business Statistics: Annual Detailed Enterprise Statistics tables. The data were available to download in batches, after having selected the sectors/industries, countries, years, and variables of interest. These selections are detailed below:

- **Sectors**: Data were available for download at the one-, two-, three-, and four-digit sector/industry level. For the main analysis, we selected and downloaded data at the two-digit sector level due to high levels of missing data at the three- and four-digit level. According to the Rev 1.1 sector classification, there were 59 possible two-digit sectors for which we might like to have data. Our process enabled us to find corresponding data as early as 1999 for 48 of these sectors. Missing sectors include A01 ("Agriculture, hunting and related service activities"), A02 ("Forestry, logging and related service activities"), B05 ("Fishing, fish farming and related service activities"), L75 ("Public administration and defence; compulsory social security"), M80 ("Education"), N85 ("Health and social work"), O90 ("Sewage and refuse disposal, sanitation and similar activities"), O91 ("Activities of membership organizations n.e.c."), O92 ("Recreational, cultural and sporting activities"), O93 ("Other service activities"), and P95 ("Activities of households as employers of domestic staff"). For each two-digit sector, data were downloaded in two pieces: firstly as per the Rev 1.1 sector classification and secondly as per the Rev 2 sector classification. Since the transition from Rev 1.1 to Rev 2 in the year 2007 was not one-to-one, fractions of two-digit sectors were carefully linked across the two classifications according to the degree of mutually overlapping three-digit sectors between any pair of two-digit sectors. More detail on the construction of this crosswalk follows in the "Creation of Crosswalk between Rev 1.1 and Rev 2 Sector Codes"

section.

- **Countries**: Data were downloaded for the following fourteen countries of interest: Bulgaria, Cyprus, the Czech Republic, Estonia, Hungary, Ireland, Latvia, Lithuania, Malta, Poland, Portugal, Romania, the Slovak Republic, and Slovenia. Data were widely available for most countries, with frequent missing values for Cyprus and Malta across the sectors, variables, and years of interest.

- **Variables**: Data were downloaded for the following five variables: Gross revenue (Turnover or gross premiums written) - million euro (V12110), Production value - million euro (V12120), Employees - number (V16130), Gross investment in tangible goods - million euro (V15110), and Enterprises - number (V11110).

- **Years**: For each sector, country, and year, relevant data were downloaded for as many years as possible, ranging from 1995 to 2016. There was low data availability prior to 1999 and after 2015 across the set of countries, industries, and variables of interest.

After having downloaded all available data for five variables from 1995 to 2016 across 14 countries and 48 sectors, there was a further selection of years, countries, and industries to include in the final analysis based on having a sufficiently high level of non-missing data to complete the analysis. These selections are detailed in the "Selection of Years, Countries, and Industries for Final Analysis" section.

Creation of Crosswalk between Rev 1.1 and Rev 2 Sector Codes. The SCM analysis compares pre-trends from before and after 2001, between Romania and synthetic Romania, for sectors with high- versus low-intensity of IT use. Since the Rev 1.1 sector classification was in effect for the treatment period of interest, our analysis seeks to classify sectors as either high- or low-intensity of IT use according to their Rev 1.1 sector classification. For this reason, the data downloaded as per Rev 2 sector classifications (for years 2007 onwards) needed to be stitched together with the data downloaded as per Rev 1.1 sector classifications (for years prior to 2007) according to a correspondence or crosswalk between the two sector classifications.

An extensive correspondence detailing equivalence between two-, three-, or four-digit sectors (as applicable and as dictated by the transition) between Rev 1.1 and Rev 2 was obtained from Eurostat and used for constructing a crosswalk that could link sectors across sector classifications at the two-digit level. The revision between Rev 1.1 and Rev 2 was in most cases not a one-to-one transition. For six of the 59 Rev 1.1 sectors, one two-digit Rev 2 sector corresponded to a single two-digit Rev 1.1 sector. Six of the 59 two-digit Rev 1.1 sectors were either split into two or more corresponding two-digit Rev 2 sectors or were part of a merge into one corresponding two-digit Rev 2 sector. The remaining 47 of 59 sectors were combinations of splits and merges, that is, some portion of the two-digit sectors as per Rev 1.1 were split into two or more sectors as per Rev 2, where they were merged with other portions of two-digit Rev 1.1 sectors that were similarly split among two or more Rev 2 sectors.

Two primary methods were used to create a crosswalk to link some fraction of each two-digit Rev 2 sector to a corresponding two-digit Rev 1.1 sector. In the first method ("gross revenue and count"), average gross revenue across all 14 countries of interest was used as a proxy for the size of each Rev 1.1

sector at the three-digit level. Data for this portion of the task was similarly downloaded from Eurostat, using the Structural Business Statistics: Annual Detailed Enterprise Statistics tables. This information on sector size was used to calculate what portion of a two-digit Rev 2 sector stemmed from distinct three-digit Rev 1.1 sectors, and in turn, to which two-digit Rev 1.1 sectors these three-digit sectors corresponded. Mechanically, this process involved calculating the average gross revenue in 2007 across all 14 countries of interest for each three-digit Rev 1.1 sector included in the correspondence table (in cases where the correspondence table indicated linkages at the four-digit level, these were aggregated up to the three-digit level so that all comparisons at this stage were made at the three-digit level). The size of corresponding three-digit sectors common to a pair of two-digit Rev 1.1 and Rev 2 sectors relative to the total size of the Rev 2 sector were used to assign fractions of that two-digit Rev 2 sector to corresponding Rev 1.1 sectors. To illustrate, suppose a particular two-digit Rev 2 sector was composed of five three-digit Rev 2 sectors that corresponded to five three-digit Rev 1.1 sectors, three of which corresponded to one two-digit Rev 1.1 sector (totaling average gross revenue of 200) and 2 of which corresponded to a second two-digit Rev 1.1 sector (totaling average gross revenue of 300). In this case, 200/500=40% of each variable value from the Rev 2 sector would be allocated to the first Rev 1.1 sector, and 300/500=60% of each variable value from the Rev 2 sector would be allocated to the second Rev 1.1 sector.

The second method, ("count only") used a similar approach. However, instead of calculating the fraction of each two-digit Rev 2 code to assign to each two-digit Rev 1.1 code using an estimate of the size of corresponding three-digit sectors, this method used the count of three-digit sectors common to a pair of two-digit Rev 1.1 and Rev 2 sectors relative to the total number of three-digit sectors that make up the two-digit Rev 2 sector to assign the fractions. In the same example as above, 3/5=60% of each variable value from the two-digit Rev 2 sector would be allocated to the first two-digit Rev 1.1 sector, and 2/5=40% of each variable value would be allocated to the second Rev 1.1 sector.

For the first method ("gross revenue and count"), in cases where gross revenue was not available in 2007 in order to make the necessary estimation of sector size, fractions were allocated based on the "count only" method.

The resulting versions of the crosswalk between the two sector classifications are similar, and there is no clear test that can be used to assess which of the two methods provides a more accurate representation of each Rev 1.1-equivalent sector from 1995 to 2016. Visual inspection of trend lines (focusing in particular on the gross revenue variable for Romania over time) seems to suggest that the "count only" method leads to a smoother correspondence across Rev 1.1 and Rev 2. While there remain some jumps in the times series according to both correspondences (which could reflect the similarity in timing of the transition from Rev 1.1 to Rev 2 along with the global financial crisis), the "count only" method appears to minimize these jumps. For this reason, the "gross revenue and count" method was disregarded in favor of moving forward with the "count only" method.

Data Cleaning. After downloading data for as many countries, variables, years, and sectors as possible and constructing a crosswalk to link the data across Rev 1.1 and Rev 2 sector classifications, the next step was to combine the data, clean the data, and select years, countries, and industries to include in the final analysis.

Data that was downloaded according to the Rev 2 sector classification (for 2007 onward) was split

fractionally to corresponding Rev 1.1 sectors as per the crosswalk and combined with the data that was downloaded according to the Rev 1.1 sector classification (for prior to 2007). In cases where data was available under both classifications for overlapping years (as was often the case for 2005, 2006, and 2007), the data from the Rev 1.1 classification was prioritized, then filled in with data from the Rev 2 classification when the Rev 1.1 data was missing.

Selection of Years, Countries, and Industries for Final Analysis. Years, countries, and industries to include in the final analysis were selected to maximize availability of data and minimize the incidence of missing observations in the resulting panel.

Less than 70% of the data across all countries and all industries were available for years prior to 1999 and after 2015, so only the years from 1999 to 2015 inclusive were retained for the analysis sample. (Note that the 70% cutoff applies to all countries and all sectors; this percentage would be higher once restricting to the final set of countries and sectors retained for analysis.)

Similarly, Malta and Cyprus were dropped from the analysis sample due to low data availability. The remaining 12 countries retained featured relatively high data availability: at least 75% of the desired data were non-missing for the 12 countries retained across all years, and at least 85% of the desired data were non-missing when restricting attention to the years 1999-2015. (Again, these percentages are for data across all sectors and would be larger once additionally restricting to the final sectors chosen for their sufficient level of data availability.)

Finally, because of the importance of estimating pre-trends for Romania, sectors were dropped if data were missing for any of the seven variables in either 1999 or 2000 for Romania. This requirement eliminated 14 sectors from the analysis sample (C10, C11, C12, C13, C14, G51, G52, I60, I61, I62, I64, J65, J66, and K67). An additional four sectors were dropped for reasons of low data availability: considering only 1999-2015 for all seven variables and the 12 retained countries, these were four industries for which more than 10% of the desired data was missing (D15, D16, D19, and D23).

After restricting the analysis sample to the years 1999-2015, dropping Cyprus and Malta, and restricting to the 48 sectors where data was available for all seven variables for Romania in 1999 and 2000 and where there was a sufficiently high level of non-missing data, remaining missing values were imputed as follows. In cases where there was a missing observation for a particular variable within a time series, the missing observation's value was estimated by a simple average of the value for the most recently available and next available year, including for cases where the gaps were greater than one year. For cases where 1999 and potentially several consecutive years were missing, all of those observations were set to match the value for the earliest available year. Similarly, in the rare situation where there were potentially several consecutive years missing data ending in 2015, all of those variable's values were set to match that variable's value for the latest available year.

The ultimate level of imputation of missing values required was minimal, ranging from 0.8% to 1.9% of desired year by country by sector observations across the seven variables.

Appendix D.2 UN Comtrade Data

Downloading the Data. For for the exports of goods and services, the data was downloaded from the UN Comtrade official website. We downloaded the data for the same set of twelve countries that

constitute the pool of donor countries for the SCM exercise using Eurostat data: Bulgaria, the Czech Republic, Estonia, Hungary, Ireland, Latvia, Lithuania, Poland, Portugal, Romania, the Slovak Republic, and Slovenia.

For goods exports, we downloaded the data disaggregated by two-digit SITC Rev 1 commodity codes. For service exports, we downloaded the data disaggregated by two-digit (X.X) EBOPS2002 codes (except Financial Services which is just classified as 6).[69] The variable of interest was Trade Value (US$) for Trade Flow=Exports.

Data Cleaning and Selection of Years and Codes for Final Analysis. The dataset for goods exports had the following features:

- There was a high rate of missing data for BG in 1995 and HU, PT in 2015
- We had to remove some codes to make sure that the codes entering the aggregates (high- versus low-intensity of IT inputs use) are consistent across countries:
  - 34. Missing: BG (1995, 1998-2000, 2003), LV (1998), RO (1996-1997), SK (2005)
  - 35. Missing: BG (1995-2006), HU (1995-1996, 2001-2002), IE (2010-2011), PT (1997-1998, 2015), RO (1996), SK (1995-1996, 2005)
  - 41. Missing: RO (2003)
  - 52. Missing: HU (1998-1999, 2006, 2009, 2015), IE (2001, 2003, 2005), LV (2009), SK (2007)
  - 93. Missing: CZ (1996-2001, 2003) , EE (1995-1998, 2002), HU (1995-1996, 2015), IE (1996), PL (1996-1997), RO (2006), SK (1995-2003)
  - 95. Missing: BG (1995, 1997-2001, 2003, 2007-2013), HU (1995, 2003, 2009, 2015), IE (1997, 1998, 2000), RO (1998-2005), SI (1996, 1999)

Removing the above sectors results in a (close to) consistent panel for 1996-2014. Note that in 2008, SK does not have sector 12.

The download of the service export data allowed us to notice that the data is available only from 2000 to 2016. Even during these years, most countries present a lot of missing values across codes. In particular, for Romania the treated NACE Rev 1 sector (72) starts reporting data only in 2005, not allowing us to study the service exports of this sector. The resulting (unbalanced) panel does not allow for an SCM, as we only observe one year before the event in 2001. Moreover, the large extent of missing values delivers noisy estimates of treatment effects.

Creation of Crosswalk between EBOPS2002 Codes and NACE Rev 1.1 Sector Codes. Our categories of high- and low-intensity (of IT usage) sectors were constructed based on Romania's input-output Table for 2000, which is at the two-digit NACE Rev 1 sector level. Hence, we had to construct a mapping between two-digit SITC Rev. 1 commodity codes and two-digit (X.X) EBOPS2002 service codes and the two-digit NACE Rev 1 codes. After commodity or service codes were mapped to NACE Rev 1 codes, we aggregated export values at the NACE Rev 1 level.

[69] EBOPS2002 includes the standard items, memorandum items, and supplementary items. Only standard items were used. For codebook, visit link.

Appendix E Context on Policies Relevant to the IT Sector in Romania

Appendix E.1 Details on Romania's Income Tax Break for Workers in IT

Table E1: 2001 Income Tax Break and Its Subsequent Amendments: Eligibility Criteria for the Income Tax Break (1/2)

Order Number	4079/ 268/ 1480/ 2001	661/ 444/ 2196/ 2001	250/ 189/ 748/ 2004	539/ 225/ 1479/ 2013	217/ 4172/ 1348/ 835/ 2015	872/ 5932/ 2284/ 2903/ 2016
Effective date	08/05/2001	11/22/2001	06/29/2004	09/10/2013	07/30/2015	01/09/2017
CAEN Code	7220 (Rev 1)	7220 (Rev 1)	**7221, 7222 (Rev 1.1)**	**5821, 5829, 6201, 6202, 6209 (Rev 2)**	5821, 5829, 6201, 6202, 6209 (Rev 2)	5821, 5829, 6201, 6202, 6209 (Rev 2)
Occupations	Analyst, Programmer, Computer Systems Designer, Engineer or Programmer of IT Systems, Database Manager, Software Engineer and Manager of IT Projects	Analyst, Programmer, Computer Systems Designer, Engineer or Programmer of IT Systems, Database Manager, Software Engineer and Manager of IT Projects	Analyst, Programmer, Computer Systems Designer, Engineer or Programmer of IT Systems, Database Manager, Software Engineer and Manager of IT Projects	Analyst, Programmer, Computer Systems Designer, Engineer or Programmer of IT Systems, Database Manager, Software Engineer and Manager of IT Projects	Analyst, Programmer, Computer Systems Designer, Engineer or Programmer of IT Systems, Database Manager, Software Engineer and Manager of IT Projects **Computer System Programmer**	Analyst, Programmer, Computer Systems Designer, Engineer or Programmer of IT Systems, Database Manager, Software Engineer and Manager of IT Projects **Computer System Programmer**
Unit within the Firm	Unit Specialized in IT	Unit Specialized in IT	Unit Specialized in IT	Unit Specialized in IT	Unit Specialized in IT	Unit Specialized in IT

Extra Eligibility Criteria for the Income Tax Break, in Addition to Those from Table E1 (2/2)

Order Number: 4079/ 268/ 1480/ 2001
Eligible Major during Higher Education for Exempted Worker: Automation, Computers, Computer Science, Cybernetics, Mathematics, Electronics.
Minimum Annual Revenue from Software Development: Annual income of at least 10,000 U.S. dollars per employee benefiting from the income tax break.
Balance for Software Development Income: A balance is required, in which the income from software development needs to be explicitly reported.

Order Number: 661/444/2196/2001
Eligible Major during Higher Education for Exempted Worker: Automation and Industrial Computer Science; Computers, Electrical Engineering and Computers; Electronics; Applied Electronics, Electronics and Telecommunications, Communications; Mathematics, Mathematical Computer Science; Computer Science, Computer Science and Economics, Applied Computer Science, Cybernetics and Computer Science and Economics, Cybernetics and Economic Prediction, Accounting and Computer Science and Management.
Minimum Annual Revenue from Software Development: Annual income, in preceding year, of at least 10,000 U.S. dollars per employee benefiting in a given year from the income tax break.
Balance for Software Development Income: A balance is required, in which the income from software development needs to be explicitly reported.

Order Number: 250/189/748/2004
Eligible Major during Higher Education for Exempted Worker: Automation and Industrial Computer Science; Computers, Electrical Engineering and Computers; Electronics; Applied Electronics, Electronics and Telecommunications, Communications; Mathematics, Mathematical Computer Science; Computer Science, Computer Science and Economics, Applied Computer Science, Cybernetics and Computer Science and Economics, Cybernetics and Economic Prediction, Accounting and Computer Science and Management.
Minimum Annual Revenue from Software Development: Annual income, in preceding year, of at least 10,000 U.S. dollars per employee benefiting in a given year from the income tax break.
Balance for Software Development Income: A balance is required, in which the income from software development needs to be explicitly reported.

Order Number: 539/225/1479/2013
Eligible Major during Higher Education for Exempted Worker: Automation and Industrial Computer Science; Computers, Electrical Engineering and Computers; Electronics; Applied Electronics, Electronics and Telecommunications, Communications; Mathematics, Mathematical Computer Science; Computer Science, Computer Science and Economics, Applied Computer Science, Cybernetics and Computer Science and Economics, Cybernetics and Economic Prediction, Accounting and Computer

Science and Management. **Newly eligible majors:** Industrial Computer Science, Applied Computer Science in Electrical Engineering, Applied Computer Science in Material Engineering, Mathematics and Applied Computer Science in Engineering; Cybernetics and Economics; Physics and Computer Science; Chemistry and Computer Science; Automation and Applied Computer Science, Equipment for Modeling, Simulation and Computerized Warfare, Engineering of Multimedia Systems; Technologies and Telecommunication Systems, Remote Controls and Electronics in Transportation; Transmissions and Military Electronic Equipment.
Minimum Annual Revenue from Software Development: Annual income, in preceding year, of at least 10,000 U.S. dollars per employee benefiting in a given year from the income tax break.
Balance for Software Development Income: A balance is required, in which the income from software development needs to be explicitly reported.

Order Number: 217/4172/1348/835/2015 and 872/5932/2284/2903/2016
Eligible Major during Higher Education for Exempted Worker: A diploma issued after a form of higher education, irrespective of major.
Minimum Annual Revenue from Software Development: Annual income, in preceding year, of at least 10,000 U.S. dollars per employee benefiting in a given year from the income tax break.
Balance for Software Development Income: A balance is required, in which the income from software development needs to be explicitly reported. **New:** New firms or firms undergoing a restructuring during that fiscal year are exempted from this requirement.

Table E2: Correspondence Table between NACE Sector Codes Rev 1.1 and Rev 2

NACE Rev 1.1	Description Rev 1.1	NACE Rev 2	Description Rev. 2	Comments
7221	Publishing of software	5821	Publishing of computer games	Publishing of computer games
7221	Publishing of software	5829	Other software publishing	All software publishing, except computer games publishing
7221	Publishing of software	6201	Computer programming activities	Software programming
7222	Other software consultancy and supply	6201	Computer programming activities	Includes: Analysis, design and programming of systems ready to use: development, production, supply and documentation of made-to-order software based on orders from specific users writing of programs following directives of the user web page design
7222	Other software consultancy and supply	6202	Computer consultancy activities	Analysis, design and programming of systems ready to use: - analysis of the user's needs and problems, consultancy on the best solution
7222	Other software consultancy and supply	6209	Other information technology and computer service activities	Software installation services

Notes: Source Eurostat.

Table E3: Correspondence Table between NACE Sector Codes Rev 2 and Rev 1.1

NACE Rev 2	Description Rev 2	NACE Rev 1.1	Description Rev 1.1	Comments
5821	Publishing of computer games	7221	Publishing of software	Publishing of computer games
5821	Publishing of computer games	724	Database activities	On-line computer games publishing
5829	Other software publishing	7221	Publishing of software	All software publishing, except computer games publishing
5829	Other software publishing	724	Database activities	All on-line software publishing, except computer games on-line publishing
6201	Computer programming activities	7221	Publishing of software	Software programming
6201	Computer programming activities	7222	Other software consultancy and supply	Includes: Analysis, design and programming of systems ready to use: - development, production, supply and documentation of made-to-order software based on orders from specific users - writing of programs following directives of the user - web page design
6201	Computer programming activities	724	Database activities	Designing of structure and content of database
6202	Computer consultancy activities	721	Hardware consultancy	All
6202	Computer consultancy activities	7222	Other software consultancy and supply	Analysis, design and programming of systems ready to use: - analysis of the user's needs and problems, consultancy on the best solution
6209	Other information technology and computer service activities	3002	Manufacture of computers and other information processing equipment	Installation of personal computers and peripheral equipment
6209	Other information technology and computer service activities	7222	Other software consultancy and supply	Software installation services
6209	Other information technology and computer service activities	726	Other computer related activities	NACE 1.1 class 72.60 was an "empty class".

Notes: Source Eurostat.

Appendix E.2 Other Policies Relevant for the IT Sector

Appendix E.2.1 State Aid Program

Between 2011 and 2016, several state aid programs that supported job creation and investment were implemented.[70] The most relevant for the sectors studied was the program created by Government Decision 797/2012. It supported large investments in new technologies with an IT component and job creation of at least of 200 new jobs. While firms in most manufacturing, energy and service sectors were eligible, mainly firms in high-tech knowledge intensive sectors benefited from it. We drop from our sample of analysis firms that have benefited from such State Aid. The program created through Government decision 332/2014, which aimed to support large investments, job creation, and regional development, also benefited several firms in high-tech knowledge-intensive services. Firms in high-tech knowledge-intensive sectors were also eligible for several smaller programs supporting SMEs and start-ups, such as "Start-up Nation," but these programs were smaller and less likely to affect major investments or job creation (e.g., "Start-up Nation" had an upper limit of approximately 44,000 euros).

Appendix E.2.2 Other Tax Exemptions

Programmers are not the only category of workers exempted from the personal income tax in Romania. Two other categories of workers exempted from the income tax could be employed by IT firms, without being programmers: workers with serious disabilities[71] and since 2016 workers in research and development (defined broadly, with no requirement to work in software development).[72] For companies with at least 50 employees, it was compulsory that at least 4% of their workers have disabilities. When firms could not comply, they had to pay a given amount to support the inclusion of people with disabilities.

Despite these other exemptions, the vast majority of exempted employees in eligible sectors (5821, 5829, 6201, 6202 and 6209) were exempted due to the tax break for workers in IT. In October 2017,[73] workers benefiting from the law under study in this article represented 96% of the exempted employees in eligible sectors.

[70] Government decisions 797/2012, 322/2014 and 807/2014.
[71] See Law 448/2006 for details.
[72] See Order 4947/899/2018/1840/906/2016 published in September 2016 for details.
[73] In 2017 we can observe the reason of the income tax exemption, which we cannot observe in previous years.