Papers in Empirical Innovation Economics

Inaugural-Dissertation
zur Erlangung des Grades

Doctor oeconomiae publicae (Dr. oec. publ.)

an der Volkswirtschaftlichen Fakultät
an der Ludwig-Maximilians-Universität München

2019

vorgelegt von

Tobias Lohse

Papers in Empirical Innovation Economics

Inaugural-Dissertation
zur Erlangung des Grades

Doctor oeconomiae publicae (Dr. oec. publ.)

an der Volkswirtschaftlichen Fakultät
an der Ludwig-Maximilians-Universität München

2019

vorgelegt von

Tobias Lohse

Referent:	Prof. Dr. Oliver Falck
Korreferent:	Prof. Dr. Stephan Heblich
Promotionsabschlussberatung:	24. Juli 2019

Datum der mündlichen Prüfung: 11. Juli 2019

Namen der Berichterstatter: Oliver Falck, Stephan Heblich, Monika Schnitzer *für blühende Landschaften*

Acknowledgements

Over the course of the past years, I have enjoyed the tremendous support of many people. All of them deserve being acknowledged, but I can only thank a selection of them here.

If it was not for my supervisor Oliver Falck, I would not have started writing a doctoral thesis. I am very grateful for the confidence he placed in me and the invaluable guidance he provided. I cannot overstate how much I appreciate that, independent of his own schedule, he always took the time to offer his advice (and a coffee) whenever I knocked on his door. I would also like to thank Stephan Heblich for his academic support from very early on. The indefatigableness with which he approaches new and old research questions was and continues to be inspiring. I further thank Monika Schnitzer for her ever-astute reviews of the work produced at the ifo Institute and for completing my dissertation committee.

I am very grateful to Andreas Mazat, with whom I had the great pleasure to share an office for more than four years. While his contributions as a coauthor and colleague were instrumental for many parts of this dissertation, his friendship turned this endeavour into something joyful and fun. I would also like to thank Johannes Koenen. Our discussions and collaboration have not only shaped my academic thinking as an economist, but also profoundly enriched my life in many other ways. I have greatly enjoyed discussing many aspects in the field of economics, especially as they relate to the development of disadvantaged regions, with Balázs Spéder, who was never short of encouragement and always provided an open ear when needed. I further thank Valentin Lindlacher, Robert Lehmann, Simone Schüller, Michael Ebnet, Udo Ludwig, Barbara Grave, and many other colleagues and friends, both present and past, for their great support throughout the years.

Finally, yet foremost, I thank my family and my girlfriend Anne. Their unconditional support, encouragement, and patience went far beyond what I could have asked for but were the privilege I needed.

Thank you. Sincerely.

Tobias Lohse, March 2019

Contents

Preface

Over the past century, it has become a political paradigm that public interventions can promote economic growth in disadvantaged areas and pave the way for regional cohesion. Already during the aftermath of the Great Depression, the federal government of the United States initiated a large-scale infrastructure investment programme to support the economic development of the Tennessee valley [Kline and Moretti, 2014*a*]. In the second half of the 20th century, many European countries followed suit and established ambitious place-based policies of their own. Among the most prominent examples are the *Zonenrandgebiet*-policy in Germany [von Ehrlich and Seidel, 2018] and the Regional Selective Assistance programme in the United Kingdom [Criscuolo et al., 2018].

Because regional prosperity continues to differ vastly [Rosés and Wolf, 2018], spatially targeted initiatives have become an integral part of the policy toolkit in most developed countries. In the United States, the federal government spends more than 15 billion Dollars per year on programmes explicitly targeting specific regions [United States Government Accountability Office, 2012]. The regional policies conducted by the EU account for annual investments of approximately 50 billion Euros, that is, a third of its entire budget.[1]

Despite the long-standing interest of policy makers and economists in regional development programmes, empirical tests of whether place-based policies indeed promote economic growth have often faced substantial difficulties. Beyond the limited availability of data, the absence of suitable evaluation methods presented a pivotal challenge for decades. This only changed during the early 1990s when seminal advances in the econometrics of programme evaluation provided a framework for systematic investigations of *causal* relationships in the social sciences [Imbens and Wooldridge, 2009]. In combination with improved data availability, these methodological developments gave rise to a burgeoning literature in the field of regional and urban economics that studies the effectiveness of targeted infrastructure investment programmes and other types of place-based policies.

[1] The EU provides funding for regional policies through the European Regional Development Fund and the Cohesion Fund (`https://ec.europa.eu/regional_policy/en/policy/what/investment-policy/`, last accessed: 16 March 2019).

Often focussing on labour market outcomes, these evaluations show that only some place-based policies successfully create local jobs whereas others failed to boost local employment.[2] Depending on the programme considered, additionally created jobs come at varying costs [Busso et al., 2013; Criscuolo et al., 2018] and do not always persist beyond the end of the regional support programme [Becker et al., 2018]. Furthermore, even if targeted places experience a sustained increase in local employment, these improvements may come at the expense of other areas [Kline and Moretti, 2014*a*].

While these studies have made significant contributions to an improved understanding of the labour market effects of place-based interventions, the evidence base on other aspects addressed by these policies continuous to be a lot more fragmental. For example, it remains an open question whether positive labour market effects induced by a specific place-based intervention also translate into improved living conditions for the residents. As economic support for regions is typically motivated by the normative objective to help specific groups of people rather than specific geographic areas, answering this question is crucial for many place-based interventions.

Studying this question theoretically, Glaeser and Gottlieb [2008], Kline [2010] and Kline and Moretti [2014*a*] introduce the idea of place-based policies in the framework of spatial equilibrium models.[3] They show that subsidies confined to specific areas only feed through to the local population if housing supply is elastic and the mobility of workers (or firms) across regions is limited. However, if housing is scarce and mobility sufficiently cheap, the benefits of the subsidies will be entirely absorbed by increases in local rents [Kline and Moretti, 2014*b*].

Beyond pursuing equity-based motives, spatially targeted interventions may also seek to counteract localised market failures. For instance, a lack of regional economic capacity can aggravate the public goods problem associated with the provision of suitable infrastructures (e.g., because low levels of local tax revenues prohibit necessary infrastructure investments). In a similar vein, using public subsidies to attract large manufacturing firms to specific regions can constitute an important policy tool to internalise the productivity externalities generated by agglomeration [Greenstone et al., 2010].[4]

Considering the plurality of objectives that many place-based interventions (ex- or implicitly) motivate, evaluating these policies along multiple dimensions is essential to better understand their effectiveness. Furthermore, more comprehensive evaluations may

[2] For more details, see, for example, the evidence reviews on area based initiatives and transport infrastructure investments provided by the What Works Centre for Local Economic Growth [2015*c*, 2016].

[3] The ideas for these models go back to some of the earliest works in the field, i.e., Alonso [1964], Rosen [1979], and Roback [1982].

[4] For a comprehensive review of potential market failures motivating place-based policies, see Kline and Moretti [2014*b*].

also reveal through which mechanisms place-based interventions operate and, thereby, enable the design of better policies in the future. Building on these insights, this dissertation makes three contributions to the literature investigating the economic effects of (targeted) infrastructure investments and place-based policies.

In the first chapter, Oliver Falck, Johannes Koenen, and I evaluate the effects of one of the largest place-based innovation policies in Germany – the Innovative Regional Growth Cores (IRGC) programme. In contrast to many other place-based policies, the IRGC does not directly support the development of physical infrastructures like roads or research facilities. Instead, it provides targeted subsidies for collaborative research and development projects at the firm level in Eastern Germany. If these collaborations are successful, they are meant to develop into self-sustaining innovative networks that, as a specific type of "infrastructure" necessary for innovation, generate local spillovers and promote regional economic development.

Following this narrative, the IRGC may be best thought of as a programme that does not only seek to address market failures commonly associated with innovation at the firm level[5], but also failures of the regional innovation system more generally [Freeman, 1987; Lundvall, 1992; Nelson and Rosenberg, 1993]. Taking the comprehensive scope of the IRGC into account, we evaluate the effects of this programme not only along a broad range of outcomes, but also separately at the firm and regional level. Our evaluation is designed to address three central questions: (1) Did the IRGC induce additional innovation efforts by firms directly subsidised through the programme? (2) Did other innovative firms that were located in the same regions, but did not receive subsidies themselves ("indirectly treated" firms), benefit from the IRGC; that is, are there firm-level spillovers? And (3), did the IRGC cause measurable improvements to aggregate economic outcomes at the regional level?

Exploiting differences in the regional and temporal incidence of IRGC funding, our empirical strategy relies on difference-in-differences (DiD) approaches. To answer the first two questions, we compare the development of firms (directly and indirectly) treated by the IRGC to that of other innovative firms located in Eastern German regions not targeted by the programme (control group). For the firm-level analysis, we are among the first to use confidential survey data covering an extensive sample of innovative firms in Germany. At the regional level, we introduce interactive fixed effects in a DiD framework to compare aggregate economic outcomes in treated and non-treated areas [Bai, 2009; Gobillon and Magnac, 2016].

[5] In particular, the externalities of knowledge creation first described by Arrow [1962] and Spence [1984], or potential failures in the capital markets (see, e.g., Hall and Lerner [2010]).

We find that firms directly receiving IRGC subsidies have significantly increased their innovative activity. As to indirect and aggregate effects, the empirical evidence does not support the notion that the IRGC affected other innovative firms located in the same regions or prosperity at the regional level. While the point estimates for the second and third question are statistically insignificant, which may be due to relatively low statistical power, they are also "small" with regard to their economic importance. Even subject to the caveat, that we cannot interpret the statistical absence of proof as proof of absence of any economic effects, these results raise questions whether the IRGC provides an effective means to address the market failures that motivated it.

While the IRGC uses a very broad definition of the term infrastructure, the second chapter of this dissertation considers the economic effects of an infrastructure in the classical sense: the Next Generation Access (NGA) networks used for the provision of very high-speed Internet. Over the past two decades, it has often been argued that access to very high-speed Internet constitutes a pre-requisite for regional economic development. According to three ministers of the federal government in Germany, for example, a fast Internet connection is an amenity as important as fresh water or electricity.[6]

However, installing NGA networks that can provide all German households with speeds of at least 50 MBit/s is estimated to cost more than 20 billion Euros [TÜV Rheinland Consulting GmbH, 2013]. As the private provision of these investments might be subject to market failures[7], better understanding their importance for regional economic development is of crucial interest. In this chapter, I therefore investigate whether and how the availability of NGA networks affects local employment in the short run. Although local labour market effects only reflect some of the determinants motivating the (publicly subsidised) development of broadband infrastructures, they may provide an important benchmark for policy makers who must decide on local infrastructure investments.

As access to very high-speed Internet can enable the uptake or development of technologies that can both complement and substitute local labour inputs, the sign of its aggregate effect is not clear *ex ante*. Focussing on a sample of relatively rural areas in Germany between 2011 and 2016, I find no empirical evidence that NGA network availability affects the employment rate in the short run. With respect to the share of employees living in a municipality, however, my results indicate that increasing the availability of NGA networks by 5 percentage points leads to an increase in the number of people living (but not necessarily also working) by 0.09 percentage points. Taken together, NGA network

[6] `https://www.zeit.de/wirtschaft/2018-11/digitalisierung-katarina-barley-hubertus-heil-olaf-scholz`, last accessed: 16 March 2019.

[7] Concerns frequently associated with the roll-out of NGA networks are discussed in the introduction of Chapter 2.

availability does not seem to be an important determinant for the creation (or destruction) of local jobs, but affects employed individuals in their decision where to live.

This study is not only among the first to investigate the causal effects of NGA network availability, but also proposes a novel identification strategy. The idea for this strategy is based on the observation that most of the technologies currently used to provide speeds of at least 50 MBit/s require the installation of local optical fibre networks that serve specific access points in the streets. From these access points, signal transmission to and from the customers is realised via pre-existing networks such as the copper wires of the public switched telephone network or the coaxial cables of the TV network. As the roll-out of optical fibre cables in Germany is costly (i.e., because they are typically installed subsurface), installing these optical fibre networks represents the most important obstacle to the provision of NGA networks.

This technical peculiarity can be exploited to identify the causal effects of NGA network diffusion. By taking advantage of detailed geographic information on local land use and methods from graph theory, I approximate the locations of these access points and construct hypothetical networks that minimise the total length of cables required to connect them. The total length of the optical fibre cables required for these hypothetical networks denotes a cost measure for local NGA network roll-out that is solely affected by the spatial dispersion of settlement structures but not any other, potentially endogenous, factors. In an instrumental variables framework, I explain the regional diffusion of NGA network availability over time as a function of network cost (i.e., cable length) while simultaneously controlling for network size (i.e., the number of access points).

Although the benefits promised by the development of new infrastructures are often in the focus of the political debate, they may also have unintended consequences. In the third and final chapter of this dissertation, Andreas Mazat, Bastian Stockinger, and I therefore evaluate how online shopping, a retail channel only enabled by the diffusion of the Internet, has affected the development of brick-and-mortar bookstores in Germany between 1999 and 2013. By investigating this question, we contribute to a better understanding of whether and how e-commerce has shaped the stationary retail industry and the structure of cities more generally.

We focus on the book market in Germany because a specific regulation (the *Buchpreisbindungsgesetz*) prohibits sellers of books to end customers to differentiate prices between retail channels. This institutional setting allows us to identify how the availability of a convenient distribution channel affects traditional retailers. While online shopping may also affect stationary retailers through other channels (e.g., price competition, differences in product portfolios, transparency, etc.), the convenience channel is of interest as it applies to most consumption products.

For our empirical analysis, we combine administrative records from social insurance with novel geomarketing information on regional shopping behaviour at the retail market level. Employing a long difference approach, we explain retail market-specific changes in the number of local brick-and-mortar bookstores (and the employment therein) as a function of changes in online shopping. To identify causal effects, we instrument the long-run change in online shopping by a Bartik-type instrumental variable that exploits historic differences in local age structures [Bartik, 1991].

Our results indicate that regional exposure to online shopping has negative effects on brick-and-mortar bookstores. For example, relative to the levels in 1999, an increase in online shopping exposure by one standard deviation reduces the number of traditional bookstores and total employment therein by on average 14% and 13%, respectively. As stationary retailers represent an important part of urban structures, these results are important for urban planners as they highlight some of the challenges that city centres and high streets may face in the future.

Although all chapters of this dissertation contribute to the same strand of the literature in the field of regional and urban economics, they are self-contained and can be read independently.

Chapter 1

Evaluating a Place-Based Innovation Policy: Evidence from the Innovative Regional Growth Cores Programme in East Germany*

1.1 Introduction

How can economic growth be promoted in disadvantaged areas? Finding answers to this question is not only pivotal to local policy makers, but also a matter of national and international relevance. For example, the European Union spends more than a third of its budget on social and economic cohesion policies promoting the development of structurally weaker regions. In Germany, the aim of providing similar living conditions throughout the country is even enshrined in the constitution. Although reducing regional disparities is an objective of many countries and institutions, the approaches taken toward its implementation differ considerably.

A growing literature in the fields of regional science and urban economics is devoted to determining which types of place-based policies are actually effective. Ex-post evaluations of prominent schemes in the United States and Europe, as well as related theoretical contributions, have improved our understanding of whether and how these programmes affect the local economy.[1] Our study analyses a particular type of place-based policy, i.e., an innovation programme. Focusing on innovation as a driver of regional development [Romer, 1990; Aghion and Howitt, 1992] has recently (re-)gained substantial political traction, but so

* This chapter is based on joint work with Oliver Falck and Johannes Koenen.

[1] This includes Busso et al. [2013], Kline and Moretti [2014*a*], Gobillon et al. [2012], von Ehrlich and Seidel [2018], and Criscuolo et al. [2018] for several ex-post evaluations and Glaeser and Gottlieb [2008], Kline [2010], and Kline and Moretti [2014*b*] for their theoretical work on the mechanisms and channels through which place-based programmes take effect. A more detailed review of the relevant literature is included in Section 1.2.2.

far little robust empirical evidence exists in this area with the exception of cluster policies.[2] In Germany, a number of high-level place-based innovation policies have been initiated recently or are planned in the immediate future. For instance, the Federal Ministry of Education and Research (BMBF) has just established the framework programme "Innovation and Structural Transformation" to bundle its efforts in using innovation policies as an explicit tool for regional economic development. Since the empirical evidence on the effectiveness and efficiency of place-based innovation policies is still scant, we contribute to this literature by evaluating the effects of the Innovative Regional Growth Cores programme (IRGC).

The IRGC is one of the largest place-based innovation policies in Germany. It provides subsidies for collaborative innovation projects of private firms and public research institutes in East Germany. Through its focus on innovation, it complements other public funding channels for regional development in East Germany. Therefore, it is important to control for this backdrop of other public funding schemes when evaluating the effects of the IRGC. To better understand the underlying mechanisms and differentiate between individual channels of this policy, especially the role of spillovers, our evaluation is designed to address three central questions: (1) Did the IRGC induce additional innovation efforts by firms directly subsidised through the programme? With regard to this question, our analysis of the IRGC is very close to Bertamino et al. [2017], who consider the Technology Districts programme in Italy, as well as other studies that investigate the causal effects of an innovation subsidy on directly treated firms.

However, we go further in that we also consider two other questions: (2) Did other innovative firms that were located in the same regions, but did not receive subsidies themselves ("indirectly treated" firms), benefit from the IRGC; that is, are there firm-level spillovers? And (3), did the IRGC cause measurable improvements to aggregate economic outcomes at the regional level? Positive externalities are an important factor in the underlying rationale both for place-based policies as well as for innovation subsidies. Taking this into account, the IRGC could affect firm- and regional-level outcomes through multiple channels. Therefore, answering the latter two questions provides a more differentiated and comprehensive analysis of the true effects of the IRGC.

Our estimation strategy relies on difference-in-differences (DiD) approaches. At the firm level, the development of firms (directly and indirectly) treated by the IRGC is compared to the development of similar innovative firms located in East German regions not targeted by the programme. For this analysis, we use extensive survey micro-data covering R&D conducting firms in Germany since 1995. For our regional analysis, we compare developments of aggregate economic outcomes in treated and non-treated areas. Along all steps of our analysis, we address endogeneity concerns by introducing appropriate

[2] See, e.g., Duranton [2011].

fixed effects and controls; at the regional level, we employ an interactive fixed effects (IFE) specification. For each step, we conduct a series of event study estimations that demonstrate empirically under which conditions our chosen specifications are valid.

We find that firms directly receiving IRGC subsidies have significantly increased their R&D activity. As to indirect effects, the empirical evidence does not support the notion that the IRGC affected other innovative firms located in the same regions. Finally, there is no measurable effect on a set of indicators for regional prosperity. In addition to their lack of statistical significance, which could be due to relative imprecision, the point estimates for the second and third question are also "small" with regard to their economic importance. Even subject to the caveat, that we cannot interpret the absence of statistical proof as proof of absence of any economic effects, this raises some questions to the underlying rationale behind the regional development-aspects of the IRGC. Given the increasing relevance of place-based innovation policies, this issue requires further study and more detailed attention.

The paper proceeds as follows: In Section 1.2, we introduce the IRGC in detail and develop a conceptual framework for our evaluation that is based on the related literature. Section 1.3 describes the comprehensive survey data and presents the firm-level analysis. Section 1.4 focuses on the regional-level analysis featuring the IFE-approach. Section 1.5 concludes.

1.2 The IRGC Programme

1.2.1 Institutional Background

The IRGC is the flagship programme of a series of innovation policies carried out within the BMBF's "Entrepreneurial Regions" (ER) initiative in East Germany. The guiding principle of all ER policies is to overcome structural weaknesses in East German regions by "improv[ing] the framework conditions for innovative processes" [BMBF, 2016*a*]. In contrast to most other place-based policies in Germany, the ER initiative does not promote this objective by subsidising private or public capital investments in general but by explicitly supporting collaborative innovation projects within given regions.[3]

Implementing this principle, the BMBF established the IRGC in 2001 and began to issue its last wave of funding in 2017. The premise underlying the IRGC is that regions possess "unique competences [that] could for example be the command of certain technologies or applications or a specific way of processing materials. [Often, these] are based on long

[3] See, e.g., the *Zonenrandgebietsförderung*, the *Gemeinschaftsaufgabe "Verbesserung der regionalen Wirtschaftsstruktur"* (GRW), and the European Regional Development Funds (ERDF) for programmes primarily focusing on the provision of investment subsidies.

regional traditions and are anchored in companies and research institutes or patented" [BMBF, 2016*c*]. Building on this premise, the IRGC supports regional collaborations between "businesses, universities and research institutes, which either already possess a joint, specific platform technology or have the potential to develop one, [and pursue a] market-oriented strategy [...] aimed at developing innovative, economically successful products in the long run" [BMBF, 2016*c*]. For example in Rostock, a former stronghold of the German ship-building industry, the IRGC supports the development of new tanks and ships for the transport and storage of cryogenic gases.[4] Based on comparable manufacturing traditions, a growth core developing new solutions for technical textiles has been established in Chemnitz – the historical centre of Germany's textile industries.[5] Focusing on development and commercialisation projects rather than basic research, these grants are meant to provide the starting point for a cluster process that can eventually increase regional value creation.

In principle, all companies and public research institutes (including universities) that have a joint platform technology and are based in the same East German region are eligible to apply. Despite the IRGC's focus on regions, the geographic boundaries of the term "region" are not explicitly defined within the programme. Based on the observed existing projects and the wider documentation of the IRGC, in the context of our evaluation we define regions according to so-called regional labour markets (RLMs).[6] To allocate funding, the BMBF has designed an elaborate application process intended to ensure the quality of selected projects.[7] This application process does not involve explicit measurable selection criteria that are publicly observable (either *ex ante* or *ex post*). Instead, the BMBF announces the winning projects without providing any information on unsuccessful contestants. Similar selection approaches are also used for a wide range of other innovation policies in Germany – especially within the ER framework.

In our analysis of the IRGC, we consider the first 13 of 15 waves of the programme and include all 54 joint research projects that were started between 2001 and 2013.[8] Within these first 13 waves, the IRGC granted a total of 275.6 million Euros (EUR), implying an average grant of roughly 6.3 million EUR per project (see Figure 1.2.1 for a map of East German regions with at least one IRGC project and Figure 1.2.2 for a timeline of all projects). On average, a supported private firm (i.e., not a public research institute) received 330,000 EUR in subsidies.

[4] `http://www.unternehmen-region.de/de/1743.php`, last accessed: 17 March 2019.
[5] `http://www.unternehmen-region.de/de/1034.phpandhttp://www.malitec.org/`, last accessed: 17 March 2019.
[6] See Appendix A.1 for a detailed discussion of this issue.
[7] See Appendix A.1 for more details on the application process.
[8] Among these 54 projects, we also include the eleven projects that are continuations of projects started in earlier waves (i.e., fanimat.nano, Maritime Safety Assistant, noa, xmlcity:berlin, ReaWeC, BioOK, Cryo Tank Systems with POLAR, BIZYCLE, highSTICK, TeMaK, and WIGRATEC).

Figure 1.2.1: IRGC Treated Regions

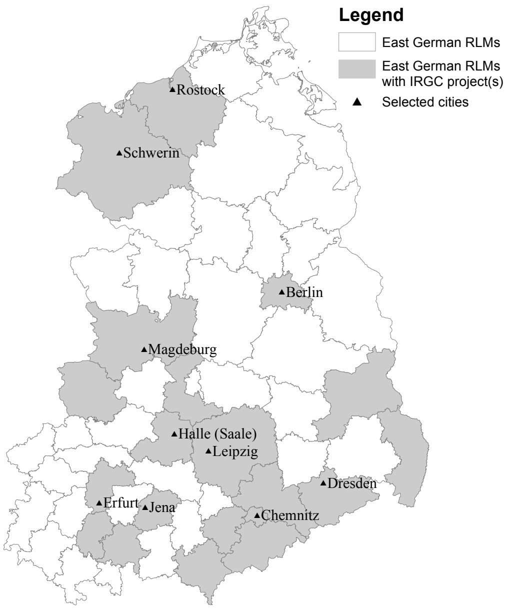

Note: This map illustrates the geographic distribution of IRGC funding in the RLMs of East Germany. Grey shaded regions indicate RLMs that were "core regions" of at least one IRGC project between 2001 and 2013.

Figure 1.2.2: Timing of IRGC Grants

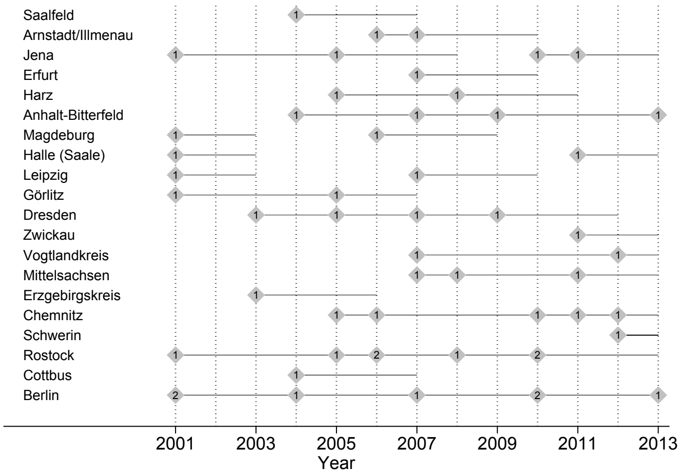

Note: This graph shows the timing of the 54 IRGC projects started between 2001 and 2013. It vertically lists all RLMs in East Germany that were "core regions" of at least one IRGC project. The number in the grey shaded diamonds indicates how many IRGC projects were started in each region that year.

1.2.2 Conceptual Framework

The IRGC provides development and commercialisation subsidies for collaborative innovation projects between firms and research institutes. These subsidies are designed to address the same two market failures commonly used to justify R&D subsidies and other innovation support schemes. First, private R&D investment tends to be lower than socially optimal since R&D activities generate externalities (e.g., in the form of spillovers) that impede firms from fully internalizing the benefits of their own R&D efforts [Arrow, 1962; Spence, 1984]. Second, R&D investment is inherently risky and often associated with severe information asymmetries. Both attributes can give rise to failures in the capital markets when they translate into financing constraints for otherwise profitable R&D projects. The latter concern is of particular relevance for smaller and younger firms as they might not be able to draw on internal funds in the same way larger and more mature firms can [Hall and Lerner, 2010]. Importantly for our study, these considerations also carry forward to commercialisation undertakings since not only the success of R&D projects, but also the demand for eventually developed products is highly uncertain *ex ante*.

Beyond these market failures at the firm level, the IRGC also seeks to address failures of the East German innovation system more generally. The idea of innovation systems dates back to Freeman [1987], Lundvall [1992], and Nelson and Rosenberg [1993], all of whom point out that institutional, geographic, and economic factors are crucial in shaping knowledge flows and, hence, impact the capacity of economies to innovate. Thus, policy interventions are not only based on market failure arguments, but also on flaws in the institutional and economic settings relevant to innovations. For example, Bertamino et al. [2017, p. 1828] note that public policies can also be justified "to overcome imperfections in the innovation systems because some essential links are missing, or the linkages within them are not working well." In fact, all ER programmes are explicitly designed to improve the framework conditions for innovative processes.

The concept of innovation systems is an integral part of most cluster and place-based policies.[9] By concentrating innovative and economic activities in specific areas, they are meant to generate positive externalities of agglomeration that are of importance to the regional innovation system [Glaeser and Gottlieb, 2008]. However, externalities caused by the IRGC are not necessarily always positive. For example, if the supply of qualified R&D workers in the regional labour pool is limited in the short run, positive employment effects on treated firms may fail to materialise or coincide with other firms in the same regions losing qualified employees. In this case, the net effect of the subsidy would be to raise demand for a given set of high-skilled workers (see, e.g., the displacement effects investigated in Einiö and Overman [2016]).

To provide an overview of the policy measures addressing possible failures of innovation systems, it is helpful to divide them into two groups based on the degree of discretion that is granted to decision makers. On the one hand, there are (mostly) rule-based measures involving few if any decisions by politicians once they are in place. This first group is comprised of programmes that are primarily designed to counteract market failures rather than improving the innovation system in general. For example, most tax incentives for R&D investment are, in principle, available to any firm incurring qualifying expenditures. R&D subsidy and loan programmes that do not have demanding eligibility requirements and are well-enough endowed fall into the same category.

On the other hand, the second group includes many large-scale (infrastructure) investment programmes that "pick" specific industries, technologies, or locations precisely because this is where decision makers expect to generate the greatest positive externalities. For example, Falck et al. [2010] evaluate a large-scale cluster policy in Bavaria, the "High-Tech-Offensive". This policy had an explicit focus on specific technological fields and helped improve public research infrastructure such that it could be used by (mostly geographically close) private firms in related industries. Hausman [2017] shows that university innovations are particularly beneficial for the local economy. In the same vein,

9 See, e.g., Chatterji et al. [2014], Kline and Moretti [2014*b*] and Neumark and Simpson [2015] for reviews on cluster and place-based policies.

the strategic establishment of new research sites, rather than simply improving existing ones, has been shown to be an important means for regional development [Andersson et al., 2004].

To this point, our discussion of market and system failures, as well as associated policy responses, has primarily focused on how the IRGC might have affected innovative firms. Focusing on these firms captures the main effects of the IRGC, but ignores the fact that the programme can affect the local economy through other channels. For instance, the commercialisation of new products could lead to an increase in the demand for marketing or accounting services.[10] Taking these possibilities into account, the fundamental logic behind the IRGC might be best thought of as the one underlying location subsidies for large manufacturers. In this context, Greenstone et al. [2010] show that attracting large manufacturing plants can generate substantial productivity spillovers in the local economy that are not necessarily limited to innovative firms (even though this group might benefit in particular; see, e.g., Agrawal et al. [2014] and Fons-Rosen et al. [2016]).

1.2.3 Outline of the Evaluation Approach

A comprehensive evaluation of the IRGC thus must encompass all effects that can occur along the supply chain as well as through other regional linkages. We follow a three-pronged evaluation approach (see Figure 1.2.3 for an illustration). The first step of our empirical analysis is concerned with evaluating the "direct effects" of the IRGC. That is, we want to determine whether firms directly subsidised through the IRGC undertook additional R&D activity. We answer this question by implementing a DiD approach in which we compare the directly treated firms with innovative firms that did not receive any IRGC subsidies and were located outside targeted regions (i.e., in the unshaded regions in Figure 1.2.1). Defining and measuring the direct effects of the IRGC in this manner is similar to the approach taken by other studies evaluating innovation policies.[11]

In the second step of our analysis, however, we go further than many previous studies in that we also evaluate whether the IRGC had "indirect effects" on other innovative firms in treated regions. We consider firms that were based in regions with at least one IRGC project but that did not receive subsidies themselves, and compare them to the same group of control firms as before. Distinguishing between these two treatment definitions at the firm level is particularly informative since localised knowledge spillovers are not merely a by-product of the IRGC but an explicitly specified goal. Even though generating positive

[10] Gans and Stern [2003] provide a framework for thinking about commercialisation settings in general.

[11] See, e.g., Zúñiga-Vicente et al. [2014] and Becker [2015] for surveys of the literature and the What Works Centre for Local Economic Growth [2015*b*] for a comprehensive review of robust evaluation studies in this field.

Figure 1.2.3: Evaluation Strategy

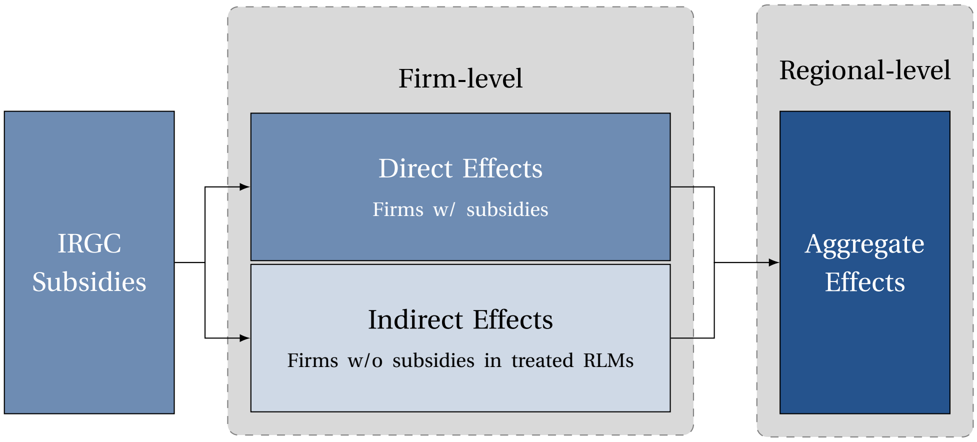

Note: Own illustration.

externalities may be an implicit goal of most conventional R&D and business support programmes, it is rarely made explicit along which dimension they are intended (i.e., within regions, industries, and/or classes of technology) and, consequently, how they should be measured. In this regard, we benefit from the strong place-based focus of the IRGC in that the relevant spillovers caused by this programme are expected to accrue within the region that the programme is embedded in.

The third step of our empirical analysis investigates the aggregate effects of the IRGC at the regional level. We again draw on the fact that the IRGC is designed to target regions that are smaller than the area for which funding is available (i.e., individual RLMs rather than East Germany as a whole) and compare treated regions with non-treated ones. Our focus on the regional dimension addresses the concern that only considering innovative firms neglects several potentially important channels of impact, for example, effects that public research institutes participating in the IRGC have on the local economy (e.g., via the links considered in Hausman [2017]). This part of our evaluation is closest to the broader literature investigating the place-based effects of subsidies for infrastructure and physical capital. For example, Ahlfeldt and Feddersen [2018] consider the agglomeration effects of a large-scale investment in high-speed rail infrastructure; von Ehrlich and Seidel [2018] and Dettmann et al. [2016] evaluate two policies that provided substantial investment subsidies for firms in disadvantaged German regions. The central place-based policy scheme within Europe, the European Structural Development Funds, is evaluated by Becker et al. [2010, 2012, 2013, 2018].

1.3 Firm-Level Analysis

1.3.1 Data

In our firm-level analysis, we are the first to use confidential survey data provided by the Wissenschaftsstatistik of the Stifterverband (WiStat) in the area of innovation policies. On behalf of the BMBF, WiStat has collected information on the R&D activity of innovative firms in Germany since the early 1970s.[12] WiStat aims to conduct a census of all R&D conducting firms in every odd-numbered year. These "full-surveys" are complemented by smaller surveys of the largest R&D conducting firms in all even-numbered years. Since participating in WiStat's surveys is not legally mandatory, the data available to us do not include complete information on the entire population of innovative firms in Germany. Nevertheless, WiStat's surveys are highly comprehensive and unique in terms of time span and detail. Designed according to the Frascati Manual [OECD, 2002], they follow international standards and are not only an integral part of national reporting systems in Germany (e.g., the regional accounts), but also incorporated internationally (e.g., the EU and the OECD). Finally, WiStat's data are particularly suitable for our analysis as they are also the central source of information about innovative activities for the BMBF – the institution that rolled out the IRGC programme.

The data available to us span the years from 1995 to 2013 and include all information WiStat collected within the full and complementary surveys during this time. They provide information on R&D expenditures by source, R&D personnel by type, and a small selection of general business indicators (i.e., total employment and turnover). WiStat matches the surveyed firms with Bureau van Dijk (BvD) identifiers if possible. We use these identifiers to merge the survey data with publicly available information on private firms participating in the first 13 waves of the IRGC.[13] Among the 389 private firms in East Germany that received IRGC grants within these waves, we are able to match 228 firms, that is, 59%. WiStat's firm panel is unbalanced, e.g., due to startups being added, panel attrition or general unit- and item non-response. To minimise related problems, WiStat complements the written surveys with telephone interviews and partly relies on imputing techniques for particularly important variables (i.e., intramural R&D expenditures and total R&D personnel).

[12] WiStat defines firms as the smallest part of a privately owned business enterprise that is required to provide balance sheet information. Since there is no administrative definition of the term "innovative firm", WiStat identifies the population of R&D conducting firms, among others, from previous R&D surveys and auxiliary variables such as a firm's industrial classification, its size and the receipt of public R&D subsidies.

[13] Information on the participants of the IRGC, i.e., their treatment status, timing of treatment, and paid out subsidies, can be downloaded from the *Förderkatalog des Bundes* (`https://foerderportal.bund.de/foekat/jsp/StartAction.do`, last accessed: 18 March 2019).

Table 1.3.1: Distribution of Firms by Industrial Sector (Percentages)

Sector	East Germany w/o Berlin (1)	Directly Treated (2)	Indirectly Treated (3)	Controls (4)
Manufacturing	70.3	68.9	70.0	71.0
Textiles & leather	3.7	6.3	4.6	2.0
Chemicals & pharmaceuticals	4.3	4.4	4.0	4.6
Non-metallic products	3.4	3.4	3.0	4.1
Basic & fabricated metals	11.5	11.7	10.5	12.9
Electronics & optics	14.5	17.0	15.2	13.2
Machinery & equipment	15.6	15.0	16.2	14.8
Cars & transport	5.3	4.4	5.3	5.4
Other manufacturing	12.1	6.8	11.3	13.9
Services	18.3	26.2	19.9	15.0
Others	11.4	4.9	10.1	15.0
Total	100.0	100.0	100.0	100.0

Note: This table shows the industrial distributions of innovative firms included in our sample. We distinguish between four different groups: all firms in East Germany, directly treated firms, indirectly treated firms, and the control group. For data confidentiality reasons, we subsume firms in industrial sections or divisions (printed in italic) with more than 1 but less than 5 firms observed between 1995 and 2013 in the "Others" section or the "*Other manufacturing*" division.

We exclude all firms with observations in less than two consecutive full surveys and interpolate values in even-numbered years.[14] Furthermore, we also exclude firms located in Berlin so as to be consistent with our regional analysis and to alleviate concerns related to the specific industry and firm structure of Berlin (see Table A.4.1 in Appendix A.4 for details). As mentioned above, one of the IRGC's objectives is to foster the clustering of (innovative) firms in targeted regions. Since actual relocations of firms are rarely observed by WiStat, however, they pose a particular challenge for our regional analysis as their treatment exposure is highly selective. Therefore, we choose to exclude relocating firms from our analysis.

After restricting the data in this way, we have a final sample of 4,448 distinct innovative firms in East Germany. Of these, 206 received direct subsidies within the IRGC (directly treated). Among the firms not directly associated with the IRGC, 2,550 are located in regions with at least one IRGC project (indirectly treated) while the remaining 1,692 firms are located outside targeted regions (control group).

Table 1.3.1 reports the industrial distribution of innovative firms included in our sample by treatment type and status (Columns (2) – (4)). Within the manufacturing sector (which is the main focus of the programme), this table shows that the industrial composition of

[14] We also interpolate values for the few firms actually surveyed in even years as small surveys only provide information on a very limited set of outcome variables. If we kept the non-interpolated values, this would prohibit consistent comparisons across outcomes.

innovative firms in East Germany is very similar across treatment types and status. The close resemblance of the industrial distribution between Columns (1) and (2), that is, between all innovative East German firms outside Berlin and the directly treated ones, indicates that the IRGC's allocation mechanism did not result in a specific industry benefiting disproportionately from the programme. Furthermore, comparing the industrial structures between treated and non-treated regions (Columns (3) – (4)), there are also no strong differences between indirectly treated and non-treated firms. This is particularly important for our regional analysis as it supports the assumption that non-treated regions are a suitable control group for IRGC regions.

1.3.2 Direct Effects

1.3.2.1 Descriptives

The first step of our evaluation is to determine whether the IRGC has effectively encouraged directly subsidised firms to undertake additional R&D that they would not have pursued otherwise. To answer this question, we employ a DiD approach in which we compare the directly treated firms with non-treated firms residing outside targeted regions (i.e., we compare firms in Columns (2) and (4) in Table 1.3.1).

Figure 1.3.1 illustrates a simple descriptive comparison of these two groups over time. The graphs depict the development of eight firm-level indicators for treated and non-treated firms between 1995 and 2013: *Total R&D expenditures* measure the overall yearly volume of expenditures classified for R&D purposes by the firm under consideration. We then further differentiate these expenditures by source of financing. *Private R&D expenditures* only include R&D spending financed from the firm's own resources, whereas the variable *public R&D expenditures* measures expenses that are financed from domestic public programmes funded by local, state or federal governments.[15] All R&D expenditures are measured in thousand Euros (TEUR).

Total R&D personnel represents the overall working hours per week (WHpW) spent on R&D activities by any type of employee of the firm. We are further able to differentiate between working hours of *scientists* (tasks mainly involve scientific research) and *technicians* (mainly applied or technical tasks, often performed under supervision of scientists). In the WiStat surveys, working time for (groups of) employees can be distributed across types of tasks, e.g., an individual's job description might involve 10% scientific and 90% technical tasks. *Turnover* and *employees* are measured in TEUR and number of employed persons, respectively.

[15] Note that Total R&D expenditures is not equal to the sum of private and public expenditures, as it also includes other funding sources such as other businesses.

Figure 1.3.1: Descriptive Statistics – Directly Treated vs Control Firms

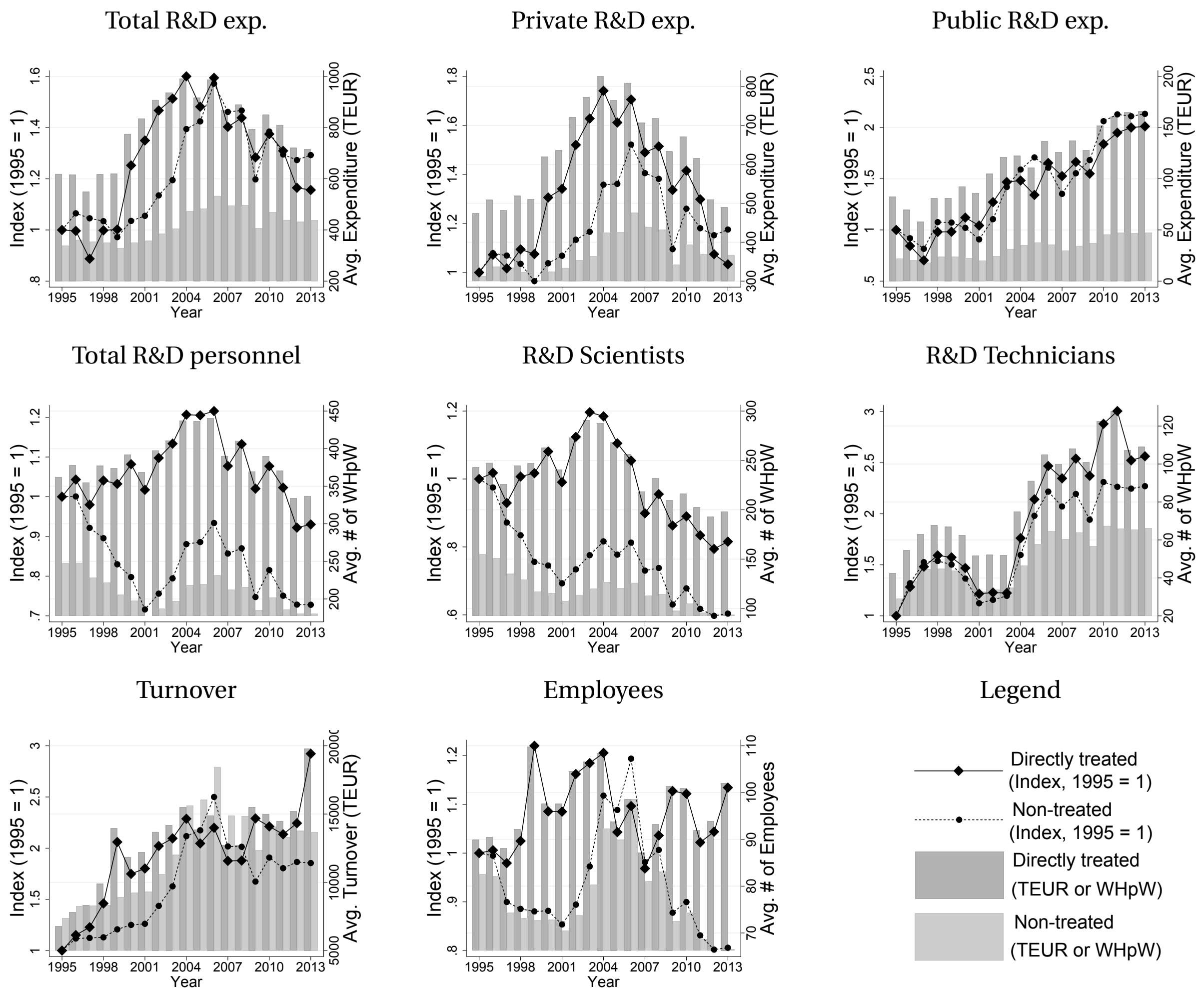

Note: This figure provides a descriptive comparison of the directly treated and non-treated firms across eight different indicators: R&D expenditures by source (TEUR), R&D personnel by type (WHpW), turnover (TEUR), and the number of employees (head count). Outcome variables are interpolated in even years. The grey bars represent yearly averages (directly treated firms are depicted in dark grey to the left of non-treated ones). The connected lines show the development of these yearly averages indexed to the year 1995 (directly treated firms are denoted by black diamonds).

For each of these variables, the graphs depict both relative/indexed developments (lines) as well as the average absolute values/levels (bars). The black diamonds and dots mark the relative development of these indicators indexed to the levels of 1995 for both groups of firms, respectively (left vertical axis). The grey bars plot the average level (right vertical axis; treated firms in dark grey, non-treated ones in light grey). Figure 1.3.1 shows that the levels of R&D activities of directly treated firms, as denoted by investments and working hours, are approximately twice as high as in the control group. This ratio is relatively constant across R&D expenditure categories (i.e., privately financed or paid for by public sources) and types of R&D personnel (i.e., scientists and technicians). Note that the treated firms differ less systematically with regard to general business indicators like turnover and overall employment.

As our identification strategy is based on a DiD approach, identification relies upon variation within firms (and regions) across time rather than on differences in levels across groups. Consequently, we are actually interested in the relative developments over time for each group. There is no obvious divergence of the relative R&D expenditure profiles over time. If the IRGC had large and persistent effects, one would expect this to be reflected in the graphs post programme start in 2001. Given that the IRGC provided relatively generous subsidies, the almost parallel development of the two groups is particularly surprising with regard to public R&D expenditures.

The relative development of R&D personnel is heterogeneous across employment types and groups. While the WHpW of R&D technicians follow a strong and mostly parallel upward trend in both groups, a substantial wedge develops for the total R&D personnel and scientists, with the indices for treated firms roughly 20 percentage points higher than those of non-treated firms. Finally, average turnover and the average number of employees develop less smoothly than the other outcome variables. With regard to turnover, treated firms appear to outperform non-treated firms both prior to and after the financial crisis. A similar pattern holds for their total number of employees, as well.

Thus, Figure 1.3.1 illustrates that despite the substantial differences in levels, the group means of treated and non-treated firms behave relatively similarly over time in most cases. One caveat is that the WiStat data are unbalanced, which can make simple comparisons of means problematic due to sample composition. Furthermore, these simple descriptives do not take into account that the IRGC issued grants at different points in time. For example, the group of "treated firms" in 2002 includes a high number of firms that were actually treated in later waves. In the next section, we derive an empirical model that, among other things, addresses both concerns and allows estimating the causal effects of the IRGC on directly treated firms.

1.3.2.2 Event Study Analysis – Testing the DiD Assumptions

A DiD approach yields consistent estimates of the treatment effect only if treatment and control groups would have followed common trends in the absence of treatment [Angrist and Pischke, 2009]. The validity of this assumption is jeopardised by the same endogeneity concerns usually associated with evaluations of innovation policies. In our context, these include, among others, the fact that the IRGC grants decision makers wide discretion in their subjective assessments, making the allocation of grants selective not only across firms but possibly also across regions and time. For example, decision makers could try to "pick winners" [Cantner and Kösters, 2012] or, on the contrary, target firms in industries with more binding financing constraints [Hyytinen and Toivanen, 2005]. Furthermore, if firms selectively targeted by the IRGC cluster in specific regions, these firm-level concerns could translate to similar selection problems at the aggregate level.

A generalization of the DiD approach – an event study – enables us to check for divergence in pre-trends. Based on this analysis, we demonstrate empirically under which conditions the common trends assumption is most likely to hold in our context and derive a standard DiD model for evaluating the direct effects of the IRGC. We begin by estimating the following two event study models:

$$\textit{TFE:} \qquad ln(y_{irt}) = \sum_{k=\underline{k}}^{\overline{k}} \gamma_k^{dir} \mathbf{1}\{K_{it} = k\} + \alpha_i + \delta_t + \epsilon_{irt} \tag{1.3.1}$$

$$\textit{TFE+LT:} \qquad ln(y_{irt}) = \sum_{k=\underline{k}}^{\overline{k}} \gamma_k^{dir} \mathbf{1}\{K_{it} = k\} + \alpha_i + \delta_t + \mu t \cdot \mathbf{1}\{D_i > 0\} + \epsilon_{irt}, \tag{1.3.2}$$

where (the natural log of) an outcome y of firm i in RLM r and year t is regressed on the "standard" set of two-way fixed effects (TFE), i.e., firm α_i and year δ_t fixed effects, as well as a full set of indicator variables depending on "relative time" K_{it}. For treated firms, K_{it} can be formally expressed as $K_{it} = t - D_i$, that is, it measures the difference between year t and the first year of treatment D_i.[16] For non-treated firms, K_{it} always assumes a value of zero. Consequently, the coefficients $\{\gamma_k^{dir}\}_{k=\underline{k}}^{\overline{k}}$ measure treatment-group specific effects that correspond to "pre-trends" for all periods $k < 0$ and dynamic treatment effects for all $k \geq 0$.[17] Note that if $\{\gamma_k^{dir}\}_{k=\underline{k}}^{-1} = 0$, this model corresponds to a standard (semi-dynamic) DiD model. In Equation (1.3.2), we follow a common approach in the literature and augment the usual TFE event study specification with a treatment group specific linear time trend (LT) to account for general differences in the growth trajectories of both groups. Given the broad scope of the IRGC, we refrain from including any additional time varying controls as they could represent "bad controls" and, therefore, should instead be considered as separate outcomes [Angrist and Pischke, 2009].

In a general investigation of event study designs, Borusyak and Jaravel [2017] show that in absence of a control group, the path of $\{\gamma_k^{dir}\}_{k=\underline{k}}^{\overline{k}}$ in Equation (1.3.1) is identified only up to a linear component in k. Intuitively, this follows from the observation that, in principle, adjustments in the firm and year fixed effects could offset an arbitrary linear trend in $\{\gamma_k^{dir}\}_{k=\underline{k}}^{\overline{k}}$ while still delivering the same predicted values as a model without any trend. In our setting, we benefit from the fact that we have a relatively large control group for our direct effects analysis (i.e., only 12% of the 1,898 firms in the sample used at this step are actually

[16] In case of multiple treatments, which occur rarely in the context of direct effects, we use the first treatment year.

[17] In our notation, $\underline{k}$ and $\overline{k}$ denote the two periods at the beginning and end of each sample indexed relative to treatment start.

treated). This large control group should ensure that the year effects are predominantly determined by the controls and, therefore, prevent arbitrary adjustments in the fixed effect components.

While the size of the control group allows us to trace out the path of $\{\gamma_k^{dir}\}_{k=\underline{k}}^{\overline{k}}$ in Equation (1.3.1), adding a group-specific trend re-introduces the same collinearity problems that would occur in absence of a control group [Borusyak and Jaravel, 2017]. To solve this problem when estimating Equation (1.3.2), we follow the authors' suggestions and drop a second period-specific indicator prior to treatment. While Borusyak and Jaravel [2017] demonstrate that dropping any two $k < 0$ is sufficient to determine non-linearities in the pre-trends, they suggest choosing a period that is as far away from the first reference period ($k = -1$) as possible. We choose $k = -6$ as it is the last pre-treatment period available in the WiStat data for firms that were treated in the first wave of IRGC funding, that is, in 2001.

Figure 1.3.2 plots $\{\hat{\gamma}_k^{dir}\}_{k=-6}^{2}$ for Equations (1.3.1) and (1.3.2). Considering the TFE estimates first (white diamonds with grey outline), there are significant pre-trends in almost all outcome variables of interest. As this indicates that the IRGC seems to "pick winners" in that directly treated firms grow significantly faster than firms in the control group even before receiving treatment, simple TFE estimates likely violate the common trends assumption. However, by estimating Equation (1.3.2), we can take advantage of the observation that many of the pre-trends in the TFE estimates are strikingly linear. After including a treatment group-specific time trend in our event study (black diamonds), most of the pre-treatment differences (i.e., for $k < 0$) become considerably smaller in magnitude and turn insignificant at the 95% significance level.[18] Based on these results, we choose the following semi-dynamic DiD model to evaluate the direct effects of the IRGC:

$$ln(y_{irt}) = \sum_{k=0}^{\overline{k}} \gamma_k^{dir} \mathbf{1}\{K_{it} = k\} + \alpha_i + \delta_t + \mu t \cdot \mathbf{1}\{D_i > 0\} + \epsilon_{irt}. \quad (1.3.3)$$

We do not replace the set of period specific post-treatment dummies with a simple post-treatment indicator since we expect treatment effects in the context of the IRGC to develop dynamically over time. In fact, given that the IRGC specifies different goals in the short and long run, a dynamic approach is per se required. Moreover, Borusyak and Jaravel [2017] emphasise that one should not expect an immediate and persistent treatment effect

[18] In Table A.4.2 in Appendix A.4, we present the results of all event studies that we conducted for the direct effects analysis. The most important take away from this table is that the estimates for TFE+LT and the specification for which we use only treated observations (TFE w/ treated) are very similar. Thereby, Equation (1.3.3) not only follows the suggestions in Borusyak and Jaravel [2017], but also takes into account that some studies do not use a control group of never treated units in an event study design.

Figure 1.3.2: Event Studies – Directly Treated Firms

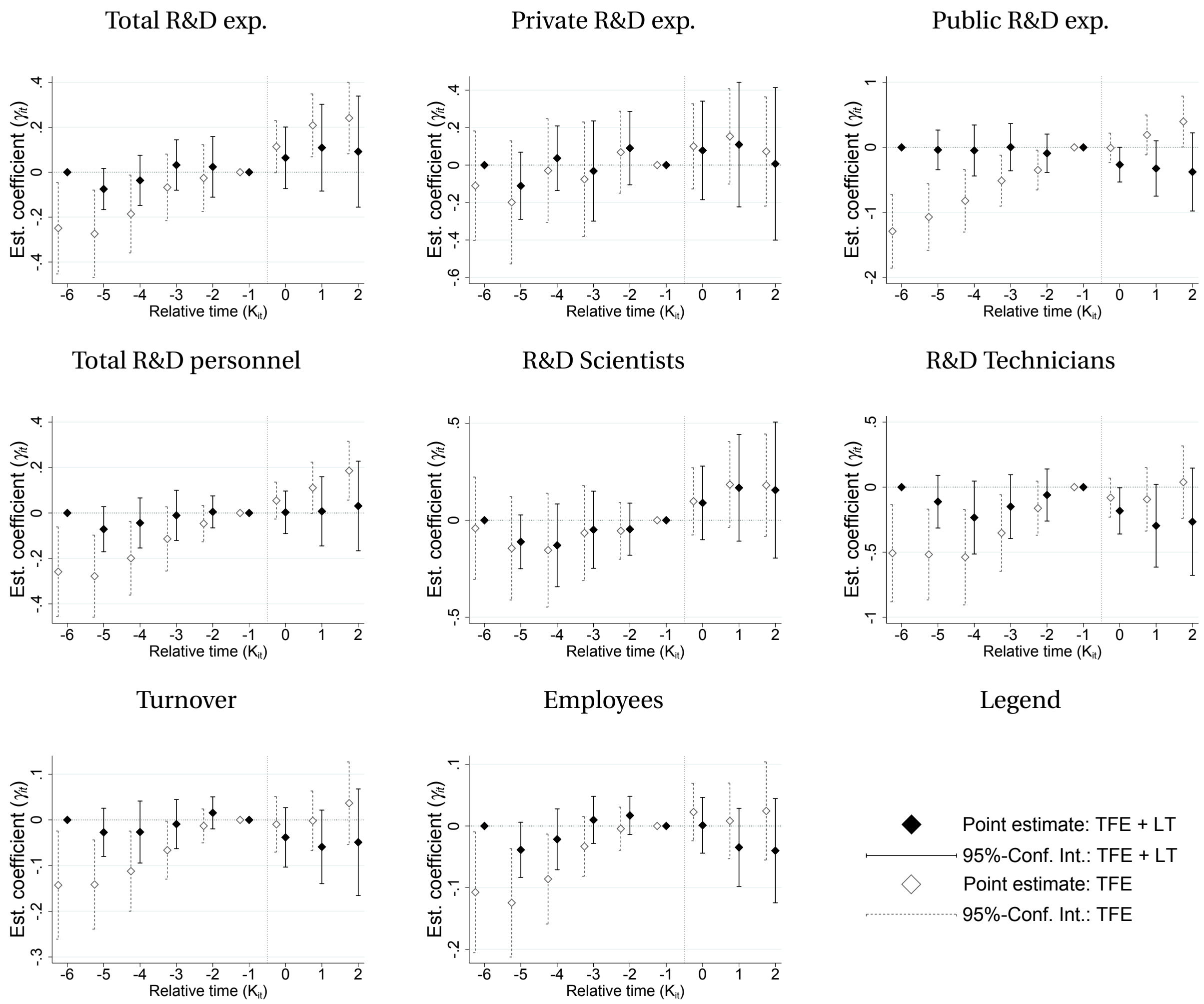

Note: This figure illustrates the results of the event study models shown in Equations (1.3.1) and (1.3.2) for $\{k\}_{-6}^{2}$, respectively. Outcome variables are interpolated in even years and expressed in natural logs. Values smaller than one have been replaced by one prior to taking logs. TFE estimates are denoted by white diamonds with black outline, TFE+LT estimates are depicted by black diamonds. Solid and dotted whiskers denote confidence intervals at the 95% level.

in the context of most event study analyses and, therefore, recommend using a more flexible specification.[19]

To keep the presentation of our results comprehensive, the following section only shows dynamic average treatment effects within four distinct intervals: years 0–1, years 2–3, years 4–5, and all years thereafter (years 6+). We calculate these effects as weighted averages of the full set of estimated coefficients $\{\hat{\gamma}_k^{dir}\}_0^{12}$. In particular, we define a weighting scheme that assigns every $\hat{\gamma}_k^{dir}$ a weight ω_k according to the share of (non-missing) observations

[19] A common alternative to Equation (1.3.3) would include firm-specific time trends rather than a single group-specific one. In our context, however, this does not seem desirable for at least two reasons. First, the group-specific trend (μ) is already sufficient to ensure the validity of the common trends assumption and, therefore, allows to estimate the treatment effects more efficiently. Second, μ is estimated using only non-treated observations. Given the strongly unbalanced nature of our data set, estimating firm-specific trends would be problematic if we observe some of the treated firms only past their actual treatment start.

recorded for firms treated in period k relative to the sum of all observations of treated firms in the relevant interval.

Formally, these weights can be defined as:

$$\omega_k = \frac{\sum_{i\in N}\sum_{t\in T}\mathbf{1}\{K_{it}=k\}}{\sum_{i\in N}\sum_{t\in T}\sum_{j=\underline{j}}^{\overline{j}}\mathbf{1}\{K_{it}=j\}}\,\forall k\in\{\underline{j},\overline{j}\}, \tag{1.3.4}$$

where N denotes the total number of firms, T the last year of our sample period (2013), and $\{\underline{j},\overline{j}\}$ the reference interval. On the one hand, this weighting scheme allows us to reduce the number of coefficients. On the other hand, it also both helps to account for the unbalanced nature of our panel data as well as the differential timing of IRGC grants, which implies that the number of observations available to estimate period-specific effects shrinks with increasing k. Our weighting scheme helps ensure that estimated average treatment effects are not influenced too strongly by these caveats.

1.3.2.3 Estimated Direct Treatment Effects

Table 1.3.2 reports the dynamic treatment effects for the directly subsidised firms. Regarding total R&D spending, we find that the IRGC induced directly treated firms to increase their total expenditures during the first four years after treatment by between 18 and 24%. We can quantify these effects in absolute terms to be able to compare them to the average subsidy received. Based on the means observed in the descriptive statistics for 1999[20], these relative increases would be equivalent to additional total R&D expenditures between 110 and 150 TEUR in each of these four years – compared to the calculated (overall, not yearly) average subsidy allocated to each private firm of about 330 TEUR. While these effects are economically and statistically significant, they do not persist beyond year 4.

Turning to the separate spending categories included in our data, we are able to differentiate the funding channels that are responsible for the overall spending observed between public and private funding. Again, there is a substantial positive effect on publicly funded R&D for the first four years, but not thereafter.[21] We do not observe statistically significant effects on privately funded R&D expenditures. As a result, these findings do not allow us to draw strong conclusions, but it is worth pointing out that most point estimates are positive. This suggests that the IRGC did not substantially crowd out private investments on average.

[20] Note that this is a conservative assumption, as average R&D expenditures were subsequently higher both for treated and non-treated firms.

[21] The absolute of the public channel accounts for about one third of the overall R&D-spending effect. It has to be remarked that the variables under consideration are survey-, not accounting data.

Table 1.3.2: Direct Effects of the IRGC

	ln R&D Expenditures			*ln R&D Personnel*			*ln Econ. Outcomes*	
Dep. Var.:	Total	Private	Public	Total	Scien-tists	Techni-cians	Turn-over	Emplo-yees
Years 0–1	0.184***	0.058	0.431**	0.089	0.164	0.026	–0.068	–0.035
	(0.068)	(0.119)	(0.212)	(0.072)	(0.124)	(0.140)	(0.043)	(0.039)
Years 2–3	0.235**	–0.010	0.586*	0.157	0.218	0.112	–0.075	–0.056
	(0.097)	(0.152)	(0.317)	(0.102)	(0.169)	(0.199)	(0.068)	(0.063)
Years 4–5	0.161	0.060	–0.094	0.077	0.215	–0.028	–0.137	–0.043
	(0.126)	(0.173)	(0.418)	(0.127)	(0.196)	(0.268)	(0.091)	(0.088)
Years 6+	0.060	0.050	–0.675	–0.060	0.214	–0.287	–0.117	–0.038
	(0.169)	(0.219)	(0.557)	(0.176)	(0.255)	(0.367)	(0.131)	(0.130)
Firm-FE	Y	Y	Y	Y	Y	Y	Y	Y
Year-FE	Y	Y	Y	Y	Y	Y	Y	Y
Lin. Trs.	Y	Y	Y	Y	Y	Y	Y	Y
Firms	1,898	1,898	1,898	1,898	1,898	1,898	1,898	1,898
N	14,787	14,787	14,787	14,774	14,774	14,774	14,787	14,787
Means in 1999	620.0 TEUR	510.1 TEUR	81.6 TEUR	374.0 WHpW	247.2 WHpW	67.0 WHpW	13,955.2 TEUR	109.8 Count

Note: This table shows dynamic treatment effects for the directly treated firms during the sub-periods "Years 0–1", "Years 2–3", "Years 4–5", and "Years 6+". The effects are weighted averages of the γ^{dir} coefficients estimated based on Equation (1.3.3). The weighting scheme is defined as: $\omega_k = \frac{\sum_{i\in N}\sum_{t\in T}\mathbf{1}\{K_{it}=k\}}{\sum_{i\in N}\sum_{t\in T}\sum_{j=\underline{j}}^{\overline{j}}\mathbf{1}\{K_{it}=j\}}\ \forall k \in \{\underline{j},\overline{j}\}$. That is, ω_k is the share of treated observations with non-missing entries in period k relative to the sum of all treated observations with valid entries in the reference interval $\{\underline{j},\overline{j}\}$. Outcome variables are interpolated in even years and expressed in natural logs. Values smaller than one have been replaced by one prior to taking logs. Monetary variables are denoted in TEUR (R&D expenditures and turnover), R&D personnel in WHpW, and employees in head counts. Standard errors are clustered at the firm level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

With regard to R&D personnel and general business indicators, we do not find that the IRGC had statistically significant effects. This may, in part, be driven by a lack of precision in measurement, therefore it can still be useful to consider the pattern of coefficients that emerges. First, we observe differential effects across different types of R&D personnel: While total R&D staff and technicians follow the same pattern as total expenditures (uptick in the first four years, followed by a reversion thereafter), the effect on scientists is stable over the entire period of observation, which corresponds to the descriptive results in Figure 1.3.1. Such an orientation towards (basic) research instead of application and implementation would clash with the aims of the IRGC, independent of the question whether a more scientific orientation proves to be a successful strategy for the treated firms in the long run. This potential risk is also reflected by the fact that we find consistently negative effects for turnover; indicating that treated firms are not successful in commercialising developments in the short or medium run.

As indicated in the previous subsection, introducing additional controls to the regression in Equation (1.3.3) as robustness checks could give rise to a bad control problem. The scope of the IRGC is deliberately broad, so many potential controls at the firm level actually have to be considered to also be potential outcomes of interest. Accordingly, we refrain from reporting robustness checks of our results involving such firm-level controls, as these might be targeted and affected by the IRGC directly. Instead, we consider all variables consistently surveyed by WiStat as outcomes of interest in our analysis. Introducing controls at the regional level, however, seems less problematic. Even though the IRGC is motivated by regional objectives, these are only addressed indirectly. Introducing two sets of regional controls therefore provides a sensible robustness check in our firm-level analysis.

The first set of regional controls includes geographic and socio-economic variables that account for many differences in structural characteristics.[22] The second set of controls is intended to address the concern that regions not hosting any IRGC projects might have been compensated with public funds from other sources. Therefore, our second set of control variables contains information on the amount of subsidies per capita that a given region received within the GRW and ERDF programmes, which represent the two most important place-based policies in Germany, as well as from the BMBF.[23] Our two measures indicating regional funding from the BMBF are particularly important as they should capture any substitution of funding that might take place within the institution that rolled out the IRGC. Our estimates for the direct effects are fully robust to the inclusion of these two sets, both qualitatively and quantitatively.[24]

1.3.3 Indirect Effects

1.3.3.1 Descriptives

The IRGC shares certain features with other "standard" firm-level innovation policies, but beyond their scope it is explicitly designed to generate regional spillovers. Arguably, among all types of regional spillovers, the one most strongly emphasised by the programme's design concerns the development of other (innovative) firms in the same regions. In the first step of our spillover analysis, we therefore focus on whether the IRGC measurably affected

[22] We include population density, the share of the working-age population, the share of elderly people (> 65 y.o.), the share of females, the share of foreigners, the share of employees with a medium level of qualification, and the share of employees with a high level of qualification.

[23] To calculate the amount of subsidies a region received from programmes carried out by the BMBF, we downloaded all information that is available within the *Förderkatalog des Bundes*. We aggregated the project-level data to the level of RLMs and calculated two measures. Our first measure is the sum of all expenditures the BMBF made toward a specific RLM, irrespective of the type of institution receiving the grant but excluding any IRGC projects, per capita. Our second measure considers only grants made to private institutions (i.e., only firms but not to universities or public research institutes).

[24] The results of these robustness tests are available from the authors upon request.

outcomes for unsubsidised innovative firms located in regions with at least one project. After investigating potential spillovers along this dimension, we turn to a regional analysis that incorporates all other types of externalities at the aggregate level (e.g., along the supply chain or through the local labour market) in the remainder of this paper.

Investigating the existence of indirect effects requires us to define “indirect” IRGC treatments of firms. Given the structure and design of the IRGC, the logical decision is to consider firms co-located in the same region as treated firms to be “indirectly treated”. Specifically, we consider a region (and the firms therein) to be treated once the BMBF declares it to be the “core region” of an IRGC project for the first time. From this point, treatment status remains unchanged until the end of our observation period, that is, our indirect treatment definition depends on the announcement of the first IRGC project rather than the disbursement of actual grants. Furthermore, we do not differentiate between regions receiving one or multiple IRGC projects. As the incidence of multiple IRGC projects is often associated with either receiving follow-up funding for the same growth core or hosting several projects that are thematically related and based on similar networks, our definition provides the advantage that it does not inflate the actual number of treatments.[25] Finally, it also allows us to remain consistent with our analysis of the direct effects.

In Figure 1.3.3, we compare the group of indirectly treated firms (grey diamonds with black outline) with the same control group as before (black dots). Again, the upper six graphs illustrate the R&D activity of both groups – using the same variables as in the previous section – whereas the lower two graphs show the development of general business indicators over time. Regarding the differences in levels, indirectly treated firms conducted roughly 1.5 times as much R&D as the firms in the control group. For instance in 2013, the indirectly treated firms spent on average 650 TEUR on R&D whereas firms in the control group spent only 440 TEUR. The most notable exception in this regard concerns publicly funded R&D expenditures. In particular, the average amount of subsidies available to indirectly treated firms drops sharply after 2003 and then remains mostly constant, whereas this amount grows steadily over time in the control group. This pattern becomes even more apparent when considering the indexed developments in these variables, relative to 1995.

With regard to both total and private R&D expenditures, there is a parallel development of the two groups until about 2007. From then on, indirectly treated firms display a substantially more robust development than the non-treated ones, with a wedge of 20 to 30 percentage points opening up. On the other hand, we observe a pronounced, opposite divergence in terms of public R&D expenditures, as noted above, with a wedge of more than 100 percentage points opening up. This development, in particular, is of interest in

[25] E.g., Rostock hosted several growth cores thematically related to the shipbuilding industry and maritime technologies.

Figure 1.3.3: Descriptive Statistics – Indirectly Treated vs Control Firms

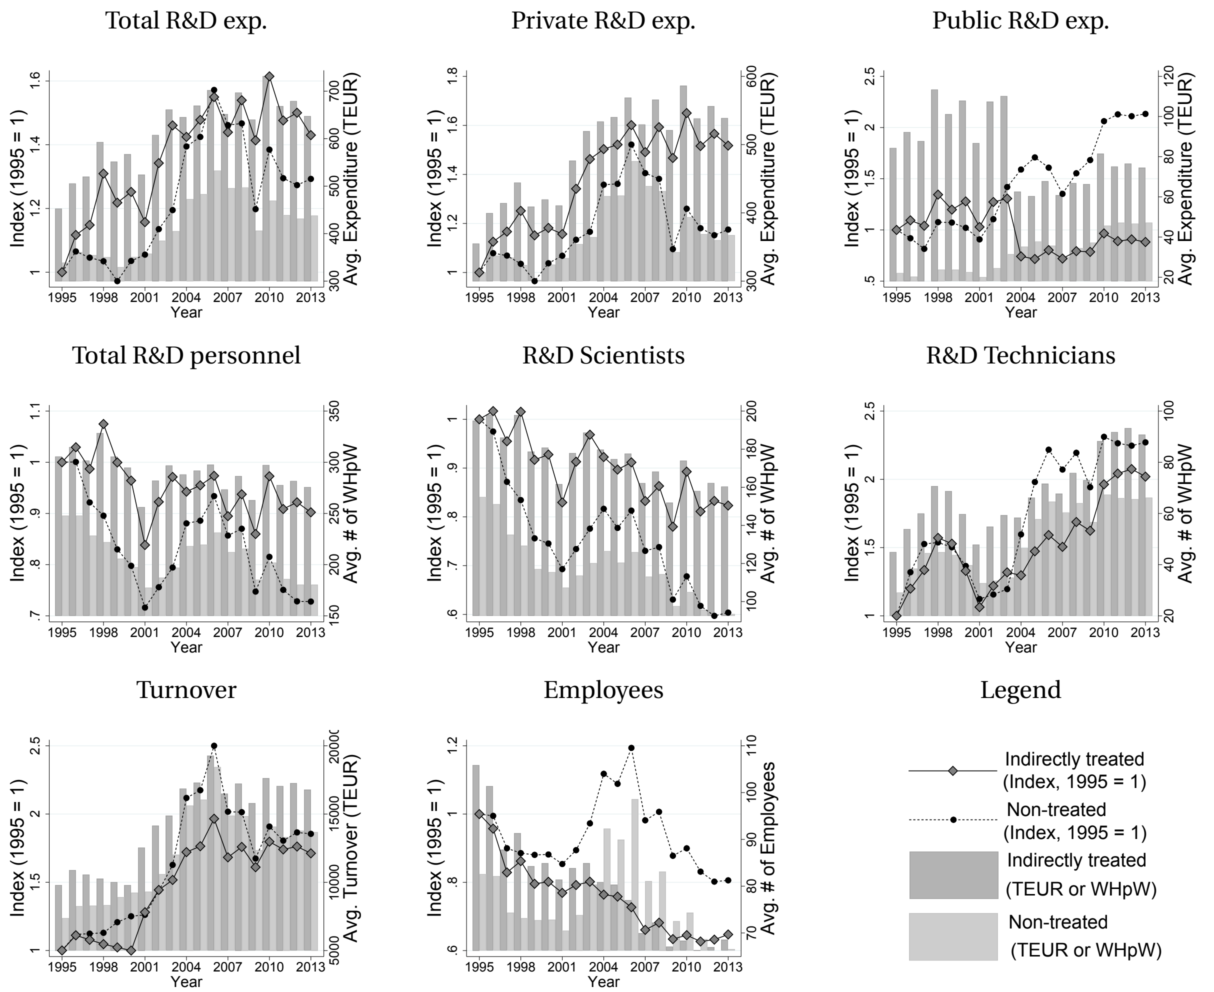

Note: This figure provides a descriptive comparison of the indirectly treated and non-treated firms across eight different indicators: R&D expenditures by source (TEUR), R&D personnel by type (WHpW), turnover (TEUR), and the number of employees (head count). Outcome variables are interpolated in even years. The grey bars represent yearly averages (indirectly treated firms are depicted in dark grey to the left of non-treated ones). The connected lines show the development of these yearly averages indexed to the year 1995 (indirectly treated firms are denoted by grey diamonds with black outline).

the context of the previous comparison of directly treated and non-treated firms (see Figure 1.3.1). There we found that the growth trajectories of directly treated firms and the control group were completely parallel for public R&D expenditures.

In particular, the divergence of directly and indirectly treated firms in the same regions with regard to publicly funded R&D expenditures raises an interesting question: Is there a discernible cause for this divergence that is potentially related to the IRGC? A possible mechanism could involve competition for public funds (in general) at the regional level. IRGC-grants to certain firms within a region would thereby make it less likely for other firms in the same region to obtain (other) public funding for projects. Our choice of control group – non-treated regions – implies that such a redistribution within treated regions, i.e., from indirectly treated firms to directly treated ones, due to the IRGC would not affect our identification strategy; the indirect treatment would only also include a negative externality via a lowered probability of obtaining (other) public funding.

With regard to our identification strategy, serious problems would only arise if there were systematic changes in the amounts and composition of public grants in the control regions associated with the IRGC programme. Investigating this concern, Figure A.3.1 in Appendix A.3 contrasts the development of publicly funded R&D expenditures in both treatment groups and the control group with the development of regionally available BMBF funds. If the IRGC led to an internal change in the BMBF's grant allocation mechanism, this should be visible in their absolute and relative developments. With regard to the overall sum of BMBF grants, however, we observe that regions not treated by the IRGC only experienced a relative increase in their total BMBF grants around the onset of the financial crisis while remaining at a substantially lower level overall. In terms of the grants the BMBF made to private firms, there is no systematic trend visible at all. As we only compare private firms in this part of our analysis, this suggests that the IRGC did not induce the BMBF to allocate substantially more (or fewer) grants from other programmes or funding lines to non-treated regions. Note in addition to this that we explicitly control for these variables in our robustness checks in our regression analyses.

Turning to average turnover and the total number of employees, the comparison of indirectly treated and non-treated firms allows for two countervailing observations. While the developments of average turnover were almost parallel between 1995 and 2013, the indirectly treated firms exhibited a substantially more negative trend with regard to the total number of employees than the non-treated group of firms.

1.3.3.2 Event Study Analysis – Testing the DiD Assumptions

Identification of indirect effects at the firm level is subject to concerns similar to those discussed in the context of direct effects. Specifically, selection might occur because industries tend to cluster in specific regions. If IRGC decision makers favour selected industries (e.g., because they are successful or, to the contrary, have recently suffered from adverse shocks), this regional clustering could cause a violation of the common trends assumption. To investigate this possibility, we pursue an event study approach parallel to the one presented above:

$$\textit{TFE}: \quad ln(y_{irt}) = \sum_{k=\underline{k}}^{\overline{k}} \gamma_k^{ind} \mathbf{1}\{K_{rt} = k\} + \alpha_i + \delta_t + \epsilon_{irt} \tag{1.3.5}$$

$$\textit{TFE+LTs}: \quad ln(y_{irt}) = \sum_{k=\underline{k}}^{\overline{k}} \gamma_k^{ind} \mathbf{1}\{K_{rt} = k\} + \alpha_i + \delta_t + \mu_r t + \epsilon_{irt}, \tag{1.3.6}$$

Figure 1.3.4: Event Studies – Indirectly Treated Firms

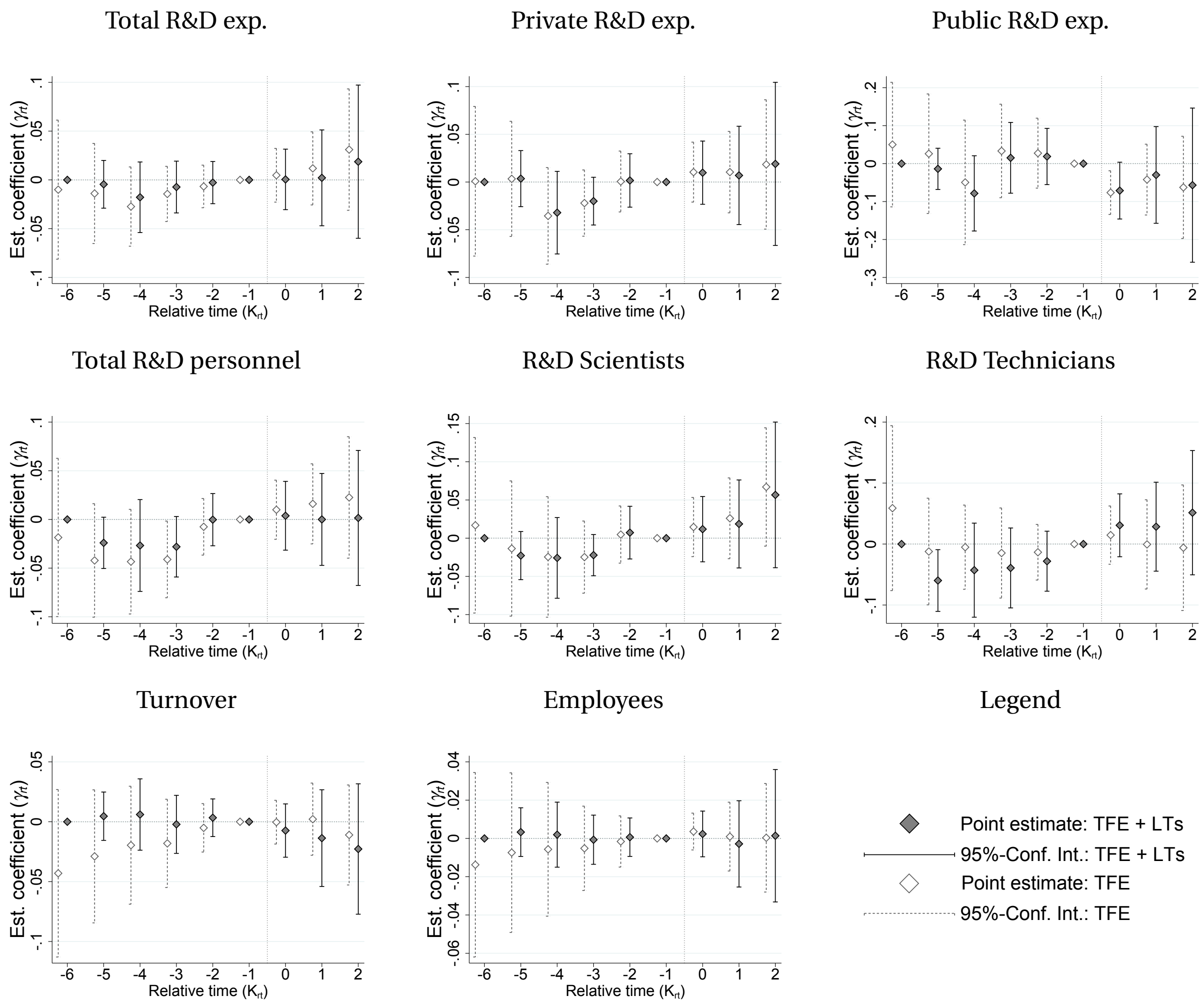

Note: This figure illustrates the results of the event study models shown in Equations (1.3.5) and (1.3.6) for $\{k\}_{-6}^{2}$, respectively. Outcome variables are interpolated in even years and expressed in natural logs. Values smaller than one have been replaced by one prior to taking logs. TFE estimates are denoted by white diamonds with black outline; TFE+LT estimates are depicted by grey diamonds with black outline. Solid and dotted whiskers denote confidence intervals at the 95% level.

where treatment now depends on K_{rt}, that is, relative time at the regional level. Note that all firms within a given region become "treated" at the same time. Therefore, we cluster standard errors within RLMs. Further, we allow for heterogeneous treatment-group specific time trends. That is, we relax the assumption of a single time trend by including a set of linear trends that vary by RLM ($\mu_r t$).

Figure 1.3.4 plots $\{\hat{\gamma}_k^{ind}\}_{k=-6}^{2}$ for Equations (1.3.5) and (1.3.6).[26] In general, it can be seen that pre-trends using the TFE specification (white diamonds with grey outline) are usually insignificant and substantially smaller in magnitude than in the direct effects analysis. This can be explained by the fact that these firms were (most likely) not explicitly targeted by the IRGC and, hence, their (unobserved) trends should not be systematically correlated

[26] In Table A.4.3 in Appendix A.4, we present the results of all event studies for the indirect effects analysis.

with the IRGC allocation decision. Furthermore, many selection concerns at the regional level should be addressed by firm-level fixed effects since they also nest all time-invariant characteristics at higher levels of aggregation. Similarly, the fact that industrial composition does not vary strongly across treated and non-treated regions gives additional confidence in our estimation strategy (see Table 1.3.1).

A closer look at Figure 1.3.4 reveals that our TFE (+LT) estimates for R&D expenditures and personnel in some cases display a slight u-shaped pre-trend – with one observation (R&D technicians, at $k = -5$) significantly different from 0 at the 95% level. That is, there are incidences and variables for which indirectly treated firms grow more quickly than firms in the control group, but only during a few years before the first IRGC project is allocated to their respective regions. Given that linear trends in our event study setting can only be identified using observations in $k = -1$ and $k = -6$, these non-linearities in the pre-trends are problematic as they make choosing a second pre-period to be dropped more problematic. For example, in case of the relatively linear TFE pre-trends for turnover and employees, dropping $k = -4$ rather than $k = -6$ likely has no effects whereas this type of choice may affect the results for the pre-trends of R&D technicians.

To ensure the comparability of the indirect effects analysis across outcomes as well as with the previous analysis, Figure 1.3.4 only includes the TFE+LTs estimates (grey diamonds with black outline) for event study specifications that drop $k = -6$. For most outcomes, including region-specific trends either improves or does not affect the pre-trends also observed in the TFE specifications at all. The only notable exception to this rule are R&D technicians as omitting $k = -6$ for this outcome leads to fitting a downward sloping trend that is not obvious in the following years. As this problem should be limited in the case of semi-dynamic specifications, i.e., because they use all observations with $k < 0$ for the identification of linear trends rather than only two, we choose the following DiD model for our evaluation of the indirect effects of the IRGC:[27]

$$ln(y_{irt}) = \sum_{k=0}^{\overline{k}} \gamma_k^{ind} \mathbf{1}\{K_{rt} = k\} + \alpha_i + \delta_t + \mu_r t + \epsilon_{irt}. \tag{1.3.7}$$

1.3.3.3 Estimated Indirect Treatment Effects

Table 1.3.3 shows the estimated treatment effects for indirectly treated firms. This part of our analysis benefits from a considerably larger sample since we can include the previously excluded group of unsubsidised firms residing in IRGC targeted regions while only losing the 206 directly treated firms. Even though the regional treatment definition implies that we are effectively relying on fewer actual treatments (i.e., 19 RLMs as opposed to 206 firms),

[27] We address potentially remaining concerns (in particular with regard to R&D technicians) by also reporting the TFE results for our indirect effects analysis in Appendix A.4.

Table 1.3.3: Indirect Effects of the IRGC

	ln R&D Expenditures			*ln R&D Personnel*			*ln Econ. Outcomes*	
Dep. Var.:	Total	Private	Public	Total	Scien-tists	Techni-cians	Turn-over	Emplo-yees
Years 0–1	–0.007	–0.002	–0.073	0.021	0.007	0.067*	–0.003	0.003
	(0.025)	(0.031)	(0.050)	(0.027)	(0.035)	(0.037)	(0.015)	(0.010)
Years 2–3	0.005	0.011	–0.062	0.031	0.044	0.108	–0.016	0.004
	(0.050)	(0.056)	(0.080)	(0.050)	(0.057)	(0.072)	(0.030)	(0.023)
Years 4–5	–0.030	–0.006	–0.175	0.021	0.027	0.144	–0.008	0.009
	(0.068)	(0.077)	(0.121)	(0.073)	(0.079)	(0.095)	(0.047)	(0.033)
Years 6+	–0.031	–0.028	–0.282	–0.019	0.015	0.255	0.001	0.051
	(0.110)	(0.135)	(0.211)	(0.125)	(0.135)	(0.181)	(0.073)	(0.054)
Firm-FE	Y	Y	Y	Y	Y	Y	Y	Y
Year-FE	Y	Y	Y	Y	Y	Y	Y	Y
RLM Trs.	Y	Y	Y	Y	Y	Y	Y	Y
RLMs	53	53	53	53	53	53	53	53
Firms	4,242	4,242	4,242	4,242	4,242	4,242	4,242	4,242
N	33,008	33,008	33,008	32,998	32,998	32,998	33,008	33,008
Means in 1999	551.7 TEUR	409.0 TEUR	101.8 TEUR	305.6 WHpW	178.7 WHpW	68.9 WHpW	10,020.6 TEUR	84.2 Count

Note: This table shows dynamic treatment effects for the indirectly treated firms during the sub-periods “Years 0–1”, “Years 2–3”, “Years 4–5”, and “Years 6+”. The effects are weighted averages of the γ^{ind} coefficients estimated based on Equation (1.3.7). The weighting scheme is defined as: $\omega_k = \frac{\sum_{i\in N}\sum_{t\in T}\mathbf{1}\{K_{rt}=k\}}{\sum_{i\in N}\sum_{t\in T}\sum_{j=\underline{j}}^{\overline{j}}\mathbf{1}\{K_{rt}=j\}}\ \forall k \in \{\underline{j},\overline{j}\}$. That is, ω_k is the share of treated observations with non-missing entries in period k relative to the sum of all treated observations with valid entries in the relevant sub-period. Outcome variables are interpolated in even years and expressed in natural logs. Values smaller than one have been replaced by one prior to taking logs. Monetary variables are denoted in TEUR (R&D expenditures and turnover), R&D personnel in WHpW, and employees in head counts. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

our estimates for the indirect treatment effects are more precise than the ones for the direct effects. Overall, the point estimates for most outcomes (i.e., total and private R&D expenditures as well as total R&D personnel and scientists) are fairly close to zero and statistically insignificant. Accordingly, these estimates are also not associated with large economic significance despite the fact that they involve regional multipliers (i.e., they apply to each innovative firm located in IRGC targeted regions).

Regarding the outcome variables measuring R&D activity, there are two notable exceptions, which we would like to discuss in detail: publicly financed R&D expenditures and the WHpW by R&D technicians. Considering R&D expenditures financed by public sources, the consistently negative point estimates suggest that research financed by public grants decreased in regions targeted by the IRGC – an observation that we also derived from the descriptive comparison above (i.e., see Figure 1.3.3). Taking the point estimates during the first six years at face value, they imply decreases in annually available funds between 6 and 17 TEUR per firm. On the one hand, it could be argued that these decreases should not affect the general R&D activities of these firms directly as they only account for 1.1% to 3.2% of total R&D expenditures per year (in 1999 levels). On the other hand, these observations may also indicate that the IRGC has served as a means to redistribute some of the public resources, that would otherwise have been available to all firms in these regions, to the winners of IRGC projects.

With regard to the activities undertaken by R&D technicians, contrasting the point estimates in Table 1.3.3 with the results of the TFE estimates in Table A.4.4 in Appendix A.4 provides interesting insights. While the estimates for all other outcome variables are qualitatively and quantitatively similar across specifications, there are pronounced differences for R&D technicians. In particular, the TFE+LTs results indicate an increasing, consistently positive effect of the IRGC whereas the TFE estimates are substantially smaller in magnitude and even change signs towards the end of the observation period. Therefore, it is only the results of the TFE+LTs model that indicate that the IRGC positively affected the activities of R&D technicians in indirectly treated firms.

As before, we also test the robustness of our results by introducing two sets of regional controls. In the context of indirect effects, controlling for regional indicators is arguably associated with more severe bad control problems. For example, if the IRGC was successful in attracting highly qualified workers to specific regions, this should result in compositional changes of the labour force that would be captured by the controls in our first set. Across all outcomes of interest, however, introducing these controls does not change the results qualitatively. If anything, our estimates tend to become smaller in magnitude.[28] Given that we may block some of the channels through which the IRGC affects the activity of innovative firms, this result might not be surprising.

1.4 Regional-Level Analysis

1.4.1 Data

We now analyse the regional effects of the IRGC. Since the programme's objectives are relatively vague in that only increasing regional value creation is specified as an explicit goal, we face the challenge of selecting a reasonable number of other outcomes to provide a comprehensive picture of the IRGC's regional impact. Apart from gross value added (GVA), which denotes a direct measure for regional value creation, we also consider four other outcomes: the number of startups, the stock of establishments, the number of employees,

[28] The results of these robustness tests are available from the authors upon request.

Figure 1.4.1: Descriptive Statistics – Treated vs Control RLMs

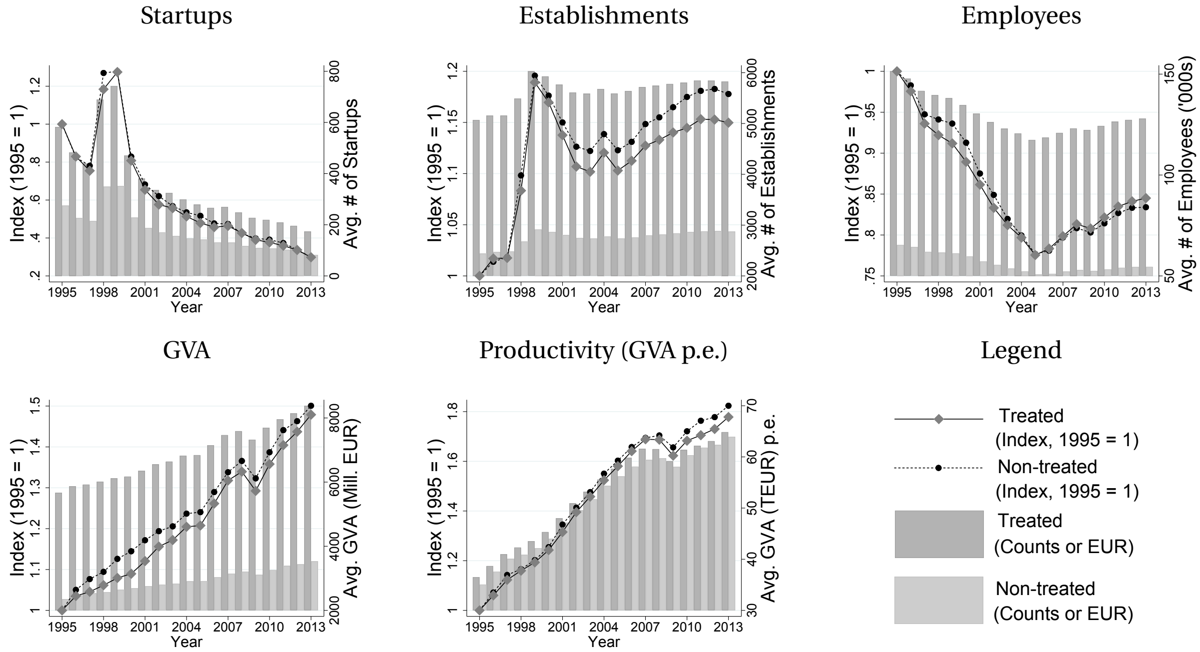

Note: This figure provides a descriptive comparison of the treated and non-treated regions across five different indicators. The grey bars represent yearly averages (treated regions are depicted in dark grey to the left of non-treated ones). The connected lines show the development of these averages indexed to the year 1995 (treated regions are denoted by grey diamonds).

and productivity measured as GVA per employee. These outcome measures allow us to evaluate a range of key policy objectives.

The regional data on startup activity and the number of establishments is aggregated from a 50% random sample of the Establishment History Panel 1975–2014 (BHP) of the Institute for Employment Research (IAB).[29] The BHP is an administrative collection of information on the employment structure of all establishments in Germany that have at least one employee covered by social security requirements. Both outcome variables provide information on the economic dynamic of a given region. Information on the total number of employees is obtained from the Federal Employment Agency. With regards to regional value creation, data on GVA is provided by the Regional Accounts (*Volkswirtschaftliche Gesamtrechnung der Länder*).

Figure 1.4.1 illustrates the average development of treated and non-treated RLMs across all five regional outcomes under consideration. In terms of differences in levels, treated RLMs are between two and three times larger than non-treated ones. For example in 2013, treated RLMs employed on average roughly 125,000 employees as compared to 55,000 employees in non-treated regions. With regard to relative development, both groups evolve

[29] The BHP is documented in Schmucker et al. [2016]. We define startups in the BHP following the methodology proposed by Hethey and Schmieder [2010].

almost in parallel over the whole sample period. The co-movement of these indicators is a strong first indication that the IRGC does not seem to have large regional effects.

1.4.2 Event Study Analysis – Testing the DiD Assumptions

The IRGC's selection mechanism potentially involves criteria at the firm as well as at the regional level. In the absence of a feature in the selection mechanism that provides credibly exogenous variation in the allocation of subsidies across regions, DiD estimators are the most commonly applied strategies [Gobillon and Magnac, 2016]. However, while the standard set of fixed effects controls for the most important determinants of the treatment allocation at the firm level, it is questionable whether this also holds true at the regional level.

Gobillon and Magnac [2016] argue that it is unlikely that all unobserved components affecting regional treatment allocation can be subsumed in additively separable parts (i.e., α_r, δ_t, and/or low dimensional polynomials of time). There are two specific concerns: First, including only a single common macro shock, δ_t, that homogeneously affects all units in the cross-section is not sufficient to model the unobserved dynamics driving regional outcomes over time. Economic variables at this level are subject to multiple macro shocks that can have heterogeneous effects, for example, the business cycle may affect regions differentially depending on their industrial structure. Second, neglecting heterogeneous effects of (potentially multiple) macro shocks in a DiD framework can lead to various forms of cross-sectional dependence. For instance, regions may be correlated spatially and therefore exhibit similar effects from specific shocks. Once these characteristics are correlated with the treatment status, they bias the estimation of the treatment effect at the regional level.

The IRGC is a prototypical example of a regional policy that is likely subject to some of these concerns. For instance, if politicians lobby for the allocation of IRGC projects because their jurisdiction is strongly affected by a current business cycle downturn, this could lead to biased estimates of the treatment effect. Similarly, technological progress can exhibit heterogeneous effects across regions if a specific invention affects individual industries more than others. In general, German innovation policy has a tendency to favour specific fields or technologies (e.g., biotechnology). If these tendencies are related to technological breakthroughs, it could give rise to biased estimates. Taking these concerns into account, we follow Gobillon and Magnac's 2016 suggestion to use IFE models.

We use the following event study models at the regional level:

$$\text{TFE+LTs:} \qquad ln(y_{rt}) = \sum_{k=\underline{k}}^{\overline{k}} \gamma_k^{agg} \mathbf{1}\{K_{rt} = k\} + \alpha_r + \delta_t + \mu_r t + \epsilon_{rt} \tag{1.4.1}$$

$$\text{TFE+LTs+IFEs:} \qquad ln(y_{rt}) = \sum_{k=\underline{k}}^{\overline{k}} \gamma_k^{agg} \mathbf{1}\{K_{rt} = k\} + \alpha_r + \delta_t + \mu_r t + \lambda_r^{'} F_t + \epsilon_{rt} \tag{1.4.2}$$

where F_t denotes an $L \times 1$ vector of common factors and λ_r an $L \times 1$ vector of factor loadings. The factor model $\lambda_r^{'} F_t$ in Equation (1.4.2) is meant to control for unobserved time trends (F_t) that are allowed to have heterogeneous effects across regions (via λ_r). Note that, in principle, $\lambda_r^{'} F_t$ could also nest the TFE+LTs model used above, namely if $F_t = (1, \delta_t, t)'$ and $\lambda_r = (\alpha_r, 1, \mu_r)'$. However, since we expect the standard fixed effects as well as the linear trends to be crucial for the IRGC's selection mechanism, we explicitly include them in all IFE-specifications.[30] As a result, the factor model in Equation (1.4.2) controls for heterogeneous effects of unobserved factors that go beyond time-invariant effects, homogeneous time shocks, and a region-specific linear time trend. Similar to the standard fixed effects notion, IFE models can account for only a given number of common factors – without the option for the researcher to explicitly distinguish which (aggregate) shocks should be absorbed into the factors and which should not.[31]

Bai [2009] derives the asymptotic properties of an estimator for γ_k^{agg} and proposes an algorithm that, given a pre-specified number of factors, allows us to estimate Equation (1.4.2). We look at two options for determining the number of factors, L, used in our estimations. First, Bai [2009] shows that L can be estimated consistently in panel data models using the information criteria approach, thus building on earlier work by Bai and Ng [2002] in the context of pure factor models. Second, Moon and Weidner [2015] show that under certain assumptions, the limiting distribution of least square estimators for models like Equation (1.4.2) is independent of the number of factors used as long as it is not underestimated. Therefore, they suggest rather using too many than too few factors, and to check whether the estimates are sensitive to changes in L.

30 Arguably, the economic reasoning for including only α_r and δ_t might be stronger than that for controlling for linear time trends, too. However, controlling for linear and often quadratic time trends is common practice in similar applications (see, e.g., Wolfers [2006], Moon and Weidner [2015], Hagedorn et al. [2016]) and keeping them in the model allows us to remain consistent with our analysis of the indirect effects.

31 An alternative to this approach is introduced by Hagedorn et al. [2016] who suggest using what they call "natural" factor models. Rather than letting λ_r and F_t be simultaneously determined by an algorithm, they replace F_t with a selection of observed aggregate factors that are deemed important by economic analysis (in their context of unemployment benefits, these include oil prices, aggregate employment figures, and an index of monetary policy). Although this seems promising in their circumstances, the lack of transparency in the IRGC's selection mechanism prohibits us from developing a strong notion about which aggregate factors might have been relevant for the allocation decision at this level.

Figure 1.4.2: Event Studies – Treated RLMs

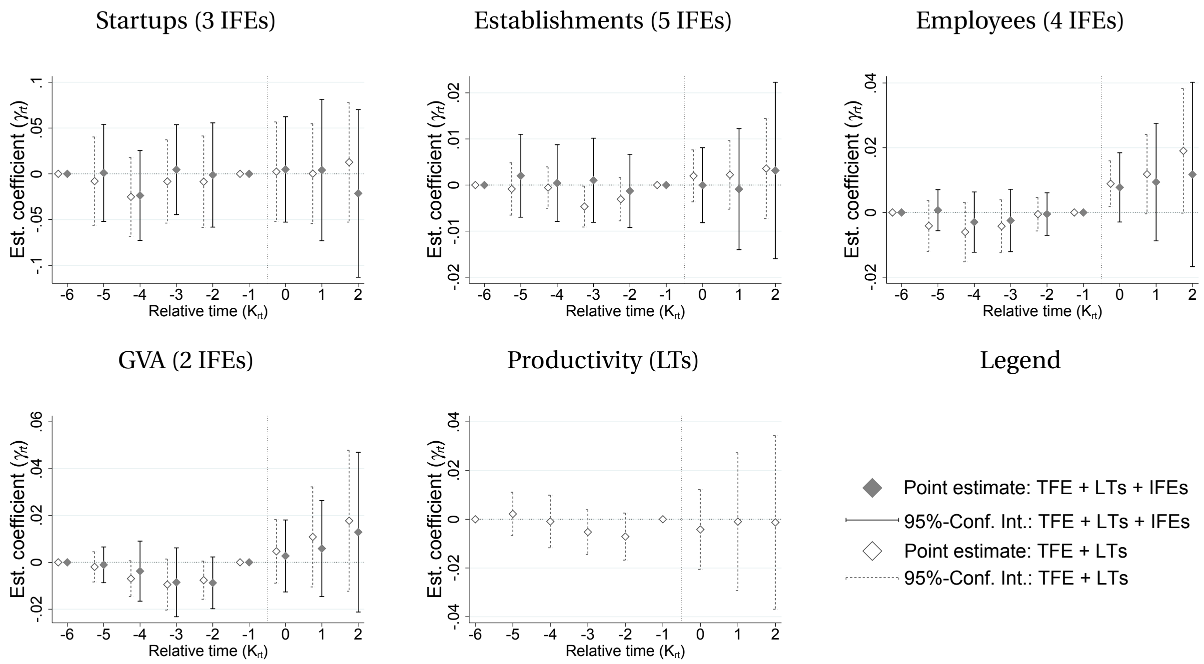

Note: This figure illustrates the results of the event study models shown in Equations (1.4.1) and (1.4.2) for $\{k\}_{-6}^{2}$, respectively. TFE+LTs estimates are denoted by white diamonds with grey outline whereas the grey diamonds depict the TFE+LTs+IFE model that is associated with the most preferable pre-trends according to our selection algorithm. If no TFE+LTs+IFE model reduces pre-trends relative to a TFE+LTs benchmark, we only show the results for the latter. Solid and dotted whiskers denote confidence intervals at the 95% level.

Combining their suggestions with the requirement of common (pre-)trends in our DiD setting, we determine our final evaluation model based on the following procedure:

1. Estimate the event study model in Equation (1.4.2) with up to six common factors (i.e., $\forall L \in \{1, 6\}$).
2. Calculate the information criteria following the suggestions made by Bai and Ng [2002] and select all TFE+LTs+IFE models that use at least as many factors as the one minimizing the information criteria.[32]
3. Compare the pre-trends of the selected TFE+LTs+IFE models with the ones obtained from a TFE+LTs benchmark. From among these models, choose the one that is associated with the smallest pre-trends.[33]

Figure 1.4.2 shows (some of) the event study estimates produced following our selection procedure. In all graphs, we present the TFE+LTs estimates (white diamonds with grey outline), which either serve as a reference point for the event study model with the "best" pre-trends (grey diamonds) or stand on their own if no IFE combination makes meaningful reductions in the pre-trends relative to this benchmark. In Figure A.3.2 in Appendix A.3, we supplement the event study graphs shown here with the results of the not selected

[32] See Appendix A.2 for more details on calculation of the information criteria.

[33] In many cases, it turns out that the number of factors minimizing the information criteria is also identical or at least very close to the one that generates the most preferable pre-trends.

specifications. That is, we show the estimates of at least three event study models – TFE, TFE+LTs, and TFE+LTs+IFEs – with the number of factors chosen following the information criterion approach. In addition, if there is an event study model that uses more common factors than the one minimizing the information criteria but is associated with better pre-trends, we include the results for that specification as well.

Across all outcomes, pre-trends at the aggregate level tend to be more heterogeneous and follow less linear patterns than the ones observed at the firm level. In most cases, however, pre-trends at the regional level are statistically insignificant at the 5% level. Nevertheless, the graphs show that adding IFEs can help improve pre-trends. For example, the stock of establishments exhibits a slight dent shortly before treatment begins that is almost statistically significant at the 5% level. For this outcome, including five IFEs flattens out this dent. Similar observations can be made for the number of employees and GVA, albeit the improvements are less pronounced. In terms of the number of startups, including only three IFEs reduces pre-trends marginally. With regard to our productivity measure, adding any number of IFEs actually leads to an exaggeration of pre-trends relative to the TFE+LTs model and, therefore, we refrain from their inclusion. Based on these results, we select one of the two following models for each outcome at the regional level for our empirical analysis of the aggregate effects of the IRGC:

TFE+LTs:
$$ln(y_{rt}) = \sum_{k=0}^{\overline{k}} \gamma_k^{agg} \mathbf{1}\{K_{rt} = k\} + \alpha_r + \delta_t + \mu_r t + \epsilon_{rt} \tag{1.4.3}$$

TFE+LTs+IFEs:
$$ln(y_{rt}) = \sum_{k=0}^{\overline{k}} \gamma_k^{agg} \mathbf{1}\{K_{rt} = k\} + \alpha_r + \delta_t + \mu_r t + \lambda_r^{'} F_t + \epsilon_{rt} \tag{1.4.4}$$

A key distinction between the models is the use of pre- and post-treatment information in estimating trends. While the effects of the linear time trends in both models are solely identified using non-treated observations, estimating the factor model in Equation (1.4.4) also requires the use of post-treatment observations since the factors F_t have to be estimated for all periods. In principle, given a sufficiently high number of factors, this could cause all (potential) effects of the IRGC to be absorbed into the factor model. As a reference point, think of a standard DiD setting in which the treatment occurs only at one point in time for all treated units. In this case, a single factor that changes at treatment start could subsume all (heterogeneous) effects related to the programme.

In our context, however, there are two reasons making it unlikely that this problem will arise. First, the start of regional IRGC treatments varies strongly across time (see Figure 1.2.1). Since there are seven different periods during which (at least) one region received its first IRGC project, it would take at least seven factors to absorb all IRGC effects. Second, in order to be fully absorbed, the IRGC effects would have to dominate all other common factors that matter for regional development. Based on our descriptive comparisons, there is very little evidence for such a conclusion. Furthermore, given that we only observe 19 treated and 34 non-treated RLMs in this part of the analysis, also using treated observations when estimating the factors might even be desirable. For instance, six of the 19 treated RLMs (without Berlin) received their first IRGC project during the first round of funding in 2001 – that is, immediately before a major reform of the German labour market (known as the *Hartz* reforms) in 2003.[34] If the Hartz reforms dominated other regional factors, using treated observations to estimate the factor model might disentangle any reform effects from post-treatment effects of the IRGC.

1.4.3 Estimated Aggregate Treatment Effects

Table 1.4.1 shows the regional effects of the IRGC across our five outcomes of interest. In contrast to the firm-level analysis, our evaluation of regional effects relies on a balanced panel data set. Therefore, the weighting scheme $\{\omega_k\}$ used for the coefficients in this table is identical across all outcomes.

The central message of Table 1.4.1 is that the IRGC did not have any statistically significant effects at the regional level. For example, the point estimates for employees, GVA, and productivity show no evidence that the IRGC had a strong effect at the regional level. Indeed, the estimates for the first six years are close to zero and, particularly for the first four years, more precisely estimated. In terms of the number of startups and the stock of establishments, there seems to be an upward trend. Taking the estimates for these variables at face value (despite their statistical insignificance) implies that the number of startups and the stock of establishments increased by 2.5% and 1.4% after three to four years post treatment, respectively.

Estimates of the long-run firm-level effects of the IRGC were accompanied by important caveats and the same is true in the regional context. Specifically, the long-run effects presented in Table 1.4.1 are estimated using only the few RLMs that are treated at the beginning of the programme, making credible inference even more problematic. We therefore focus on the short- and medium-run effects for which more data are available.

Performing the same robustness checks as before (i.e., introducing our two sets of control variables) does not meaningfully affect the results presented in Table 1.4.1.[35] Since the

[34] This could also be a concern for our analysis of indirect effects at the firm level. However, in the previous analysis we were able to control for a more fine-grained set of fixed effects, i.e., firm vs. region fixed effects, and could also take advantage of the general sample structure. Innovative firms are inherently more similar to each other than RLMs to RLMs. As a consequence, aggregate effects such as the Hartz reforms are more likely to have homogeneous effects across these firms and thus would be captured by time fixed effects.

[35] See Table A.4.5 in Appendix A.4.

Table 1.4.1: Aggregate effects of the IRGC (East Germany w/o Berlin)

Dep. Var.:	Startups	Establish-ments	Employees	GVA	Productivity (GVA p.e.)
Years 0–1	0.012	0.004	0.004	0.006	0.001
	(0.030)	(0.006)	(0.007)	(0.008)	(0.011)
Years 2–3	0.025	0.014	0.005	0.007	0.000
	(0.037)	(0.011)	(0.015)	(0.020)	(0.016)
Years 4–5	0.065	0.023	0.007	–0.004	–0.015
	(0.063)	(0.018)	(0.022)	(0.035)	(0.027)
Years 6+	0.123	0.049	0.029	–0.013	–0.027
	(0.096)	(0.036)	(0.036)	(0.066)	(0.048)
RLM-FE	Y	Y	Y	Y	Y
Year-FE	Y	Y	Y	Y	Y
RLM Trends	Y	Y	Y	Y	Y
Factors	3	5	4	2	-
RLMs	53	53	53	53	53
N	1,007	1,007	1,007	1,007	1,007
Means in 1999	742.3 Count	6,027.8 Count	138.1 Count ('000s)	6,116.4 Mill. EUR	43.5 TEUR p.e.

Note: This table shows dynamic treatment effects at the regional level (RLMs) during the sub-periods "Years 0–1", "Years 2–3", "Years 4–5", and "Years 6+". The effects are weighted averages of the γ^{agg} coefficients estimated based on Equations (1.4.3) or (1.4.4). The weighting scheme is defined as: $\omega_k = \frac{\sum_{r\in R}\sum_{t\in T}\mathbf{1}\{K_{rt}=k\}}{\sum_{r\in R}\sum_{t\in T}\sum_{j=\underline{j}}^{\overline{j}}\mathbf{1}\{K_{rt}=j\}}\ \forall k \in \{\underline{j},\overline{j}\}$. That is, ω_k is the share of treated observations with non-missing entries in period k relative to the sum of all treated observations with valid entries in the relevant sub-period. Outcome variables are expressed in natural logs. Startups, establishments and employees are measured as counts, GVA in million Euros and productivity as GVA per employee. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

bad control problems associated with including regional-level controls are arguably most pronounced in the context of this robustness check (i.e., regional adoption mechanisms should at least partly be captured by our controls), the fact that the estimates remain largely unaffected could be interpreted as an additional indication that the IRGC did not have large effects at this level.

1.5 Discussion and Conclusions

In the last few years, a growing body of theoretical and empirical research has focused on the question of whether place-based policies are effective in promoting regional growth. Even though growth theory emphasises the key role that innovation plays for (regional) development, the regional effects of innovation policies are rarely considered. Given that many innovation policies are eventually pursuing regional rather than purely firm-level objectives this paper investigates the effects of the IRGC, a major place-based innovation policy in Germany. We develop a methodological framework that allows us to analyse the causal effects of this policy not only at the firm, but also at the regional level.

We are among the first to use confidential firm-level survey data and assemble a comprehensive data set of regional indicators to answer the following three questions: (1) Did the IRGC induce additional innovation efforts by firms directly subsidised through the programme? (2) Did other innovative firms that were located in the same regions but did not receive subsidies benefit from the IRGC? (3) Did the IRGC cause measurable improvements at the regional level? Looking at the direct effects on innovative firms first, we find that firms receiving IRGC subsidies have increased their R&D activity above the levels they would have reached when following their counterfactual trends. These effects, however, persist only within the short and medium run. With regard to the indirect effects, the empirical evidence does not support the notion that the IRGC affected other innovative firms located in the same regions. Finally, we find no indication that the IRGC discernibly affected regional prosperity as measured by five different outcomes. Taken together, our results are in line with the broader literature on the effectiveness of cluster policies and suggest that, despite continued public efforts, local spillovers are not easily generated by political means.

Chapter 2

Next Generation Access Networks: Very High-Speed Internet and Regional Employment in Germany

2.1 Introduction

In its Digital Agenda for Europe (DAE), the European Commission emphasises that universal access to very high-speed Internet is of strategic importance for the promotion of social and economic cohesion. When published in 2010, the DEA formulated the explicit goal that Internet speeds of at least 30 MBit/s should become universally available in the entire EU territory until 2020 [European Commission, 2010].[1] Following suit, many member states of the EU developed national strategies to implement these, and sometimes even more ambitious, objectives. For example, in 2014, the federal government in Germany announced its aspirations to provide all German households with universal downstream speeds of at least 50 MBit/s by 2018 [German Federal Government, 2014]. However, for Germany alone, it is estimated that the installation of the Next Generation Access (NGA) networks, that are needed to provide these speeds, require investments of at least 20 billion Euros [TÜV Rheinland Consulting GmbH, 2013].[2]

Given the size of the investments needed for the deployment of NGA networks, it is often argued that market or government failures may prevent private telecommunication operators from providing the socially optimal level of NGA network availability. Three reasons for market failure are of particular concern. First, the high fixed costs required for the installation of a telecommunication network in combination with its low marginal operational costs can give rise to market power in the form of a natural monopoly [Baumol, 1977]. Second, private telecommunication operators may fail to internalise that a regionally

[1] In 2016, the European Commission updated this goal. It now states that by 2025, all Europeans should have access to at least 100 MBit/s in download speeds [European Commission, 2016].

[2] Of the 20 billion Euros, 40% are required to give access to the “last” 5% of the households, that is, the most remotely located ones [TÜV Rheinland Consulting GmbH, 2013].

unequal provision of NGA networks can lead to costly distortions of spatial equilibria. For example, if a better Internet availability in one place attracts people from surrounding areas, regional in- and out-migration can generate externalities (e.g., by creating agglomeration economies or congestion). Third, the uncertainty associated with the development of digital technologies may prevent some market participants from accurately predicting future demand for very high-speed Internet. If the capability to make accurate predictions differs systematically between banks and NGA network developers, resulting information asymmetries can give rise to failures in the financial markets.

With respect to government failures, the interplay of regulations and investment incentives plays an important role. In various countries, the legislator uses access regulations to restrict the (potential) market power of private telecommunication network owners. Beyond regulating already existing infrastructures, governments may also make the development of new networks contingent on the willingness to accept specific requirements. For example, the German government currently tenders the frequencies necessary to operate a nation-wide 5G mobile network. Among others, the award conditions of the tender include that the prospective operator of this network is obligated to negotiate with competitors technical and contractual access opportunities. If these negotiations are unsuccessful, the regulatory authority for telecommunications in Germany (*Bundesnetzagentur*) may enforce regulations that grant access to competitors [Bundesnetzagentur, 2018*b*]. As a consequence of these regulations, investments in telecommunication networks can have public good character which increases the likelihood that they are underprovided by the private sector.

Given that the provision of very high-speed Internet access faces multiple challenges, it is important to better understand what a market or government failure in this field might imply. Therefore, this paper evaluates empirically how the availability of NGA network has affected regional employment in Germany between 2011 and 2016. For the purposes of this paper, the availability of NGA networks is measured by the share of households in a region that has access to a fixed line connection with a minimum downstream data transfer rate of 50 MBit/s. As emphasised above, the regional diffusion of NGA networks is likely driven by multiple factors. In the context of an empirical evaluation, these factors can give rise to severe endogeneity concerns. Most importantly, profit maximising telecommunication operators may react to demand-side factors, such as regional differences in the willingness to pay for broadband services, and only install NGA networks in places where returns are high. As regional willingness to pay is not directly observable but likely correlated with both NGA roll-out and local labour market outcomes, this factor constitutes an omitted variables problem.

Addressing this and related concerns, I propose a novel identification approach to estimate the causal effects of NGA network diffusion. My approach is based on the observation that most of the technologies that are currently used to provide Internet speeds of at least 50 MBit/s require the installation of local optical fibre networks that serve specific access points in the streets. From these access points, signal transmission to and from the customers is typically realised via pre-existing networks such as the copper wires of the public switched telephone network (PSTN) or the coaxial cables of the TV network. As the roll-out of optical fibre cables in Germany is costly (i.e., because they are typically installed subsurface), installing these optical fibre networks represents the most important obstacle to the provision of NGA networks.

Taking advantage of detailed geographic information on local land use and methods from graph theory, I approximate the locations of these access points and construct hypothetical networks that minimise the total length of cables required to connect them. In the absence of any geographic frictions, these hypothetical networks could be used as blueprints for optical fibre networks that could theoretically serve all inhabited areas with speeds of at least 50 MBit/s. The total length of optical fibre cables required for these hypothetical networks denotes a cost measure for local NGA network roll-out that is solely affected by the spatial dispersion of local settlements but not any other, potentially endogenous, factors. Incorporating this cost measure in an instrumental variables (IV) framework, I explain the regional diffusion of NGA network availability over time as a function of network cost (i.e., cable length) while simultaneously controlling for network size (i.e., the number of access points).

In general, NGA network availability may have different effects where people work and where they live. This paper therefore distinguishes between two outcomes of interest: the employment rate and the share of employees (defined as the the number of employed individuals *living* in a given place relative to the local working age population). Overall, I find no evidence that NGA network availability affects the municipality-level employment rate in the first three years after provision. This implies that the availability of NGA networks does not seem to constitute an important determinant for the creation (or destruction) of local jobs and the location decision of firms in the short run. With respect to the share of employees, however, I find a positive and statistically significant effect. The point estimate of my preferred specification indicates that an increase in NGA network availability by 5 percentage points leads to an increase in the number of employed people living (but not necessarily also working) in a municipality by 0.09 percentage points.

By evaluating the short-term effects of NGA network availability, this paper contributes to an emerging literature investigating the differential effects of various broadband speeds on local employment outcomes. It is closest to Briglauer et al. [2019] who consider the case of a federal state aid programme in Bavaria, Germany, to evaluate the indirect labour market effects of different broadband speeds. Distinguishing between three different categories of available bandwidths (with 16 MBit/s and more being the highest category), they find that improved broadband infrastructures have no effect on the employment rate but positively affect the number of qualified jobs and the share of employees. For the US, Ford [2018] finds that higher broadband speeds do not have a different effect on county-level employment than lower bandwidths (that is, 25 MBit/s compared to 10 MBit/s) .[3]

This paper also contributes to the broader literature evaluating whether and how the Internet affects local labour markets. Among others, this includes studies analysing the effects of earlier broadband technologies in Germany [Fabritz, 2013; Czernich, 2014], the USA (see, e.g., Kandilov and Renkow [2010], Kolko [2012], Atasoy [2013], Whitacre et al. [2014*a*,*b*], and Kim and Orazem [2016]), Italy [Canzian et al., 2015], and Norway [Akerman et al., 2015]. Furthermore, the aggregate effects considered in this paper may be indicative for the effects to be expected at the firm level. Studies focussing on the firm-level broadband effects include Bertschek et al. [2013] for innovation as well as Colombo et al. [2013], De Stefano et al. [2014], Canzian et al. [2015], and Haller and Lyons [2015] for outcomes measuring employment and productivity.[4]

The remainder of this paper is organised as follows: In Section 2.2, I briefly describe the architecture of the networks used for the provision of speeds of at least 50 MBit/s. In the third section, I outline the empirical framework and illustrate the identification approach proposed. Section 2.4 summarises the data and describes sample selection. The fifth section reports the baseline estimation results as well as two additional analyses investigating the composition of compliers and effect heterogeneity. Section 2.6 discusses the main findings and concludes.

2.2 Architecture of NGA Networks: A Brief Overview

To empirically assess the labour market effects of NGA network availability, I take advantage of the fact that most of the technologies currently used for the provision of speeds of at least 50 MBit/s require the installation of a local optical fibre network. This section illustrates the structure of these networks and provides a brief overview of the technical features that are most relevant for the identification approach described in the next section.

In Germany, telecommunication operators provide access to speeds of at least 50 MBit/s in downstream data transfer rates via different types of fixed and mobile technologies.

[3] In a correction of a paper by Bai [2017], Whitacre et al. [2018] come to similar conclusions.

[4] For more comprehensive reviews of the literature, see Niebel et al. [2016] and the Evidence Review "Broadband" by the What Works Centre for Local Economic Growth [2015*a*].

According to the Bundesnetzagentur, the vast majority of connections in this speed category is currently realised via very high-speed digital subscriber line technologies (VDSL) and signal transmission over the cable TV network (CATV) [Bundesnetzagentur, 2018*a*].[5] In 2016, these technologies accounted for a combined market share of 45% among all fixed line broadband connections (i.e., independent of bandwidth) whereas fibre to the home/building (FTTH/B) only made up less than 2% [Bundesnetzagentur, 2017]. At least partly, these differences in take-up also reflect disparities in regional availabilities. In 2016, only 7.1% of all German households had access to FTTH/B whereas VDSL and CATV were technically feasible for 28.3% and 63.5%, respectively [TÜV Rheinland Consulting GmbH, 2016].

In terms of NGA roll-out, VDSL and CATV offer significant cost advantages. In contrast to FTTH/B, the provision of speeds of at least 50 MBit/s based on both technologies does not require that each home or building is connected via individual optical fibre lines. Instead, VDSL and CATV facilitate signal transmission along the "last mile", that is, between the households and the interface to the next layer of the internet architecture, based on pre-existing networks. While VDSL uses the copper wires of the public switched telephone network (PSTN), CATV is based on the coaxial cables installed for cable TV. As optical fibre lines usually have to be rolled out subsurface in Germany, employing these pre-existing infrastructures implies significant cost reductions. According to a report commissioned by the Federal Ministry for Economics and Energy, covering all German households with FTTH/B is estimated to cost between 85 and 94 billion Euros whereas using a mix of VDSL, CATV, and mobile technologies would only amount to 20 billion Euros [TÜV Rheinland Consulting GmbH, 2013].

Figure 2.2.1 illustrates the architecture of the two most commonly used types of NGA networks. To provide VDSL, the PSTN has to be changed in two central ways. First, the copper wires between the main distribution frame (MDF) and the service area interfaces (SAI) have to be replaced with optical fibre cables. Second, the digital subscriber line access multiplexer (DSLAM) has to be transferred from the MDF to the SAI. In combination, these two changes substantially reduce the usage of copper wires within the access network. Thereby, line loss is reduced and higher bandwidths become technically feasible. The bottom half of the picture illustrates a hybrid optical fibre-coaxial-network (HFC). Within this type of network, a process called "segmentation" applies a similar principle to increase the diffusion of speeds of at least 50 MBit/s. To segment the cable TV network, the cable company first installs a so-called "bypass", that is, it supplies an amplifier with an optical fibre cable from the hub.[6] Consequently, this amplifier is converted into a fibre node. As

[5] VDSL is defined according to the ITU-T G.993.2 standard, that is, "VDSL2".

[6] Hubs are often located in the same buildings as MDFs.

Figure 2.2.1: Network Architecture Used for VDSL and CATV Provision

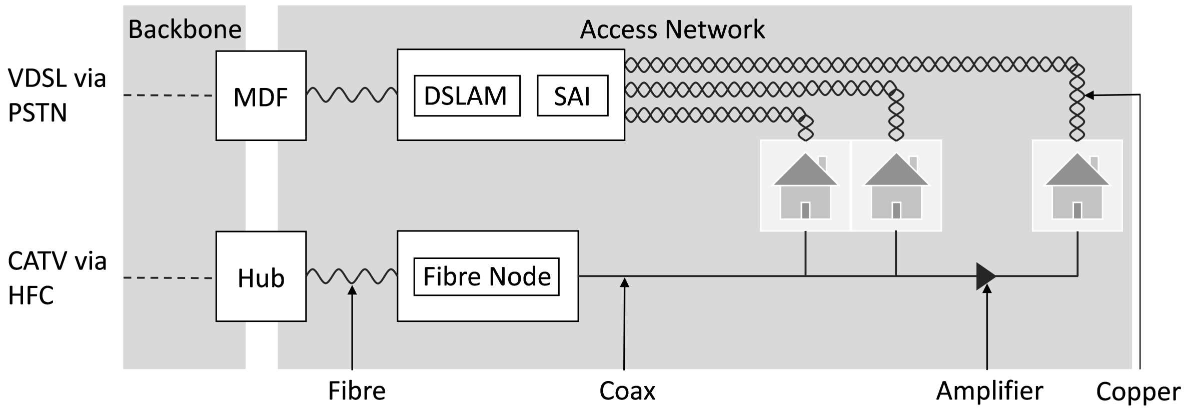

Note: This figure illustrates the architectures of two types of NGA networks. In the upper panel, speeds of at least 50 MBit/s are provided via VDSL based on a hybrid network that combines optical fibre cables with copper wires. In the lower panel, speeds of at least 50 MBit/s are provided via CATV based on a hybrid optical fibre-coaxial-network.

the bandwidth of coaxial cables has to be shared among all connected users (that is, it is a "shared medium"), adding another branch to the tree structure of the cable TV network increases available bandwidths as less users have to share the same coaxial cable [Schnabel, 2015].

In general, segmentation does not represent the only way to increase the availability of speeds of at least 50 MBit/s via cable TV networks. Once the amplifiers of a local cable TV network have been upgraded to support a return channel, the provision of Internet services becomes technically feasible as data can be transferred up- and downstream. Depending on the transmission technologies and standards[7] employed, the small line loss associated with coaxial cables allows cable companies to provide download speeds of up to 100 MBit/s and more without necessarily having to segment the network. However, if the number of customers in a branch surpasses a given threshold and the increased demand for bandwidth cannot be compensated by expanding the frequency range, segmentation may be required.

2.3 Empirical Framework and Identification Strategy

2.3.1 Hypotheses, Baseline Specification, and Endogeneity Concerns

Considering the employment rate and the share of employees as outcomes of interest, this paper evaluates the labour market effects of NGA network availability in the short run. With respect to the employment rate, sign and magnitude of the effects to be expected

[7] In broadband cable networks, the data over cable service interface specification (DOCSIS) represents the internationally accepted standard for data transfer. In Europe, the standards EuroDOCSIS 3.0 and EuroDOCSIS 3.1 are currently used and permit data transfer at speeds of at least 100 MBit/s.

depend on whether the availability of NGA networks predominantly serves as a substitute or complement for local labour inputs. For instance, NGA network availability may help local firms to develop or adopt technologies that substitute local workers. However, it may also lower entry barriers for businesses [Kim and Orazem, 2016] and therefore increase the dynamic of the local labour market.

Similarly, it is *ex ante* not clear whether or how the availability of NGA networks affects the share of employees. On the one hand, having access to very high-speed Internet at home may offer the opportunity to work remotely. As home office constitutes an effective means to reduce commuting costs, this opportunity may allow workers to live in areas where they do not work [Briglauer et al., 2019]. Additionally, NGA network availability may enable job seekers to participate in online training courses that can help them find (better) matching jobs in their regional labour markets [Kuhn and Skuterud, 2004; Kuhn and Mansour, 2014]. On the other hand, considering the availability of NGA networks from a consumer's perspective, it may also be argued that access to very high-speed Internet can have negative effects on local productivity as it provides new entertainment opportunities [Falck et al., 2014]. Even if the time spent on entertainment does not change overall, meeting this demand online rather than offline can reduce the residual demand available to local producers and retailers. Given the data at hand, this paper investigates the employment effects of NGA network availability in the aggregate as these effects may be of immediate relevance for local policy makers who must decide on infrastructure investments.

To estimate how NGA network availability affects the two outcomes of interest, I combine a stacked difference model with a novel IV approach. The baseline specification is of the following form:

$$\Delta Y_{mrt} = \beta \Delta NGA_{mrt} + \phi' \Delta \mathbf{X}_{mrt} + \omega_m + \Delta(\omega_r \times \omega_t) + \Delta \epsilon_{mrt}, \tag{2.3.1}$$

where Y_{mrt} denotes labour market outcome Y in municipality m, regional labour market r, in year t.[8] NGA_{mrt} is the central explanatory variable and measures the share of households with fixed line access to download speeds of at least 50 MBit/s. $\mathbf{X}_{mrt}$ is a vector of time-varying control variables. ω_m, ω_r, and ω_t are fixed effects at the levels of municipalities, regional labour markets, and years, respectively. ϵ_{mrt} denotes the error term. Δ represents a time series operator that calculates annualised changes between the years of interest. For the estimation of Equation (2.3.1), I stack the differences for the two time periods from 2011 to 2013 and from 2013 to 2016. That is, I rely on a balanced panel data set that contains

[8] This size of municipalities varies considerably within and across federal states. If a group of municipalities is organised in a municipality association, I aggregate the municipality-level observations to the level of municipality associations to ensure that observed units are comparable. As such, the term "municipality" is used interchangeably with the term "municipality association".

two observations per municipality with each observation representing a time-differenced value.

For the empirical analysis, I restrict my attention to a sample of 1,625 municipalities that gained access to NGA networks *after* 2013. That is, I exclude those municipalities for which NGA networks were available *before* 2014. By focussing on these municipalities, Equation (2.3.1) effectively implements a difference-in-differences (DiD) design with varying treatment intensity and municipality-specific linear time trends. The most important advantage of this design is that it allows to identify the municipality-level fixed effects ω_m (i.e., the linear time trends) from variation in the first period (i.e., 2011–2013) that cannot be driven by NGA network roll-out. Therefore, I can separate actual effects of NGA network availability from various types of (unobserved) underlying trends, including lagged effects of the introduction of other Internet speeds.[9]

Furthermore, by introducing the interacted fixed effects $\omega_t \times \omega_r$, I also control for arbitrary shocks at the level of regional labour markets. To see why this may be important, consider the case in which a local administration decides to provide public subsidies for NGA network roll-out as a counter-cyclical measure during a business cycle downturn. In this case, the business cycle shock is positively correlated with NGA network diffusion but negatively associated with the current labour market conditions. As a consequence, the estimated effects for NGA network availability could be downward biased if this correlation was not captured by the interacted fixed effects. Additionally, as most (subsidy) programmes in Germany are designed and financed by either the national government or the federal states, controlling for these effects at the level of regional labour markets should be sufficient to capture general changes in public policy at the national and sub-national level.

However, while the fixed effects in Equation (2.3.1) account for linear trends and aggregate shocks, they leave room for important endogeneities. For example, one might be concerned that profit maximising telecommunication operators react to regional differences in the willingness to pay for broadband services and, therefore, only install NGA networks in places where returns are high. If the willingness to pay varies over time and is also correlated with local labour market outcomes, omitting this factor could bias the estimated coefficient for NGA network availability. Similarly, the uptake of subsidy programmes at the municipality level may depend on the absorptive capacity of the local administration. Even if the regional labour market fixed effects account for changes in the availability of these programmes, differential uptake may still render NGA network roll-out endogenous. Considering these two endogeneity concerns jointly, it can be seen that OLS estimates of

[9] Note that by including municipality-level fixed effects in a stacked difference model, I estimate a model that is very similar to a three-period model in levels employing unit specific fixed effects and time trends. That is, a model of the following form: $Y_{mrt} = \beta NGA_{mrt} + \phi' \mathbf{X}_{mrt} + \omega_m \times t + \omega_r \times \omega_t + \epsilon_{mrt}$. The main difference between these two models regard the assumptions made on the error term [Wooldridge, 2002].

Equation (2.3.1) may not only be biased, but it is also hard to determine in which direction the bias should be expected. While the willingness to pay is likely positively correlated with local income levels (and, hence, labour market outcomes), the uptake of subsidy programmes may only be necessary in regions that are financially less well off. Addressing these and related concerns, the next section describes a novel identification strategy that exploits differences in the spatial dispersion of local settlements to identify and use plausibly exogenous variation in NGA network diffusion.

2.3.2 Minimum Spanning Trees and Identifying Variation

As described in Section 2.2, most of the currently used types of NGA networks require the installation of local optical fibre networks that serve specific access points in the streets (e.g., within the PSTN, these are often the SAIs). From these access points, signal transmission to and from the customers is typically realised via pre-existing networks such as the copper wires of the PSTN or the coaxial cables of the TV network. As the roll-out of optical fibre cables in Germany is costly (i.e., because they are typically installed subsurface), deploying these new optical fibre cables constitutes the most important impediment for the diffusion of NGA networks. Furthermore, due to differences in population density, the per capita costs for upgrading the PSTN and the cable TV networks are typically smaller in urban areas than in the country side [Hafner et al., 2015]. Using geographic information systems (GIS), I employ a novel three step procedure to consistently measure these settlement-structure induced differences in NGA network deployment costs.

In the first step, I obtain and combine geographic data from Corine Land Cover (CLC) and OpenStreetMaps (OSM) to comprehensively map economic and residential land use.[10] As OSM is not solely based on satellite images but also incorporates data from manual surveys, GPS devices, and other sources, combining these two data sets allows mapping "inhabited areas" in a very detailed manner. In the second step, I approximate the locations of roughly 360,000 SAIs. According to TÜV Rheinland Consulting GmbH [2013], this is the total number of SAIs currently installed in the PSTN. My approximation is based on the following assumption: If all economic and residential activities were evenly distributed across inhabited areas, then one SAI would on average serve an inhabited area of $0.1km^2$.[11] Correspondingly, I intersect all inhabited areas in Germany with a raster of quadratic tiles that have an edge length of $316.23m$ ($=\sqrt{100,000m^2}$). From the resulting intersections, I only keep those that cover at least $0.05km^2$ of inhabited areas and calculate the geometric

[10] I use the CLC wave from 2018 and data extracted from OSM in October 2018. According to CLC and OSM, 11% of the German territory are used for economic or residential purposes.

[11] As the sum of all inhabited areas in Germany equals $39{,}187.4km^2$, this result follows from: $\frac{39{,}187}{360{,}000} \approx 0.1$. The length of the edge of a $0.1km^2$ square is approximately $316.23m$ long.

centroids within each remaining tile to approximate the SAI locations. In the following, I will refer to these centroids as "nodes". To the degree that the structure of the PSTN resembles local HFC networks, this approximation should also be correlated with the (potential) locations of fibre nodes or amplifiers. Note that even if the actual locations of the SAIs were publicly available, approximating the SAI sites only based on land use offers the advantage that the locations determined by this procedure should be less affected by any other, potentially unobserved and endogenous, factors.

Finally, using Prim's [1957] algorithm, I calculate a minimum spanning tree (MST) within each municipality to derive a measure for local NGA network roll-out costs.[12] In graph theory, a MST describes a subset of edges that connects all nodes within a weighted undirected graph such that no cycles are created and the weighted sum of all edges is minimised. For each municipality, the weighted undirected graph used for the calculation of the MST consists of all the nodes approximated above and the geodesic distances between them (i.e., so-called edges). By using geodesic distances, I abstract from all geographic features (e.g., rivers, mountains, trees) as well as the availability of economic infrastructures (e.g., streets, rail ways, sewage systems). Although all of these aspects can and likely do affect the deployment costs for new optical fibre lines, disregarding these factors offers the virtue that differences in the resulting MSTs are solely driven by variation in the spatial dispersion of local settlements and not any other, potentially endogenous, factors. In other words, this approach trades off the practical relevance of the MSTs against the number of factors that determines them. For the regression analyses below, I aggregate both the nodes and the edges within each municipality.

Incorporating these hypothetical networks in an IV framework, I explain the regional diffusion of NGA network availability over time as a function of network cost (i.e., total edge length) while simultaneously controlling for network size (i.e., the number of nodes). The first stage equation is of the following form:

$$\begin{aligned}\Delta NGA_{mrt} =& \Gamma\Delta[ln(edges_m) \times \mathbf{1}\{t = 2016\}] + \Phi_1\Delta[ln(nodes_m) \times \mathbf{1}\{t = 2016\}] \\ &+ \Phi_2'\Delta\mathbf{X}_{mrt} + \Omega_m + \Delta(\Omega_r \times \Omega_t) + \Delta\xi_{mrt},\end{aligned} \tag{2.3.2}$$

where the variables $edges_m$ and $nodes_m$ denote total edge length and number of nodes, respectively. As Equation (2.3.2) again describes a stacked difference model, I interact these NGA network-specific measures for settlement structures with period-specific fixed effects. Considering the time interval from 2011 to 2013 as the period of reference, the coefficients

[12] While this problem could also be solved with other algorithms, one advantage of this approach is that Prim's algorithm is greedy. As such, it makes locally optimal decision at each step rather than optimising the final outcome. This may well reflect the approach of local telecommunication providers if administrative frictions such as municipality borders or local planning permissions prohibit more global approaches.

Figure 2.3.1: MST for Stockstadt am Rhein

Note: This map illustrates a MST for Stockstadt am Rhein. The MST consists of 25 nodes and has a total edge length of 8.8km. Data: CLC and OSM. Map: Borders of municipalities as of December 31, 2015, obtained from the Service Centre of the Federal Government for Geo-Information and Geodesy.

Figure 2.3.2: MST for Ranstadt

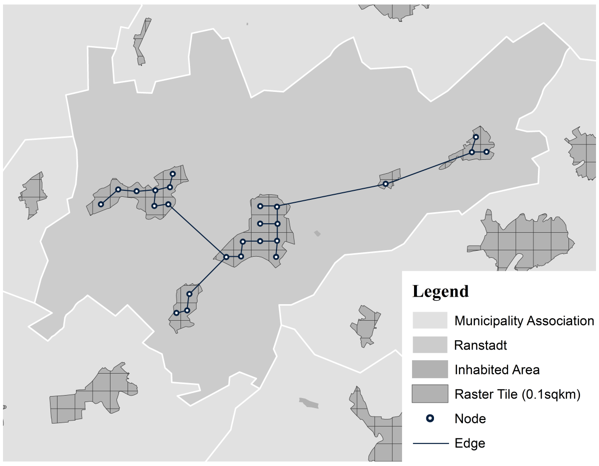

Note: This map illustrates a MST for Ranstadt. The MST consists of 25 nodes and has a total edge length of 11.8km. Data: CLC and OSM. Map: Borders of municipalities as of December 31, 2015, obtained from the Service Centre of the Federal Agency for Geo-Information and Geodesy.

Γ and Φ_1 describe the effects of $ln(edges_m)$ and $ln(nodes_m)$ relative to the period when NGA networks were not available. My identification approach builds on the idea that the interaction including the total edge length serves as an excluded instrument whereas the one including the number of nodes constitutes an included instrument (i.e., a control variable). That is, the effect of NGA network availability is only identified from variation in the period-specific effect of spatial dispersion but not the size of the local settlements. While one might be concerned that my NGA network-specific measure for spatial dispersion is correlated with general agglomeration tendencies, the stacked difference model alleviates these concerns as long as the effects of agglomeration remain relatively stable over the short time period considered (i.e., are constant or follow a linear trend) . Furthermore, I address this concern directly by including annualised changes in population density, separately measured within inhabited areas and across the entire municipality, as baseline controls in all specifications reported in the next section. In a robustness check, I also test whether the regression results are sensitive to the inclusion of a comprehensive set of variables measuring agglomeration in various dimensions.

Figures 2.3.1 and 2.3.2 illustrate the idea underlying the identification approach graphically. The figures show two municipalities – Stockstadt am Rhein and Ranstadt – that are both located in the same regional labour market in Hesse, represent small towns, have between 5,000 and 6,000 residents, and, given the assumption of evenly distributed economic and residential activities, require the same number of nodes (25) to provide full coverage with speeds of at least 50 MBit/s. With respect to my identification approach, the most important difference between these two municipalities is that their inhabited areas are unequally dispersed across space. While the total length of edges required to connect all nodes only amounts to 8.8km in Stockstadt am Rhein, it is 3km longer in Ranstadt. As the deployment costs for optical fibre networks are directly related to the total length of optical fibre cables installed, this relative difference implies a significant impediment for the diffusion of speeds of at least 50 MBit/s in Ranstadt. Until 2016, only 9% of the households in Ranstadt gained access to NGA networks whereas this became technically feasible for 37% of the households in Stockstadt am Rhein.

Furthermore, these figures also help clarifying that the exact number of nodes used in Equation (2.3.2) does not affect the validity of my identification approach (as long as both measures are based on the same MST). To see why, assume I would cut the size of the tiles in half, that is, double the number of resulting nodes. As the geographic extent of the local settlements remains constant, this change would affect all municipalities proportionally and therefore not alter the comparisons made in Equation (2.3.2). In Appendix B.1, Figure B.1.2 combines six histograms that illustrate the cross-sectional distributions of the number of nodes and the total edge length for Germany overall and the municipalities included in the sample, respectively. In the sample, the average municipality consists of 59.07 nodes and requires 30.91km of optical fibre cables to span a network that serves all inhabited areas with speeds of at least 50 MBit/s. The minimum (maximum) number of nodes and total edge length are 3 (436) and 0.51km (213.07km), respectively. Figure B.1.2 shows that the distributions of the number of nodes and the total edge length are highly skewed to the right and only approach the shape of a normal distribution after a log transformation.

2.4 Data

2.4.1 Diffusion of NGA Networks in Germany between 2011 and 2016

To measure the regional diffusion of NGA network availability between 2011 and 2016, I use municipality-level data from the German broadband atlas (*Breitbandatlas*). The Breitbandatlas is an administrative collection of self-reported information of telecommunication operators that provides detailed measures for regional Internet coverage in Germany.[13] If municipalities are organised in municipality associations, I use population weights to aggregate the municipality information. By aggregating, I ensure that observed units are comparable across federal states as the size of municipalities differs strongly across state-level borders.

[13] Until 2018, the Breitbandatlas was produced by TÜV Rheinland Consulting GmbH on behalf of the Federal Ministry of Transport and Digital Infrastructure (BMVI). See, for example, TÜV Rheinland Consulting GmbH [2018] for a description of the methodology.

For the sake of brevity, I use the term "municipality" to refer to both administrative entities.[14]

The Breitbandatlas does not distinguish between the availabilities of individual access technologies (e.g., ADSL, VDSL, FTTH/B, or CATV), but only provides superordinate information on fixed and mobile technology groups. With respect to fixed technologies, the Breitbandatlas measures the shares of households that are covered at minimum downstream data transfer rates of 1 MBit/s, 16 MBit/s, and 50 MBit/s. The availability of mobile technologies is measured by the share of households that can access the Internet with at least 2 MBit/s via Long Term Evolution (LTE). The share of households with fixed line access to download speeds of at least 50 MBit/s serves as my measure for the availability of NGA networks. As this variable is measured in terms of household rather than firm coverage, it is a direct measure for NGA network availability at home but may not fully capture access possibilities at place of work. However, taking into account that not all firms can afford to lease or build private (optical fibre) lines, this variable also provides a useful approximation for businesses that have to rely on the same infrastructures also available to households.

Figure 2.4.1 illustrates the change in the regional availability of NGA networks between 2011 and 2016 at the municipality level. The blue-coloured municipalities represent approximately 36% of all German municipalities (1,625 out of 4,538) and constitute the sample described in Section 2.3.1. They are selected based on the following four restrictions: First, I exclude 2,550 municipalities that received access to NGA networks before 2014 as late deployment constitutes a necessary condition for the validity of my identification strategy. For the purposes of this paper, a municipality is considered to have no access to NGA networks if the share of households covered with at least 50 MBit/s in download speed is smaller than 5%.[15] Second, I drop 166 municipalities that have missing values in at least one of the two outcomes of interest due to data privacy restrictions. Third, I exclude 76 municipalities that represent unincorporated areas ("Gemeindefreie Gebiete"). Finally, I exclude 121 further municipalities. Of these, 87 can be characterised as outliers[16] and 34 constitute singleton groups. While the exclusion of outliers is necessary to account for the high degree of skewness in both outcomes of interest, I follow Correia's

[14] As of December 31, 2015, the Federal Republic of Germany was organised in 11,165 municipalities. Of these, 3,269 were independent whereas 7,896 municipalities belonged to one of 1,269 municipality associations. Therefore, the sample of municipalities considered in this paper is selected from a population of 4,538 municipalities/municipality associations.

[15] In Table B.2.4 in Appendix B.2, I conduct a robustness check to show that this assumption does not affect the qualitative results of the regression analyses.

[16] To identify outliers, I take long differences (i.e., from 2011 to 2016) of the two employment related outcomes of interest. If the long difference of a municipality is below (above) the 1^{st} (99^{th}) percentile with respect to at least one of the corresponding distributions, it is considered to be an outlier.

Figure 2.4.1: Change in NGA Network Availability

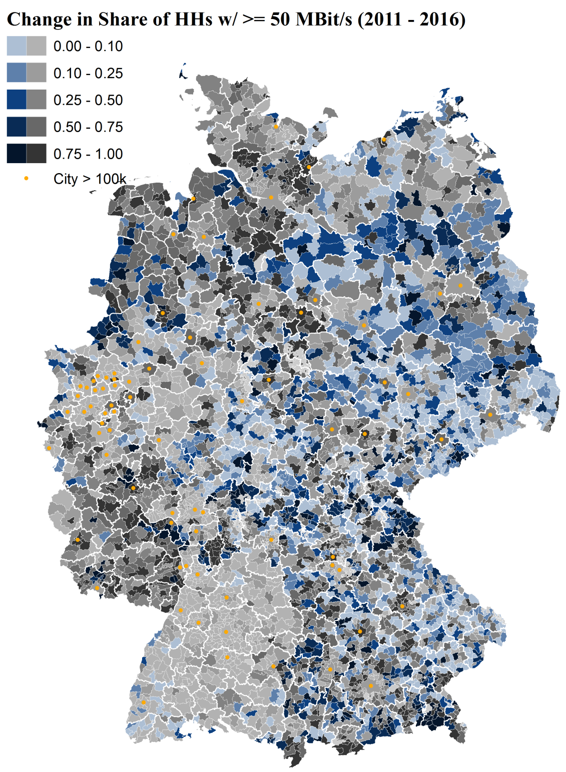

Note: This map illustrates the change in the regional availability of NGA networks between 2011 and 2016 at the level of municipalities. The availability of NGA networks is measured by the share of households for which a downstream data transfer rate of at least 50 MBit/s via fixed lines is technically feasible. Broad white lines define regional labour markets while thin white lines delineate municipalities. Orange dots mark cities with at least 100,000 residents. Gray shaded areas are not included in the estimation sample. Data: Breitbandatlas (2017). Map: Borders of municipalities as of December 31, 2015, obtained from the Service Centre of the Federal Agency for Geo-Information and Geodesy.

Figure 2.4.2: Availability of NGA Networks over Time

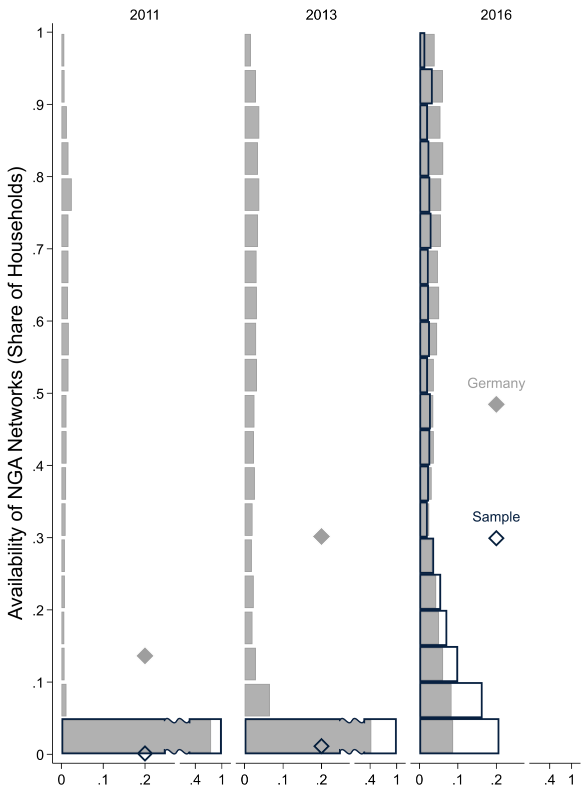

Note: These graphs show frequency distributions that illustrate the regional availability of NGA networks in Germany in 2011, 2013, and 2016. The availability of NGA networks is measured by the share of households for which a downstream data transfer rate of at least 50 MBit/s via fixed lines is technically feasible. Grey and blue bordered bars span 5 percentage point bins of NGA network availability and denote frequencies for Germany overall and the municipalities included in the sample, respectively. Grey and blue bordered diamonds represent averages for Germany overall and the municipalities included in the sample, respectively. Data: Breitbandatlas (2017).

[2015] recommendation to drop singletons as maintaining them could give rise to invalid inference.

In Figure 2.4.1, it is striking to see that the availability of NGA networks in large cities as well as surrounding areas only changed relatively little between 2011 and 2016 (e.g., the light grey areas in the west and the south-west). For example, considering the highly agglomerated Ruhr area, the NGA network availability in most municipalities changed by less than 10 percentage points over the sample period. These relatively small changes are explained by the fact that high rates of NGA network diffusion in 2011 were almost exclusively restricted to agglomerated areas. In contrast, NGA networks only became available in provincial towns and rural areas during the following years. Focussing on the blue shaded areas, this map also shows that the municipalities included in the sample are primarily concentrated in Bavaria, Hesse, and the federal states of former East Germany. In total, more than 80% of the municipalities included in the sample are located in these areas.

Figure 2.4.2 shows the diffusion of NGA networks in Germany (grey) and the sample municipalities (blue bordered) in terms of frequency distributions over time. It focusses on the years 2011, 2013, and 2016 as these years confine the time intervals considered in the regression analyses.[17] As the vast majority of German municipalities did not have access to the Internet at speeds of at least 50 MBit/s in 2011, the horizontal axes in all three vertical panels are cut off at 25 percent and subsequently compressed to enable a detailed illustration of availabilities at all points of the distributions. With respect to the average municipality in Germany (grey diamonds), this figure shows that the availability of NGA networks increased by more than 30 percentage points between 2011 and 2016. While the municipalities included in the sample (blue bordered diamonds) did not have access to NGA networks before 2014, they experienced an increase in availability of similar magnitude until the end of 2016.

Turning to the shape of the availability distributions, it can be seen that both groups approach a u-shape over time. This pattern seems to be driven by strong decreases at the lower ends of the distributions as well as strong increases at the upper ends. In other words, once NGA networks become available in a region with previously low coverage rates, very high-speed Internet tends to become technically feasible for larger shares of the local population. This observation is consistent with the technical feature of NGA networks that most of the technologies currently used for the provision of speeds of at least 50 MBit/s do not require the installation of new optical fibre lines to every home or building, but only to the specific access points (e.g., VDSL). From these points, clusters of households can be served via pre-installed networks.

2.4.2 Labour Market Outcomes and Regional Control Variables

To quantify the labour market effects of NGA network availability, I obtain data on employment from the Federal Employment Agency (*Bundesagentur für Arbeit*). As described before, I separately consider the employment rate and the share of employees. Distinguishing between these two outcomes is vital because the availability of NGA networks may not only constitute an infrastructure that can affect the creation (or destruction) of jobs where people work, but also impact where and how they live. On the one hand, it may help employees to work remotely from home or provide access to educational content and, therefore, serve productive purposes. On the other hand, NGA network availability may

[17] Figure B.1.1 in Appendix B.1 shows these frequency distribution for all years to provide a complete overview on the NGA network deployment timing.

Table 2.4.1: Descriptive Statistics

	2011		2013		2016	
Panel I: Levels	Mean	SD	Mean	SD	Mean	SD
Outcome Variables (in percent)						
Employment rate	36.45	19.23	37.88	19.87	40.26	20.76
Share of employees	56.84	4.60	58.33	4.52	61.15	4.63
NGA Network Availability (share of households)						
Fixed line access to ≥ 50 MBit/s	0.00	0.00	0.01	0.01	0.30	0.29
	$\Delta^1_{2013-2011}$		$\Delta^2_{2016-2013}$		Avg. of (Δ^1, Δ^2)	
Panel II: Annualised Differences	Mean	SD	Mean	SD	Mean	SD
Outcome Variables (in percent)						
Employment rate	0.72	1.19	0.79	0.96	0.75	1.08
Share of employees	0.75	0.54	0.94	0.51	0.84	0.54
NGA Network Availability (share of households)						
Fixed line access to ≥ 50 MBit/s	0.01	0.01	0.10	0.10	0.05	0.08
Regional Labour Markets (RLMs)	184		184		184	
Municipalities	1,625		1,625		1,625	

Note: This table reports descriptive statistics for the two labour market outcomes of interest and the availability of NGA networks. The statistics refer to the municipalities included in the sample (see Section 2.4.1). The definitions of the variables are summarised in Table B.2.1 in Appendix B.2.

constitute a consumptive amenity. For both outcomes, employees are measured relative to the local working age population, that is, the number of residents between 15 and 64 years old (in percent).

Table 2.4.1 reports descriptive statistics for the two employment outcomes of interest and NGA network availability. As before, I restrict my attention to the years 2011, 2013, and 2016 as these years confine the time periods considered in the empirical analysis. Looking at the averages in Panel I, it can be seen that the employment rate and the share of employees increased steadily between 2011 and 2016. These labour market improvements reflect the continuous upswing of the German economy following the financial crisis in 2008 and 2009. In terms of annualised differences (Panel II), both outcomes changed by similar magnitudes. Considering the average year-to-year changes in the third column, the employment rate increased by 0.75 percentage points per year whereas it was slightly higher (0.84) for the share of employees. Focussing on the changes in the availability of NGA networks among the municipalities included in the sample, it can be seen that the share of households covered by NGA networks increased on average by 10 percentage points per year after 2013. Averaging across both differences, this increase translates to a change of 5 percentage points per year across the entire sample period.

I complement the data on employment and NGA network availability with characteristics of the local population as well as measures indicating the "reachability" of various infrastructure facilities in the region.[18] Data on population characteristics comes from the Genesis database of the Federal Statistical Office. It includes information on total population size, population density (separately measured within inhabited areas as well as the entire area of the municipality and expressed in terms of residents per km^2), the share of female residents, the share of the working age population, the share of people older than 64 years, and the net migration rate which is defined as the difference between regional immigrants and emigrants relative to the number of residents. To account for regional differences with respect to the availability of various infrastructure facilities, I obtain data from the Federal Institute for Research on Building, Urban Affairs and Spatial Development (BBSR). At the municipality level, the BBSR provides reachability indicators that measure the area-weighted average travel times (in minutes) that it takes by car to reach the next motorway, larger train station, airport, hospital, and economic centre.[19] All of these measures are time invariant and based on 2015.

Table B.2.2 in Appendix B.2 provides summary statistics for all variables considered and compares the estimation sample (Columns (3) – (5)) to the general population of municipalities across all years (Column (1)) and the three sample years (Column (2)). Overall, the estimation sample includes 1,625 municipalities with complete information on both outcomes and all explanatory variables of interest. These municipalities are located in 184 out of 258 regional labour markets (RLMs) in Germany.[20] With respect to the characteristics of the local population, the municipalities included in the estimation sample are smaller (e.g., they have on average 8,000 residents of which 5,300 are at working age), less densely populated, and have lower shares of net migration. This difference in size is also reflected in the shares of different settlement types: the estimation sample includes few large or medium-sized towns and over-represents rural areas as compared to Germany overall. Furthermore, almost half of the municipalities included in the sample are "shrinking" whereas this is only true for 34% in the general population.[21] This discrepancy is

[18] Table B.2.1 in Appendix B.2 provides a complete overview of all variables used as well as their sources.

[19] The BBSR defines economic centres as municipalities that are classified as middle- or higher-order centres according to the system of central places in Germany (`https://www.bbsr.bund.de/BBSR/DE/Raumentwicklung/RaumentwicklungDeutschland/Projekte/Archiv/ZentraleOrte/ZentraleOrte.html?nn=411742`, last accessed: 15 March 2019).

[20] The BBSR distinguishes 258 RLMs in Germany. These are geographic areas that are economically integrated and characterised by a high degree of within-region commuting (`https://www.bbsr.bund.de/BBSR/DE/Raumbeobachtung/Raumabgrenzungen/AMR/Arbeitsmarktregionen.html?nn=443270`, last accessed: 15 March 2019).

[21] The BBSR's classification is based on the long-run development of a set of socio-economic characteristics (`https://www.bbsr.bund.de/BBSR/DE/Raumbeobachtung/Raumabgrenzungen/wachsend-schrumpfend-gemeinden/Begleitinformation.pdf?__blob=publicationFile&v=5`, last accessed: 15 March 2019).

mainly driven by the fact that the sample includes a relatively large share of East German municipalities. In summary, the sample considered in this paper describes a group of relatively rural and economically disadvantaged areas. As sample selection is mainly driven by the degree of NGA network penetration in 2013, however, these sample municipalities are also the ones in the focus of policy makers discussing the provision of broadband deployment subsidies as a means for economic development.

2.5 Results

2.5.1 Baseline Results and Robustness

Table 2.5.1 reports the baseline results for both outcomes of interest. To better understand potential biases, Columns (1) and (2) report OLS and IV results for the most parsimonious specifications considered (i.e., beyond $\Delta[ln(nodes_m) \times \mathbf{1}\{t = 2016\}]$, these specifications only control for annualised changes in population density, separately measured within inhabited areas and across the entire municipality). With respect to the employment rate, it can be seen that the OLS estimate is greater in magnitude than the corresponding IV coefficient; suggesting that the OLS estimate is upward biased. For the share of employees, the bias goes in the opposite direction. An explanation that can be reconciled with this pattern is linked to unobserved (uptake of) digital innovations. On the one hand, these innovations are likely positively correlated with the availability of NGA networks as they increase the demand for improved broadband services. On the other hand, the productivity gains generated by them may be largest where absorptive capacity is high. If absorptive capacity is higher in urban areas, then the fact that population density is positively (negatively) correlated with the employment rate (the share of employees) could give rise to an upward (downward) bias.[22]

Considering the first stage estimates in the lower panel, it can be seen that the coefficients for $\Delta[ln(edges_m) \times \mathbf{1}\{t = 2016\}]$ and $\Delta[ln(nodes_m) \times \mathbf{1}\{t = 2016\}]$ are statistically significant and of opposite signs. The signs of these point estimates are consistent with the intuition that it is more attractive to deploy NGA networks in larger markets (measured by the number of nodes), but more costly in spatially dispersed areas (measured by total edge length). With respect to total edge length, a coefficient of –0.09 implies that doubling the spatial dispersion decreases NGA network availability by 9 percentage points on average. In distributional terms, a doubling of the total edge length corresponds approximately to a move from the 20^{th} percentile to the median.

[22] Pooling across all years, Pearson's correlation coefficient for the relationship of population density and the employment rate is 0.232. For population density and the share of employees, the corresponding coefficient is –0.057.

Beyond the direction of the biases, Table 2.5.1 shows that the IV estimates do not vary strongly across specifications. In Column (3), I introduce annualised changes in the share of the female population, the share of the working age population, the share of people older than 64 years, and the net migration share as additional controls. Accounting for these population specific controls relaxes the conditional independence assumption that the number of nodes in combination with the fixed effects and measures for population density are sufficient to control for all regional differences that simultaneously affect local labour market outcomes and NGA network diffusion. Recalling that the measures derived in Section 2.3.2 are solely based on sheer geographic extent, it is not surprising that accounting for these qualitative aspects also affects the IV coefficients to some degree. Although including them does not affect the results qualitatively, Column (3) represents my preferred baseline IV specification.

Given an average annual increase in the availability of NGA networks by 5 percentage points, an estimated effect of 1.867 implies an annual increase in the share of employees by 0.09 percentage points. Note that I do not find statistically significant effects on the employment rate. As a consequence, these results do not imply that NGA network availability leads to a creation of jobs at place of residence but that municipalities with a higher availability of NGA networks attract more employed individuals as residents than other areas or that they have relatively lower levels of out-migrating employees. In the sample considered, this effect explains roughly 14% of the annual change in the share of employees. Compounded over the whole sample period, this coefficient implies that the availability of NGA networks increased the number of employed individuals living (but not necessarily also working) in a municipality by roughly 29.71.[23]

With respect to the robustness of the results, one might be concerned that the NGA network-specific measure for spatial dispersion is confounded with other types of agglomeration that directly affect economic outcomes (e.g., through pooling, returns to scale, or spillovers). Therefore, Column (4) of Table 2.5.1 investigates whether the total length of edges has explanatory power that goes beyond a rich set of measures for reachability (i.e., the average travel time by car that it takes to reach the next motor way, larger train station, airport, hospital, and economic centre in minutes). As the reachability measures are time-invariant, they enter the regressions as interactions with a period-specific fixed effect. Both quantitatively and qualitatively, the estimated effects are robust to this check. Furthermore, I test whether other types of Internet access affect the results for NGA network availability. For example, this could be the case if current effects of NGA network availability are confounded with the effects of lower speeds. To this end, I run a specification that includes changes

[23] In 2011, the average municipality in the sample had a working age population of 5,304 people. Given a total increase in NGA network availability by 30 percentage points, the average increase in the number of employed individuals can be calculated as: $\frac{0.3\times1.867\times5,304}{100} \approx 29.72$.

Table 2.5.1: Baseline OLS and IV Results

	OLS	IV				
	(1)	(2)	(3)	(4)	(5)	(6)
Outcomes						
ΔEmployment rate	0.706	–0.143	0.379	0.336	0.386	0.377
	(0.476)	(1.558)	(1.544)	(1.949)	(1.588)	(2.005)
ΔShare of employees	0.429	1.484*	1.867**	1.933*	1.862**	1.970*
	(0.235)	(0.895)	(0.820)	(0.993)	(0.837)	(1.014)
First stage						
$\Delta[ln(edges_m) \times$ $\mathbf{1}\{t=2016\}]$	–	–0.092***	–0.092***	–0.081***	–0.090***	–0.079***
		(0.009)	(0.009)	(0.008)	(0.009)	(0.009)
$\Delta[ln(nodes_m) \times$ $\mathbf{1}\{t=2016\}]$	–	0.134***	0.134***	0.114***	0.133***	0.114***
		(0.009)	(0.009)	(0.010)	(0.009)	(0.009)
Kleibergen-Paap rk Wald F statistic	–	107.831	105.395	84.093	101.111	80.350
Controls						
Density	Yes	Yes	Yes	Yes	Yes	Yes
Population	–	–	Yes	Yes	Yes	Yes
Reachability	–	–	–	Yes	–	Yes
Other Internet	–	–	–	–	Yes	Yes
Municipalities	1,625	1,625	1,625	1,625	1,625	1,625
N	3,250	3,250	3,250	3,250	3,250	3,250

Note: This table reports OLS (Column (1)) and IV (Columns (2) – (6)) results for Equation (2.3.1). In all IV regressions, $\Delta[ln(edges_m)\times\mathbf{1}\{t=2016\}]$ serves as the excluded instrument while $\Delta[ln(nodes_m)\times\mathbf{1}\{t=2016\}]$ is an included instrument. Furthermore, all regressions include annualised changes in population density, separately measured within inhabited areas and across the entire municipality ("Density"), and municipality-level as well as period-specific fixed effects as baseline controls. The specification in Column (3) additionally controls for annualised changes in the shares of the female population, the working age population, the population aged 65 and older, and net migration ("Population"). Columns (4) – (6) present robustness checks. Beyond the control variables in Column (3), the specification in Column (4) controls for five reachability measures that are interacted with period-specific fixed effects ("Reachability") whereas the one in Column (5) accounts for annualised changes in the shares of households that have access to 1 MBit/s via fixed lines and 2 MBit/s via LTE ("Other Internet"). The specification in Column (6) includes all sets of control variables simultaneously. All Kleibergen-Paap rk Wald F statistic are above the critical values calculated by Stock and Yogo [2005], rejecting the null hypothesis that the excluded instrument is weak. Standard errors are clustered within regional labour markets and shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

in the shares of households that have access to at least 1 MBit/s via fixed lines as well as to 2 MBit/s via LTE (Column (5)). The results are robust to the inclusion of these variables and a specification in which both the controls for reachability and other Internet speeds are included simultaneously (Column (6)).

In an additional set of robustness tests, I include different sets of fixed effects to control for potential confounders in three important dimensions. First, I generate three dummy variables that indicate whether a municipality is a large or medium-sized city (1), a small town (2), or a rural area (3). By interacting these dummy variables with period-specific fixed effects, I account for arbitrary shocks at the level of these settlement types. This could be important because absorptive capacity may vary at the level of local administrations. Comparing Columns (1) and (2) of Table B.2.3 in Appendix B.2, it can be seen that both point estimates decrease in magnitude and are statistically insignificant. However, for the share of employees, the point estimate is only marginally insignificant (p-value = 0.113) and qualitatively very similar to the baseline specification. In Column (3), I classify how important a municipality is for the supply of goods and services to surrounding areas, that is, whether it represents an economic centre. The effect of NGA network availability on the share of employees is also robust to this check. Although the point estimate for the employment rate decreases, it remains statistically insignificant. Column (4) includes controls for shocks at the level of growing, stagnating, and shrinking municipalities. This could be important if the municipality-level fixed effects fail to adequately account for growth trends. However, this does not appear to be the case. In Column (5), I control for all sets of fixed effects simultaneously and find qualitatively similar effects.

Finally, I investigate whether the sample selection based on a 5% threshold for NGA network diffusion in 2013 affects the results. To this end, Table B.2.4 in Appendix B.2 reports the IV results for my preferred baseline specification (i.e., Column (3) of Table 2.5.1) when using different exclusion thresholds. Although the qualitative results are again similar across specifications, the magnitude of the effect for the employment rate in Column (2) indicates that the results for this outcome are more sensitive to the sample restrictions imposed. In general, it may not be too surprising that the sample used in Column (2) produces the largest differences as imposing a threshold of 2.5% also leads to the largest relative change in sample size (i.e., a reduction of 17.8%.)

As compared to the effects of other broadband speed upgrades in Germany, these estimates are larger than the ones found for lower speeds but still comparable. For example, Briglauer et al. [2019] find that an increase of approximately 25 percentage points in the availability of at least 6 MBit/s increases the share of employees in the average municipality by 6.35. While the difference between their and my point estimates may reflect genuine differences in the types of services and technologies that can be accessed at very high Internet speeds, it can also be the result of methodological differences between both studies. For example, this paper evaluates the effects of NGA network availability during an early phase. If the effects of NGA network availability are non-linear in the degree of diffusion (e.g., because NGA networks are especially important for municipalities with a high share of early adopters), this could imply that I estimate a local average treatment effect (LATE) for early adopters. In contrast, Briglauer et al. [2019] study Bavarian municipalities that participate in a state aid programme and may therefore rather represent late adopters. In the following section, I conduct a complier analysis to evaluate the role of local average treatment effects in determining the magnitude of my estimates.

2.5.2 Complier Analysis and Heterogeneity of Effects

While the robustness tests presented in the previous section may limit concerns related to the validity of the identification strategy, it is *ex ante* not clear whether the IV estimates represent an average treatment effect (ATE) or one that is local to a specific subgroup (i.e., a LATE). Employing the terminology introduced by Imbens and Angrist [1994], the central question is whether the population of municipalities that "complies" with the instrument, that is, delays NGA network roll-out due to a spatially more dispersed settlement structure, is identical with the average municipality. To evaluate this possibility, I conduct a complier analysis that is similar to the one suggested by Akerman et al. [2015]. In a first step, I generate an identifier that distinguishes four mutually exclusive groups of municipalities and classifies whether a municipality is located in an urban or rural area. Among the group of urban municipalities, I include large and medium-sized cities as well as small towns. In a second step, I subdivide each of these two groups in two equally sized subgroups depending on how remote they are. For the purposes of this paper, a municipality is considered to be remote if the travel time needed to reach the next economic centre by car is above the median within the subgroup considered.

To conduct a complier analysis as in Akerman et al. [2015], I would have to split the sample and estimate separate first stage regressions for each group considered. However, using split samples in my application is problematic as this would not only allow the effects of the NGA network-specific measures to vary by group (i.e., Γ and Φ_1), but also all other coefficients, including the fixed effects. As Equation (2.3.2) employs high-dimensional fixed effects, this can play a crucial role because it may affect the number of observations that can be used in the estimation. For example, by only considering relatively central municipalities in urban areas, controlling for regional labour market-by-year fixed effects requires to drop all regions that have only one municipality of this type. Therefore, I employ an approach that is closely related to Akerman et al. [2015] but does not require splitting the sample. In particular, I estimate two first stage regressions of the following form:

$$\begin{aligned}\Delta NGA_{mrt} = {} & \Gamma^{\bar{g}}\Delta[ln(edges_m) \times \mathbf{1}\{t=2016\}] + \Phi_1^{\bar{g}}\Delta[ln(nodes_m) \times \mathbf{1}\{t=2016\}] \\ & + \Phi_2'\Delta\mathbf{X}_{mrt} + \Omega_m + \Delta(\Omega_r \times \Omega_t) + \Delta(\Omega_g \times \Omega_t) + \Delta\xi_{mrt} \end{aligned} \tag{2.5.1}$$

$$\begin{aligned}
\Delta NGA_{mrt} = \Gamma^{1}\Delta[ln(edges_{m}) \times \mathbf{1}\{t = 2016\}] \\
+ \sum_{G=2}^{4} \Gamma^{G}\Delta[ln(edges_{m}) \times \mathbf{1}\{t = 2016\} \times \mathbf{1}\{g = G\}] \\
+ \Phi_{1}^{1}\Delta[ln(nodes_{m}) \times \mathbf{1}\{t = 2016\}] \\
+ \sum_{G=2}^{4} \Phi_{1}^{G}\Delta[ln(nodes_{m}) \times \mathbf{1}\{t = 2016\} \times \mathbf{1}\{g = G\}] \\
+ \Phi_{2}'\Delta\mathbf{X}_{mrt} + \Omega_{m} + \Delta(\Omega_{r} \times \Omega_{t}) + \Delta\xi_{mrt},
\end{aligned} \tag{2.5.2}$$

where $g \in \{1,2,3,4\}$ indicates group membership, $\bar{g}$ denotes estimates for the full sample, and Ω_g is a group-specific fixed effect. Note that by including group-by-year fixed effects, $\Gamma^{\bar{g}}$ in Equation (2.5.1) is only identified from variation within group-by-year clusters. In Equation (2.5.2), I interact the two NGA network-specific measures with group-specific dummies ($\mathbf{1}\{g = G\}$). That is, I estimate an interaction model in which the group of centrally located urban municipalities ($g = 1$) constitutes the group of reference. Based on these estimates, the first-stage coefficient Γ^{g} for group $g \in \{2,3,4\}$ follows from a linear combination of Γ^{1} and Γ^{G} (i.e., $\Gamma^{g} = \Gamma^{1} + \Gamma^{G} \forall G \in \{2,3,4\}$). Finally, I calculate the share of compliers as $Complier^{g} = \frac{\hat{\Gamma}^{g}}{\hat{\Gamma}^{\bar{g}}} \times \frac{N^{g}}{N^{\bar{g}}}$, where N^{g} and $N^{\bar{g}}$ represent the total number of observations in group g and the full sample, respectively.

The results of the complier analysis are reported in Table 2.5.2. As can be seen in Column (4), urban municipalities are over-represented among the population of compliers as compared to their shares in the sample overall. Furthermore, the degree of over-representation is especially large for urban municipalities that are relatively remote. With respect to the employment rate in 2011 (i.e., before the considered municipalities gained access to NGA networks), the complier municipalities in urban but relatively remote areas are smaller than their more central counterparts. If this lower average employment rate reflects a relative lack of productive firms and infrastructures in these areas, an interaction between the availability of NGA networks and these potential shortcomings could explain relatively sizeable effects in the IV regressions.

Finally, I investigate whether NGA network availability has heterogeneous effects with respect to settlement type and relative location. After generating two dummy variables that classify whether a municipality is rural or remote, I estimate a variant of Equation (2.3.1)

Table 2.5.2: Characterising Complier Municipalities

Municipality type	Obs. (1)	Share in Sample (2)	$\hat{\Gamma}$ (3)	Share of Compliers (4)	Employment rate (5)	Share of employees (6)
Urban & central	804	0.25	–0.075	0.26	48.70	56.28
Urban & remote	804	0.25	–0.094	0.33	37.24	56.57
Rural & central	826	0.25	–0.063	0.22	29.46	57.49
Rural & remote	816	0.25	–0.056	0.19	30.69	56.99
Overall	3,250	1.00	–0.072	1.00	36.45	56.84

Note: This table reports the results of a complier analysis based on Equations (2.5.1) and (2.5.2). The share of compliers in Column (4) is calculated as $Complier^g = \frac{\hat{\Gamma}^g}{\hat{\Gamma}^{\bar{g}}} \times \frac{N^g}{N^{\bar{g}}}$, where N^g and $N^{\bar{g}}$ represent the total number of observations in group g and the full sample, respectively, and $\hat{\Gamma}$ are the corresponding estimates listed in Column (3). Columns (5) and (6) report the average employment rates at place of work and at place of residence in 2011, respectively.

in which NGA network availability is not only included as a main effect but also in an interaction term with one of the dummy variables defined.[24]

The results for these interaction models are reported in Table 2.5.3. With respect to the employment rate (Columns (1) – (2)), it can be seen that the point estimates for NGA network availability vary considerably, though not statistically significantly, with the degree of rurality and remoteness. Although the large standard errors prohibit drawing strong conclusions, the negative point estimate for the interaction term for rural municipalities suggests that the availability of NGA networks rather substitutes than complements local workers in these areas.

In contrast, the positive effect for remote municipalities indicates that NGA network availability may act as a substitute for other infrastructures in remote areas. In combination, these results are in line with previous findings in the literature suggesting that (very high-speed) Internet access may not lead to the "death of the city" [Kolko, 2000], but reduce the cost of distance. With respect to the share of employees, there is no evidence that the effect observed before strongly varies between rural or remote municipalities. This implies that the benefits associated with the availability of NGA networks (such as the opportunities to do home office or train for jobs more conveniently) are not restricted to areas where commuting costs are high but apply to all regions equally.

[24] The definition for rurality follows the settlement types classified by BBSR [2018]. The definition for remoteness depends on the travel time that it takes to reach the next economic centre. Travel times are measured by a traffic model of the BBSR (`https://www.bbsr.bund.de/BBSR/DE/Raumbeobachtung/UeberRaumbeobachtung/Kompo-nenten/Erreichbarkeitsmodell/erreichbarkeitsmodell_node.html`, last accessed: 15 march 2019). A municipality is considered to be remote, if its travel time is above the 50^{th} percentile of the corresponding distribution.

Table 2.5.3: Heterogeneity – Rural vs. Urban and Remote vs. Central

	ΔEmployment rate		ΔShare of employees	
	(1)	(2)	(3)	(4)
Explanatory Variables				
Δ NGA	1.143	0.490	2.062**	2.084**
	(1.947)	(2.013)	(0.889)	(1.047)
Δ NGA × Rural	–1.555		–0.314	
	(3.735)		(1.926)	
Δ NGA × Remote		1.128		–0.210
		(3.997)		(1.960)
First stage				
Kleibergen-Paap rk Wald F statistic	12.487	18.406	12.487	18.406
Controls				
Density	Yes	Yes	Yes	Yes
Population	Yes	Yes	Yes	Yes
Municipalities	1,625	1,625	1,625	1,625
N	3,250	3,250	3,250	3,250

Note: This table reports IV results for two interaction models in which NGA network availability is not only included as a main effect (as in Equation (2.3.1)) but also as an interaction with two dummy variables indicating whether a municipality is rural or remote. Across all specifications, $\Delta[ln(edges_m) \times \mathbf{1}\{t = 2016\}]$ and the interaction of this term with the corresponding dummy serve as excluded instruments while $\Delta[ln(nodes_m) \times \mathbf{1}\{t = 2016\}]$ and the interaction of this term with the corresponding dummy denote included instruments. Furthermore, all regressions include annualised changes in population density, separately measured within inhabited areas and across the entire municipality ("Density"), the shares of the female population, the working age population, the population aged 65 and older, net migration ("Population"), and municipality-level as well as period-specific fixed effects as baseline controls. That is, the control variables correspond to the ones used in the preferred baseline specification shown in Column (3) of Table 2.5.1. All Kleibergen-Paap rk Wald F statistic are above the critical values calculated by Stock and Yogo [2005], rejecting the null hypothesis that the excluded instruments are weak. Standard errors are clustered within regional labour markets and shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

2.6 Summary and Conclusions

In summary, the empirical evidence provided in this paper does not support the hypothesis that NGA networks affect the employment rate within the first three years after deployment. This implies that NGA network availability does not seem to constitute an important determinant for the creation (or destruction) of local jobs and the location decision of firms in the short run. This is an important information for policy makers who are concerned about the immediate effects of NGA network availability – or a lack thereof. With respect to the share of employees, I show that the availability of NGA networks has a positive and statistically significant effect. According to my preferred specification, an increase in NGA network availability by 5 percentage points (the average annual increase in the municipalities considered) leads to an increase in the share of employees at place of residence by 0.09 percentage points. Compounding this effect over the sample period, NGA network availability has increased the number of employees living (but not necessarily also working) in the average municipality by 29.72 individuals. A complier analysis reveals that these effects are likely driven by urban, but relatively remotely located municipalities.

The results presented in this paper only denote short-term effects of NGA network availability in the aggregate. As a consequence, they may mask considerable heterogeneity at the firm or individual level and more research is needed to identify how the availability of NGA networks affects local labour market outcomes in the medium and long run. Disentangling these effects provides a promising avenue for future research as it may help to clarify the ambiguity of some of the previous findings in the literature. For example, it is of crucial interest to policy makers whether positive employment effects of faster Internet are only confined to specific regions [Forman et al., 2012] or contingent on firm-level factors [Colombo et al., 2013].

Employing the identification approach pursued in this paper also provides a means to investigate whether and how NGA network availability affects other outcomes of interest. Given that many of the applications requiring the availability of NGA networks are primarily used at home, investigating the short-term effects of NGA network availability on social capital [Bauernschuster et al., 2014] and house prices [Ahlfeldt et al., 2017] may be of particular interest. Finally, it seems possible to translate the identification approach described in this paper to the case of FTTH/B networks. That is, regional differences in the spatial dispersion of local settlement structures are likely not only a determinant for the deployment of the hybrid networks that are currently used in most places, but also for complete optical fibre networks serving connected premises directly.

Chapter 3

Online Shopping and Retail Employment: Evidence from Brick-and-Mortar Bookstores*

3.1 Introduction

The retail industry has undergone substantial change over the past two decades. While prospering malls have been a symbol of successful retailing in the United States for years, they now often fail to lease out their vacant spaces or even have to shut down entirely. In various European regions, the decline of brick-and-mortar retailers has alerted local politicians and real estate developers who are worried about the prospect of deserted city centres and high streets.

Although this situation is often blamed on the rise of online shopping, rigorous empirical evaluations of this relationship are scant. We contribute to closing this gap in the literature by investigating the causal effects of online shopping exposure on brick-and-mortar bookstores in Germany between 1999 and 2013. Our analysis is based on a novel data set that combines detailed geomarketing information on regional shopping behaviour with administrative records from social insurance.

We find that online shopping exposure has substantial negative effects on both the stock of local bookstores and the number of employees therein. For example, an increase in regional exposure to online shopping by one standard deviation reduces the stock of brick-and-mortar bookstores by about 0.11 bookstores per 10,000 residents. Relative to the density of bookstores in 1999, this corresponds to a reduction of 14%. With respect to overall employment, we find effects of similar magnitude. A closer examination further reveals that the effect on overall employment is mainly driven by bookstores in local retail markets where the density of bookstores in 1999 was relatively high. In contrast, overall employment in more concentrated retail markets remains unchanged. This implies that bookstores in

* This chapter is based on joint work with Andreas Mazat and Bastian Stockinger.

rather concentrated retail markets reduce the number of hours worked by their employees while being able to hold overall employment constant.

Our results quantify the effects of the emergence of a new distribution channel in combination with the complementary services it offers. The institutional setting of the German book market allows us to abstract from differences in price, quality, and the set of books that are sold on- and offline for two reasons: First, the German fixed book price law (*Buchpreisbindungsgesetz*) regulates prices for books that are sold to end customers and therefore prohibits price differentiation between retail channels. Second, both a low market share of electronic books (3.9% in 2013) as well as an even smaller share of self-published books (below 2% in 2014) indicate that, at least until the end of our sample period in 2013, the German book market only provided very limited opportunities to differentiate the qualities or sets of books offered via different retail channels.

Identifying the causal effects of online shopping exposure on brick-and-mortar bookstores in a regional context gives rise to endogeneity concerns. For example, availability and usage of the Internet is likely influenced by the same individual and regional characteristics that also drive online shopping and which themselves may have a direct effect on reading and the consumption of books. More generally, any unobserved local shock or regional differences in the composition of the population, as well as agglomeration structures, may affect both brick-and-mortar retailers and the uptake of online shopping. Furthermore, reverse causality may be a problem if a poor local shopping infrastructure pushes consumers to buy online.

To address endogeneity concerns, we estimate a long difference model in combination with an instrumental variables (IV) strategy. The central idea of our IV approach is to construct a Bartik-type instrument that exploits differences in local age structures before the Internet was introduced to explain regional variation in the uptake of online shopping. Within our long difference model, this approach controls for the effects of all time-invariant factors, unobserved shocks that are uncorrelated with (pre-Internet) local age structures, and measurement error. While other endogeneity concerns may persist, we argue that our IV approach is valid, conditional on controlling for agglomeration dynamics and changes in Internet usage over time. Robustness checks, including a comprehensive set of socio-demographic control variables, support the conclusion that our results are valid.

This paper relates to the empirical literature that investigates how supermarket entry affects local retail employment, as well as the survival and size of incumbent retailers. For example, Neumark et al. [2008], Jia [2008], and Basker [2005] consider county-level data in the United States to investigate local retail employment effects of Wal-Mart entry. Even though Basker [2005] finds positive employment effects in the short run, all three papers show that long-run effects on employment as well as the survival of incumbent retail stores are negative. A similar paper by Zhang and Lei [2015] uses city-level data. Borraz et al. [2014] and Igami [2011] instead define very small geographic markets within cities to study the effect of supermarket entry on existing stores. They mostly confirm the previously discovered negative effects on employment and establishment survival. Igami [2011] documents that, unlike large and medium-sized supermarkets, small ones actually profit from entry of large supermarkets as their product portfolios are sufficiently different from larger stores and can even serve as complements.

Online shopping likely distorts spatial equilibria even more than the entry of big box stores because online retailers are not as locally attached as these traditional retailers. While big box stores create local employment, online retailers may generate employment in an area far away from the retail markets served. This possibility is investigated by Goolsbee [2000], Ellison and Ellison [2009], Goolsbee et al. [2010], and Einav et al. [2014], who find that online demand is less dependent on geographic proximity to sellers. In particular, factors such as the price may divert demand from the place of residence to other regions. They further show that customers tend to shop more often in federal states in which sales taxes are low.

We also contribute to the literature that analyses local labour market effects of technological change and digitisation. For example, Acemoglu and Restrepo [2017] investigate the effect of robot usage on local industrial employment, using a Bartik-type instrument that combines variation in historic industry shares across commuting zones in the United States and aggregate trends in robot adoption by industry. Similarly, Graetz and Michaels [2018] and Dauth et al. [2017] construct a Bartik-type instrument to investigate labour market effects of industry-level robot adoption in a panel of 17 countries as well as in regional labour markets in Germany. Similar to other studies, such as Akerman et al. [2015] or Atasoy [2013], they investigate how (different groups of) workers in manufacturing are affected by technological change. In contrast, we focus on employment effects in retail, a sector that is rarely studied. Among the few studies dealing with the retail sector are Basker [2012] and Doms et al. [2004] who examine how the introduction of barcode scanners and other IT equipment affected (labour) productivity in the retail sector. In summary, our results extend the existing literature with respect to multiple aspects and will be particularly useful for urban planners, policy makers, brick-and-mortar retailers, and real estate developers, who are interested in understanding how high streets and city centres will evolve in future.

The remainder of the paper proceeds as follows. Section 3.2 reviews the development of the German book market between 1999 and 2013. Section 3.3 describes the data used in the empirical analysis. Section 3.4 outlines our empirical approach and discusses the validity of our identification strategy. Section 3.5 reports the results. Section 3.6 summaries the main findings and concludes.

3.2 The German Book Market Between 1999 and 2013

Figure 3.2.1 illustrates the market shares of different book distribution channels between 1999 and 2013. Over the whole period, brick-and-mortar retail was the primary distribution channel through which private customers obtained their books. In 1999, sales in traditional bookstores accounted for 59% of total market turnover, whereas not even 1% of German book sales took place online.[1] At the same time, roughly 6% of all books were sold via telephone or catalogues. All other books were sold by warehouses that do not specialise in selling books (9%), other selling points like kiosks or gas stations (5%), book clubs (4%), and publishers that sell directly to institutional customers or private companies (16%).

Due to the rise of online retailers, neither brick-and-mortar bookstores nor companies selling books via telephone or catalogues were able to maintain their initial market shares. The revenues of brick-and-mortar retailers decreased by more than 10 percentage points between 1999 and 2013; sales made via telephone or catalogues saw even sharper decreases and only accounted for 2.3% of the market in 2013. In contrast, online retailers gained substantially over the same period. The first online bookseller to enter the German market was ABC Bücherdienst, which began its online services in 1996. Over the next few years, others followed. However, the online market soon consolidated as many of the initial entrants had already disappeared by 2003. Among those surviving was Amazon, which grew to be the most important player in the German online market for books during the following decade. In 2013, Amazon accounted for almost three-quarter of the 19% market share for online retailers [Schrape, 2011].

With respect to competition between online and offline retailers, the economic literature so far mostly focuses on price differences and product differentiation. Investigating price differences across channels, Brynjolfsson and Smith [2000] and Cavallo [2017] show that, depending on the product considered, online prices tend to be lower. For the computer industry, these price differences are particularly important as Goolsbee [2001] finds that online computer demand is highly elastic to offline price changes. In our context, however, differences between online and offline prices play a negligible role for cross-channel competition in the German book market as resale prices of books are regulated by the Buchpreisbindungsgesetz.

The Buchpreisbindungsgesetz requires anyone who publishes or imports non-foreign-language books to set a common retail price to be paid by the final customers. These prices have to be made publicly available and cannot be changed by the outlets eventually selling the books (Buchpreisbindungsgesetz 2002, §5). Following a series of book price agreements since 1888, this law was enacted in 2002 to safeguard books as cultural goods

[1] This also includes the books that brick-and-mortar retailers sold online.

Figure 3.2.1: Book Sales by Retail Channel

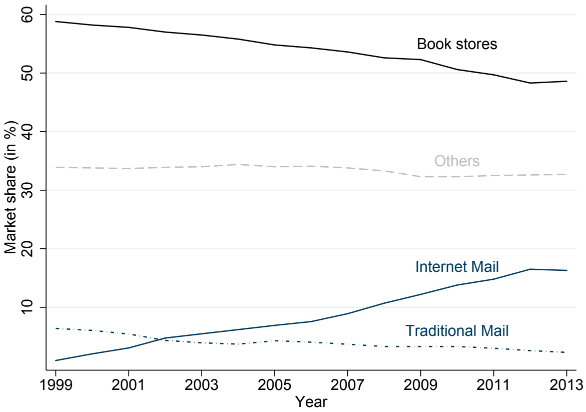

Note: This figure illustrates shares of sales for four different book distribution channels between 1999 and 2013. Data: Langendorf [2003], Cronau [2006, 2011, 2013, 2014].

Figure 3.2.2: Online Shopping Preferences

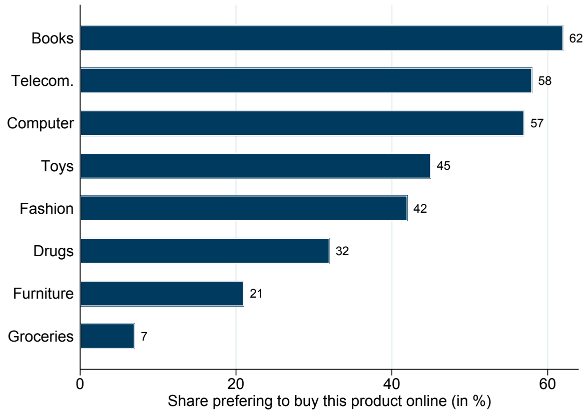

Note: This figure presents the results of a survey among 1,003 Germans aged 16–69 who were asked whether they prefer to buy a given product online or offline. Data: Bundesverband E-Commerce und Versandhandel and Boniversum [2017] .

of particular value that should be ubiquitously available to the general public. To this end, the Buchpreisbindungsgesetz also makes explicit that it is in favour of a large number of retail outlets (Buchpreisbindungsgesetz 2002, §1). While this law does not affect the resale of used books, it applies to all other books, including cartographic products, music sheets, and foreign-language books that are primarily sold in Germany (i.e., dictionaries).[2] As a result, online and offline prices for new books are mostly identical in Germany.

Publishers are allowed to set different prices for different types of books. For example, they may charge different prices for hard- and soft-cover copies, or the electronic version of a book (e-book). If online retailers primarily sold e-books, there would be price differences across online and offline retail channels (at least if e-books are considered to be suitable substitutes for printed books). However, market penetration of e-books in Germany is small, accounting for only 0.5% (3.9%) of the overall consumer market for books in 2010 (2013). Figure C.1.1 in Appendix C.1 further illustrates that this puts Germany on the bottom of an international ranking that is led by the USA, with an e-book penetration of 29%, followed by Great Britain and Japan with 23% and 14%, respectively. Overall, differences in price, quality, and the set of books offered between retail channels should thus have only limited impact on cross-channel competition in the German book market [Börsenverein des Deutschen Buchhandels, 2018].

In terms of product differentiation, it is important to note that books are standardised goods that typically come in a single form. Beyond the design of different book types, most publishers do not face quality competition from other publishers as they are the only ones

[2] Until a reform of the Buchpreisbindungsgesetz in 2016, cross-border book sales were also exempt from the law unless books were exported solely for the purpose of re-import. As cross-border sales of foreign-language books only represent a small share of the overall book market, we do not further discuss this exception.

distributing a particular book. Similarly, retailers have little opportunity to alter the quality of a book (e.g., by bundling the book with other products or services). As a result, the quality of a particular book does not differ across different distribution channels and is thus easily verifiable. These properties may also contribute to the fact that books were among the first goods to be frequently bought online (Figure 3.2.2).

The economic literature also discusses whether online and offline product portfolios differ and how this affects customer demand. For example, Brynjolfsson et al. [2009] show that online retailers are virtually immune from brick-and-mortar competition when selling niche products. This finding is the so-called the long tail phenomenon: online retailers can sell a wider range of products than stationary retailers due to lower distribution and inventory costs [Brynjolfsson et al., 2009]. As Zentner et al. [2013] show for the case of video rental patterns, a greater product portfolio may also affect customers' choices in that people renting videos online rather than in brick-and-mortar stores are more likely to watch niche films.

In our context, however, differences between the book portfolios of stationary and online retailers should be limited. Even if a bookstore does not showcase a specific book, customers can typically order any book available on the market at a local bookstore at no or very low cost. If brick-and-mortar bookstores can achieve economies of scale in ordering books or have better access to specific delivery services, books ordered by local stores may even arrive earlier at the customer than the ones ordered privately online. Another reason why the product supply between online and offline channels could differ involves self-published books. As self-published books are mainly sold online, local stores might not be able to order them. Yet, self-published books accounted for less than 2% of the general book market in 2014 [Behrens, 2017]. Consequently, any differences between traditional and online sale channels are likely not substantial and should not affect our empirical analysis.

In summary, our empirical analysis of the relationship between online shopping and brick-and-mortar bookstores in Germany should neither be driven by differences between online and offline prices nor by product differentiation. Instead, we are primarily identifying the joint effects of having a new, convenient *distribution channel* available as well as the complementary services associated with this channel (e.g., better information based on customer reviews as discussed in Friberg and Grönqvist [2012]). As these joint effects are increasingly important for the retail of other products as well, our results provide important insights for understanding how retail industries may develop in the future.

3.3 Data

3.3.1 Geomarketing Information

Our empirical analysis is based on a data set that combines novel information on regional shopping behaviour with administrative records. To measure the regional uptake of online shopping, we obtained information from Nexiga GmbH (Nexiga). As a geomarketing company, Nexiga gathers micro-geographical data for more than 22 million addresses in Germany and offers marketing services to businesses that want to target specific groups. Nexiga's extensive database contains information on more than 500 characteristics, ranging from the structure of a building and the quality of the surrounding infrastructure to detailed socio-demographic attributes of the residents.

Importantly for us, Nexiga calculates a building-level index of the *affinity to shop online* (*OSA*). Its *OSA* index is based on an extensive consumer survey conducted by Schober Information Group Deutschland GmbH (Schober) between 1997 and 2012. With more than 5 million respondents, this survey was the largest written consumer survey over that time period and covered more than 10% of all German households. As responses to this survey can be geocoded via the postal addresses of the participants, Nexiga is able to match the responses with additional address- and building-level data. Using these additional data, Nexiga applies multivariate regression techniques to extrapolate survey responses across buildings and time. Eventually, it assigns each building b an ordinal value between 1 (low) and 9 (high) that denotes the residents' affinity to shop online (OSA_b).

We obtained an aggregation of OSA_b that measures the *regional exposure to online shopping* at the municipality level (EOS_m) in 2015. To that end, Nexiga weighted OSA_b by the number of households in building b (HH_b) and calculated the share of households per municipality that are associated with an OSA_b of 7 or higher. Formally:

$$EOS_m = \frac{1}{\sum_{b\in m} HH_b} \sum_{b\in m} (OSA_b \times HH_b)\mathbf{1}\{OSA_b \geq 7\},$$

where $\mathbf{1}\{OSA_b \geq 7\}$ is an indicator function assuming the value of 1 if OSA_b is greater than or equal to 7 and 0 otherwise. That is, EOS_m measures the share of households per municipality that have a high affinity to shop online.

Figure 3.3.1 illustrates the regional distribution of EOS_m. Considering online shopping preferences at this fine spatial scale reveals the striking insight that high values of EOS_m do not have to coincide with the location of large agglomerations (i.e., "cities" are marked with orange dots). For example, most of the municipalities in the Ruhr region[3], the largest

[3] In the maps, the Ruhr region encompasses the cluster of large cities located close to the western German border.

Figure 3.3.1: Regional Distribution of EOS_m

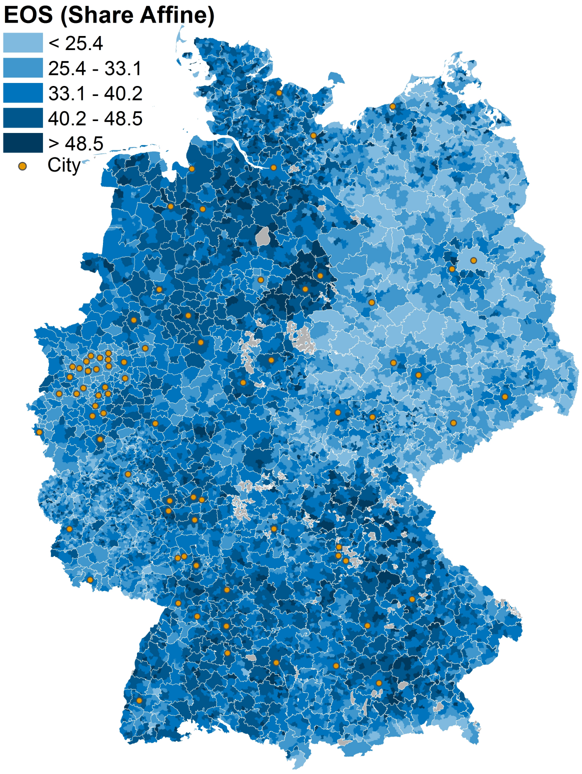

Note: This map illustrates the regional distribution of EOS (share of highly online shopping affine population) at the municipality level. White lines define regional retail markets. Orange dots mark municipalities with at least 100,000 residents. Grey shaded areas are unincorporated territories (*Gemeindefreie Gebiete*) or water areas. Data: Nexiga, Federal Statistical Office. Map: Borders of municipalities as of December 31, 2014, obtained from the Service Centre of the Federal Agency for Geo-Information and Geodesy.

Figure 3.3.2: Regional Distribution of EOS_{rm}

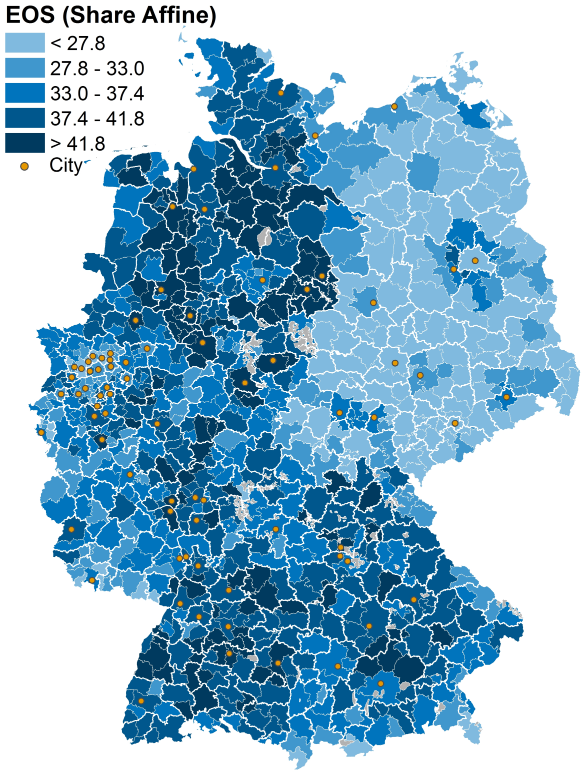

Note: This map illustrates the regional distribution of EOS (share of highly online shopping affine population) at the retail-market level. Thin (thick) white lines define regional retail markets (regional labour markets). Orange dots mark municipalities with at least 100,000 residents. Grey shaded areas are unincorporated territories (*Gemeindefreie Gebiete*) or water areas. Data: Nexiga, Federal Statistical Office. Map: Borders of retail markets as of December 31, 2014, are aggregated based on municipality-level information obtained from the Service Centre of the Federal Agency for Geo-Information and Geodesy.

agglomeration area in Germany, falls in the second lowest quintile of EOS_m whereas many of the highest values can be found around or inbetween larger cities. In the area of the New German Länder, however, this relationship is less pronounced. This may be at least partly explained by the general difference in the average levels of EOS_m between East and West Germany. Taking these observations together, the figure suggests on the one hand, that regional exposure to online shopping is positively correlated with purchasing power at higher levels of aggregation. On the other hand, the lower uptake of online shopping in densely populated municipalities might be explained by the fact that these areas typically provide better access to traditional shopping infrastructure.

In 1999, roughly 83.2% of all municipalities in Germany did not have a brick-and-mortar bookstore, implying that we have to assign these municipalities to those local retail markets in which their residents were most likely to buy books before they were able to shop online. In the absence of comprehensive data on individual shopping behaviour, we address this challenge by relying on an administrative definition of the Federal Institute for Building, Urban Affairs and Spatial Development (BBSR). For regional planning purposes, the BBSR defines middle regions (*Mittelbereiche*) to describe areas in which residents are able to meet their demand for high(er) quality goods and services. That is, these areas are explicitly designed to account for retail-specific linkages between regions. In our book-specific context, using this definition for local retail markets passes an important robustness check in that only 31 of the 878 retail markets in Germany did not have at least one traditional bookstore in 1999.

Since municipalities are nested within retail markets, we can calculate our measure for regional exposure to online shopping at the retail market level (EOS_{rm}) as a population-weighted average of EOS_m. Figure 3.3.2 illustrates EOS_{rm} across all 878 retail markets in Germany. On average, 35% of the households in a retail market are highly affine to online shopping. Even though the average retail market nests 13 municipalities, the patterns observed at the municipality level are also visible after aggregation. For example, the difference between East and West Germany even becomes more pronounced when considering retail markets.

It is important to note that regional variation in exposure to online shopping is largely driven by differences in local population characteristics rather than by technological factors. For example, as online book retailers offered their services throughout the whole country soon after market entry, their services were not only available to a limited set of selected municipalities. Furthermore, local availability of Internet infrastructure did not restrict the opportunity to shop online as integrated services digital networks (ISDN) have been universally available to all German households since the end of the 1990s.[4] Even though access to higher Internet bandwidths may have fostered shopping online, we argue that this is a negligible factor for the uptake of online shopping as compared to personal characteristics such as age, education, and income. As a result, EOS_{rm} can best be thought of as a regional aggregate of the individual decisions made by people living in retail market rm.

3.3.2 Administrative Data

Our main source of administrative data is the Establishment History Panel (*Betriebs-Historik-Panel*, BHP) of the Institute for Employment Research (IAB). The BHP provides detailed information on the universe of all German establishments that have at least one employee covered by social security on June 30th in each year. Although our version of the BHP spans the years from 1975 to 2014 (East German establishments are included

[4] Large online retailers optimise their websites to make efficient use of available bandwidths.

since 1992), the detailed industry classifications necessary to identify brick-and-mortar bookstores are available only from 1999 onward.[5] Prior to this year, the BHP did not include any information on establishments solely employing so-called marginal workers, i.e., part-time employees not subject to social security requirement and with very low working hours and wages (currently up to 450 Euro per month). Since these workers may be of particular interest in the retail sector, our empirical analysis uses administrative data only from 1999 onward. One of the largest book publishing houses in Germany (the Verlagsgruppe Weltbild GmbH) went bankrupt in 2014. As this publishing house also operated numerous traditional brick-and-mortar stores, we exclude this year from our analysis to avoid identifying short-term fluctuations. Our final sample covers the years from 1999 to 2013.

In terms of brick-and-mortar bookstores, we clean the BHP by excluding all establishments that had more than 100 employees in 2013 (i.e., that are above the 99.5% percentile). This is necessary as some large online retailers for books are misclassified. To match our cleaned establishment-level information with the regional geomarketing data, we aggregate the BHP to the retail market-by-year level. For each retail market, this aggregation provides us with the number of brick-and-mortar bookstores, the number of employees (also distinguished by full-time, part-time, and marginal), and the average wage of full-time employees. By considering all industries at once (i.e., rather than focusing on traditional bookstores only), the BHP also allows us to define a comprehensive set of control variables that measure employment-related characteristics of the retail markets. To this end, we count the number of employees by qualification, nationality, and gender, as well as measure the average wage across all industries.

Figure 3.3.3 shows the development of the number of brick-and-mortar bookstores in Germany in levels and the number of employees therein (indexed to the year 1999) between 1999 and 2013. Over this time period, the number of bookstores decreased by roughly 10%, which matches the decline in the share of sales observed in Figure 3.2.1. In terms of retail employment, there is strong divergence across types of workers. While the number of full-time employees has decreased by 40% since 1999, the number of regular and marginal part-time workers increased by roughly 25% and 10%, respectively. Overall, the number of employees working in traditional bookstores decreased by around 15%. At least in the aggregate, these developments imply that the dramatic change in full-time employees has been partly compensated for by increases in the employment of part-time workers.

We complement our aggregation of the BHP and the data provided by Nexiga with a series of regional variables that are available from public sources. These data include information on population (by age and gender) and the geographic area of municipalities

[5] We define a consistent measure of stationary bookstores using five-digit industry classifications from the German classification systems of 1993 and 2008 (*Wirtschaftszweigsklassifikation 2008*).

Figure 3.3.3: Brick-and-Mortar Bookstores and their Employees

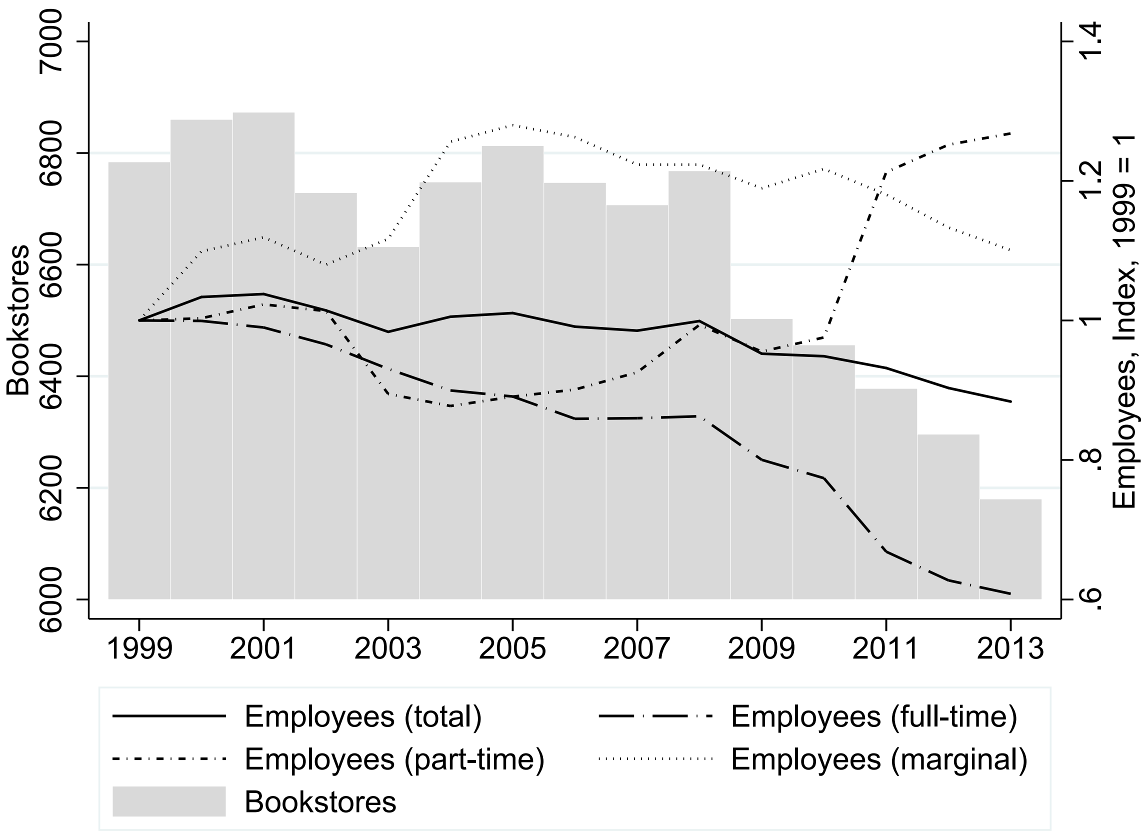

Note: This graph illustrates the development of brick-and-mortar bookstores in Germany (levels, left axis) and the employees, distinguished by type, therein (indexed to year 1999, right axis). Data: BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

(as well as all aggregations thereof). Table 3.3.1 reports descriptive statistics for all variables of interest. In 1999, the average retail market had 0.77 traditional bookstores per 10,000 residents. By 2013, this number had decreased by 0.06 or 7.8%, which is close to the change in aggregate numbers illustrated in Figure 3.3.3. All dependent variables are measured per 10,000 residents, i.e., denote densities, to account for size differences between retail markets. With regards to employment, it is worth noting that the national developments in the composition of employees (shown in Figure 3.3.3) are also reflected in the average numbers rather than being limited to a few selected places like large cities.

Table 3.3.1: Descriptive Statistics

	1999		2013		$\Delta_{2013-1999}$	
	Mean	SD	Mean	SD	Mean	SD
Dependent Variables						
Bookstores	0.77	0.40	0.71	0.36	–0.06	0.39
Employees (total)	3.48	2.35	3.19	2.12	–0.29	1.87
Employees (full-time)	1.44	1.11	0.86	0.80	–0.57	0.90
Employees (part-time)	0.60	0.66	0.81	0.76	0.22	0.73
Employees (marginal)	1.25	1.10	1.38	1.08	0.13	1.08
Explanatory Variables						
EOS	0.00	0.00	0.35	0.06	0.35	0.06
EOS^{bartik}	0.00	0.00	0.33	0.01	0.33	0.01
Share Population <20	0.22	0.02	0.18	0.02	–0.04	0.01
Share Population 20–29	0.11	0.01	0.11	0.02	–0.00	0.01
Share Population 30–39	0.16	0.01	0.11	0.01	–0.05	0.01
Share Population 40–49	0.15	0.01	0.15	0.01	0.00	0.01
Share Population 50–59	0.12	0.01	0.16	0.01	0.05	0.02
Share Population ≥60	0.24	0.02	0.28	0.03	0.04	0.02
Control Variables						
Population Density	0.15	0.15	0.15	0.15	0.00	0.00
Internet Usage	0.00	0.00	0.29	0.09	0.29	0.09
Share Female	0.48	0.04	0.49	0.04	0.02	0.02
Share German	0.94	0.04	0.93	0.04	–0.01	0.01
Share Medium Skilled	0.73	0.04	0.74	0.06	0.00	0.03
Share High Skilled	0.07	0.03	0.12	0.04	0.04	0.02
Share Employed	0.34	0.07	0.40	0.09	0.06	0.04
Average Hourly Wages	58.64	19.67	73.01	27.95	14.36	17.66
N	792		792		792	

Note: This table provides descriptive statistics for the variables of interest. The dependent variables are measured per 10,000 residents. *EOS* denotes online shopping exposure and measures the share of residents within a retail market who have a high preference for shopping online. EOS^{bartik} is a Bartik-type instrument for *EOS* (for more details on its computation, see Section 3.4.2). Population density is measured in terms of 100 residents per km^2. Average hourly wages are calculated for all full-time workers within a retail market. Data: Nexiga, Federal Statistical Office, BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

3.4 Empirical Framework

3.4.1 Baseline Specification

Identifying the causal effects of online shopping exposure on the development of local brick-and-mortar retailers is not straight forward. In the context of a simple cross-sectional regression, our measure for the regional exposure to online shopping may be endogenous for various reasons. First, regional differences in the structural composition of the local population may constitute an omitted variables concern. For example, if younger and better educated individuals are more likely to shop online and also differ in their reading behaviour from the average population, omitting these differences would likely bias the estimate. Second, if consumers choose to shop online because the local shopping infrastructure is poorly developed, e.g., bookstores are too small, far away, or cannot offer specific complementary services, the estimated effects may be affected by reverse causality. Finally, our main explanatory variable (EOS_{rm}) might be subject to measurement error.

Taking (some of) these concerns into account, we estimate a long difference model to determine whether regional exposure to online shopping affects the development of brick-and-mortar bookstores. The baseline specification is of the following form:

$$\Delta Y_{rm,r} = \beta \Delta EOS_{rm} + \Delta \mathbf{X}'_{rm}\theta + \gamma_r + \epsilon_{rm,r}, \tag{3.4.1}$$

where $\Delta Y_{rm,r}$ denotes a change in outcome Y of retail market rm and regional labour market r between 1999 and 2013. Our main outcomes of interest are the local stock of brick-and-mortar bookstores and the number of employees working in them (both measured per 10,000 residents). Our central explanatory variable, ΔEOS_{rm}, measures the change in regional exposure to online shopping over the same time period. In 1999, we assume that EOS_{rm} equals zero as only a negligible share of books was offered and purchased online at that time. Consequently, any variation observed in ΔEOS_{rm} is equal to the regional variation in 2013. Since this variable describes long-term changes in shopping behaviour, we use Nexiga's data for 2015 as a proxy for 2013. $\Delta \mathbf{X}_{rm}$ is a vector of changes in socio-economic and labour market specific characteristics of the local retail market. γ_r is a fixed effect at the level of regional labour markets (RLMs). $\epsilon_{rm,r}$ denotes the error term.

In Equation (3.4.1), taking long differences addresses many of the endogeneity concerns associated with a cross-sectional regression. In particular, we account for all types of time-invariant differences across retail markets. Among others, these include persistent differences in the composition of the local population, e.g., with respect to age, education, and income levels. Furthermore, the long differences also remove general differences in the quality of shopping infrastructures and, thereby, help to alleviate concerns related to reverse causality.

However, as we continue to be concerned about unobserved regional shocks (e.g., in the form of heterogeneous agglomeration trends), we also include fixed effects for RLMs (γ_r). We use the BBSR's definition of RLMs, which describes 258 geographic areas in Germany that are economically integrated and characterised by a high degree of within-region commuting. Empirically, using RLM fixed effects implements a *geographic* conditional independence assumption (CIA). That is, by only comparing retail markets within the same

RLM, we account for unobserved shocks that vary smoothly across space within confined geographical regions (i.e, RLMs in this case). Among others, this allows us to control for changes in average local income levels due to the closing of a large plant or the decline of an industry clustered in a specific region.

Even though these fixed effects eliminate unobserved heterogeneity in various dimensions, they may still fail to adequately capture other changes in the local population structure that are related to online shopping. Our greatest concern is that using the Internet in general may have direct effects on the development of brick-and-mortar bookstores. For example, Internet usage may be positively correlated with the number of books purchased at traditional stores as the information available online helps customers to identify books they want to buy. Furthermore, there could also be a negative relationship if spending time online causes people to read less [Falck et al., 2014] or replace traditional media like encyclopedias, dictionaries, and tourist guides with similar online services. As Internet usage is likely determined by the same personal characteristics as online shopping, our estimates in Equation (3.4.1) might be biased if we neglect these effects.

Therefore, we include the change in Internet usage across retail markets between 1999 and 2013 as a control variable in our regressions. By further including changes in population density, we account for differences in agglomeration across retail markets. Finally, we also control for differences in population shares by employment status, skill level, gender, and nationality as they might be correlated with the adoption of new technologies. With respect to income levels, we include the change in the average hourly wage of full-time workers.

3.4.2 Identification

Even though our combination of fixed effects and time-varying control variables alleviates some endogeneity concerns related to ΔEOS_{rm}, others remain. For example, our measure for regional exposure to online shopping might be subject to measurement error. On the one hand, measurement error could be introduced by Nexiga's extrapolation of the Schuber survey data across space and time. On the other hand, our population-weighted aggregation of municipality-level information could fall prey to similar problems. Furthermore, reverse causality could affect our OLS estimates of Equation (3.4.1) if the availability or quality of traditional shopping infrastructure (e.g., the number of stores) affects the uptake of online shopping and varies over time. Similarly, unobserved shocks to local retail markets, for example, the opening or closing of a large mall, could cause biases in the same direction.

Figure 3.4.1: Expenditures for Books by Channel and Age

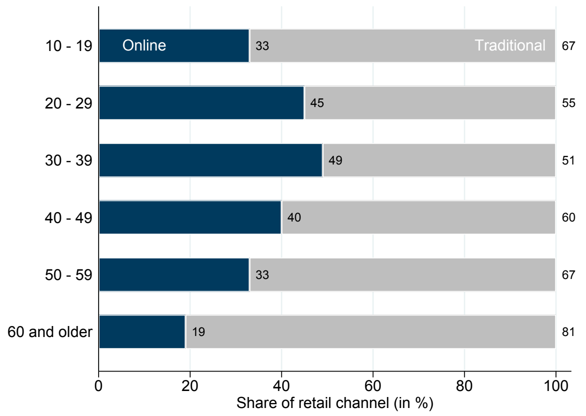

Note: This graph illustrates the expenditure shares for online and stationary retail channels, distinguished by age group. The expenditure shares are based on a representative survey among German households that was conducted in 2015. The stationary retail channel comprises purchases made in brick-and-mortar stores as well as other types of traditional retailers. Textbooks are excluded. Data: Börsenverein des Deutschen Buchhandels [2015], $n = 60{,}070$.

Figure 3.4.2: First Stage Regression of ΔEOS_{rm} on ΔEOS_{rm}^{bartik}

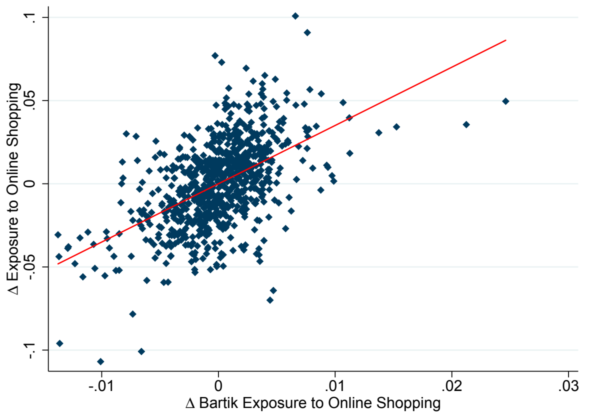

Note: This graph illustrates an added variables plot of ΔEOS_{rm} and ΔEOS_{rm}^{bartik}, accounting for RLM-level fixed effects, and changes in population density, online usage, the share of female employees, the share of foreign employees, the share of medium skilled employees, share of high skilled employees, the employment rate, and average hourly wages. Data: Nexiga, Federal Statistical Office, BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

In the absence of technological frictions or political incentives that could lead to exogenous variation in ΔEOS_{rm}, we employ a Bartik-type identification approach to address these endogeneity concerns. First introduced by Bartik [1991], this type of identification approach has been used in many fields of economics. Most closely related to this paper are Acemoglu and Restrepo [2017] and Dauth et al. [2017], who use Bartik-type instruments to identify the effects of industrial robots usage on regional employment. We adapt their empirical framework to our retail-specific context and construct a Bartik-type instrument that exploits regional variation in the uptake of online shopping based on different age structures. Specifically, we take advantage of the fact that the share of book expenditures online differs considerably across age groups (see the survey data illustrated in Figure 3.4.1).[6]

Figure 3.4.1 shows that almost half the people in the age group from 30 to 39 preferred buying their books online. In contrast, only 19% of the people who were 60 and older used this retail channel. Assuming that all books at the beginning of our sample period were bought in traditional stores, and keeping everything else constant, this figure implies that traditional bookstores have lost, on average, 31.6% of their business since the introduction of online shopping. However, since the share of online expenditures varies strongly across age groups ($OnlShare_a$), we do not expect that all brick-and-mortar bookstores are affected equally by this aggregate development. Instead, traditional bookstores that are located in

[6] Acemoglu and Linn [2004] exploit similar variation to identify the effect of market size on pharmaceutical innovations.

retail markets with a population more prone to online shopping should also be subject to stronger effects of online shopping.

Formally, we exploit this relationship by constructing a Bartik-type instrument:

$$\Delta EOS_{rm}^{bartik} = \sum_a \underbrace{\frac{Pop_{a,rm,2001}}{Pop_{rm,2001}}}_{\text{Historic Share}} \underbrace{\left(\frac{Pop_{a,2013} \times OnlShare_{a,2013}}{Pop_{a,1999}} - \frac{Pop_{a,1999} \times \overbrace{OnlShare_{a,1999}}^{=0}}{Pop_{a,1999}} \right)}_{\text{Aggregate Change}}, \tag{3.4.2}$$

where the first factor in the product denotes the share of people in age group a, living in retail market rm, in year 2001, and the second factor indicates the aggregate change in "online shoppers" standardised by the population in 1999. Since online shopping only started to take off in the early 2000s, the share of online expenditures for books in 1999 was negligible (i.e., we assume a value of zero). The age shares in 2001 are intended to approximate the regional age structure before online shopping was introduced. We use 2001 as it is the first year for which we have comprehensive information on age shares across all retail markets. The persistence in regional age structures should ensure the validity of this choice.

Using this Bartik-type instrument, we estimate the following first stage relationship:

$$\Delta EOS_{rm} = \delta \Delta EOS_{rm}^{bartik} + \Delta \mathbf{X}'_{rm}\phi + \omega_r + \eta_{rm}. \tag{3.4.3}$$

Figure 3.4.2 illustrates this relationship graphically. This scatter plot is based on the most comprehensive specification, that is, it includes RLM-level fixed effects (ω_r) and all of the control variables described above. The solid red line is an OLS estimate of the slope coefficient (δ) in Equation (3.4.3). The coefficient δ is greater than 3, indicating a strong and positive correlation between ΔEOS_{rm} and our Bartik-type instrument ΔEOS_{rm}^{bartik}.

To explain why this coefficient is greater than 1, recall that both measures for online shopping differ within two dimensions. On the one hand, ΔEOS_{rm} reflects a general change in the local population's *affinity* to shop online that is not specific to certain products or services. On the other hand, ΔEOS_{rm}^{bartik} measures *actual online expenditures* for books as a fraction of overall spending. Therefore, it may not be surprising that a high share of actual expenditures requires an even greater share of online shopping affine people as even highly affine people occasionally shop for books offline. Furthermore, general online shopping affinity may be greater than that related to a specific product (in our case books).

As illustrated in Equation (3.4.1), Bartik-type instruments are an interaction of two components: a (historic) local share and and aggregate shock. Most of the empirical literature using this identification approach argues that the aggregate shock (and, therefore, the instrument) is exogenous in its specific context. In some cases, this also involves instrumenting the shock (as, for example, in the studies using the “China shock” based on the work by Autor et al. [2013]). Goldsmith-Pinkham et al. [2018] recently challenged this line of reasoning by arguing that identification using Bartik-type instruments should be discussed in terms of the (historic) local shares rather than the aggregate shocks. To formalise their argument, they show that using a Bartik-type instrument is numerically equivalent to using the local shares directly and that aggregate shocks only contribute to instrument relevance.

To cater both approaches pursued in the literature, we discuss the validity of our Bartik-type instrument from both angles. With respect to the exogeneity of the aggregate shock in our context, it can be argued that online shopping only became possible due to the invention of the Internet. Since this was an international-level technological advance, it should not be correlated with regional, book market-specific developments in Germany. In a similar vein, using an aggregate shock also allows us to abstract from selection problems caused by regional migration as the aggregate population in Germany is only affected by demographic factors and international migration.

With respect to our age shares in 2001, it is not straightforward to argue that they are neither directly related to developments in local book markets nor affected by any omitted factors. For example, a high share of young residents in 1999 may have given rise to dynamic economic developments at the level of retail markets that cannot be captured by RLM-specific trends. If these economic developments are reflected in both a higher uptake of online shopping and a higher demand for education and books, this might cause an upward bias in our IV estimates. To assess this possibility, we follow a suggestion by Goldsmith-Pinkham et al. [2018] and conduct a simple empirical test. Specifically, we investigate whether our regional age shares in 2001 are balanced across the set of control variables available to us. This test is inspired by the idea that, ideally, a valid instrument should not be affected too strongly by observable characteristics as this suggests that unobservables may also be important (see, e.g., Oster [2018] for an elaboration on this argument). Even though one may argue that controlling for observable confounders solves any problems that can be detected with this test (in fact, we also explicitly build on a CIA above), it still provides important insights in our context.

As an example, Figure 3.4.3 shows a selection of four scatter plots that indicate how balanced the share of 30 to 39 year old people is across four important control variables (i.e., the change in Internet usage, average hourly wages of full-time workers, and the share

Figure 3.4.3: Relationship between Age Shares in 1999 and Changes in Control Variables

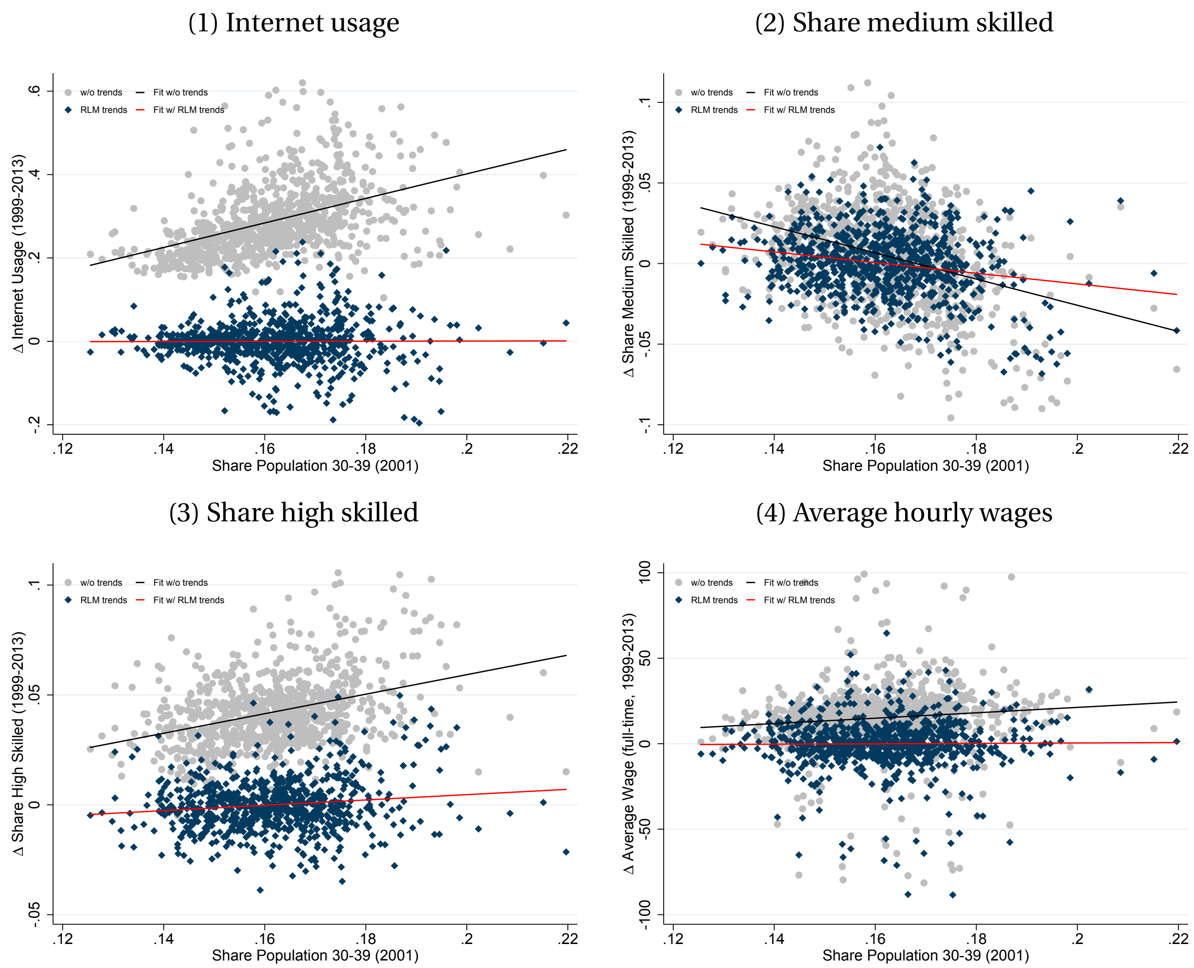

Note: This figure presents four scatter plots that illustrate the relationship between the share of 30 to 39 year old people in 2001 and the change in a selection of control variables between 1999 and 2013. Grey dots (black lines) indicate scatter plots (OLS regression coefficients) of the unconditional relationship between the change in the control variable of interest and the pre-Internet share of people between 30 and 39. Blue diamonds (red lines) represent added variable scatter plots (OLS regression coefficients) of the same relationships after accounting for RLM-level fixed effects. Data: Nexiga, Federal Statistical Office, BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

of medium as well as high skilled employees).[7] The grey dots indicate the unconditional relationship between the change in the control variable of interest and the (pre-Internet) share of the 30 to 39 year old residents (black lines denote corresponding coefficients of OLS regressions).

As can be seen in the upper left graph, there is a strong (unconditional) positive correlation between the share of 30 to 39 year old people in 1999 and the change in Internet usage. However, if we compare retail markets only within RLMs, this relationship vanishes entirely. Qualitatively similar observations can also be made in the other graphs in that adding RLM-level fixed effects at least considerably reduces the relationship observed in the

[7] In Appendix C.1, we show the corresponding graphs for the remaining control variables used in the regressions with respect to the share of 30 to 39 year old people. Conducting the same visual inspections for all other age shares leads to similar conclusions (results are available from the authors upon request).

unconditional case. Combining these observations with the result that we can draw similar conclusions for all other age shares, we conclude that our Bartik-type instrument is valid only when we include RLM-level fixed effects to account for underlying regional trends. Conditional on doing so, however, our approach addresses endogeneity concerns related to omitted variables and reverse causality. Furthermore, as ΔEOS_{rm}^{bartik} provides a measure for ΔEOS_{rm} that should be subject to different types of measurement error, using it as an instrument should reduce measurement error in our survey-based information.

3.5 Results

3.5.1 Baseline OLS and IV Results

Table 3.5.1 reports the results of our baseline OLS and IV estimations for the number of brick-and-mortar bookstores (measured per 10,000 residents). Across all specifications, we include RLM-level fixed effects and cluster standard errors at the same level. We exclude retail markets below the bottom and above the top percentile (in terms of bookstore density) as we do not want our effects to be driven by outliers and also drop retail markets that do not have at least one bookstore in 1999. As we are estimating long difference models, the number of observations corresponds to the number of retail markets considered (i.e., $N = 792$).

Our OLS estimates show that regional online shopping exposure has robust negative, though statistically insignificant, effects on the number of brick-and-mortar bookstores. As expected, including the change in Internet usage (i.e., in Column (2)) decreases the estimated effect. This supports the idea that online shopping and Internet usage are positively correlated and should both affect traditional retailers. Neglecting Internet usage therefore leads to an upward bias of the coefficient for ΔEOS in Column (1). In our most comprehensive OLS specification (Column (3)), we also include a set of socio-demographic as well as labour market-specific control variables. Among these, only the change in the average hourly wage of full-time workers has a statistically significant effect; suggesting that changes in income levels affect the development of local bookstores.

With respect to the IV results reported in Columns (4) to (6), all regressions are associated with strong first stage regression results. Across specifications, a 1 percentage point increase in ΔEOS^{bartik} leads to an increase in online shopping exposure by more than 3 percentage points. The corresponding Kleibergen-Paap F-statistics are either close to or above 100, rejecting the null hypothesis that our Bartik-type instrument is weak. Regarding the second stage, we find that our IV estimates for ΔEOS are considerably smaller than their OLS counterparts and become statistically significant at the 10% level (Column (6)).

Table 3.5.1: Baseline OLS and IV Results (Bookstores)

ΔBookstores	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ*EOS*	–0.614	–0.773	–0.908	–1.453	–1.702	–1.771*
	(0.563)	(0.581)	(0.574)	(0.922)	(1.084)	(1.018)
ΔInternet Usage		0.257	0.308		0.347	0.389*
		(0.205)	(0.212)		(0.219)	(0.229)
ΔPopulation Density		–9.213	–21.59		66.93	42.03
		(127.8)	(106.3)		(153.7)	(130.8)
ΔShare Female			–0.528			–0.553
			(0.781)			(0.784)
ΔShare German			1.939			1.879
			(1.435)			(1.433)
ΔShare Medium Skilled			–0.286			–0.267
			(0.930)			(0.933)
ΔShare High Skilled			–1.504			–1.347
			(1.502)			(1.535)
ΔShare Employed			0.238			0.310
			(0.527)			(0.503)
ΔAverage Hourly Wages			0.008***			0.008***
			(0.001)			(0.001)
First Stage						
ΔEOS^{bartik}				3.619***	3.470***	3.517***
				(0.323)	(0.352)	(0.354)
Kleibergen-Paap F-Stat.				125.40	97.19	98.67
RLM-FE	Yes	Yes	Yes	Yes	Yes	Yes
N	792	792	792	792	792	792

Note: This Table reports baseline OLS and IV results for Equation (3.4.1). The dependent variable is the change in the number of bookstores per 10,000 residents between 1999 and 2013. All specifications include RLM-level fixed effects. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. Data: Nexiga, Federal Statistical Office, BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

There are (at least) two reasons that may help to explain the differences between our OLS and IV estimates. On the one hand, an upward bias in our OLS regressions could result from unobserved shocks that are positively correlated with ΔBookstores and Δ*EOS*, respectively. For example, if our control variables fail to adequately account for changes in local income levels (which are likely positively correlated with *EOS*, see the maps in Section 3.3.1), this would explain the pattern observed. On the other hand, Nexiga constructs our measure of *EOS* based on interpolations across space and time. If these interpolations induce classical measurement error, this would lead to an attenuation bias that could also explain our smaller IV estimates.

Table 3.5.2: Baseline IV Results (Employees)

ΔEmployees	Total	Full-Time	Part-Time	Marginal
	(1)	(2)	(3)	(4)
Δ*EOS*	–7.297*	–2.896	–1.953	–2.318
	(4.261)	(2.232)	(1.644)	(2.912)
ΔInternet Usage	0.532	0.436	–0.331	0.604
	(1.594)	(0.851)	(0.573)	(0.776)
ΔPopulation Density	452.9	203.7	29.44	205.3
	(663.1)	(359.8)	(260.6)	(340.5)
ΔShare Female	2.148	1.578	0.251	0.366
	(3.490)	(1.839)	(1.899)	(2.075)
ΔShare German	–1.784	–0.174	–5.830*	3.340
	(7.282)	(4.627)	(3.235)	(4.266)
ΔShare Medium Skilled	6.133	1.273	4.188	–0.710
	(5.186)	(2.261)	(2.753)	(2.827)
ΔShare High Skilled	–1.596	–9.298***	6.268*	1.849
	(7.253)	(3.198)	(3.734)	(4.395)
ΔShare Employed	1.727	–0.006	0.871	1.007
	(2.126)	(1.040)	(0.970)	(1.154)
ΔAverage Hourly Wages	0.020***	0.004***	0.005***	0.001***
	(0.004)	(0.002)	(0.002)	(0.002)
First Stage				
ΔEOS^{bartik}	3.517***	3.517***	3.517***	3.517***
	(0.354)	(0.354)	(0.354)	(0.354)
Kleibergen-Paap F-Stat.	98.67	98.67	98.67	98.67
RLM-FE	Yes	Yes	Yes	Yes
N	792	792	792	792

Note: This tables reports baseline OLS and IV results for Equation (3.4.1). The dependent variables are measured per 10,000 residents and indicate changes the number of brick-and-mortar bookstore employees (also distinguished by type) between 1999 and 2013, respectively. All specifications include RLM-level fixed effects. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. Data: Nexiga, Federal Statistical Office, BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

Considering Column (6) as our preferred specification, the point estimate of –1.771 implies that an increase in regional online shopping exposure by one standard deviation (i.e., 6 percentage points) reduces the local stock of brick-and-mortar bookstores per 10,000 residents by 0.11. Given that the average retail market in 1999 had 0.77 bookstores per 10,000 residents, this effect corresponds to a reduction by approximately 14%.

In Table 3.5.2, we report the IV results for changes in employment (i.e., the total number of employees as well as employees distinguished by type). As for bookstores, the employment variables are measured per 10,000 residents. Our IV estimates for these outcomes fit into the picture observed for bookstores. For example, Column (1) shows that an increase in the regional exposure to online shopping substantially reduces overall employment in brick-and-mortar bookstores. This effect is statistically significant at the 10% level. The point estimate of –7.297 implies that an increase in regional online shopping exposure by one standard deviation leads to a reduction in total brick-and-mortar retail employment by 0.44 employees per 10,000 residents. In terms of 1999 employment levels, this corresponds to a reduction by approximately 13%. With respect to our results in Columns (2) – (4), the degree of precision associated with our IV estimates prohibits us from concluding that ΔEOS has heterogeneous effects across employment types. However, considering the point estimates, our results suggest that full-time employees may be affected most severely by the introduction of online shopping and contribute the largest part to the overall effect.

3.5.2 Robustness and Heterogeneity

In Section 3.4.2, we have argued that our estimation strategy is only valid conditional on a set of fixed effects. While we have also motivated the importance of our time-varying control variables, our estimates have been relatively robust toward their inclusion. So far, one important limitation of our controls is that they only account for general characteristics of the local population and the labour market, but lack a book market specific component. Incorporating this concern, we test the robustness of our results by including the number of bookstores (per 10,000 residents) in 1999 as a control. Arguably, this should provide a good approximation of book market specific aspects of the local retail market before the Internet (and, hence, online shopping) was introduced.

Table 3.5.3 reports the results for all outcomes of interest when controlling for this lag. Except for part-time employees, we find that the density of bookstores in 1999 has a negative and statistically highly significant influence on the development of local book markets. This result suggests that retail markets with a previously high number of brick-and-mortar bookstores are associated with more severe decreases between 1999 and 2013.

With respect to our estimates for ΔEOS, we find that, after including this lag, our estimates become slightly smaller and are more precisely estimated. For example, our point estimate for the change in bookstores decreases from –1.771 (in our preferred IV specification of Table 3.5.1) to –2.139. While these two point estimates are not statistically significant different from each other, the latter is associated with substantially lower standard errors. Similarly, our estimates for changes in the total number of employees as well as the number of full-time workers (per 10,000 residents) do not change much in magnitude, but are now statistically significant at the 5% and 10% level, respectively.

Table 3.5.3: Robustness Test of IV Results

	ΔBookstores	ΔEmployees			
		Total	Full-Time	Part-Time	Marginal
	(1)	(2)	(3)	(4)	(5)
Δ*EOS*	–2.139***	–8.266**	–3.388*	–2.009	–2.689
	(0.764)	(4.208)	(2.224)	(1.649)	(2.861)
ΔInternet Usage	0.353*	0.439	0.789	–0.337	0.568
	(0.195)	(1.483)	(0.851)	(0.573)	(0.736)
ΔPopulation Density	46.90	465.7	210.2	30.18	210.2
	(119.6)	(668.9)	(349.1)	(261.8)	(348.0)
ΔShare Female	–0.295	2.827	1.922	0.290	0.626
	(0.630)	(3.496)	(1.945)	(1.895)	(2.047)
ΔShare German	–0.069	–6.903	–2.771	–6.125*	1.380
	(1.184)	(7.529)	(4.747)	(3.206)	(4.489)
ΔShare Medium Skilled	–1.215	3.639	0.008	4.044	–1.665
	(0.759)	(4.899)	(2.178)	(2.747)	(2.728)
ΔShare High Skilled	–0.925	–0.487	–8.736***	6.332*	2.274
	(1.193)	(6.787)	(2.814)	(3.767)	(4.196)
ΔShare Employed	0.039	1.016	–0.354	0.830	0.735
	(0.485)	(2.149)	(0.995)	(0.977)	(1.173)
ΔAverage Hourly Wages	0.008***	0.019***	0.004*	0.005***	0.010***
	(0.001)	(0.004)	(0.002)	(0.002)	(0.002)
L.Bookstores	–0.573***	–1.506***	–0.764***	–0.087	–0.577***
	(0.050)	(0.221)	(0.125)	(0.086)	(0.151)
First Stage					
ΔEOS^{bartik}	3.504***	3.504***	3.504***	3.504***	3.504***
	(0.347)	(0.347)	(0.347)	(0.347)	(0.347)
Kleibergen-Paap F-Stat.	102.11	102.11	102.11	102.11	102.11
RLM-FE	Yes	Yes	Yes	Yes	Yes
N	792	792	792	792	792

Note: This table reports IV results for Equation (3.4.1), but includes the share of bookstores in 1999 (measured per 10,000 residents) as an additional control variable. The dependent variables are measured per 10,000 residents and indicate changes in the number of bookstores and the number of employees therein (also distinguished by type) between 1999 and 2013, respectively. All specifications include RLM-level fixed effects. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. Data: Nexiga, Federal Statistical Office, BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

An additional benefit of including the lagged density of bookstores as control variable is that it enables us to investigate whether our estimated effects for Δ*EOS* are homogeneous. More specifically, it allows to ask whether regional exposure to online shopping has heterogeneous effects along the distribution of the number of local bookstores in 1999. Table C.2.1 in Appendix C.2 provides summary statistics for a selection of variables within six intervals to illustrate how retail markets differ along this distribution. Apart from the average bookstore density in 1999, we also include the pre-Internet number of bookstore employees (by type) per establishment to characterise local book markets. In the second panel, we list most of our control variables (i.e., we omit those that are zero in 1999) and the total numbers of residents and employees.

Regarding book market specific aspects in 1999, this table reveals that retail markets with higher bookstore densities in 1999 have, on average, “smaller” bookstores (i.e., lower ratios of employees per bookstore). This relationship does not only hold for the total number of employees, but across all types. Maybe surprisingly, the retail markets with the highest bookstore densities are not the largest ones, e.g., in terms of total number of residents or employees. Instead, there is an inverted u-shape relationship between absolute retail market size and bookstore density. In terms of our control variables, there are no stark differences between retail markets in the intervals considered. The only exception is that higher initial bookstore densities are also associated with higher average wages of full-time workers.

To estimate heterogeneous effects with respect to bookstore density in 1999, we standardise the number of bookstores per 10,000 residents in 1999 at five different percentiles ($p \in \{10, 25, 50, 75, 90\}$) and estimate the following model:

$$\begin{aligned} \Delta Y_{rm,r} = & \beta_1 \Delta EOS_{rm} + \beta_2 (\Delta EOS_{rm} \times L.Bookstores^{p}_{rm}) + \beta_3 L.Bookstores^{p}_{rm} \\ & + \Delta \mathbf{X}^{'}_{rm} \theta + \gamma_r + \epsilon_{rm,r}, \end{aligned} \tag{3.5.1}$$

where $L.Bookstores^{p}_{rm}$ denotes the number of bookstores per 10,000 residents in 1999, standardised at percentile p. As this standardised measure also enters the interaction term, the coefficients β_1 and β_3 represent “main effects” that reflect conditional relationships. In other words, β_1 denotes the effect of ΔEOS_{rm} conditional on $L.Bookstores^{p}_{rm}$ being zero, whereas β_3 is the effect of the lagged (and standardised) measure for pre-Internet bookstore density in absence of a change in ΔEOS_{rm}.

In Table 3.5.4, we report our results for all outcomes (panels) and percentiles (columns) of interest. For the sake of brevity, we omit the point estimates for all control variables as well as for the lagged value of bookstore density. In Panel A, we depict how changes in regional online shopping exposure heterogeneously affect the development of the local stock of brick-and-mortar bookstores. Considering Column (1), an increase in ΔEOS by one standard deviation leads to a decrease of approximately 0.11 bookstores per 10,000 residents (i.e., -1.895×0.06) at the 10th percentile of the lagged density distribution. In contrast, an identical increase at the 90th percentile has an effect of 0.14. Although the difference between these two point estimates has some economic meaning, we lack the statistical power to conclude that it is also statistically significant.

Table 3.5.4: Heterogeneity of IV Results

Percentile	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
	(1)	(2)	(3)	(4)	(5)
Panel A:	ΔBookstores				
ΔEOS	–1.895**	–1.979**	–2.077***	–2.203***	–2.367**
	(0.934)	(0.828)	(0.767)	(0.804)	(1.017)
$\Delta EOS \times L.$Bookstores	–0.202	–0.202	–0.202	–0.202	–0.202
	(0.517)	(0.517)	(0.517)	(0.517)	(0.517)
Panel B:	ΔEmployees (Total)				
ΔEOS	–0.388	–3.091	–6.250	–10.30***	–15.580***
	(5.265)	(4.690)	(4.203)	(3.975)	(4.427)
$\Delta EOS \times L.$Bookstores	–6.491***	–6.491***	–6.491***	–6.491***	–6.491***
	(2.308)	(2.308)	(2.308)	(2.308)	(2.308)
Panel C:	ΔEmployees (Full-Time)				
ΔEOS	–3.083	–3.188	–3.310	–3.466	–3.670
	(2.754)	(2.410)	(2.206)	(2.334)	(3.038)
$\Delta EOS \times L.$Bookstores	–0.251	–0.251	–0.251	–0.251	–0.251
	(1.607)	(1.607)	(1.607)	(1.607)	(1.607)
Panel D:	ΔEmployees (Part-Time)				
ΔEOS	–0.033	–0.908	–1.580	–2.442	–3.565
	(2.163)	(1.865)	(1.670)	(1.738)	(2.284)
$\Delta EOS \times L.$Bookstores	–1.381	–1.381	–1.381	–1.381	–1.381
	(1.271)	(1.271)	(1.271)	(1.271)	(1.271)
Panel E:	ΔEmployees (Marginal)				
ΔEOS	2.321	0.602	–1.406	–3.981	–7.339**
	(3.473)	(3.133)	(2.878)	(2.833)	(3.254)
$\Delta EOS \times L.$Bookstores	–4.127***	–4.127***	–4.127***	–4.127***	–4.127***
	(1.572)	(1.572)	(1.572)	(1.572)	(1.572)
First Stage					
ΔEOS^{bartik}	3.205***	3.293***	3.396***	3.529***	3.702***
	(0.471)	(0.400)	(0.351)	(0.365)	(0.492)
$\Delta EOS^{bartik} \times L.$Bookstores	3.867***	3.779***	3.675***	3.543***	3.370***
	(1.182)	(1.098)	(1.009)	(0.913)	(0.831)
Kleibergen-Paap F-Stat.	13.332	13.332	13.332	13.332	13.332
RLM-FE	Yes	Yes	Yes	Yes	Yes
Controls (incl. *L.*Bookstores)	Yes	Yes	Yes	Yes	Yes
N	792	792	792	792	792

Note: This table reports IV results of Equation (3.5.1). The dependent variables denote standardised changes in the number of bookstores and the number of employees therein (distinguished by type) per 10,000 residents between 1999 and 2013. All specifications include RLM-level fixed effects. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. Data: Nexiga, Federal Statistical Office, BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

Turning to the total number of employees in Panel B, we find an economically and statistically more distinct pattern. While retail markets at the bottom of the density distribution do not reduce their total number of bookstore employees significantly, there is a strong and steady increase in the absolute size of the coefficients at higher percentiles. At the 75^{th} and 90^{th} percentile, an increase in ΔEOS by one standard deviation leads to a decrease by approximately 0.62 and 0.93 employees per 10,000 residents, respectively. These estimates are statistically significant at the 1% level.

To better understand where these heterogeneous effects come from, we further investigate whether our results differ across types of bookstore employees (Panels C to E). With respect to full-time workers, we find a pattern that is qualitatively similar to the one observed for changes in the number of bookstores, but lack the statistical power to draw strong conclusions. Nevertheless, if we take the point estimates for this outcome at face value, retail markets across the whole density distribution respond to the introduction of online shopping in a similar manner, i.e., by reducing full-time employment. Consequently, the small differences in the effects for full-time workers cannot explain the heterogeneity observed for the total number of bookstore employees.

In terms of part-time and marginal workers, our heterogeneity analysis reveals that retail markets at higher percentiles of the bookstore density distribution experience larger decreases. Again, most of our point estimates are not statistically significant (with the only exception being the estimate for marginal employees in Column (5)), but the general patterns associated with them are in line with the heterogeneity observed for the total number of employees. This implies that our estimates for the total number of employees mask heterogeneous effects among employment categories. For example, considering low density retail markets, the very small (and statistically insignificant) reduction in the total number of employees seems to be the result of a replacement of full-time workers with marginal employees. In contrast, at the upper end of the distribution, the absolute large effects stem from a joint decline in all employment categories.

Summing up, our heterogeneity analysis indicates that retail markets with a higher density of bookstores in 1999 are more strongly affected by the introduction of online shopping. While retail markets at the lower end of the density distribution seem to substitute different types of employees (i.e., full-time vs marginal), we observe that bookstores in retail markets at the high end decrease employment across all types of workers. One possible explanation for this heterogeneity is that bookstores in less concentrated retail markets (i.e., with a higher number of competitors) may face stronger competition from other bookstores and are therefore more vulnerable to the additional competition induced by online shopping.

3.6 Summary and Conclusions

A local decline of brick-and-mortar retailers is often associated with online shopping. Considering the case of traditional bookstores in Germany, we combine novel geomarketing data with administrative information on traditional bookstores to investigate this relationship empirically. Based on a long difference framework, we implement a Bartik-type IV strategy that exploits local variation in pre-Internet age structures to identify the causal impact of regional online shopping exposure on the development of brick-and-mortar bookstores between 1999 and 2013.

In summary, we find that regional exposure to online shopping has a robust negative effect on the development of brick-and-mortar bookstores. For example, an increase in the uptake of online shopping by one standard deviation leads to a reduction in the local stock of bookstores by approximately 0.11 establishments per 10,000 residents. Relative to the levels in 1999 (i.e., the pre-Internet era), this corresponds to a decline by 14%. In terms of total employment, we find effects of similar relative size.

Across different types of employment, we find that our effects are heterogeneous with respect to the density of bookstores in 1999. Although parts of our heterogeneity analysis lack statistical power, our point estimates suggest that all retail markets reduce full-time workers in a similar manner. While retail markets at the lower end of the density distribution seem to substitute these full-time workers by marginal employees, we find consistently negative effects for all types of employees at higher percentiles. Taking into account that retail markets at the upper end of this density distribution were less concentrated in 1999 (i.e., had a higher degree of competition), one possible explanation for our results may be that the additional competition induced by online shopping affected those retail markets the most that were also more competitive in the past.

Finally, we have analysed the impact of online shopping on the development of brick-and-mortar bookstores in Germany because the institutional setting is particularly well suited to identify how the availability of a convenient distribution channel affects traditional retailers. While the convenience aspect is arguably important for most consumption products and, therefore, also retailers, it may not be the dominant channel in every industry. For example, in some industries, limited regional supply or the demand for niche products may be more important. It would therefore be promising to know whether our effects also persist in other industries and, in comparison, how much convenience matters relative to other channels. Furthermore, we view our paper as a step towards explaining the development of brick-and-mortar retailers as this constitutes an important aspect for city planners and local policy makers. As general equilibrium effects are therefore missing from our paper (i.e., we do not investigate what happens to former employees of bookstores who lose their jobs), this topic provides further promising avenues for future research.

Appendices

Appendix A

Appendix to Chapter 1

A.1 Additional Programme Details

Defining the Term 'Region'

Despite the IRGC's regional focus, the geographic boundaries of the term "region" are not defined within the programme itself. The only reference made in this regard is that federal states (*Bundesländer*) are not considered to be regions and that a functional connection within the collaborations has to be made explicit in the application process [BMBF, 2016*c*]. However, a pragmatic approach towards narrowing down the scope of this term can be based on two remarks: First, every IRGC funded collaboration has to declare a so called *core region* in which most of the significant activities of the growth core have to take place. In many cases, the core regions match the names or centres of German counties. Second, the IRGC has been complemented by the Growth Cores Potential (GCP) programme in 2007 [BMBF, 2016*b*]. The GCP aims to prepare the grounds for an IRGC funded project by bridging public research and R&D conducting companies residing in the same geographic areas. Importantly, the programme motivates scientists from universities and public research institutes to collaborate with companies located within *50km* distance to explore the opportunities for shaping joint platform technologies. Even though this 50km distance threshold has most likely not been strictly enforced, it is relatively close to the average dimension of a typical regional labour market (RLM) as defined by the BBSR. Since RLMs nest the administrative entities of counties under consideration of regional commuter flows, we consider them to be a sensible choice for a definition of regions in the context of the IRGC. Furthermore, on average more than 50% of the private firms subsidised within a given project also reside in the RLM assigned to this project.

The Application Process

The application process begins with a consultation of the applicants and the project executing organisation Projektträger Jülich (PtJ). After this consultation, the applicants have to submit an idea sketch that is subsequently examined by PtJ. If this examination is successful, the applicants are invited to an interview with the BMBF and have to prepare a so-called innovation concept, which is essentially a business plan for their collaborative innovation project, as well as a formal grant proposal. After submitting both documents to PtJ, the applicants are invited to an assessment centre conducted by the BMBF in collaboration with a jury of external experts which makes the final funding decision.

A.2 Estimating the Number of Factors

We estimate the number of factors following the information criterion approach described by Bai and Ng [2002] and Bai [2009]. A concise description of this approach can be found in the appendix of Hagedorn et al. [2016]. Borrowing their notation, the decision criterion CP can be expressed as a function of the number of factors L:

$$CP(L) = \hat{\sigma}^2(L) + \hat{\sigma}^2(\bar{L})[L(N+T) - L^2]\frac{\log(NT)}{NT} \tag{A.2.1}$$

where $\bar{L}$ denotes the maximum number of factors (which is six in our case), N is the number of RLMs, and T is the number of years. $\hat{\sigma}^2$ represents the mean squared error of a semi-dynamic DiD model that only includes the post-treatment indicators, firm and year fixed effects, region-specific linear time trends, and a factor model. Formally, this can be expressed as:

$$\hat{\sigma}^2(L) = \frac{1}{NT}\sum_{r=1}^{N}\sum_{t=1}^{T}\Big(ln(y_{rt}) - \sum_{k=0}^{\bar{k}}\gamma_k^{agg}\mathbf{1}\{K_{rt}=k\} - \alpha_r - \delta_t - \mu_r t - \lambda^{'}(L)_r F_t(L)\Big), \tag{A.2.2}$$

where $\lambda^{'}(L)_r F_t(L)$ denotes the estimates for a factor model with L common factors. The number of factors can then be estimated by:

$$\hat{L} = argmin_{L\leq\bar{L}} CP(L). \tag{A.2.3}$$

A.3 Additional Figures

Figure A.3.1: Descriptive Statistics – Publicly Funded R&D Expenditures and BMBF Grants

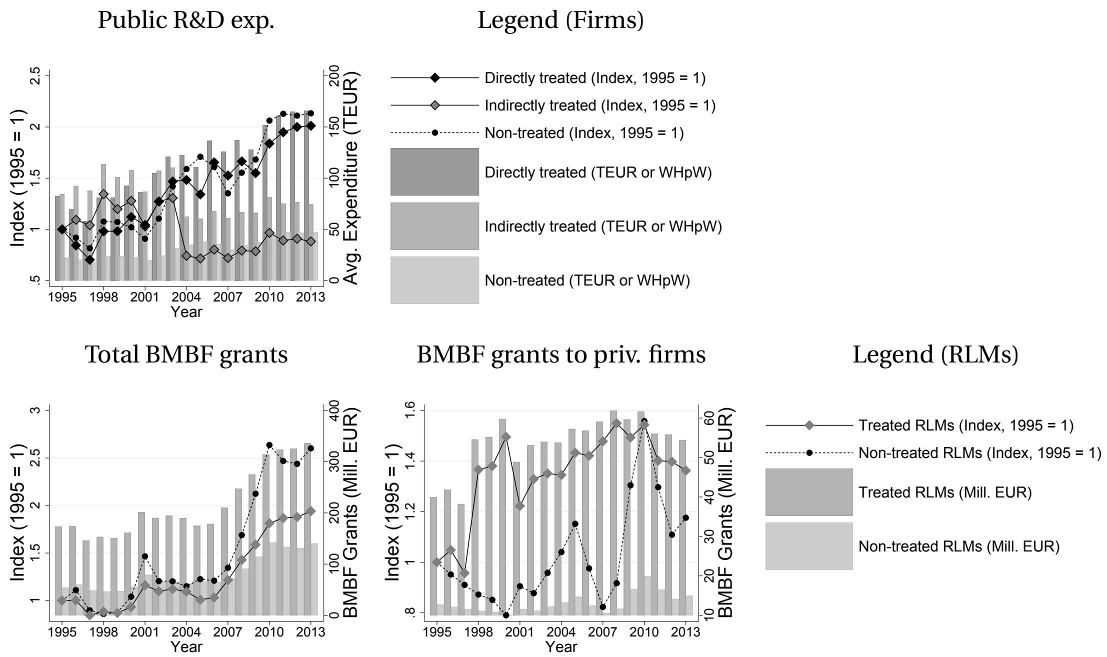

Note: The first panel shows the descriptive development of publicly financed R&D expenditures for the directly and indirectly treated firms, as well as for the control group. The second panel denotes similar developments for the sum of grants made by the BMBF to treated and non-treated regions. The left graph depicts the sum of all grants whereas the right graph only shows the sum of all BMBF grants made to private firms. The grey bars in the first panel represent yearly averages (directly treated firms are depicted in dark grey to the left of indirectly treated and non-treated ones). The connected lines indicate the development of these yearly averages indexed to the year 1995 (directly treated firms are denoted by black diamonds, indirectly treated firms by grey diamonds with black outline, and the control group by black dots). In the second panel, grey bars indicate the sum of BMBF grants (treated regions are to the left of non-treated ones). The connected lines indicate the development of these yearly sums indexed to the year 1995 (treated regions are denoted by grey diamonds and non-treated ones by black dots).

Figure A.3.2: Event Studies – Treated RLMs (All Specifications)

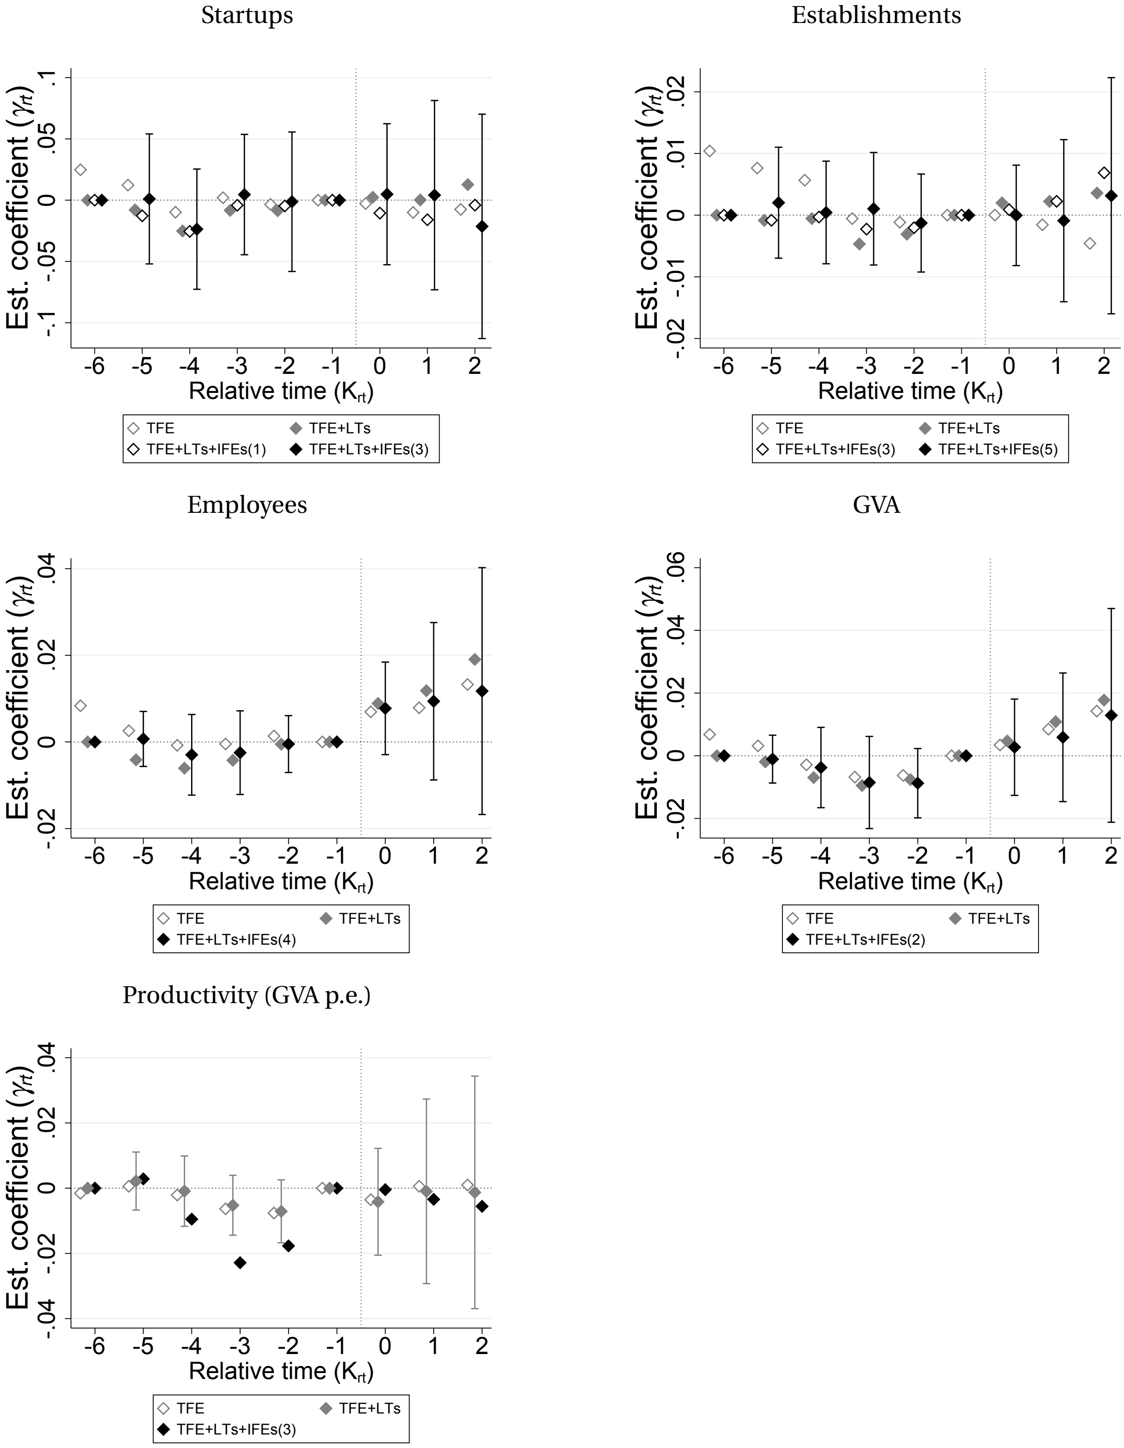

Note: This figure illustrates the results of the event study models shown in Equations (1.4.1) and (1.4.2) for $\{k\}_{-6}^{2}$, respectively. TFE (TFE+LTs) estimates are denoted by white diamonds with grey outline(grey diamonds). White diamonds with black outline (block diamonds) depict the TFE+LTs+IFE model that use the number of factors that is associated with the lowest information criterion (provides the most preferable pre-trends according to our selection algorithm). Solid whiskers denote confidence intervals for our preferred specification at the 95% level.

A.4 Additional Tables

Table A.4.1: Distribution of Firms by Industrial Sector (Percentages)

Sector	East Germany w/o Berlin (1)	Berlin (2)
Manufacturing	70.3	44.3
Textiles & leather	3.7	0.0
Chemicals & pharmaceuticals	4.3	3.9
Non-metallic products	3.4	0.9
Basic & fabricated metals	11.5	3.5
Electronics & optics	14.5	21.1
Machinery & equipment	15.6	8.7
Cars & transport	5.3	2.7
Other manufacturing	12.1	3.6
Services	18.3	46.2
Others	11.4	9.5
Total	100.0	100.0

Note: This table shows the industrial distributions of innovative firms included in our sample. We distinguish between two different groups: all firms in East Germany and firms located in Berlin. For data confidentiality reasons, we subsume firms in industrial sections or divisions (printed in italic) with more than 1 but less than 5 firms observed between 1995 and 2013 in the "Others" section or the "*Other manufacturing*" division.

Table A.4.2: Robustness of Event Studies – Direct Effects of the IRGC

Dep. Var.	*ln R&D Expenditures*				*ln R&D Personnel (Total)*			
Model	TFE	TFE	TFE+LT	TFE w/ treated	TFE	TFE	TFE+LT	TFE w/ treated
$k = -6$	–0.249**	.	.	.	–0.259**	.	.	.
	(0.104)	.	.	.	(0.101)	.	.	.
$k = -5$	–0.274***	–0.162***	–0.075	–0.076	–0.278***	–0.160***	–0.071	–0.071
	(0.099)	(0.061)	(0.047)	(0.047)	(0.092)	(0.059)	(0.051)	(0.049)
$k = -4$	–0.186**	–0.075	–0.036	–0.046	–0.199**	–0.083	–0.044	–0.058
	(0.088)	(0.060)	(0.057)	(0.058)	(0.083)	(0.058)	(0.056)	(0.057)
$k = -3$	–0.068	0.041	0.032	0.023	–0.114	0.000	–0.011	–0.034
	(0.076)	(0.057)	(0.057)	(0.058)	(0.072)	(0.056)	(0.056)	(0.055)
$k = -2$	–0.026	0.080	0.024	0.029	–0.047	0.065	0.005	–0.021
	(0.076)	(0.068)	(0.069)	(0.072)	(0.041)	(0.043)	(0.036)	(0.036)
$k = -1$	.	.	.	.	.	.	.	.
	.	.	.	.	.	.	.	.
$k = 0$	0.114*	0.218***	0.064	0.062	0.055*	0.164***	0.003	0.005
	(0.059)	(0.058)	(0.070)	(0.073)	(0.041)	(0.057)	(0.048)	(0.050)
$k = +1$	0.209***	0.313***	0.109	0.109	0.111*	0.221***	0.007	0.014
	(0.071)	(0.063)	(0.099)	(0.103)	(0.058)	(0.063)	(0.078)	(0.080)
$k = +2$	0.241***	0.348***	0.092	0.079	0.186***	0.298***	0.031	0.035
	(0.081)	(0.068)	(0.126)	(0.132)	(0.066)	(0.069)	(0.101)	(0.102)
Interpol.	Y	Y	Y	Y	Y	Y	Y	Y
Firm-FE	Y	Y	Y	Y	Y	Y	Y	Y
Year-FE	Y	Y	Y	Y	Y	Y	Y	Y
Lin. Trs.	N	N	Y	N	N	N	Y	N
Firms	1,898	1,898	1,898	206	1,898	1,898	1,898	206
N	14,787	14,787	14,787	2,017	14,774	14,774	14,774	2,012

Note: This table shows additional event study results for the directly treated firms (i.e., Equations (1.3.1) and (1.3.2) for $\{k\}_{-6}^{2}$). Dependent variables (measured in natural logs): monetary variables are denoted in TEUR (R&D expenditure and turnover), R&D staff in WHpW and employees in head counts. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Table A.4.3: Robustness of Event Studies – Indirect Effects of the IRGC

Dep. Var.	*ln R&D Expenditures*				*ln R&D Personnel (Total)*			
Model	TFE	TFE	TFE+LT	TFE w/ treated	TFE	TFE	TFE+LT	TFE w/ treated
$k = -6$	–0.010	.	.	.	–0.019	.	.	.
	(0.036)	.	.	.	(0.041)	.	.	.
$k = -5$	–0.014	–0.008	–0.005	.	–0.042	–0.032**	–0.024*	.
	(0.026)	(0.012)	(0.012)	.	(0.030)	(0.013)	(0.013)	.
$k = -4$	–0.027	–0.022	–0.018	.	–0.043	–0.033	–0.027	.
	(0.021)	(0.018)	(0.018)	.	(0.027)	(0.024)	(0.024)	.
$k = -3$	–0.014	–0.009	–0.007	.	–0.041**	–0.032*	–0.028*	.
	(0.014)	(0.015)	(0.014)	.	(0.020)	(0.018)	(0.016)	.
$k = -2$	–0.007	–0.002	–0.003	.	–0.008	0.002	0.000	.
	(0.011)	(0.018)	(0.011)	.	(0.015)	(0.019)	(0.014)	.
$k = -1$	.	.	.	.	.	.	.	.
	.	.	.	.	.	.	.	.
$k = 0$	0.005	0.010	0.000	.	0.010	0.019	0.004	.
	(0.014)	(0.024)	(0.016)	.	(0.015)	(0.025)	(0.018)	.
$k = +1$	0.012	0.016	0.002	.	0.016	0.024	0.000	.
	(0.019)	(0.026)	(0.025)	.	(0.021)	(0.031)	(0.024)	.
$k = +2$	0.031	0.036	0.019	.	0.023	0.031	0.002	.
	(0.032)	(0.038)	(0.040)	.	(0.032)	(0.043)	(0.035)	.
Interpol.	Y	Y	Y	.	Y	Y	Y	.
Firm-FE	Y	Y	Y	.	Y	Y	Y	.
Year-FE	Y	Y	Y	.	Y	Y	Y	.
Lin. Trs.	N	N	Y	.	N	N	Y	.
RLMs	53	53	53	.	53	53	53	.
Firms	4,242	4,242	4,242	.	4,242	4,242	4,242	.
N	33,008	33,008	33,008	.	32,998	32,998	32,998	.

Note: This table shows dynamic treatment effects for the indirectly treated firms. Dependent variables (measured in natural logs): monetary variables are denoted in TEUR (R&D expenditure and turnover), R&D staff in WHpW and employees in head counts. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Table A.4.4: Robustness of Indirect Effects of the IRGC (TFE)

	ln R&D Expenditures			*ln R&D Personnel*			*ln Econ. Outcomes*	
Dep. Var.:	Total	Private	Public	Total	Scien-tists	Techni-cians	Turn-over	Emplo-yees
Years 0–1	0.020	0.019	–0.068	0.031	0.025	0.009	0.014	0.006
	(0.021)	(0.028)	(0.050)	(0.027)	(0.036)	(0.042)	(0.018)	(0.013)
Years 2–3	0.049	0.042	–0.036	0.051	0.080*	0.012	–0.003	–0.002
	(0.038)	(0.044)	(0.059)	(0.040)	(0.046)	(0.068)	(0.025)	(0.018)
Years 4–5	0.036	0.043	–0.131*	0.058	0.090	0.012	0.007	–0.003
	(0.045)	(0.049)	(0.077)	(0.048)	(0.054)	(0.079)	(0.034)	(0.022)
Years 6+	0.068	0.046	–0.185	0.058	0.125	–0.010	0.008	0.008
	(0.058)	(0.072)	(0.113)	(0.070)	(0.100)	(0.108)	(0.053)	(0.033)
Firm-FE	Y	Y	Y	Y	Y	Y	Y	Y
Year-FE	Y	Y	Y	Y	Y	Y	Y	Y
RLM Trs.	N	N	N	N	N	N	N	N
RLMs	53	53	53	53	53	53	53	53
Firms	4,242	4,242	4,242	4,242	4,242	4,242	4,242	4,242
N	33,008	33,008	33,008	32,998	32,998	32,998	33,008	33,008
Means in 1999	551.7 TEUR	409.0 TEUR	101.8 TEUR	305.6 WHpW	178.7 WHpW	68.9 WHpW	10,020.6 TEUR	84.2 Count

Note: This table shows dynamic treatment effects for the indirectly treated firms during the sub-periods "Years 0–1", "Years 2–3", "Years 4–5", and "Years 6+". The effects are weighted averages of the γ^{ind} coefficients estimated based on Equation (1.3.7). The weighting scheme is defined as: $\omega_k = \frac{\sum_{i \in N} \sum_{t \in T} \mathbf{1}\{K_{rt}=k\}}{\sum_{i \in N} \sum_{t \in T} \sum_{j=\underline{j}}^{\overline{j}} \mathbf{1}\{K_{rt}=j\}} \forall k \in \{\underline{j}, \overline{j}\}$. That is, ω_k is the share of treated observations with non-missing entries in period k relative to the sum of all treated observations with valid entries in the relevant sub-period. Outcome variables are interpolated in even years and expressed in natural logs. Values smaller than one have been replaced by one prior to taking logs. Monetary variables are denoted in TEUR (R&D expenditures and turnover), R&D personnel in WHpW, and employees in head counts. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Table A.4.5: Robustness of Aggregate Effects of the IRGC (East Germany w/o Berlin)

Dep. Var.:	Startups	Establish-ments	Employees	GVA	Productivity (GVA p.e.)
Years 0–1	0.019	0.007	0.005	0.003	–0.001
	(0.029)	(0.006)	(0.007)	(0.008)	(0.010)
Years 2–3	0.035	0.017	0.007	0.007	–0.003
	(0.036)	(0.011)	(0.016)	(0.019)	(0.017)
Years 4–5	0.069	0.029	0.012	–0.005	–0.018
	(0.063)	(0.020)	(0.025)	(0.032)	(0.027)
Years 6+	0.143	0.059	0.030	–0.024	–0.034
	(0.105)	(0.039)	(0.042)	(0.060)	(0.051)
Controls	Y	Y	Y	Y	Y
Firm-FE	Y	Y	Y	Y	Y
Year-FE	Y	Y	Y	Y	Y
RLM Trends	Y	Y	Y	Y	Y
Factors	3	5	4	2	-
RLMs	53	53	53	53	53
N	1,007	1,007	1,007	1,007	1,007
Means in 1999	742.3 Count	6,027.8 Count	138.1 Count ('000s)	6,116.4 Mill. EUR	43.5 TEUR p.e.

Note: This table shows dynamic treatment effects at the regional level (RLMs) during the sub-periods "Years 0–1", "Years 2–3", "Years 4–5", and "Years 6+". The effects are weighted averages of the γ^{agg} coefficients estimated based on Equations (1.4.3) or (1.4.4). The weighting scheme is defined as: $\omega_k = \frac{\sum_{r\in R}\sum_{t\in T}\mathbf{1}\{K_{rt}=k\}}{\sum_{r\in R}\sum_{t\in T}\sum_{j=\underline{j}}^{\overline{j}}\mathbf{1}\{K_{rt}=j\}}\ \forall k \in \{\underline{j},\overline{j}\}$. That is, ω_k is the share of treated observations with non-missing entries in period k relative to the sum of all treated observations with valid entries in the relevant sub-period. Outcome variables are expressed in natural logs. Startups, establishments and employees are measured as counts, GVA in million Euros and productivity as GVA per employee. Standard errors are clustered at the RLM level and presented in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Appendix B

Appendix to Chapter 2

B.1 Additional Figures

Figure B.1.1: Technical Availability of NGA Networks in Germany and the Sample Municipalities over Time

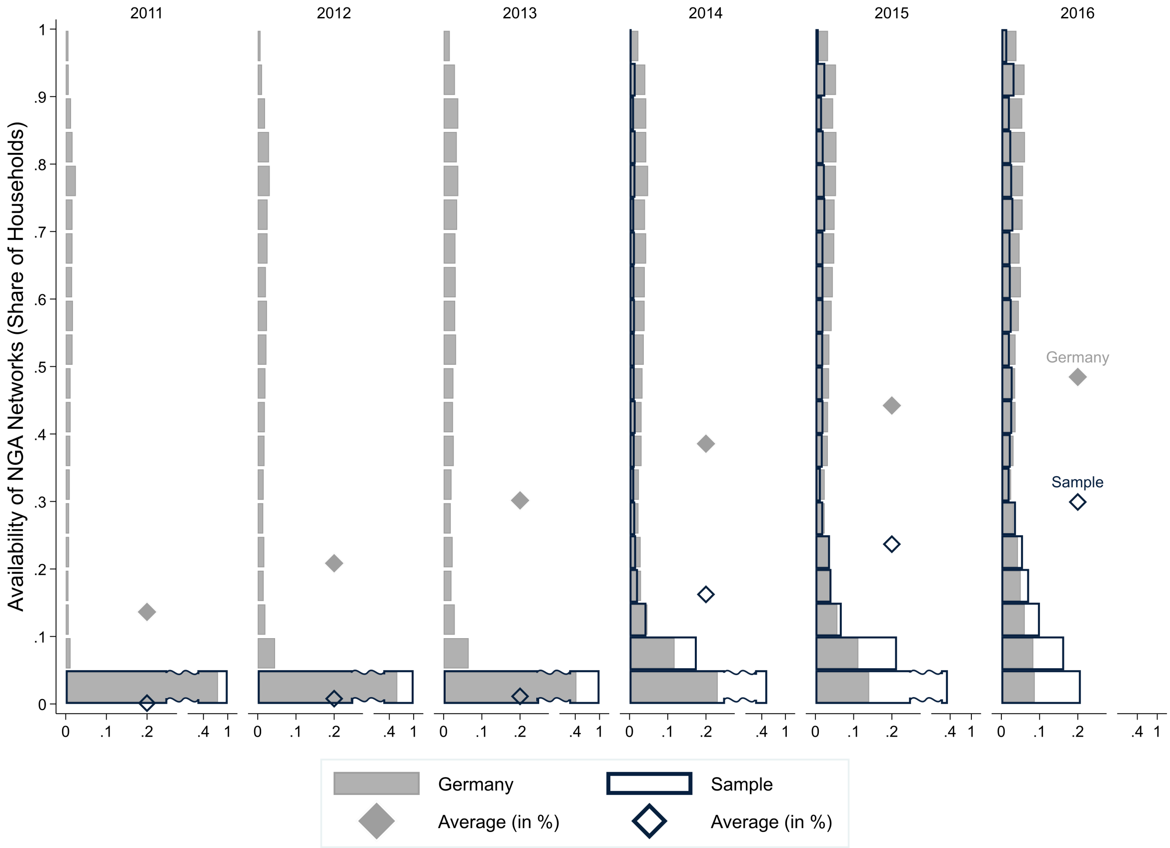

Note: These graphs show frequency distributions that illustrate the availability of NGA networks in Germany for each year between 2011 and 2016. Grey and blue bordered bars span 5 percentage point bins of NGA network availability and denote frequencies for Germany overall and the municipalities included in the sample, respectively. Grey and blue bordered diamonds represent averages for Germany overall and the municipalities included in the sample, respectively. NGA network availability is measured as the share of households in a municipality for which a downstream data transfer rate of at least 50 MBit/s via fixed lines is technically feasible. Data: Breitbandatlas (2017).

Figure B.1.2: Cross-sectional Distributions of the Number of Nodes and Total Edge Length

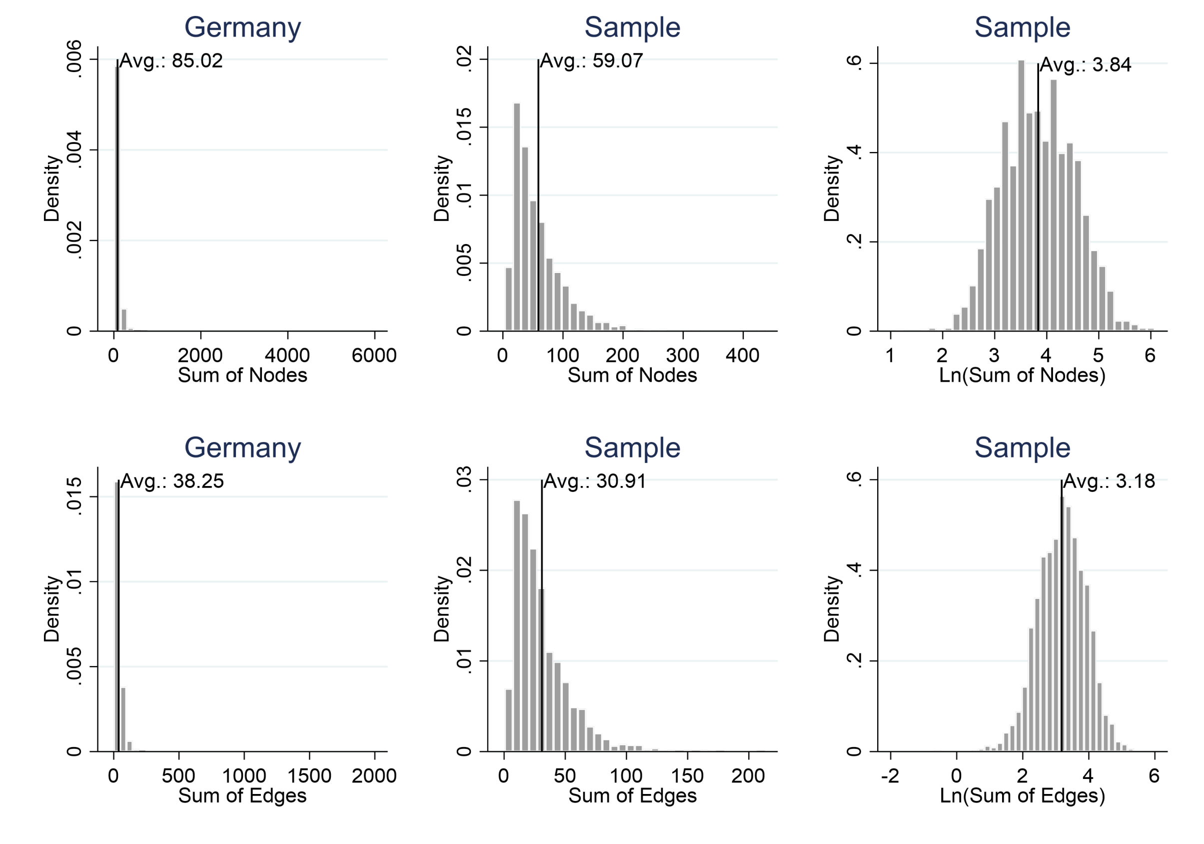

Note: These graphs show cross-sectional histograms that illustrate the number of nodes and total edge length in Germany overall and the municipalities included in the sample, respectively.

B.2 Additional Tables

Table B.2.1: Variable Definitions and Sources

Variable	Description	Source
Employment rate	Number of employed individuals at place of work divided by the working age population (in percent)	Federal Employment Agency ("revidierte Gemeindebände")
Share of employees	Number of employed individuals at place of residence divided by the working age population (in percent)	Federal Employment Agency ("revidierte Gemeindebände")
Fixed line access to ≥ X MBit/s	Share of households covered at a minimum downstream data transfer rate of at least X MBit/s via fixed lines $\forall X \in \{1, 16, 50\}$, where $X \geq 50$ MBit/s measures NGA	Breitbandatlas/TÜV Rheinland
Mobile access to ≥ 2 MBit/s	Share of households covered by Long Term Evolution at a minimum downstream data transfer rate of 2 MBit/s	Breitbandatlas/TÜV Rheinland
Total population	Number of residents per municipality	Federal Statistical Office
Population density (inhabited areas)	Number of residents per km^2 in inhabited areas	Federal Statistical Office
Population density (total area)	Number of residents per km^2 in total area	Federal Statistical Office
Share of female population	Number of female residents divided by the number of all residents	Federal Statistical Office
Share of working age population	Number of residents between 15 and 64 years old divided by the number of all residents	Federal Statistical Office
Share of population aged 65 and older	Number of residents older than 64 years divided by the number of all residents	Federal Statistical Office
Share of net migration	Number of immigrants minus number of emigrants divided by the number of all residents	Federal Statistical Office
Reachability of next X	Travel time it takes by car to reach next X (in minutes) $\forall X \in$ {motorway, train station, airport, hospital, middle-order centre, higher-order centre,}	BBSR
Dummy for settlement type X	Dummy equals 1 if municipality is of settlement type X $\forall X \in$ {Large- & medium sized, small, rural}	BBSR
Dummy for centre type X	Dummy equals 1 if municipality is an economic centre of type X $\forall X \in$ {middle-order, higher-order}	BBSR
Dummy for growth type X	Dummy equals 1 if municipality is of growth type X $\forall X \in$ {shrinking, stagnating, growing}	BBSR

Note: This table describes the definitions of the variables and their sources. The data set is at the level of municipalities/municipality associations in the administrative borders as of December 31, 2015. Information on administrative boundaries were obtained from the Service Centre of the Federal Agency for Geo-Information and Geodesy.

Table B.2.2: Summary Statistics

	All municipalities		Sample		
	'11–'16	'11,'13,'16	2011, 2013, 2016		
	Mean (1)	Mean (2)	Mean (3)	Median (4)	SD (5)
Outcome Variables (in percent)					
Employment rate	42.80	42.67	38.20	33.76	20.02
Share of employees	58.23	58.10	58.77	58.83	4.92
Internet Availability (share of households)					
Fixed line access to ≥ 50 MBit/s (NGA)	0.33	0.31	0.10	0.01	0.22
Fixed line access to ≥ 16 MBit/s	0.58	0.56	0.40	0.33	0.30
Fixed line access to ≥ 1 MBit/s	0.90	0.89	0.85	0.92	0.18
Mobile access to ≥ 2 MBit/s	0.61	0.56	0.56	0.70	0.39
Population Characteristics					
Total population ('000s)	18.21	18.20	8.12	6.16	7.19
Population density in inhabited areas (people per km^2)	1,778.48	1,778.25	1,451.05	1,363.62	565.39
Population density (people per km^2)	279.40	279.28	151.51	100.35	156.10
Share of female population	50.42	50.44	50.26	50.28	1.10
Share of working age population (15–64 years old)	65.38	65.35	65.14	65.47	2.65
Share of population aged 65 and older	21.28	21.25	21.76	21.49	3.76
Share of net migration	0.42	0.24	0.08	0.04	1.03
Reachability (in minutes by car, 2015)					
Travel time to next motorway	15.12	15.12	17.08	14.00	12.76
Travel time to next train station	28.84	28.84	33.75	30.90	17.98
Travel time to next airport	56.82	56.82	62.53	62.70	24.63
Travel time to next hospital	12.14	12.14	13.83	13.86	6.56
Travel time to next centre	10.47	10.47	12.00	12.20	7.40
Settlement Type (shares)					
Large & medium-sized towns	0.16	0.16	0.05	–	0.22
Small towns	0.47	0.47	0.44	–	0.50
Rural areas	0.37	0.37	0.51	–	0.50
Higher-order centre	0.03	0.03	0.01	–	0.08
Middle-order centre	0.17	0.17	0.12	–	0.33
Growth Type (shares)					
Shrinking	0.34	0.34	0.48	–	0.50
Constant	0.14	0.14	0.13	–	0.34
Growing	0.52	0.52	0.39	–	0.49
Municipalities	4,462	4,462	1,625		
Regional Labour Markets (RLMs)	258	258	184		
Observations	26,772	13,386	4,875		

Note: This table reports summary statistics for all variables included in the empirical analysis. It distinguishes between three different samples: Column (1) reports average values for the universe of municipalities in Germany pooled over the years from 2011 to 2016. The average values in Column (2) are based on a pooled sample including the universe of German municipalities during the years 2011, 2013, and 2016. Columns (3) – (5) report three summary statistics for the municipalities included in the estimation sample. As in Column (2), the estimation sample only includes the years 2011, 2013, and 2016.

Table B.2.3: Robustness I – Additional Types of Fixed Effects

	IV				
	(1)	(2)	(3)	(4)	(5)
Dep. Var.					
ΔEmployment rate	0.379	0.244	–0.774	0.265	–0.648
	(1.544)	(2.131)	(2.161)	(1.637)	(2.752)
ΔShare of employees	1.867**	1.735	1.976*	2.179***	2.315*
	(0.820)	(1.094)	(1.126)	(0.850)	(1.398)
First stage					
$\Delta[ln(edges_m)\times$ $\mathbf{1}\{t=2016\}]$	–0.092***	–0.073***	–0.070***	–0.089***	–0.059***
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
$\Delta[ln(nodes_m)\times$ $\mathbf{1}\{t=2016\}]$	0.134***	0.092***	0.092***	0.130***	0.068***
	(0.009)	(0.011)	(0.010)	(0.010)	(0.011)
Kleibergen-Paap rk Wald F statistic	105.395	63.912	64.570	99.017	43.015
Controls					
Density	Yes	Yes	Yes	Yes	Yes
Population	Yes	Yes	Yes	Yes	Yes
Fixed Effects					
Municipalities	Yes	Yes	Yes	Yes	Yes
RLM×year	Yes	Yes	Yes	Yes	Yes
Settlement×year	–	Yes	–	–	Yes
Centre×year	–	–	Yes	–	Yes
Growth×year	–	–	–	Yes	Yes
Municipalities	1,625	1,625	1,625	1,625	1,625
N	3,250	3,250	3,250	3,250	3,250

Note: This table reports IV results for Equation (2.3.1) using additional sets of fixed effects. Across all specifications, $\Delta[ln(edges_m)\times\mathbf{1}\{t=2016\}]$ serves as the excluded instrument while $\Delta[ln(nodes_m)\times\mathbf{1}\{t=2016\}]$ is an included instrument. Furthermore, all regressions include annualised changes in population density, separately measured within inhabited areas and across the entire municipality ("Density"), the shares of the female population, the working age population, the population aged 65 and older, net migration ("Population"), and municipality-level as well as period-specific fixed effects as baseline controls. That is, the control variables correspond to the ones used in the preferred baseline specification shown in Column (3) of Table 2.5.1 (repeated for convenience in Column (1)). Additional to these baseline controls, the following specifications separately control for arbitrary shocks at the settlement type-by-year (Column (2)), type of centre-by-year (Column (3)), and growth type-by-year (Column (4)) level. Column (5) controls for all of these fixed effects simultaneously. All Kleibergen-Paap rk Wald F statistic are above the critical values calculated by Stock and Yogo [2005], rejecting the null hypothesis that the excluded instrument is weak. Standard errors are clustered within regional labour markets and shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table B.2.4: Robustness II – Different Thresholds for "No NGA" in the Pre-Period

	IV			
Threshold	5%	2.5%	7.5%	10%
	(1)	(2)	(3)	(4)
Dep. Var.				
ΔEmployment rate	0.379	1.815	–0.491	–0.004
	(1.544)	(1.652)	(1.600)	(1.535)
ΔShare of employees	1.868**	1.509*	1.444*	1.568**
	(0.818)	(0.903)	(0.772)	(0.726)
First stage				
$\Delta[ln(edges_m) \times$	–0.092***	–0.094***	–0.092***	–0.095***
$\mathbf{1}\{t = 2016\}]$	(0.009)	(0.010)	(0.008)	(0.008)
$\Delta[ln(nodes_m) \times$	0.131***	0.134***	0.127***	0.135***
$\mathbf{1}\{t = 2016\}]$	(0.010)	(0.011)	(0.009)	(0.009)
Kleibergen-Paap rk Wald F statistic	105.341	93.150	117.432	128.812
Controls				
Density	Yes	Yes	Yes	Yes
Population	Yes	Yes	Yes	Yes
Municipalities	1,625	1,335	1,779	1,880
N	3,250	2,670	3,558	3,760

Note: This table reports IV results for Equation (2.3.1) using different definitions for "No NGA" in the base period, that is, the threshold values referring to the share of households with access to NGA networks in 2013. Across all specifications, $\Delta[ln(edges_m) \times \mathbf{1}\{t = 2016\}]$ serves as the excluded instrument while $\Delta[ln(nodes_m) \times \mathbf{1}\{t = 2016\}]$ is an included instrument. Furthermore, all regressions include annualised changes in population density, separately measured within inhabited areas and across the entire municipality ("Density"), the shares of the female population, the working age population, the population aged 65 and older, net migration ("Population"), and municipality-level as well as period-specific fixed effects as baseline controls. That is, the control variables correspond to the ones used in the preferred baseline specification shown in Column (3) of Table 2.5.1 (repeated for convenience in Column (1)). All Kleibergen-Paap rk Wald F statistic are above the critical values calculated by Stock and Yogo [2005], rejecting the null hypothesis that the excluded instruments are weak. Standard errors are clustered within regional labour markets and shown in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Appendix C

Appendix to Chapter 3

C.1 Additional Figures

Figure C.1.1: Turnover Share of E-Books in 2013 across Countries

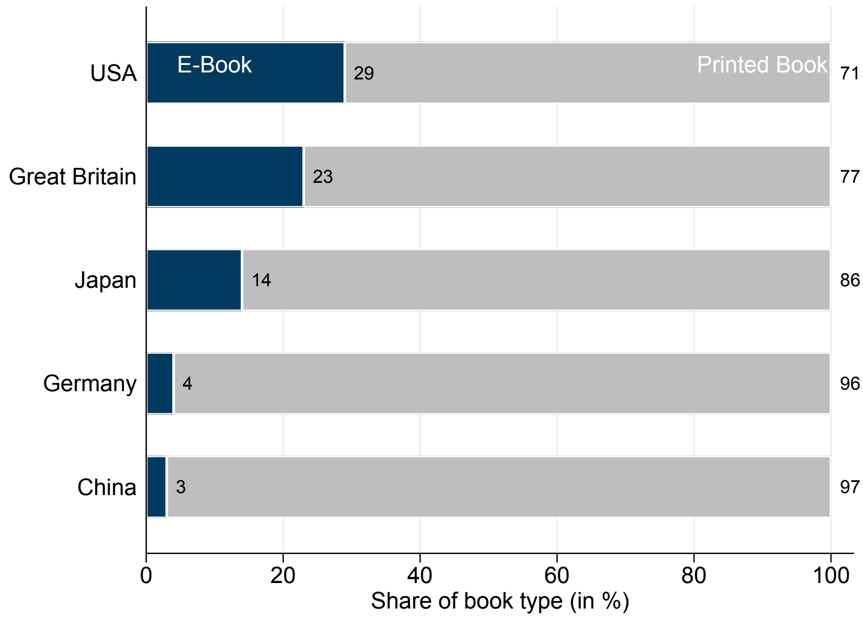

Note: This figure depicts the turnover shares of traditional books (including audio books) and e-books in 2013 across different countries. Data: Richter [2014].

Figure C.1.2: Relationship between (pre-Internet) Age Shares and Changes in Control Variables (Continued)

(1) Population density

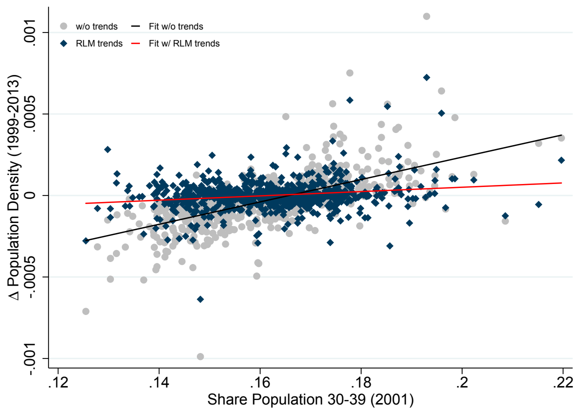

(2) Share female

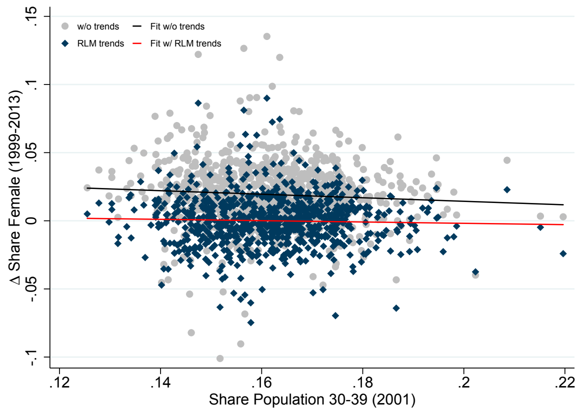

(3) Share German

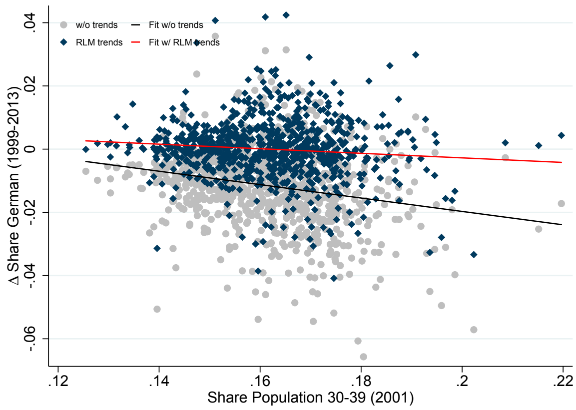

(4) Share employed

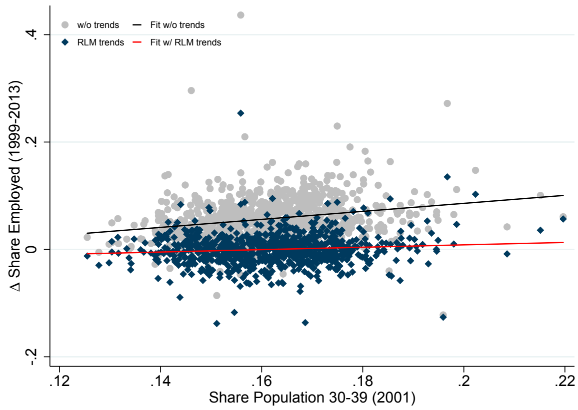

(5) Lagged share of establishments

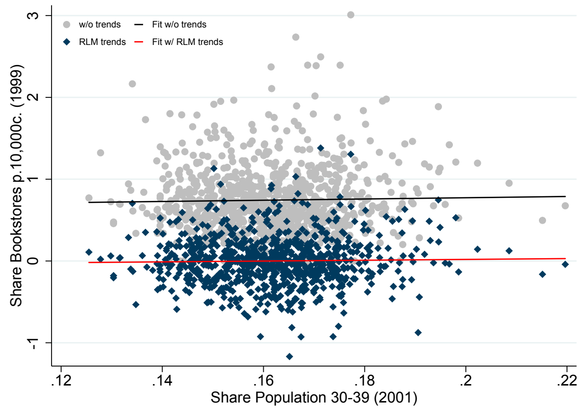

Note: This figure presents five scatter plots that illustrate the relationship between the share of 30 to 39 year old people in 2001 and the change in a selection of control variables between 1999 and 2013. Gray dots (black lines) indicate scatter plots (OLS regression coefficients) of the unconditional relationship between the change in the control variable of interest and the pre-Internet share of people between 30 and 39. Blue diamonds (red lines) represent added variable scatter plots (OLS regression coefficients) of the same relationships after accounting for RLM-fixed effects. Data: Nexiga, Federal Statistical Office, BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

C.2 Additional Tables

Table C.2.1: Summary Statistics Along the Distribution of Bookstore Density in 1999

Percentiles	(0, 10]	(10, 25]	(25, 50]	(50, 75]	(75, 90]	(90, 100]
	(1)	(2)	(3)	(4)	(5)	(6)
Local Book Markets in 1999						
Bookstores	0.27	0.43	0.60	0.81	1.10	1.63
Employees (total) p.b.	5.76	5.54	4.70	4.73	4.40	3.63
Employees (full-time) p.b.	2.46	2.34	1.92	2.03	1.82	1.36
Employees (part-time) p.b.	0.93	1.00	0.82	0.81	0.80	0.56
Employees (marginal) p.b.	1.92	1.85	1.64	1.62	1.56	1.50
Local Population in 1999						
Population (total, tsd.)	49.17	72.92	87.39	121.27	97.81	77.60
Population Density	0.11	0.14	0.16	0.17	0.15	0.11
Employees (total, tsd.)	15.60	25.2	30.71	48.81	41.30	31.20
Share Employed	0.32	0.33	0.33	0.35	0.36	0.36
Share Female	0.47	0.47	0.47	0.48	0.48	0.48
Share German	0.94	0.95	0.95	0.95	0.94	0.94
Share Medium Skilled	0.73	0.74	0.74	0.74	0.72	0.73
Share High Skilled	0.06	0.07	0.07	0.08	0.08	0.07
Average Hourly Wages	48.93	53.50	56.90	60.40	64.72	66.97
N	80	118	198	198	119	79

Note: This table provides additional summary statistics for variables describing the structure of the local book market. These are either measured per 10,000 residents (bookstores) or per bookstore (total number of employees as well as full-time, part-time, and marginal employees). The total number of residents and employees is expressed in thousands. Population density is measured in terms of 100 residents per km^2. Average hourly wages are calculated for all full-time workers within a retail market. Data: Nexiga, Federal Statistical Office, BHP (for a documentation of the BHP, see Schmucker et al. [2016]).

List of Figures

List of Tables

Bibliography

Acemoglu, D. and Linn, J. [2004], 'Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry', *Quarterly Journal of Economics* **119**(3), 1049–1090.

Acemoglu, D. and Restrepo, P. [2017], 'Robots and Jobs: Evidence from US Labor Markets', *NBER Working Paper* (23285).

Aghion, P. and Howitt, P. [1992], 'A Model of Growth Through Creative Destruction', *Econometrica* **60**(2), 323–351.

Agrawal, A., Cockburn, I., Galasso, A. and Oettl, A. [2014], 'Why Are Some Regions More Innovative Than Others? The Role of Small Firms in the Presence of Large Labs', *Journal of Urban Economics* **81**, 149–165.

Ahlfeldt, G., Koutroumpis, P. and Valletti, T. [2017], 'Speed 2.0: Evaluating Access to Universal Digital Highways', *Journal of the European Economic Association* **15**(3), 586–625.

Ahlfeldt, G. M. and Feddersen, A. [2018], 'From Periphery to Core: Measuring Agglomeration Effects using High-Speed Rail', *Journal of Economic Geography* **18**(2), 355–390.

Akerman, A., Gaarder, I. and Mogsatd, M. [2015], 'The Skill Complementarity of Broadband Internet', *Quarterly Journal of Economics* **130**(4), 1781–1824.

Alonso, W. [1964], *Location and Land Use: Toward a General Theory of Land Rent*, Harvard University Press, Cambridge, Massachusetts.

Andersson, R., Quigley, J. M. and Wilhelmson, M. [2004], 'University Decentralization as Regional Policy: The Swedish Experiment', *Journal of Economic Geography* **4**(4), 371–388.

Angrist, J. D. and Pischke, J.-S. [2009], *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, Princeton and Oxford.

Arrow, K. J. [1962], Economic Welfare and the Allocation of Resources for Invention, *in* R. R. Nelson, ed., 'The Rate and Direction of Inventive Activity: Economic and Social Factors', Princeton University Press, Princeton, New Yersey, pp. 609–628.

Atasoy, H. [2013], 'The Effect of Broadband Internet Expansion on Labor Market Outcomes', *Industrial and Labor Relations Review* **66**(2), 315–345.

Autor, D. H., Dorn, D. and Hanson, G. H. [2013], 'The China Syndrome: Local Labor Market Impacts of Import Competition in the United States', *American Economic Review*

103(6), 2121–2168.

Bai, J. [2009], 'Panel Data Models with Interactive Fixed Effects', *Econometrica* **77**(4), 1229–1279.

Bai, J. and Ng, S. [2002], 'Determining the Number of Factors in Approximate Factor Models', *Econometrica* **70**(1), 191–221.

Bai, Y. [2017], 'The Faster, the Better? The Impact of Internet Speed on Employment', *Information Economics and Policy* **40**, 21–25.

Bartik, T. J. [1991], *Who Benefits from State and Local Economic Development Policies?*, W. E. Upjohn Institute for Employment Research, Kalamazoo, Michigan.

Basker, E. [2005], 'Job Creation or Destruction? Labor Market Effects of Wal-Mart Expansion', *Review of Economics and Statistics* **87**(1), 174–183.

Basker, E. [2012], 'Raising the Barcode Scanner: Technology and Productivity in the Retail Sector', *American Economic Journal: Applied Economics* **4**(3), 1–27.

Bauernschuster, S., Falck, O. and Woessmann, L. [2014], 'Surfing Alone? The Internet and Social Capital: Evidence from an Unforeseeable Technological Mistake', *Journal of Public Economics* **117**, 73–89.

Baumol, W. J. [1977], 'On the Proper Cost Tests for Natural Monopoly in a Multiproject Industry', *American Economic Review* **67**(5), 809–288.

BBSR [2018], *INKAR 2018: Erläuterungen zu den Raumbezügen*, Bundesinstitut für Bau-, Stadt- und Raumforschung im Bundesamt für Bauwesen und Raumordnung (BBSR, Federal Institute for Research on Building, Urban Affairs and Spatial Development), Bonn. **URL:** *https://www.inkar.de/documents/Erlaeuterungen Raumbezuege.pdf, last accessed: 18 March 2019*

Becker, B. [2015], 'Public R&D Policies and Private R&D Investment: A Survey of the Empirical Evidence', *Journal of Economic Surveys* **29**(5), 917–942.

Becker, S. O., Egger, P. H. and von Ehrlich, M. [2010], 'Going NUTS: The Effect of EU Structural Funds on Regional Performance', *Journal of Public Economics* **94**(9-10), 578–590.

Becker, S. O., Egger, P. H. and von Ehrlich, M. [2012], 'Too Much of a Good Thing? On the Growth Effects of the EU's Regional Policy', *European Economic Review* **56**(4), 648–668.

Becker, S. O., Egger, P. H. and von Ehrlich, M. [2013], 'Absorptive Capacity and the Growth and Investment Effects of Regional Transfers', *American Economic Journal: Economic Policy* **5**(4), 29–77.

Becker, S. O., Egger, P. H. and von Ehrlich, M. [2018], 'Effects of EU Regional Policy: 1989–2013', *Regional Science and Urban Economics* **69**, 143–152.

Behrens, J. [2017], *Book on Demand: Auswirkungen auf den Deutschen Buchmarkt*, BoD Books on Demand, Norderstedt.

Bertamino, F., Bronzini, R., De Maggio, M. and Revelli, D. [2017], 'Regional Policies for Innovation: the Case of Technology Districts in Italy', *Regional Studies* **51**(12), 1826–1839.

Bertschek, I., Cerquera, D. and Klein, G. J. [2013], 'More Bits – More Bucks? Measuring the Impact of Broadband Internet on Firm Performance', *Information Economics and Policy* **25**, 190–203.

BMBF [2016*a*], *Entrepreneurial Regions: The Guiding Principles*, Bundesministerium für Bildung und Forschung (BMBF, Federal Ministry of Education and Research), Berlin.
URL: *https://web.archive.org/web/20150307045625/http://www.unternehmen-region.de:80/en/54.php#, last accessed: 4 January 2018*

BMBF [2016*b*], *Growth Cores Potential: Programme Description*, Bundesministerium für Bildung und Forschung (BMBF, Federal Ministry of Education and Research), Berlin.
URL: *https://web.archive.org/web/20150307045556/http://www.unternehmen-region.de:80/en/4578.php, last accessed: 4 January 2018*

BMBF [2016*c*], *The Innovative Regional Growth Cores Programme: Programme Description*, Bundesministerium für Bildung und Forschung (BMBF, Federal Ministry of Education and Research), Berlin.
URL: *https://web.archive.org/web/20150307045729/http://www.unternehmen-region.de:80/en/4577.php, last accessed: 4 January 2018*

Borraz, F., Dubra, J., Ferrés, D. and Zipitría, L. [2014], 'Supermarket Entry and the Survival of Small Stores', *Review of Industrial Organization* **44**(1), 73–93.

Börsenverein des Deutschen Buchhandels [2015], *Buchkäufer und -leser: Strukturen, Kennzahlen und Einstellungen*, Börsenverein des Deutschen Buchhandels e.V., Frankfurt am Main.
URL: *https://web.archive.org/web/20171106224843/http://www.boersenverein.de/sixcms/media.php/1117/Buchkäufer_und_leser_2015_Pressemappe.pdf, last accessed: 18 March 2019*

Börsenverein des Deutschen Buchhandels [2018], Umsatzanteil von E-Books im Publikumsmarkt in Deutschland in den Jahren 2010 bis 2018, *in* 'Statista - Das Statistik-Portal', Börsenverein des Deutschen Buchhandels e.V., Frankfurt am Main.
URL: *https://de.statista.com/statistik/daten/studie/303339/umfrage/umsatzanteil-von-e-books-im-buchmarkt/, last accessed: 6 January 2018*

Borusyak, K. and Jaravel, X. [2017], 'Revisiting Event Study Designs, with an Application to the Estimation of the Marginal Propensity to Consume', *Working Paper* .
URL: *https://scholar.harvard.edu/files/borusyak/files/event_studies_may8_website.pdf, last accessed: 18 March 2019*

Briglauer, W., Dürr, N. S., Falck, O. and Hüschelrath, K. [2019], 'Does State Aid for Broadband Deployment in Rural Areas Close the Digital and Economic Divide?', *Information Economics and Policy, forthcoming* .

Brynjolfsson, E., Hu, Y. and Rahman, M. S. [2009], 'Battle of the Retail Channels: How Product Selection and Geography Drive Cross-Channel Competition', *Management Science* **55**(11), 1755–1765.

Brynjolfsson, E. and Smith, M. D. [2000], 'Frictionless Commerce? A Comparison of Internet and Conventional Retailers', *Management Science* **46**(4), 563–585.

Bundesnetzagentur [2017], *Annual Report 2016: Markets in the Digital Revolution*, Bundesnetzagentur für Elektrizität, Gas, Telekommunikation, Post und Eisenbahnen (Bundesnetzagentur), Bonn.
URL: *https://www.bundesnetzagentur.de/SharedDocs/Downloads/EN/BNetzA/PressSection/ReportsPublications/2017/AnnualReport2016.pdf?__blob=publicationFile&v=2, last accessed: 18 March 2019.*

Bundesnetzagentur [2018*a*], *Annual Report 2017: Networks for the Future*, Bundesnetzagentur für Elektrizität, Telekommunikation, Gas, Post und Eisenbahnen (Bundesnetzagentur), Bonn.
URL: *https://www.bundesnetzagentur.de/SharedDocs/Downloads/EN/BNetzA/PressSection/ReportsPublications/2018/AnnualReport2017.pdf?__blob=publicationFile&v=2, last accessed: 18 March 2019.*

Bundesnetzagentur [2018*b*], *Decision of the President's Chamber of the Bundesnetzagentur on the Determinations and Rules in Detail (Award Rules) and on the Determinations and Rules for Conduct of the of the Proceedings (Auction Rules) to Award Spectrum in the 2 GHz and 3.6 GHz Bands*, Bundesnetzagentur für Elektrizität, Gas, Telekommunikation, Post und Eisenbahnen (Bundesnetzagentur), Bonn.
URL: *https://www.bundesnetzagentur.de/SharedDocs/Downloads/EN/Areas/ Telecommunications/Companies/TelecomRegulation/FrequencyManagement/ElectronicCommunicationsServices/FrequencyAward2018/20181214_Decision_III_IV.pdf?__blob=publicationFile&v=2, last accessed: 18 March 2019.*

Bundesverband E-Commerce und Versandhandel and Boniversum [2017], Würden Sie die folgenden Produkte lieber im Online- und Versandhandel oder lieber im klassischen Einzelhandel kaufen?, *in* 'Statista - Das Statistik-Portal', Bundesverband E-Commerce und Versandhandel Deutschland, Berlin.
URL: *https://de.statista.com/statistik/daten/studie/201914/umfrage/einkaufsverhalten-im-onlinehandel-vs-einzelhandel-nach-produktgruppen/, last accessed: 10 September 2017.*

Busso, M., Gregory, J. and Kline, P. [2013], 'Assessing the Incidence and Efficiency of a Prominent Place Based Policy', *American Economic Review* **103**(2), 897–947.

Cantner, U. and Kösters, S. [2012], 'Picking the Winner? Empirical Evidence on the Targeting of R&D Subsidies to Start-Ups', *Small Business Economics* **39**(4), 921–936.

Canzian, G., Poy, S. and Schüller, S. [2015], 'Broadband Diffusion and Firm Performance in Rural Areas: Quasi-Experimental Evidence', *IZA Working Paper* (9429).

Cavallo, A. [2017], 'Are Online and Offline Prices Similar? Evidence from Large Multi-Channel Retailers', *American Economic Review* **107**(1), 283–303.

Chatterji, A., Glaeser, E. and Kerr, W. [2014], 'Clusters of Entrepreneurship and Innovation', *Innovation Policy and the Economy* **14**(1), 129–166.

Colombo, M. G., Croce, A. and Grilli, L. [2013], 'ICT Services and Small Businesses' Productivity Gains: An Analysis of the Adoption of Broadband Internet Technology', *Information Economics and Policy* **25**, 171–189.

Correia, S. [2015], 'Singletons, Cluster-Robust Standard Errors and Fixed Effects: A Bad Mix', *Working Paper* .
URL: *http://scorreia.com/research/singletons.pdf, last accessed: 18 March 2019.*

Criscuolo, C., Martin, R., Overman, H. G. and Reenen, J. V. [2018], 'Some Causal Effects of an Industrial Policy', *CEP Discussion Paper* (1113).

Cronau, S. [2006], *Buch und Buchhandel in Zahlen 2006*, Börsenverein des Deutschen Buchhandels e.V., Frankfurt am Main.

Cronau, S. [2011], *Buch und Buchhandel in Zahlen 2011*, Börsenverein des Deutschen Buchhandels e.V., Frankfurt am Main.

Cronau, S. [2013], *Buch und Buchhandel in Zahlen 2013*, Börsenverein des Deutschen Buchhandels e.V., Frankfurt am Main.

Cronau, S. [2014], *Buch und Buchhandel in Zahlen 2014*, Börsenverein des Deutschen Buchhandels e.V., Frankfurt am Main.

Czernich, N. [2014], 'Does Broadband Internet Reduce the Unemployment Rate? Evidence for Germany', *Information Economics and Policy* **29**, 32–45.

Dauth, W., Findeisen, S. and Suedekum, J. [2017], 'Trade and Manufacturing Jobs in Germany', *American Economic Review* **107**(5), 337–342.

De Stefano, T., Kneller, R. and Timmis, J. [2014], 'The (Fuzzy) Digital Divide: Universal Broadband Access and Firm Performance', *University of Nottingham Discussion Paper* **6**.

Dettmann, E., Brachert, M. and Titze [2016], 'Identifying the Effects of Place-Based Policies – Causal Evidence from Germany', *CESifo Working Paper* (5901).

Doms, M. E., Jarmin, R. S. and Klimek, S. D. [2004], 'Information Technology Investment and Firm Performance in US Retail Trade', *Economics of Innovation and New Technology* **13**(7), 595–613.

Duranton, G. [2011], 'California Dreamin': The Feeble Case for Cluster Policies', *Review of Economic Analysis* **3**, 3–45.

Einav, L., Knoepfle, D., Levin, J. and Sundaresan, N. [2014], 'Sales Taxes and Internet Commerce', *American Economic Review* **104**(1), 1–26.

Einiö, E. and Overman, H. G. [2016], 'The (Displacement) Effects of Spatially Targeted Enterprise Initiatives: Evidence from UK LEGI', *CEPR Discussion Paper* (DP11112).

Ellison, G. and Ellison, S. F. [2009], 'Tax Sensitivity and Home State Preferences in Internet Purchasing', *American Economic Journal: Economic Policy* **1**(12), 53–71.

European Commission [2010], 'A Digital Agenda for Europe', *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions* **COM(2010)**(2).

European Commission [2016], 'Connectivity for a Competitive Digital Single Market – Towards a European Gigabit Society', *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions* **COM(2016)**.

Fabritz, N. [2013], 'The Impact of Broadband on Economic Activity in Rural Areas: Evidence from German Municipalities', *ifo Working Paper* (166).

Falck, O., Gold, R. and Heblich, S. [2014], 'E-lections: Voting Behavior and the Internet', *American Economic Review* **104**(7), 2238–2265.

Falck, O., Heblich, S. and Kipar, S. [2010], 'Industrial Innovation: Direct Evidence from a Cluster-Oriented Policy', *Regional Science and Urban Economics* **40**(6), 574–582.

Fons-Rosen, C., Scrutinio, V. and Szemeredi, K. [2016], 'Colocation and Knowledge Diffusion: Evidence from Million Dollar Plants', *CEP Discussion Paper* (1447).

Ford, G. S. [2018], 'Is Faster Better? Quantifying the Relationship between Broadband Speed and Economic Growth', *Telecommunications Policy* **42**, 766–777.

Forman, C., Goldfarb, A. and Greenstein, S. [2012], 'The Internet and Local Wages: A Puzzle', *American Economic Review* **102**(1), 556–575.

Freeman, C. [1987], *Technology and Economic Performance: Lessons from Japan.*, Pinter, London.

Friberg, R. and Grönqvist, E. [2012], 'Do Expert Reviews Affect the Demand for Wine?', *American Economic Journal: Applied Economics* **4**(1), 193–211.

Gans, J. S. and Stern, S. [2003], 'The Product Market and the Market for "Ideas": Commercialization Strategies for Technology Entrepreneurs', *Research Policy* **32**(2), 333–350.

German Federal Government [2014], *Digitale Agenda 2014 – 2017*, Bundesministerium für Wirtschaft und Energie (Federal Ministry for Economics and Energy), Berlin.

Glaeser, E. L. and Gottlieb, J. D. [2008], 'The Economics of Place-Making Policies', *Brookings Papers on Economic Activity* (1), 155–253.

Gobillon, L. and Magnac, T. [2016], 'Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls', *The Review of Economics and Statistics* **98**(3), 535–551.

Gobillon, L., Magnac, T. and Selod, H. [2012], 'Do Unemployed Workers Benefit from Enterprise Zones? The French Experience', *Journal of Public Economics* **96**(9-10), 881–892.

Goldsmith-Pinkham, P., Sorkin, I. and Swift, H. [2018], 'Notes on Bartik Instruments', *NBER Working Paper* (24408).

Goolsbee, A. [2000], 'In a World Without Borders. The Impact of Taxes on Internet Commerce', *Quarterly Journal of Economics* **115**(2), 561–576.

Goolsbee, A. [2001], 'Competition in the Computer Industry: Online Versus Retail', *Journal of Industrial Economics* **49**(4), 487–499.

Goolsbee, A., Lovenheim, M. F. and Slemrod, J. [2010], 'Playing with Fire: Cigarettes, Taxes, and Competition from the Internet', *American Economic Journal: Economic Policy* **2**(1), 131–154.

Graetz, G. and Michaels, G. [2018], 'Robots at Work', *Review of Economics and Statistics* **100**(5), 753–768.

Greenstone, M., Hornbeck, R. and Moretti, E. [2010], 'Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings', *Journal of Political Economy* **118**(3), 536–598.

Hafner, J., Lehmann, C. and Windolph, A. [2015], *Kurzstudie: Ausbau Ultraschneller Breitbandnetze: Potentiale von Ausbaustrategien in unterschiedlichen Siedlungsräumen für eine Versorgung mit 100 Mbit/s*, TÜV Rheinland Consulting GmbH, Berlin.
URL: *https://www.tuv.com/content-media-files/master-content/services/ict-business-solutions/d02-telco-solutions-consulting/1317-tuv-rheinland-broadband-consulting/tuv-rheinland-ausbau-ultraschneller-breitbandnetze-studie-de.pdf, last accessed: 18 March 2019*

Hagedorn, M., Manovskii, I. and Mitman, K. [2016], 'The Impact of Unemployment Benefit Extensions on Employment: The 2014 Employment Miracle?', *NBER Working Paper* (20884).

Hall, B. H. and Lerner, J. [2010], Chapter 14 – The Financing of R&D and Innovation, *in* B. H. Hall and N. Rosenberg, eds, 'Handbook of the Economics of Innovation', Vol. 1, Elsevier, Amerstam and Oxford, pp. 609–639.

Haller, S. A. and Lyons, S. [2015], 'Broadband Adoption and Firm Productivity: Evidence from Irish Manufacturing Firms', *Telecommunications Policy* **39**, 1–13.

Hausman, N. [2017], 'University Innovation and Local Economic Growth', *Working Paper* . **URL:** *https://en.falk.huji.ac.il/sites/default/files/falk/files/paper_17-05.pdf, last accessed: 18 March 2019.*

Hethey, T. and Schmieder, J. F. [2010], 'Using Worker Flows in the Analysis of Establishment Turnover. Evidence from German Administrative Data', *FDZ-Methodenreport* **6**.

Hyytinen, A. and Toivanen, O. [2005], 'Do Financial Constraints Hold Back Innovation and Growth?: Evidence on the Role of Public Policy', *Research Policy* **34**(9), 1385–1403.

Igami, M. [2011], 'Does Big Drive Out Small? Entry, Exit, and Differentiation in the Supermarket Industry', *Review of Industrial Organization* **38**(1), 1–21.

Imbens, G. W. and Angrist, J. D. [1994], 'Identification and Estimation of Local Average Treatment Effects', *Econometrica* **62**(2), 467–475.

Imbens, G. W. and Wooldridge, J. M. [2009], 'Recent Developments in the Econometrics of Program Evaluation', *Journal of Economic Literature* **47**(1), 5–86.

Jia, P. [2008], 'What Happens When Wal-Mart Comes to Town? An Empirical Analysis of the Discount Retailing Industry', *Econometrica* **76**(6), 1263–1316.

Kandilov, I. T. and Renkow, M. [2010], 'Infrastructure Investment and Rural Economic Development', *Growth and Change* **41**(2), 165–191.

Kim, Y. and Orazem, P. F. [2016], 'Broadband Internet and New Firm Location Decisions in Rural Areas', *American Journal of Agricultural Economics* **99**(1), 285–302.

Kline, P. [2010], 'Place Based Policies, Heterogeneity, and Agglomeration', *American Economic Review: Papers and Proceedings* **100**(2), 383–387.

Kline, P. and Moretti, E. [2014*a*], 'Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority', *Quarterly Journal of Economics* **129**(1), 275–331.

Kline, P. and Moretti, E. [2014*b*], 'People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs', *Annual Review of Economics* **6**, 629–662.

Kolko, J. [2000], The Death of Cities? The Death of Distance? Evidence from the Geography of Commercial Internet Usage, *in* I. Vogelsang and B. M. Compaine, eds, 'The Internet Upheaval: Raising Questions, Seeking Answers in Communications Policy', MIT Press and TPRC, Cambridge, Massachusetts, pp. 73–98.

Kolko, J. [2012], 'Broadband and Local Growth', *Journal of Urban Economics* **71**(1), 100–113.

Kuhn, P. and Mansour, H. [2014], 'Is Internet Job Search Still Ineffective?', *Economic Journal* **124**(581), 1213–1233.

Kuhn, P. and Skuterud, M. [2004], 'Internet Job Search and Unemployment Durations', *American Economic Review* **94**(1), 218–232.

Langendorf, B. [2003], *Buch und Buchhandel in Zahlen 2003*, Börsenverein des Deutschen Buchhandels e.V., Frankfurt am Main.

Lundvall, B.-A. [1992], *National Systems of Innovation: Towards a Theory of Innovation and Interactive Learning*, Pinter, London.

Moon, H. R. and Weidner, M. [2015], 'Linear Regression for Panel with Unknown Number of Factors as Interactive Fixed Effects', *Econometrica* **83**(4), 1543–1579.

Nelson, R. R. and Rosenberg, N. [1993], Technical Innovation and National Systems, *in* 'National Innovation Systems: A Comparative Analysis', Oxford University Press, New York, pp. 3–28.

Neumark, D. and Simpson, H. [2015], Place-Based Policies, *in* G. Duranton, V. Henderson and W. Strange, eds, 'Handbook of Regional and Urban Economics', Vol. 5, Elsevier, Amsterdam and Oxford, pp. 1197–1287.

Neumark, D., Zhang, J. and Ciccarella, S. [2008], 'The Effects of Wal-Mart on Local Labor Markets', *Journal of Urban Economics* **63**(2), 405–430.

Niebel, T., Briglauer, W., Kauf, B., Hüschelrath, K. and Bertschek, I. [2016], 'The Economic Impacts of Broadband Internet: A Survey', *Review of Network Economics* **14**(4), 201–227.

OECD [2002], *Frascati Manual: Proposed Standard Practice for Surveys of Research and Experimental Development*, Organisation for Economic Co-operation and Development (OECD), Paris.
URL: *https://www.oecd-ilibrary.org/docserver/9789264199040-en.pdf?expires=1552935374&id=id&accname=ocid56012597&checksum=D178A337CFB3C7E98CDF7A3D296B4CA9, last accessed: 18 March 2019.*

Oster, E. [2018], 'Unobservable Selection and Coefficient Stability: Theory and Evidence', *Journal of Business and Economic Statistics, forthcoming* .

Prim, R. C. [1957], 'Shortest Connection Networks And Some Generalizations', *Bell System Technical Journal* **36**(6), 1389–1401.

Richter, F. [2014], Deutschland hängt (noch) am gedruckten Buch, *in* 'Digitales Bild. Statista Infografiken'.
URL: *https://de.statista.com/infografik/2859/umsatz-mit-buechern-und-e-books/, last accessed: 10 October 2018.*

Roback, J. [1982], 'Wages, Rents, and the Quality of Life', *Journal of Political Economy* **90**(6), 1257–1278.

Romer, P. M. [1990], 'Endogenous Technological Change', *Journal of Political Economy* **98**(5), 71–102.

Rosen, S. [1979], Wage-Based Indexes of Urban Quality of Life, *in* P. Mieszkowski and M. Straszheim, eds, 'Current Issues in Urban Economics', Johns Hopkins University Press, Baltimore, pp. 74–104.

Rosés, J. R. and Wolf, N. [2018], 'Regional Economic Development in Europe, 1900-2010: A Description of the Patterns', *Economic History Working Papers* (278).

Schmucker, A., Seth, S., Ludsteck, J., Eberle, J. and Ganzer, A. [2016], 'Establishment History Panel 1975 – 2014', *FDZ-Datenreport* **3**.

Schnabel, P. [2015], *Kommunikationstechnik-Fibel*, 4 edn, Self-published, Ludwigsburg.

Schrape, J.-F. [2011], 'Der Wandel des Buchhandels durch Digitalisierung und Internet', *SOI Discussion Paper* (2011-01).

Spence, M. [1984], 'Cost Reduction, Competition, and Industry Performance', *Econometrica* **52**(1), 101–122.

Stock, J. H. and Yogo, M. [2005], Testing for Weak Instruments in Linear IV Regression, *in* D. W. K. Andrews and J. H. Stock, eds, 'Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg', Cambridge University Press, Cambridge, Massachusetts, pp. 80–108.

TÜV Rheinland Consulting GmbH [2013], *Szenarien und Kosten für eine kosteneffiziente flächendeckende Versorgung der bislang noch nicht mit mindestens 50 Mbit/s versorgten Regionen*, TÜV Rheinland Consulting GmbH, Berlin.
URL: *https://www.bmwi.de/Redaktion/DE/Publikationen/Studien/kostenstudie-zum-breitbandausbau.pdf?__blob=publicationFile&v=5, last accessed: 18 March 2019.*

TÜV Rheinland Consulting GmbH [2016], Bericht zum Breitbandatlas Mitte 2016 im Auftrag des Bundesministeriums für Verkehr und digitale Infrastruktur (BMVI). Teil 1: Ergebnisse, *in* 'Bericht zum Breitbandatlas', TÜV Rheinland Consulting GmbH, Berlin.
URL: *https://www.bmvi.de/SharedDocs/DE/Anlage/Digitales/bericht-zum-breitbandatlas-mitte-2016-ergebnisse.pdf?__blob=publicationFile, last accessed: 18 March 2019.*

TÜV Rheinland Consulting GmbH [2018], Bericht zum Breitbandatlas Mitte 2018 im Auftrag des Bundesministeriums für Verkehr und digitale Infrastruktur (BMVI). Teil 2: Methode, *in* 'Bericht zum Breitbandatlas', TÜV Rheinland Consulting GmbH, Berlin.
URL: *https://www.bmvi.de/SharedDocs/DE/Anlage/Digitales/bericht-zum-breitbandatlas-mitte-2018-methode.pdf?__blob=publicationFile, last accessed: 18 March 2019.*

United States Government Accountability Office [2012], 'The Distribution of Federal Economic Development Grants to Communities with High Rates of Poverty and Unemployment', *Report* (GAO-04-306), 1–35.
URL: *https://www.gao.gov/assets/650/648367.pdf*

von Ehrlich, M. and Seidel, T. [2018], 'The Persistent Effects of Place-Based Policy: Evidence from the West-German Zonenrandgebiet', *American Economic Journal: Economic Policy* **10**(4), 344–374.

What Works Centre for Local Economic Growth [2015*a*], *Broadband*, Vol. 6, London School of Economics and Political Science, Centre for Cities, and Arup, London.

What Works Centre for Local Economic Growth [2015*b*], *Innovation: Grants, Loans and Subsidies*, Vol. 9, London School of Economics and Political Science, Centre for Cities, and Arup, London.

What Works Centre for Local Economic Growth [2015*c*], *Transport*, Vol. 7, London School of Economics and Political Science, Centre for Cities, and Arup, London.

What Works Centre for Local Economic Growth [2016], *Area Based Initiatives*, Vol. 10, London School of Economics and Political Science, Centre for Cities, and Arup, London.

Whitacre, B. E., Alam, M. R. and Lobo, B. J. [2018], 'Econometric Error Nullifies Finding of the Impact of Broadband Speed on County-level Employment', *Information Economics and Policy* **44**, 58–60.

Whitacre, B., Gallardo, R. and Strover, S. [2014*a*], 'Broadbands Contribution to Economic Growth in Rural Areas: Moving Towards a Causal Relationship', *Telecommunications Policy* **38**, 1011–1023.

Whitacre, B., Gallardo, R. and Strover, S. [2014*b*], 'Does Rural Broadband Impact Jobs and Income? Evidence from Spatial and First-Differenced Regressions', *Annals of Regional Science* **53**, 649–670.

Wolfers, J. [2006], 'Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results', *American Economic Review* **96**(5), 1802–1820.

Wooldridge, J. M. [2002], *Econometric Analysis of Cross Section and Panel Data*, MIT Press, Cambridge, Massachusetts.

Zentner, A., Smith, M. and Kaya, C. [2013], 'How Video Rental Patterns Change as Consumers Move Online', *Management Science* **59**(11), 2622–2634.

Zhang, C. and Lei, X. [2015], 'Short-Term Jobs and Long-Term Job Loss of Large Supermarket Entry', *China Economic Journal* **8**(1), 74–85.

Zúñiga-Vicente, J. Á., Alonso-Borrego, C., Forcadell, F. J. and Galán, J. I. [2014], 'Assessing the Effect of Public Subsidies on Firm R&D Investment: A Survey', *Journal of Economic Surveys* **28**(1), 36–67.