UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

Essays on Political Economy

Inauguraldissertation zur Erlangung des Doktorgrades
des Fachbereichs Wirtschaftswissenschaften der
Ruprecht-Karls-Universität Heidelberg

vorgelegt von
Matthias Quinckhardt

Heidelberg
Juni 2022

Acknowledgements

This dissertation is the product of many hours of work, fruitful conversations and helpful feedback from many people.

First and foremost, I want to thank my adviser, Christina Gathmann, for her continued support throughout this dissertation project. She gave me great freedom to pursue my research but offered helpful comments and feedback at all points along the way. Furthermore, I thank Axel Dreher for being the second referee on my dissertation committee.

This thesis also profited greatly from discussions with many colleagues who became good friends and created the necessary social environment for a successful PhD. I was very lucky to be immediately welcomed at the chair by Eduard, Ole and Franziska and benefited greatly from their advice throughout all stages of my dissertation. Anna and Felix were wonderful additions to the team and great companions in the final two years, when Covid made all of our work much more challenging.

Two of the papers in this thesis are the product of productive collaborations with coauthors. I am very grateful to Francesco, Christina, Tobias, and Lukas for their cooperation and many hours of work that went into our projects.

From the chair in Heidelberg, I would like to thank Sabine for many engaging and stimulating conversations as well as endless support in navigating the pitfalls of the university administration.

This thesis also benefited from generous support of the German National Academic Foundation. I was very lucky to have been supported by the Foundation since my first semester of undergraduate studies up to the end of my PhD.

Several friends, most importantly Nils, Marie, Anna, Viktoria and Cäcilia, had to endure many rants that go along with a project like this. Thank you for always being there.

I also want to thank my family, my parents - who supported me all these years -, siblings, and especially my nieces and nephews, for taking my mind off this project during many calls and visits.

Lastly, I want to dedicate this thesis to my grandmother. From a very early age on, she taught me to follow my passions and interests. She was a constant source of emotional support and inspiration through our many conversations during this journey.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Geographic differences in economic performance are an omnipresent phenomenon. They range from large cross-country disparities in wealth and income to within-city gaps between the rich and poor parts of town. Figure 1.1 shows the within-country dispersion of unemployment rates at the NUTS-2 level for six European countries over time and provides graphical evidence of these differences. There exist large within-country inequalities. For instance, the range between the highest and lowest unemployment rate among Italian regions was nearly 20 percentage points in 2015 while the country-wide unemployment rate was 11%. In Germany, the dispersion decreased markedly between 2006 and 2018 from around 14 to less than 5 percentage points in 2018. Over the same period, it nearly doubled in Turkey from 10 to 20 percentage points.

Such regional inequality has a multitude of societal and economic consequences. It encourages migration from poorer to wealthy regions (Borjas, 1989), is associated with negative health effects for the population (Case and Deaton, 2020) and leads to higher support for right-wing or populist political movements in poorer areas (Autor et al., 2020). For instance, support among inhabitants of left-behind areas contributed to the victories of Donald Trump in the United States (Autor et al., 2020) or of the Brexit movement in the United Kingdom (Fetzer, 2019). Overall, regional inequality leads to a decrease in societal cohesion within countries.

Figure 1.1: Dispersion of NUTS-2 Level Unemployment Rates in six Countries

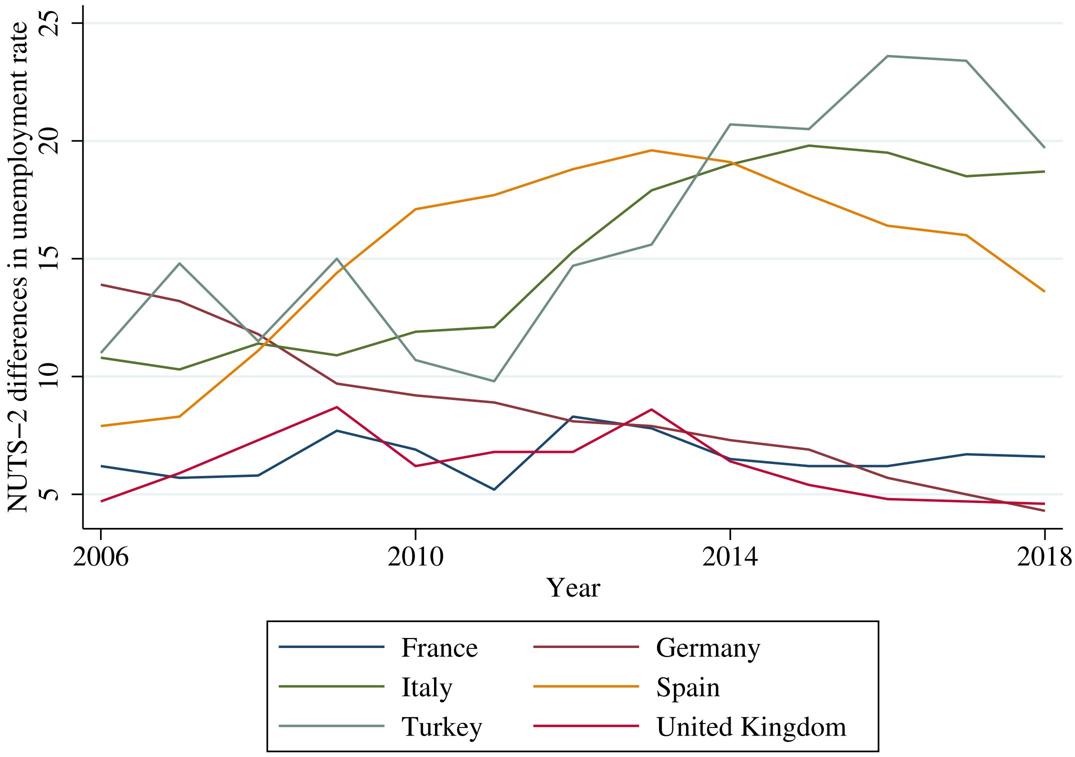

Note: This graph reports the development of the percentage point range in the unemployment rate at the NUTS-2 level for six European countries between 2006 and 2018. The range was calculated as the difference between the highest and lowest unemployment rate of a NUTS-2 region in a given country and year using data from Eurostat (2022).

For these reasons, governments make great efforts to ensure equal living standards and stimulate economic growth in regions that are struggling economically. For instance, the EU plans to spend €392 bn towards regional development funds between 2021 and 2027 (European Commission, 2021). The main tool for German regional policy, the Joint Federal/Länder Task for the Improvement of Regional Economic Structures (GRW), spent €3.9 bn between 2017 and 2021 on regional development alone (BmWK, 2022). These efforts are aimed at providing citizens with equal standards of living and economic opportunities as well as easing societal tensions stemming from regional inequalities. To underline the importance of such policies, for instance the German Basic Law specifically provides the Federal Government with the right to enact legislation "to the extent that the establishment of equivalent living conditions throughout the federal territory or the maintenance of legal or economic unity renders federal regulation necessary in the national interest" (Federal Republic of Germany, 1949).

In addition to policymakers, economists have long studied the origins of regional inequality and the effectiveness of place-based policies to counter it. In the economic growth literature, several factors have been identified which may explain differences in economic activity between locations. Among these are time-invariant characteristics, such as geographic location, or temporary shocks to the population, such as natural disasters or wars (for an overview, see for instance Desmet and Henderson, 2015). At the same time, the effectiveness of place-based policies, such as grants for private and public investments or the provision of specific public infrastructure, to address regional inequality has received increased attention in the literature as well (for an overview, see for instance Neumark and Simpson, 2015 or Südekum, 2021).

The focus of this thesis is to make a contribution to the economic literature on regional differences within countries. It includes three essays, which are all related to either the debate over the origins of regional inequality or the effectiveness of place-based policies. The second chapter studies the political determinants of intergovernmental transfers from states to municipalities in Germany. It investigates the so-called 'alignment bias', i.e., distortions in the distribution of public funds which favor regions which are politically aligned with the central government. Therefore, it relates both to the debate on the origins of regional inequality and the effectiveness of place-based policies. On the one hand, a large alignment bias, favoring one region at the expense of others, can lead to regional inequality. On the other hand, targeted intergovernmental transfers can be viewed as a place-based policy aimed at decreasing regional inequality. An alignment bias in the distribution of these transfers may imply that the policy is inefficient, if political rather than economic considerations determine the allocation of the funds. The chapter establishes that even in a developed country with strong democratic institutions, such as Germany, there exists a sizable alignment bias which is comparable in size to the effects found in some developing countries.

The third chapter of this thesis, which is joint work with Francesco Berlingieri and Christina Gathmann, analyzes the local labor market effects of opening a college. Providing access to higher education by opening a college or university is a popular place-based policy with the goal of creating economic growth through a more educated local labor force and spurring innovation which benefits local firms. We evaluate whether such a policy has a lasting positive effect on the local economy. We find that German counties which saw a college opening exhibit a large and persistent increase in the student population, but that the labor market effects are limited. Wages remain stable and apart from an increase in the number of young workers with a college degree, the composition of the labor force remains unchanged. Furthermore, this increase in young high-skilled workers stems mostly from incumbent firms, suggesting that a college does not lead to large changes in the industry structure. Therefore, the chapter provides evidence that the effectiveness of college openings as a place-based policy to stimulate the local economy is limited.

The fourth and final chapter of this thesis, which is joint work with Tobias Korn and Lukas Wellner, examines how a large and negative shock to the population affects the local economy. We study a historical setting, namely Germany before and after World War I, and exploit regional variation between counties in the number of soldiers which were wounded or killed in action. We show that regions which were more affected by this casualty shock had a worse economic performance after the war, such as lower wages and tax revenues, than less affected counties. Hence, the casualty shock may be responsible for an increase in regional inequality. The results further suggest that such differences have negative societal and political consequences. We find that counties with a higher casualty rate also had higher vote shares for extreme right wing parties and the Nazi party in the interwar period.

In addition to contributing to existing economic research on regional inequality and place-based policies, each of the papers in this thesis also has a high relevance for economic policy. The first chapter shows that when distributing intergovernmental transfers, policymakers need to account for political considerations which may bias the allocation of funds. Further results in this paper show that politicians benefit from better information about transfers if they are a member of a political party. Hence, policymakers ideally also need to address information asymmetries between local and state governments about the availability of intergovernmental grants. The second chapter shows that a policy of opening colleges in itself will not necessarily decrease regional inequality. Instead, the results suggest that such a policy is only successful if there already exists an industrial base which can absorb the increase in high-skilled labor. The fourth chapter provides insights into the economic effects of a negative population shock. While the specific setting is historical, the findings may still translate into more recent settings, for instance the economic effects of outmigration in developing countries or the economic effects of pandemics or natural disasters.

In the following, I will briefly introduce each of the three chapters individually.

1.1 The Value of a Party: Local Politics and the Allocation of Intergovernmental Transfers

The second chapter investigates the effect of party connections on fiscal policies. Specifically, I investigate whether party alignment between mayors and state governments in Germany influences the allocation of investment grants - an important tool for regional policy - from states to municipalities. In developing countries, such an alignment bias has been frequently documented (see, e.g., Brollo and Nannicini, 2012; Hodler and Raschky, 2014; Asher and Novosad, 2017). For developed countries, the evidence is more mixed (see, e.g., Tavits, 2009; Albouy, 2013; Baskaran and Hessami, 2017), with only some studies finding an effect of partisan alignment on the distribution of public funds. In the literature, such favoritism has been interpreted as a sign of partisanship and blamed on central governments which use public funds for partisan purposes.

The empirical analysis employs newly-collected data for mayoral elections results for eight German states for the period from 2000 to 2008. I combine this data with a detailed account-level data set on municipal finances from the German State Statistical Offices. In order to address the endogeneity of party alignment, I employ a regression discontinuity design for close local elections, comparing municipalities which closely elected an aligned mayor, i.e., a mayor who is member of one of the parties in the state government, with those that elected an unaligned mayor.

I find that party alignment between a local mayor and the state government increases transfers to a municipality by about 20% per year. This effect is more pronounced in the years leading up to a local election. Furthermore, the effect disappears when the incumbent mayor does not run for reelection, suggesting that mayors' own electoral incentives matter for the alignment bias. In a further analysis of the mechanisms behind the alignment bias, I show that independent mayors who are not members of any political party obtain less transfers from the state government than mayors who are members of one of the opposition parties in the state legislature. I interpret this as a sign that party connections and better access to information increase the amount of transfers to a municipality. This acts as an additional mechanism to partisanship, where the central government systematically favors municipalities which elected an aligned politician.

Overall, the chapter shows that German municipalities with a mayor who is politically aligned with the state government receive significantly higher transfers. Hence, the results expose some inefficiencies in the allocation of transfers which are meant to decrease regional inequalities. Furthermore, local political conditions, such as reelection incentives and information transmission, provide an additional explanation for the alignment effect, as opposed to partisanship.

1.2 College Openings and Local Economic Development

In the third chapter we examine the labor market effects of college openings in Germany. We exploit 20 college openings in the period between 1988 and 1998 which occurred predominantly in counties that did not have colleges before. The openings were meant to increase access to education and stimulate local economic growth. This chapter evaluates whether this policy was successful and adds to a larger literature which examines the effects of proximity to institutions of higher learning (Beeson and Montgomery, 1993; Carneiro et al., forthcoming; Fuest and Immel, 2021). The chapter also contributes to ongoing debates on the effects of positive shocks to local labor supply, for instance through migration (Card, 2001; Borjas, 2003; Glitz, 2012). In our setting, the supply of young high-skilled labor exogenously increases due to the college opening. Therefore, we investigate the effects of the opening on the wages and employment of workers in different skill and age groups.

We use administrative social security data which covers the vast majority of German workers and contains detailed information on wages and employment by skill, age and industry. In order to account for the endogeneity of a college opening, we use Mahalanobis matching on rich pre-opening covariates to find suitable control regions for each region with a college opening. We then estimate an event study for a variety of different outcomes.

A new college substantially increased the student population and share of high-skilled workers in a region. Yet, we find no effect on regional wages or employment indicating that the local economies did not experience additional growth through skill-biased technological change, for instance. Instead, there is sizable heterogeneity in the local gains: high-tech firms in manufacturing absorb most of the new college graduates, especially in engineering professions. We find little impact on the low- or high-skilled service sector or employment in managerial professions. Finally, we show that local labor market conditions prior to the opening matter: in regions with a more dynamic labor market, the opening encourages firm creation and a permanent upskilling of the workforce. Areas with a less dynamic labor market experience little sustained growth in high-skilled workers who are absorbed by incumbent firms.

Our findings strongly suggest that college openings had only a limited positive effect on the local economy. Most importantly, incumbent and high-tech firms benefited most from the opening. Hence, the positive effects of the opening accrued only to a small number of local firms. The colleges did not improve labor market outcomes across all industries. Taken together, the findings call into question the effectiveness of college openings as a place-based policy to improve local market conditions.

1.3 The Local Consequences of War: Evidence from Germany after World War I

The final chapter of this thesis examines the economic effects of German casualties in World War I. During the war an estimated 15% percent of the male working-age population perished. At the same time, the capital stock of the German economy was largely unaffected since the war was mostly fought outside of German borders. Hence, the casualties can be viewed as a large negative shock to the local population. Such a shock can affect the local economy in two ways: First, it can act as a negative shock to local labor supply, thereby increasing the value of labor and inducing higher investments in labor-saving technologies. Second, the shock can depress local demand for goods and services due to smaller consumer base. This would lead to a decrease in firm profits, lower output and wages. We empirically test which channel dominates in our setting. Therefore, this chapter relates to a larger literature which estimates the effect of such exogenous shocks to the local population, for instance pandemics (Brainerd and Siegler, 2003; Voigtländer and Voth, 2013a; Karlsson et al., 2014), natural disasters (Vigdor, 2008; Boustan et al., 2012) or wars (Voigtländer and Voth, 2013b).

In order to measure regional exposure to the casualty shock, we geocoded the birth places of 8.5 million wounded and killed German soldiers during World War I from official casualty records. We linked these casualties to different county-level statistics from before and after the war. Our main results are based on difference-in-differences estimations for newly-digitized wage data across genders and age groups in order to establish whether the casualty shock had a negative effect on the local labor supply. We further use data on post-war tax revenues to measure local economic activity as well as detailed agricultural census data to examine the industry-specific effects of the casualty shock.

The findings show clearly that the casualty shock had a negative impact on the local economy. Wage growth is lower in counties with a higher casualty rate and more exposed counties encountered significant out-migration, predominantly of women. Overall economic activity is also lower, evidenced by a decrease in corporate and payroll tax revenues. Furthermore, we find that counties with a higher casualty rate economically transformed away from industrial production and towards medium-scale farming activities. Lastly, the casualty shock also had negative political consequences: voters in more affected counties were more likely to support radical right-wing parties in interwar elections.

This chapter provides evidence that negative shocks to the local population can have significant negative economic effects and thereby increase regional inequality. In the setting studied, the population shock leads to a large decrease in local demand which negatively affects wages, especially for women. The population shock also did not have a positive effect on the use of technology, suggesting that the decrease in the labor supply was negligible and firms did not substitute for labor through technology. The results in this chapter therefore run counter to other studies which have argued that negative population shocks have a positive on economic development due to favorable shifts in the labor-to-capital ratio and increased technology usage (Young, 2005; Voigtländer and Voth, 2013b).

Chapter 2

The Value of a Party: Local Politics and the Allocation of Intergovernmental Transfers[1]

In federal systems, central governments frequently channel public funds to politically aligned local governments or politicians. In the theory of distributive politics by Grossman (1994), a central government aims to buy political capital from local politicians in exchange for transfers. His model predicts higher transfers to regions which are aligned politically with the central government, as it places higher value on the political capital of aligned politicians. This prediction has been empirically tested in a range of different environments, for instance Brollo and Nannicini (2012) for Brazil, Curto-Grau et al. (2018) for Spain or Asher and Novosad (2017) for India.

In this literature, responsibility for party favoritism is assigned to the central government which 'punishes' local politicians of the opposing party and 'rewards' those of their own party respectively by denying or providing them public funds. For instance, Bracco et al. (2015, p. 89) find that "the central government will tend to divert resources toward aligned jurisdictions for electoral purposes." Consequently,

[1]I thank Christina Gathmann for invaluable advice during this project. Furthermore, I am grateful to Manuel Bagues, Eduard Brüll, Anne Burton, Kai Gehring, Ole Monscheuer, James M. Snyder, participants of the European Public Choice Society Meeting 2019, the ZEW Public Finance Conference 2019, 1st Genoa Summer School on Political Economy, the Ammersee Workshop 2019, the ZEW Summer Workshop 2019 in Public Economics, the Verein für Socialpolitik Annual Meeting 2019 and seminar participants at Heidelberg University for valuable comments and suggestions. Maximilian Brien, Felix Grimm, Jongoh Kim, Nora Mac and Davis Schweinberg provided excellent research assistance.

an alignment effect is viewed as a sign of partisanship in the political system and a distortion in the allocation of public funds.

However, there exist potential alternative explanations. First, local political conditions may explain the alignment effect more successfully than assigning blame to the central government. This implies that the central government does not distribute funds based on a systematic electoral strategy but instead local politicians react to their own electoral incentives. Second, politicians' connectedness and access to information are potential explanations for the alignment effect, aside from partisanship. For instance, local politicians of the governing party may be more successful at obtaining transfers from the state government because they are better informed by fellow party members about funding opportunities.

In this paper, I investigate party favoritism in a setting where it was previously believed to play only a minor role in the allocation of public funds, namely Germany. The expectation that the German political system leaves little room for party favoritism builds on the observation that institutional factors, such as the electoral system and the role of political parties, can limit politicians' ability to use public funds for partisan purposes (for an overview see, e.g., Golden and Min, 2013). Hence, party favoritism is viewed as less common in settings with a proportional representation system, strong national parties and transparent institutions.[2] In line with this assessment, Baskaran and Hessami (2017) find no effect of alignment between municipal councils and state governments on the allocation of intergovernmental transfers in Germany.

Using a new, and largely hand-collected data set for mayoral elections and municipal finances in eight of Germany's sixteen states, I analyze the effect of party alignment between state governments and mayors on the allocation of state infrastructure transfers, which constitute an important source of income for municipalities.

[2]For Germany, the mixed electoral system sets incentives for parties to conduct national campaigns, as opposed to ones which favor only very specific groups of voters or regions. Nonetheless, Stratmann and Baur (2002) find some evidence that directly-elected German MPs become members of parliamentary committees which may dispense pork-barrel funds. However, they do not conclude that this behavior actually distorts the allocation of public funds to favor individual regions.

I employ regression discontinuity methods to estimate the effect of alignment in municipalities with close elections for mayor to address the endogeneity of alignment.

Contrary to the expectation, I document a sizable, positive effect of alignment on the amount of transfers allocated to a municipality. Splitting the sample based on the state and local electoral cycle shows that the alignment effect is most prominent in the years leading up to local elections, whereas it is unaffected by the state electoral cycle. Furthermore, the alignment effect disappears when the local incumbent is not running for reelection.

In a last step, I split the sample of unaligned mayors into those who are members of an opposition party and those who are unaffiliated with any political party. The results show that opposition party mayors receive more funding than unaffiliated ones. Furthermore, a large part of the initial alignment effect is driven by differences between municipalities with aligned and unaffiliated mayors.

In addition to the literature on alignment effects, this paper also relates to the empirical literature which studies the effects of electoral incentives, such as term limits or electoral accountability, on politicians' behavior (for instance Besley and Case 1995; Ferraz and Finan 2011; Curto-Grau et al. 2018). In my setting, aligned mayors appear to enjoy a larger incumbency bonus than unaligned ones. Furthermore, the alignment effect disappears in those municipalities where mayors are not seeking reelection, suggesting that aligned mayors do not take advantage of their connections to the state government if they cannot personally derive electoral benefit from it.

The role of politicians' connectedness has been shown to influence the allocation of public resources as well. For instance, there exists ample empirical evidence that politicians direct public funds to their home regions (Hodler and Raschky 2014; Gehring and Schneider 2018) or to their own ethnic groups (Burgess et al., 2015). Typically, these results are interpreted as signs of partisanship and favoritism. However, information transmission about local demand or central government policy can also be influenced by politicians' networks. Azulai (2017) finds evidence for this in a study of connectedness between local politicians and federal ministers in Brazil. I exploit a unique feature of the German setting, the existence of independent mayors who are not aligned with any political party, and show that they obtain less transfers from state governments than mayors who are members of one of the opposition parties in the state government. Independent mayors are unlikely to be part of partisan conflicts between political parties but at the same time cannot rely on a network of fellow party members for information. Hence, their disadvantage in obtaining transfers underlines the importance of information transmission as opposed to partisanship.

Lastly, a large literature building on Nordhaus (1975) examines the role of the electoral cycle on fiscal policies. Most studies find strong signs of a political business cycle in developing countries, with more mixed evidence for developed ones (Akhmedov and Zhuravskaya 2004; Brender and Drazen 2005; Shi and Svensson 2006). For Germany, Schneider (2010) and Engelmaier and Stowasser (2017) find evidence of electoral cycle effects in state budgets and public bank lending. My findings underline that, even in a developed country such as Germany, politicians adjust fiscal policies close to elections. In my case, state transfers to municipalities increase in the run-up to local elections in those municipalities where the incumbent mayor is aligned. However, the state electoral cycle does not affect the allocation of transfers to aligned municipalities.

The rest of this paper is organized as follows. Section 2.1 describes the institutional background, Section 2.2 the data set and identification strategy. The results are presented in Section 2.3. Section 2.4 shows results for the mechanisms of the alignment effect and Section 2.5 offers a conclusion.

2.1 Institutional Background

2.1.1 Municipal Finances and Intergovernmental Transfers

The German Basic Law divides competencies between the federal, state and local governments. As such, municipalities are responsible for the majority of local infrastructure, such as school buildings, sewage systems, hospitals, cultural institutions or minor roads. However, their ability to raise revenues is confined to adjustments of a business tax levied on local companies and local property taxes, both of which are subject to tax competition with neighboring municipalities. Hence, taxes make up only around half of municipal revenues, with the rest consisting mainly of transfers from the respective state governments.[3]

These transfers can be classified into three groups. First, all states run fiscal equalization schemes which redistribute tax revenues across municipalities.[4] Rule-based transfers make up the largest share of transfers received by municipalities and nearly all of them receive some amount in a given year. Second, municipalities can receive transfers to their general budget, awarded at the discretion of the state government.[5] However, such transfers are rare and only a low share of municipalities receive any type of discretionary budget transfers in a year. Third, state governments also provide project-specific infrastructure transfers ('Investitionszuweisungen') for a host of different purposes, such as the modernization and construction of public buildings or investments in road infrastructure. These transfers require an interested municipality to write a grant proposal to the state government and cover a certain percentage of expenses with their own funds. The state laws regarding infrastructure transfers only contain vague criteria for the evaluation of grant applications and the decision making process.[6] Hence, state governments have discretion over the allocation and generosity of infrastructure transfers, the size of which can be substantial for individual municipalities. They average 80 euros per capita per year and account for roughly a third of total transfers per year. As municipalities face fiscal constraints and carry responsibility for large parts of public infrastructure, they are very interested in obtaining investment transfers. In fact, nearly all municipalities apply for and receive at least some amount of infrastructure transfers in a year.

[3]Due to legal constraints in the Basic Law, even federal transfers to municipalities have to be allocated by the states. Also, fines and other types of income are of minor importance for municipalities.

[4]These rule-based transfers ('Schlüsselzuweisungen') go into the general budget of municipalities and are allocated according to pre-defined formulas which rarely change and are hard to manipulate to favor individual municipalities. Typical variables included in this formula are for instance population, population density, the local tax base and employment.

[5]These can be in response to some extraordinary circumstances affecting a particular municipality, such as a natural disaster or unexpected budget shortfalls, but there exist no deterministic rules governing the allocation of these transfers.

[6]For instance, sometimes the decisions are made by individual career public servants, in other cases by committees which include elected officials or political appointees such as the respective state ministers or their deputies.

These infrastructure transfers provide an ideal setting to study whether political factors shape their allocation for several reasons. First, nearly all municipalities receive some amount of infrastructure transfers in a year. Hence, there is no obvious sample selection as to which municipalities receive transfers. Second, state politicians can easily influence the allocation of transfers to favor politically aligned municipalities, as funding decisions are made on a case-by-case basis and the criteria for providing transfers are not transparent. Third, local politicians will likely benefit electorally from obtaining transfers, as they enable municipalities to undertake infrastructure investments, which are salient to local voters.

2.1.2 Mayors and Local Politics

Local politics in Germany are dominated by two institutions, the mayor and the city council. It has been argued that mayors have a comparatively weak position in the German political system (e.g., Baskaran and Hessami 2017), as their main role lies in leading the city administrations.[7] Most legislative competencies are assigned to the city council in which the party of the mayor may not necessarily hold a majority of seats and where formal coalitions between parties are uncommon, except in larger cities. However, mayors are usually professional, full-time politicians, whereas city council positions are part-time and filled by citizens with fairly different professional backgrounds. Hence, mayors enjoy high visibility among the local population and can make use of this position in achieving their policy goals, even in the absence of a majority in the city council (Gehne, 2012). Most importantly for my setting, mayors are responsible for the formulation of transfer applications to the state government, giving them a potent tool for shaping policy by controlling the access to potential sources of intergovernmental transfers. Hence, the focus of this paper lies on the role of mayors and parties in attracting infrastructure transfers to their municipality.

In contrast to the leaders of the state and federal governments, who are voted into office by the respective parliaments, German mayors are directly elected by the

[7] The role of mayors varies to some degree by state. In the southern German states of Baden-Württemberg and Bavaria, the mayor is viewed to have strongest position versus the city council in comparison to other German states. However, these two states are not part of my data set.

local population.[8] Mayors are typically elected by simple majority with a runoff between the two candidates with the most votes, in the case that no candidate captures a majority in the first round. Terms of office vary between four and eight years depending on the state and elections do not typically coincide with those at the state and federal level, which isolates local politics from overall political trends (Holtmann et al., 2017). Furthermore, mayors are very likely to be reelected, if they decide to stand for office again. The salary of a mayor is high when compared to that of other elected officials in Germany, and provides good pension benefits based on the time spent in office.[9]

Even though all citizens can run for mayor, elected officials typically have a background in either law or public administration. This may be seen as evidence that voters prefer professional managers over partisan politicians as the head of their city. Gehne (2012) argues that this may be one of the reasons why independent mayors who are unaffiliated with a political party enjoy a certain amount of success in local elections. Furthermore, he points out that voters tend to assign more importance to the personality of candidates when voting for local than for state or federal offices. However, about two thirds of mayors are members of and nominated by a political party, suggesting that party politics also plays an important role in their election.

2.1.3 Parties

Parties play a dominant role in the German political system by shaping public opinion, their ability to elect members of parliament through lists, and selecting candidates. The two dominant forces in German politics – at least for the time covered in this study – are the right-of-center Christian Democratic Union (CDU) and the left-of-center Social Democratic Party of Germany (SPD). Furthermore,

[8] Most German states introduced the direct election of mayors in the 1990s. The states of Schleswig-Holstein, Bavaria and Baden-Württemberg had direct elections of mayors since 1949. However, they are not part of the data set. The switch from an indirect election of mayors through the city council to a direct one by the citizenry aimed to increase voter engagement and accountability of local politicians.

[9] For instance a mayor of a municipality with 5,000 inhabitants earns approximately 7,400 Euros per month in North Rhine-Westphalia.

there exist several smaller parties which typically become coalition partners of one of the major parties to form governments at both the state and federal level.[10]

Despite differences in the political positions and core constituencies, party organizations are similar and mirror the administrative levels in Germany: local chapters in cities and municipalities form the basis of each party. From the ranks of their members, the local party conventions nominate candidates for municipal office, as well as for direct elections in state and federal parliament districts. Personnel decisions for party leadership are made by state and federal congresses, to which local party chapters send delegates. In order to be nominated for higher office by a party congress, successful candidates have typically worked their way up through the local party chapter. Hence, there exist deep ties between party officials at the local and state levels and they have motivation to aid their fellow party members whom they rely on for advancement in their own careers.

Candidates not aligned with one of the major parties play only a limited role in German politics, especially at the state or federal level. However in local elections, independent candidates or local voter initiatives enjoy more success, but are unlikely to have access to the same kind of networks as career party officials.

Taken together, the institutional setting offers an interesting laboratory to examine the role of party favoritism on the allocation of intergovernmental transfers. The peculiar position of the mayor in the German political system as a directly elected official, and the comparatively nontransparent procedures governing the allocation of infrastructure transfers, suggest that there exists at least some potential for pork-barrel politics.

[10]For the states and the time period observed in this study, all state governments consisted of either the SPD or CDU in a coalition with one or two of the smaller parties.

2.2 Data and Identification Strategy

2.2.1 Data Set

The data set encompasses the universe of all direct elections for full-time mayors in eight German states.[11] Since these data are not centrally collected in Germany, I obtained information from various sources, for instance from the respective State Statistical Offices, local newspaper archives and contacting individual municipalities. I observe the name, party membership and vote share of the two candidates who received the most votes in an election. A detailed description of the data sources and construction of the data set for the individual states is provided in Appendix 2.A.

The key variable for the empirical analysis of party effects is the margin of victory of the aligned candidate in a mayoral election. A candidate is aligned if she belongs to one of the parties in the state government. Hence, the alignment status of a municipality can change either with the election of a new mayor or the formation of a new state government. As the electoral cycles at the state and local levels differ, alignment can change during a mayor's term. To each municipality-year pair, I assign the margin of victory of the aligned candidate in the last local election.[12] Consider the following example: The city of Cologne in the state of North-Rhine Westphalia held a mayoral election in the year 2000 for a term of eight years. In the runoff between the CDU and SPD candidate, the CDU candidates narrowly prevailed with a margin of 5 percentage points. At the time, the state government consisted of a coalition of the SPD and Green party. Consequently, the margin of

[11]These are Lower Saxony, North-Rhine Westphalia, Hesse, Rhineland-Palatine and the Saarland in West Germany and the East German states of Saxony, Saxony-Anhalt and Thuringia. The city states of Bremen, Hamburg and Berlin were excluded due to their special status. Furthermore, no information on mayoral election results were found for Schleswig-Holstein, Baden-Württemberg, Mecklenburg-Vorpommern and Brandenburg. Another special case is Bavaria, where the Christian Social Union (CSU), a local party aligned with the CDU, has taken a dominant role and controlled the governor's office since 1957. This dominance is also evident at the local level where around 75% of mayors are either members of the CSU or independents who govern with CSU approval. The lack of variation in the parties running the state or local governments makes Bavaria infeasible to use for a statistical analysis.

[12]In election years, the alignment status of a municipality is assigned based on the month in which a local or state election occurred. I exclude elections in which only one candidate competes, a total of 1,588 observations. For elections where two unaligned or two aligned candidates compete, the margin of victory cannot be calculated. However, I run additional analyses with some of these cases later.

victory takes the value -5, as the aligned candidate for mayor lost by five percentage points. This value is assigned to all observations for Cologne until 2005. That year, the state government led by SPD and Greens was succeeded by a government of CDU and Liberal Democrats. Hence, the margin of victory changes to 5 for the rest of the mayor's term, as the CDU mayor is now aligned with the state government. The margin of victory changes again with the next mayoral election in 2009.

For further analyses, I construct several other variables on local politicians' characteristics. First, I generate a dummy variable indicating whether an incumbent mayor is running for reelection in the next local election in order to assess the electoral incentives of the mayor. Second, I use the margin of victory of independent candidates over candidates whose party has seats in the state legislature in a given year, regardless of whether the party is part of the state government or not. The intent is to capture whether the mayor is independent of the political establishment in a state and consequently has fewer contacts to the state government.

The electoral data is merged with a detailed administrative data set on municipal finances from the joint research center of the German State Statistical Offices. The data set contains information on the amount of any type of transfer a municipality receives in a year. Furthermore, I am able to distinguish for what kind of infrastructure transfers were earmarked for, for instance investments in education or traffic infrastructure. In addition to transfers, I also have information on municipal tax revenues and expenditures. I further add the population, area and number of full-time employees at the municipal level as controls for local economic conditions.

The final data set is an unbalanced panel and covers the years 2000 to 2008 with a total of 17,669 municipality-year pairs for 2,322 municipalities.[13] Of all pairs, 10,917 are assigned to elections between an aligned candidate and an unaligned one and 6,752 to elections where either two aligned or unaligned candidates compete.

[13]There are two reasons for unbalancedness: First, some municipalities only begin direct elections in the early 2000s, as the direct elections of mayors was staggered within states. Second, municipalities may be dissolved or amalgamated. This applies especially to the three East German states in the sample.

The data are constrained by two factors, the introduction of direct elections for mayors in the 1990s and reforms in the accounting rules for municipalities beginning in the late 2000s.[14] As both of these changes occurred in different years across states, I restrict the analysis to those years in which all states had already implemented the direct election of mayors and not yet reformed municipal finances. However, in a robustness check I also use all available years which yields similar results.

Table 2.1 presents descriptive statistics for the data set used in the baseline estimation divided between aligned and unaligned observations.

Table 2.1: Descriptive Statistics

	Aligned	*Unaligned*	*P-Value*
Population	28,848	27,899	0.44
	(66,420)	(58,101)	
Population density	408.12	391.56	0.06
	(465.58)	(429.37)	
Employment share	0.3274	0.3237	0.00
	(0.0304)	(0.0298)	
Tax revenue	327.16	321.51	0.37
	(315.45)	(260.56)	
Infrastructure transfers	87.16	79.53	0.00
	(107.21)	(119.03)	
Rule-based transfers	193.38	191.92	0.57
	(120.56)	(110.61)	
East	0.23	0.14	0.00
	(0.42)	(0.35)	
Observations	4,555	5,738	

Note: Mean, standard deviation in parentheses, and p-value for a two-sided t-test between aligned and unaligned municipalities. The employment variable is calculated by dividing the number of individuals in full-time employment by population. Tax revenue and transfers are in per capita terms. East is a dummy indicating whether a municipality is in East Germany.

As is evident, there exist several differences between aligned and unaligned municipalities, especially in terms of population density and local employment.

[14] The reform in accounting standards makes municipal finance data not comparable before and after the reform in individual states and also makes comparisons across states impossible, as each state introduced different accounting standards.

Furthermore, more East German municipalities appear to have an aligned mayor than an unaligned one. Most importantly, the unconditional comparison of aligned and unaligned municipalities reveals significant differences in the amount of infrastructure transfers received, with aligned municipalities receiving slightly more funding per capita than unaligned ones on average. However, this result cannot be interpreted as a sign of party favoritism, as differences in other characteristics may be responsible for higher infrastructure transfers to aligned municipalities. In the next chapter, I will briefly elaborate on the empirical challenges of estimating the causal effect of alignment.

2.2.2 Identification Strategy

There are two empirical challenges when estimating the causal effect of alignment on the allocation of transfers: First, unobserved socio-economic factors may influence both the amount of transfers going to and the alignment status of a municipality. For instance, rural or urban areas may receive more transfers. This challenge can be overcome by adding relevant controls, however some of these socio-economic factors may be unobservable. Second, the alignment status of a municipality is likely endogenous, even when other characteristics are controlled for. Voters may be more likely to vote for aligned mayors who have been successful previously at obtaining transfers. Hence, the current alignment status of a municipality may depend on its past alignment status as well on the past allocation of transfers.

In order to address both of these issues, I follow the identification strategy proposed by Brollo and Nannicini (2012) and use a regression discontinuity design to compare municipalities with close elections for mayors between aligned and unaligned candidates. Focusing on these municipalities addresses the endogeneity of alignment, as its assignment is quasi-random in a close local election where small swings in vote shares can change the alignment status. If the sample size is large enough, this identification strategy also addresses the first challenge: Municipalities which barely elect an unaligned mayor are unlikely to differ systematically in terms of socio-economic factors from those which barely elect an aligned one. Chapter 2.2.3 provides evidence that the identification strategy is successful in doing so.

The main specification estimates the effect of party alignment between the mayor and state government on the outcome in question in a municipality following a close local election. To identify the influence of electoral incentives and connectedness, I split the sample along several dimensions and rerun the main specification. In the baseline model, I run parametric specifications according to

$$Y_{it} = \alpha + \tau align_{it} + \gamma f(mov_{it}) + (\sigma - \gamma) align_{it} * f(mov_{it}) + \epsilon_{it}, \qquad (2.1)$$

where Y_{it} is the dependent variable in question for municipality i in year t. $align_{it}$ is a dummy indicating whether the mayor of a municipality is a member of one of the parties forming the respective state government in year t. $f(mov_{it})$ is a 4th order polynomial of the margin of victory of the aligned candidate as the running variable. As demonstrated by Gelman and Imbens (2019), a parametric specification may bias my estimates and is sensitive to outliers. To mitigate these concerns, I run nonparametric local linear regressions with optimal bandwidths, a uniform kernel and inference according to Calonico et al. (2014). In the nonparametric specification, I include state and year dummies as covariates following Calonico et al. (2019). Hence, I estimate equation 2.2 only for observations in the bandwidth $[-h^*, +h^*]$, where h^* is the optimal bandwidth.

$$Y_{it} = \alpha + \tau align_{it} + \gamma mov_{it} + (\sigma - \gamma) align_{it} * mov_{it} + \epsilon_{it}, \qquad (2.2)$$

In all specifications, standard errors are clustered at the state-year level. Additional robustness checks with different levels of clustering leave the results largely unchanged. The coefficient of interest in the parametric and nonparametric specifications is τ, the effect of having an aligned mayor on the outcome in question. In the baseline, this is the log of discretionary infrastructure transfers per capita which a municipality receives in a given year.

To investigate the influence of the electoral cycle, I provide results both for the full sample of all observations and for a sample containing only observations for the two years leading up to a local election. A comparison of the estimates for both samples will yield insight into whether the electoral cycle plays a role in the distribution of the outcome in question. In addition to the baseline specifications for the full and pre-election sample, I conduct a variety of robustness checks by changing the relevant RD parameters, such as bandwidth and kernel, to establish the robustness of the estimated effects.

2.2.3 Test of RD Assumptions

Before presenting the main results I investigate whether the assumptions for the use of regression discontinuity methods are fulfilled in my setting. Regression discontinuity analysis requires fairly weak assumptions in order to identify local treatment effects (Lee and Lemieux, 2010). First, subjects may not be able to manipulate their treatment status. In my case, the treatment is alignment between the party of the mayor and one of the parties in the state government in a given year for a municipality. Hence, in order to sort into treatment, aligned candidates need to disproportionately win in close races. Eggers et al. (2015) show that incumbents tend to win a larger share of close races than non-incumbents. In my setting, this would pose a problem only if incumbents are disproportionately aligned or unaligned. I test for this directly through placebo tests in the robustness checks and find no evidence that this is the case. Another potential threat to the assumption is manipulation of election results by individual candidates or different parties across states.[15] Given the strong institutional framework and democratic traditions in Germany, this seems unlikely.[16] Nonetheless, I perform a density test of the running variable according to Cattaneo et al. (2018). The results and a histogram of the running variable are shown in Figure 2.1.

Both graphs suggest that there exist no discontinuities in the density of observations around the RD cutoff, supporting the view that there exists no widespread

[15]Note that alignment differs across states, as the eight states in my sample have different governing parties. Hence, systematic manipulation of election results by one of the major parties would not be problematic, as it affects alignment differently across states.

[16]Mayors could also switch parties in order to manipulate their alignment status. However, I find no evidence of such behavior in the data.

Figure 2.1: Density Tests of Running Variable

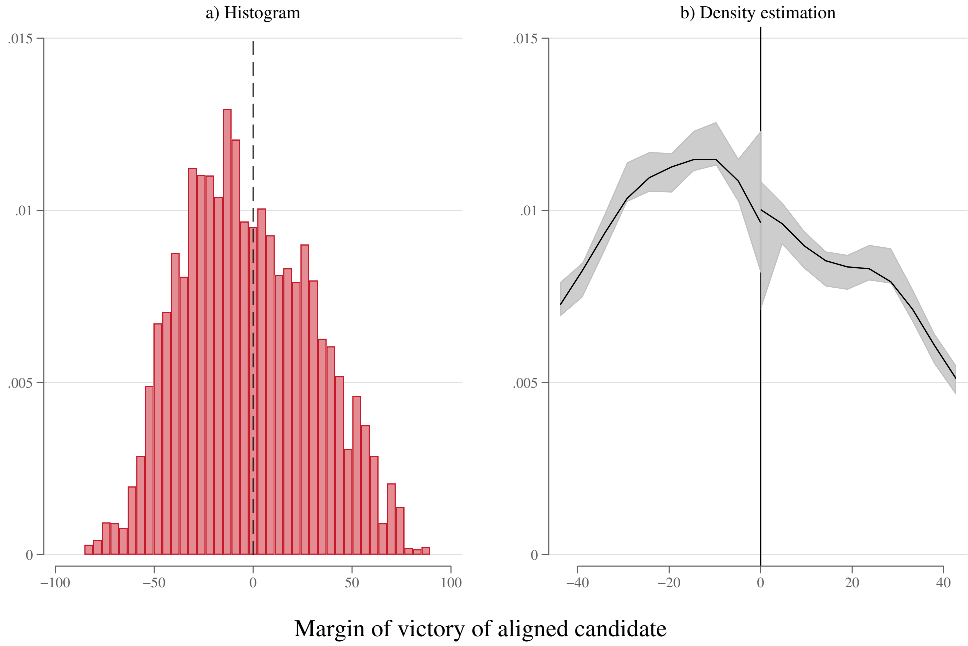

Note: Histogram and results of local density estimation according to Cattaneo et al. (2018). The black line in Graph b) represents the point estimates and the grey shaded area the 95% confidence interval.

manipulation of election results in favor of aligned candidates. Notably, the largest density of observations is slightly to the left of the RD cutoff, indicating that unaligned candidates may fare better than aligned ones on average.

The second assumption necessary for a valid RD identification is that the selection of subjects close to the cutoff needs to be quasi-random. That is, having a close local election in which the aligned candidate wins needs to be uncorrelated to other municipal characteristics which may influence the size of infrastructure transfers to it. There are several such potential factors, like population, population density, employment of the local population and municipal tax revenues. While there exists no obvious relationship between these characteristics and the success of aligned candidates in close local elections, I check whether they change discontinuously around the cutoff. If that were the case, the RD point estimate may be biased. I conduct placebo estimations with different dependent variables for both the full sample and the sample restricted to two years before the next local election to test the assumption empirically. For these regressions, I use the optimal bandwidth from the nonparametric baseline regression results shown in Table 2.2. Table 2.B.2 in the Appendix presents the results of these regressions.

It is evident that there exists a discontinuous effect of alignment on only one of the variables in question, population. One potential explanation is that local politics in large cities tend to be more partisan than in smaller municipalities. Hence, voters may vote for the same party in local and state elections, leading to higher probability of having an aligned mayor. This may explain the positive and significant coefficient for the population variable, which could lead to biased estimates. I deal with this problem in two ways: First, I conduct several robustness checks in which I drop municipalities with the largest population and those which receive the highest amounts of transfers per capita. Furthermore, all outcomes used for transfers are in per capita terms in order to control for the fact that larger municipalities will receive more transfers in absolute terms. Using the logarithm further reduces the weight given to outliers, for instance large cities with very large infrastructure projects. In both robustness checks, the main results remain unchanged.

2.3 Results

I now estimate the effect of having an aligned mayor on the amount of infrastructure transfers received. Figure 2.2 provides first evidence of the main result and shows the amount of transfers received by municipalities around the alignment cutoff.

According to the graphical evidence, there exists no noticeable discontinuity in the amount of transfers received for municipalities with an aligned mayor in the full sample. However, transfers to aligned municipalities appear to increase in the years leading up to local elections. This result is also confirmed in the estimation, where the alignment effect is most prominent in the pre-election sample. One caveat with the graphical evidence is the high variation in transfers to individual municipalities, even when observations are binned as in Figure 2.2. This fact is also evident in the comparably high standard deviation of infrastructure transfers in Table 2.1. In order to ascertain that any alignment effect is not contingent on a specific bandwidth

Figure 2.2: RD Plot Infrastructure Transfers

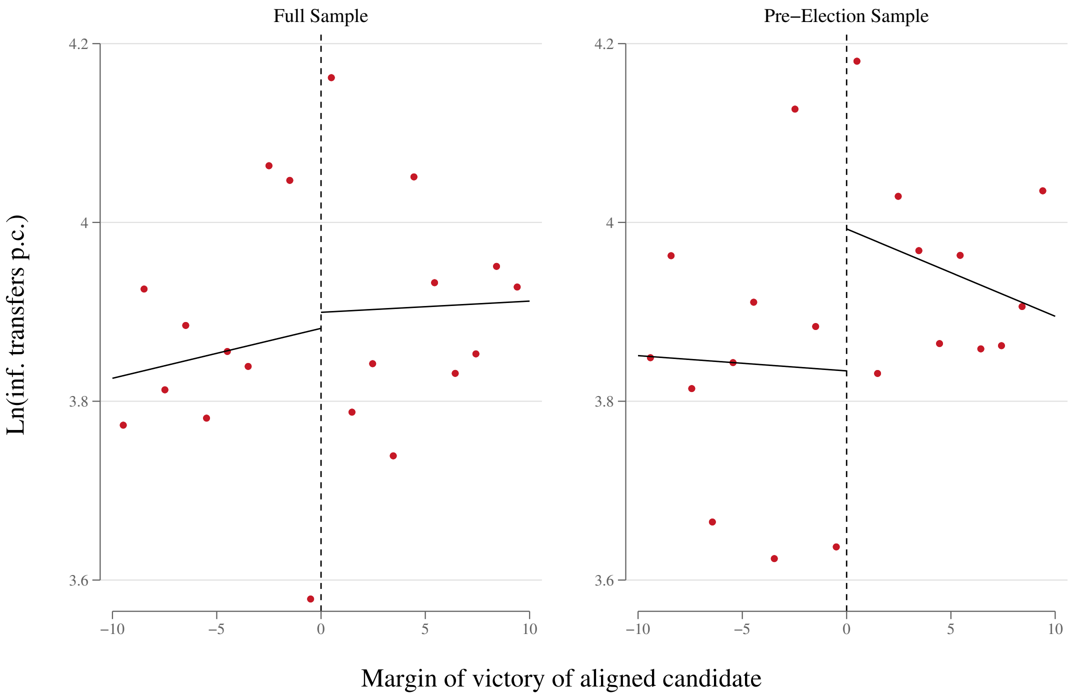

Note: RD plot for the baseline RD result. The x-axis is the margin of victory of the aligned candidate in the last local election and the y-axis the log of infrastructure transfers received by a municipality in a given year. Observations were binned in 10 evenly-spaced bins to the left and right of the cutoff. The line represents a nonparametric local linear fit of the data. The figure on the left contains all observations in the sample, the figure on the right only those for the two years leading up to the next local election.

choice or driven by individual outliers, I conduct several robustness checks to vary the bandwidth or exclude observations with particular high infrastructure transfers. All of these robustness checks leave the main results unchanged.

Table 2.2 presents the estimation results for alignment effect for parametric and nonparametric specifications. The first two columns show estimates with no covariates and the last a specification with state and year dummies following Calonico et al. (2019) to increase the efficiency of the RD estimator.

The point estimates are positive in every specification which points to an increase of infrastructure transfer to municipalities with aligned mayors. When covariates are included, the size of the coefficient grows and the optimal bandwidth and standard errors decrease. In line with the graphical evidence, the effect of alignment

Table 2.2: Baseline Regression Results

Ln(inf. transfers p.c.)	*Parametric* (1)	*Nonparametric* (2)	(3)
Panel A: Full Sample			
Alignment effect	0.149	0.168	0.215***
	(0.110)	(0.105)	(0.068)
Observations	10293	2632	2510
Bandwidth	-	12.57	11.99
Panel B: Pre-Election Sample			
Alignment effect	0.302*	0.327**	0.342***
	(0.165)	(0.152)	(0.112)
Observations	4566	1249	1179
Bandwidth	-	13.08	12.49
State Dummies	No	No	Yes
Year Dummies	No	No	Yes

Results of parametric RD with a fourth order polynomial (Column 1) and nonparametric local linear RD with a uniform kernel (Columns 2 and 3). The running variable is the margin of victory of the aligned candidate in the last mayoral election. Panel A contains the results for the full data set and Panel B for the pre-election sample with only observations for up to two years before a local election. Column (3) adds state and year dummies as covariates. Standard errors are clustered at the state-year level, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

is most pronounced and nearly doubles in the lead-up to local elections. The point estimate suggests about 30% higher infrastructure transfers per capita per year for aligned municipalities in the years leading up to a local election, which underlines the importance of the local electoral cycle.

The optimal bandwidth for the nonparametric specifications seems reasonably small so as to not include observations which are too far away from the cutoff. Overall, the parametric and nonparametric estimates are fairly similiar both in terms of size and significance, but due to the concerns expressed in Gelman and Imbens (2019), I will focus on the nonparametric specifications with covariates in the following analyses.

There exists a strong, positive effect of alignment on the amount of infrastructure transfers received. This result runs contrary to the existing literature on the alignment effect in Germany (Baskaran and Hessami, 2017). In terms of coefficient size, the effect is comparable to the one found by Brollo and Nannicini (2012) for Brazil, which is surprising given Germany's institutional setting. Hence, before I examine the channels responsible for the alignment effect, I briefly conduct a variety of robustness checks of the main result.

One potential caveat to the main result is that state governments may offer infrastructure funding for specific types of infrastructure which align with their political values. For instance, a left government may offer more funding for public transport, whereas a right government may focus on roads. If left and right mayors are more likely to apply for funding in line with these priorities, the estimated alignment effect may reflect the fact that political priorities of local and state governments are aligned and cannot be interpreted as a sign of favoritism. Luckily, the municipal finance data also contains information on the allocation of other types of discretionary transfers which are not earmarked for infrastructure investments.

Table 2.B.1 in the Appendix shows results for nonparametric RD estimations where different types of transfers are used as the dependent variable. The results show positive coefficients for most other types of discretionary transfers, suggesting that party alignment also has a positive effect for other types of discretionary transfers. The generally smaller magnitude of the point estimates compared to those in the baseline can be attributed to the fact that the discretionary transfers used here are less frequently awarded to municipalities. Hence, the vast majority of observations reports values of zero for any of these transfers. Notably, there is no effect for rule-based transfers. This in line with expectations, as rule-based transfers are allocated based on municipal characteristics, such as size and fiscal capacity.

I also check robustness to the RD specification by changing the bandwidth, kernel and polynomial of the baseline identification. Table 2.B.3 in the Appendix presents these results. Changing the order of the polynomial leads to significantly larger optimal bandwidths and slightly reduces the size of the coefficients. Using a triangular kernel leaves the results virtually unchanged. For the bandwidth, the point estimates seem to decrease for larger bandwidths. The significance of the point estimates in the full sample seems to be somewhat sensitive to changes in the RD parameters. However, the effects for the pre-election sample all remain strongly statistically significant.

Another concern for the robustness of the estimates is the level of clustering for standard errors. Recent contributions in the debate on clustering, most notably Abadie et al. (2017), have argued that applied researchers either cluster at levels that lead to clusters which are too small or too few in number. I chose to cluster standard errors at the state-year level, which leads to a total of forty clusters with a minimum size of twenty-nine observations per cluster.[17] In addition to resulting in an acceptable size and number of clusters, the structure of clusters allows for the control of interdependence of transfer payments in municipalities in the same state and year, for example due to effects of the local business cycle. Nonetheless, in order to ascertain the robustness of the baseline, I present the results with clustered standard errors at different levels for the nonparametric specification in Table 2.B.4 in the Appendix.

The significance of the estimates varies slightly with the level of clustering. For the state, district or municipality level, the coefficients for the full sample become insignificant. However, results for the pre-election sample are robust with only slightly lower levels of statistical significance. Clustering at the state, district or municipal level also increases the optimal bandwidth.

I further investigate whether the size of the coefficients is dependent on a particular group of observations, such as states or outliers in terms of population and transfers received. Table 2.B.5 in the Appendix presents results for the baseline estimation with each of the eight states excluded. In subsequent specifications I drop outliers – the 1% of observations with the highest infrastructure transfers per capita[18]

[17]This is for the Saarland in the year 2000, as not all municipalities had yet adopted direct election of mayors. The number of observations for the Saarland increases to 52 in subsequent years.

[18]This corresponds to a €536 per capita.

– or observations for large cities, i.e., those which do not belong to a district ('kreisfreie Städte'), as differences in their size and legal status may affect transfer payments.

There are several things to note in the results. First, the results do not appear sensitive to the exclusion of large cities. Second, dropping outliers decreases the size of the coefficient by about two thirds and the significance in the full sample. For the pre-election sample the point estimate decreases slightly but remains statistically significant. The results for only municipalities without cities show a similar pattern. Last, the effect does not seem to be driven by one particular state, as the point estimate all remain positive and significant in the pre-election sample.

One further potential confounding factor of the alignment effect estimated here is that alignment between the party of the mayor and the state government may be correlated with alignment of other elected officials who have similar incentives to obtain transfers for a municipality. For example, municipalities that narrowly elect a conservative mayor may be more likely to have a conservative representative in the state legislature, a conservative county executive or a conservative majority in the city council who also benefit electorally from obtaining transfers for the municipality. If that were the case, the effect of alignment estimated here may falsely be attributed to mayors. Hence, I use information on the party membership of the local MP, county executive and the majority in the city council and generate dummy variables, indicating whether they are aligned with the state government.[19] I use these dummies as dependent variables and rerun the baseline nonparametric specification. Any positive effect found here may point to a potentially problematic correlation between having an aligned mayor and an aligned state MP, county executive or city council.

In order to address the concerns raised by Eggers et al. (2015) that candidates who win close elections may differ systematically from those who lose, I also estimate the RD effect of alignment on mayors' characteristics following the same procedure that was used for the placebo estimation of municipal characteristics. As dependent

[19]For the county executives, the sample consists of information for all eight states in the main sample. For state MPs the sample includes all states except the Saarland where the electoral system slightly differs. For the city councils, the data set includes information for all states from the main sample except Saxony, where data for seat allocations was not available.

variable I use the dummy for runoffs, i.e., whether the mayor won election in the first round or in the runoff. The estimation results for both alignment at different levels of government and mayors' characteristics are presented in Table 2.B.6 in the Appendix.

The results show that electing an aligned mayor in a close election is not associated with electing an aligned state MP, county executive or city council. The point estimates are all negative, suggesting that electing an aligned mayor is negatively correlated with having aligned representatives at other levels of government. In fact, the negative coefficient for alignment of the city council is strongly significant, suggesting that municipalities with an aligned mayor are about 11% less likely to have an aligned city council in the pre-election sample. Furthermore, there is no effect on runoffs, suggesting that aligned mayors are not more likely to be elected in a runoff.

2.4 Mechanisms

If one follows the dominant interpretation in the literature, the alignment effect estimated here is a sign of partisanship in the political system and the result of state governments using intergovernmental transfers for electoral purposes. Both of these interpretations are only of limited suitability for Germany, where partisanship is generally considered weak and there exists no conclusive evidence of widespread central government favoritism. Hence, I investigate whether there exist alternative explanations for the alignment effect estimated here. First, I examine whether local or state politics drives the alignment effect. Second, I test whether politicians' connectedness, as opposed to partisanship provides an additional mechanism for the alignment effect.

2.4.1 State or Local Politics?

The baseline results have already established that the local electoral cycle strongly influences the size of the alignment effect. In the two years leading up to a local election, the size of the coefficient doubles. One aspect to note here is that there exists no local electoral cycle within states. Election dates for mayors in a specific state vary based on unforeseen shocks, for instance resignations or death of previous mayors while in office. Hence, the electoral effect measured in the main result cannot be attributed to an overall increase in intergovernmental transfers. Instead, it stems from higher transfers to municipalities with an upcoming local election. I now test differences between the state and local electoral cycle explicitly by splitting the data set into years before and after local as well as state elections. Table 2.3 presents the results of these regressions.

Table 2.3: RD Results for State and Local Electoral Cycle

	Local Electoral Cycle		State Electoral Cycle	
Ln(inf. transfers p.c.)	*First Two* *(1)*	*Last Two* *(2)*	*First Two* *(3)*	*Last Two* *(4)*
Alignment Effect	-0.002	0.342***	0.168**	0.160**
	(0.053)	(0.112)	(0.077)	(0.070)
Observations	2276	1179	1792	1931
Bandwidth	19.67	12.49	13.21	15.14
State Dummies	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes

Note: Results for local linear nonparametric RD regressions for sample splits based on the local and state electoral cycle. The running variable is the margin of victory of the aligned candidate in the last mayoral election. Columns (1) and (2) contain results for the local, column (3) and (4) for the state electoral cycle. Odd-numbered columns show results for the sample containing only the first two years after the last local or state election. Even-numbered column contain only observations for the last two years before the next local or state election. All specifications use a uniform kernel, optimal bandwidths and clustered standard errors at the state-year level, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

It is immediately clear that the size of the alignment coefficient depends strongly on the local electoral cycle and appears to be driven by years before local elections, when a mayor may expect electoral benefits from obtaining higher transfers. The coefficient does not react to splitting the sample based on the state electoral cycle, suggesting only a limited role of state politics.

Do aligned mayors benefit electorally from their ability to obtain higher transfers from the state government? In order to disentangle the electoral benefit of alignment from other factors, such as previous incumbency, I focus on first term mayors who won in a close election at time t and compare whether their probability to run for reelection, their margin of victory in the next election and probability of being reelected at time t+1 is larger if the mayor was aligned when she was elected at time t.[20] Table 2.4 shows the results of these regressions.

Table 2.4: RD Results for Electoral Effects of Alignment

	Running again (1)	*MoV next election* (2)	*Won next election* (3)
Alignment Effect	-0.067	20.2**	0.177*
	(0.121)	(8.631)	(0.094)
Observations	280	176	242
Bandwidth	16.58	13.6	19.88
State Dummies	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes

Note: Results for local linear nonparametric RD regressions for data set containing only election observations. The sample is restricted only to mayors in their first term. The running variable is the margin of victory of an aligned candidate at time t. Dependent variables are a dummy indicating whether the incumbent is running for reelection in the next election (1), the margin of victory of the incumbent in the next election (2) and a dummy indicating whether the incumbent won reelection in the next election (3). All specifications use a uniform kernel, optimal bandwidths and clustered standard errors at the state-year level, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The results suggest that there exist no significant differences in the probability to seek reelection between mayors who were aligned or unaligned at the time of their first election. However, aligned mayors are more successful electorally than unaligned ones. Their margin of victory is about 20% higher. The magnitude of this effect is significant as the incumbency advantage for German mayors has been estimated to range between 15 and 17 percentage points (Freier, 2015). The electoral advantage also translates into a higher probability of reelection, with aligned mayors enjoying a 15% higher probability of being reelected.[21] It is likely

[20] For this analysis, I move away from the previous panel structure of the data to one which contains only observations for election years. Hence, I cannot account for changes in alignment which occur during a mayor's term due to a change in the state government and focus on mayors who were aligned at the time of election.

[21] The study by Freier estimates the effect of incumbency on reelection probability to range between 30 and 40%.

that this effect is at least partially caused by the ability of aligned mayors to obtain additional funds for their municipalities.

Together, these results indicate that local politics plays a more important role than those at the state level.

2.4.2 Reelection Incentives

To underline the importance of local politics, I use the names of the winner and runner-up to identify cases in which an incumbent mayor is a 'lame duck' and not running for reelection in the next local election. Arguably, incumbents in their final years in office have fewer incentives to obtain transfers than those who have to compete again at the ballot box. However if the alignment effect is driven by a partisan state government instead of local conditions, the alignment effect should either persist or even increase if the state government aims to convince voters to continue electing a mayor of their party. Hence, by splitting the sample between mayors running for reelection and those who are not, I can determine whether the alignment effect is driven more by the local electoral incentives of the mayor or by partisanship and patronage on the part of the state government.

One potential concern with this sample split is that the decision to run for reelection may be endogenous and related to the amount of transfers a mayor was able to obtain for a municipality during her previous term. In these cases, unsuccessful mayors may be less likely to seek reelection in the first place. However in 75% of cases, eligible incumbent mayors run for reelection (Gehne and Holtkamp, 2002) and the most common reason for retirement is an age limit which varies by state but specifies either an age at which mayors have to resign or a maximum age at which they can run for reelection (see for instance Bertelsmann Stiftung, 2008 for the age distribution of mayors). Hence, whether a mayor runs for reelection again is mostly determined by her age.

Figure 2.B.1 in the Appendix shows graphical evidence for the sample split based on mayors' electoral incentives. In the graphs for the full sample, there appears to be no difference in the amount of transfers going to aligned or unaligned municipalities,

regardless of whether the incumbent is seeking reelection. In the pre-election sample, the graph shows a small discontinuity for those cases where the incumbent is running for reelection but none for those where the incumbent retires at the end of term.

Table 2.5 presents the estimation results for the sample split for mayors' electoral incentives. It has to be noted that due to limitations in the size of the data set in the three East German states, the covariate structure had to be adjusted slightly. Instead of three state dummies for the East German states, the estimates use only one dummy for East Germany and state dummies for the West German ones. Otherwise, the optimal bandwidth cannot be calculated.[22]

Table 2.5: RD Results for Reelection Incentives

	Incumbent not running		Incumbent running	
Ln(inf. transfers p.c.)	*Full* *(1)*	*Pre-Election* *(2)*	*Full* *(3)*	*Pre-Election* *(4)*
Alignment Effect	0.256**	0.088	0.203***	0.384***
	(0.120)	(0.148)	(0.073)	(0.127)
Observations	839	474	1725	792
Bandwidth	12.26	14.08	12.52	12.72
State/East Dummies	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes

Note: Results for local linear nonparametric RD regressions for sample split on incumbents seeking reelection. The running variable is the margin of victory of the aligned candidate in the last mayoral election. Columns (1) and (2) contain results for the data set when the incumbent mayor is not running for reelection. The results for an incumbent running again are presented in columns (3) and (4). Odd-numbered columns are for the full and even ones for the pre-election sample for observations in the two years leading up to an election. All specifications use a uniform kernel, optimal bandwidths and clustered standard errors at the state-year level, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The positive and significant effect of having an aligned mayor stems mostly from cases when the incumbent mayor is running for reelection. When the incumbent is not running for reelection, the alignment effect disappears in the pre-election sample. This result suggests that mayors' electoral incentives play a significant role in the allocation of infrastructure transfers. It also points to a limited role of

[22] Additional results with the previous covariates and sensible bandwidths specified by hand lead to similar results.

state politics. If the allocation of transfers were based on partisan politics, the state government would allocate transfers especially to municipalities with an upcoming open mayoral election in order to support their own candidate and secure that the office of mayor remain in the hands of a fellow party member.

Taken together, the results on mechanisms indicate that the dominant interpretation of alignment effects as signs of partisan politics by the central government may be incomplete. In fact, local politics, such as the electoral cycle and mayors' reelection incentives, is driving the alignment effect.

2.4.3 Partisanship or Connectedness?

In a next step, I will evaluate the role of partisanship directly and provide an alternative channel for the alignment effect, namely local politicians' connectedness. Attributing the alignment effect to partisanship is a logical explanation when the effect is caused by a central government which distributes funds based on political considerations to aligned local politicians. However in my case, local political conditions seem to play a more important role.

An alternative explanation is a politician's connectedness to state politics. Politicians with better connections may be able to obtain transfers for their municipality more successfully than less connected colleagues. Such connectedness can be independent of membership in one of the governing parties as it only reflects better access to information, for instance on suitable transfers programs by the state government. Most importantly, connectedness does not imply favoritism in the allocation of transfers by the state but instead a higher efficiency in the process of applying for and obtaining transfers by local politicians.

To test this empirically, I exploit the fact that candidates unaffiliated with a party enjoy a certain amount of success in municipal elections in Germany. These independent mayors differ from others in that they typically are not members of any political party and therefore do not have the same connections to state politics as members of a political party, regardless of whether their party is currently in the government or opposition.

From the viewpoint of a partisan state government, opposition party mayors and independent ones do not differ significantly. If the allocation of transfers by state officials is driven by partisan considerations, both groups of mayors will be at a disadvantage compared to aligned mayors. However, opposition party mayors and independents differ in that opposition party mayors typically have connections to fellow party members at the state level, even if those are not currently serving in the state government. Hence, they may enjoy access to more accurate information on transfer opportunities and successful strategies in writing grant proposals.[23].

To test whether differential access to information is responsible for the alignment effect, I run three additional RD regressions. First, I compare cases in which an independent and opposition party candidate compete. The running variable is the margin of victory of the opposition party candidate. Note that these cases were previously excluded from the analysis, as there were no aligned candidates running. If party mayors enjoy better access to information, one would expect a positive effect of electing an opposition party mayor on the amount of transfers a municipality receives. If partisanship is responsible for the effect, one would expect no differences between opposition and independent mayors. Second, I compare cases in which a government party, i.e., an aligned, and an independent candidate compete with each other. Third, I compare cases where opposition party and government party mayors compete. The size of the coefficient in these cases sheds light on whether the baseline effect estimated earlier is driven by partisanship or connectedness. If the effect stems mainly from differences between independent and government party mayors, it is plausible that connectedness is the dominating mechanism. If the effect comes from differences between government party and opposition party mayors, partisanship may be more important.

Table 2.6 presents the results of these regressions. These estimates rely on a smaller sample because the East German states and the Saarland have an insufficient

[23]One potential problem with this estimation strategy is that independent and party mayors may differ in other characteristics which could influence the ability to obtain transfers, such as their education. The available evidence suggests, however, few differences in the level of education of mayors of different parties (Gehne and Holtkamp, 2002)

number of case to estimate the effect of electing an independent.

Table 2.6: RD Results for Information Transmission

Ln(inf. transfers p.c.)	*Oppo/Indep* *(1)*	*Gov't/Indep* *(2)*	*Gov't/Oppo* *(3)*
Panel A: Full Sample			
Alignment effect	0.200*	0.447***	0.114*
	(0.102)	(0.163)	(0.068)
Observations	637	361	1784
Bandwidth	18.47	15.13	11.54
Panel B: Pre-Election Sample			
Alignment effect	0.023	0.849***	0.302**
	(0.192)	(0.259)	(0.119)
Observations	235	143	881
Bandwidth	15.39	14.15	11.9
State Dummies	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes

Note: Results for nonparametric RD regressions for races between independent and party candidates. Columns (1) contains results for the data set on elections between opposition party and independent candidates. The results for races between aligned and independent candidates are presented in column (2). In column (3), the results are for elections between government party mayors and opposition party mayors. Panel A contains the results for the full data set and Panel B for the pre-election sample with only observations for up to two years before a local election. All specifications use a uniform kernel, optimal bandwidths and standard errors clustered at the state-year level, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column (1) shows a positive, weakly significant effect on transfers to municipalities with an opposition party mayor when compared to those with an independent mayor. This effect disappears in the pre-election sample. Running the baseline estimation only for cases in which aligned and independent candidates competed in the last local election (columns 3 and 4) results in an alignment effect more than twice as large as in the main specification with a significant electoral cycle effect. Furthermore, the comparison between government and opposition party mayors yields a positive and significant alignment effect with slightly smaller coefficient size than in the baseline.

Taken together, the results suggest that municipalities with independent mayors are comparatively worse off in terms of the amount of transfers they receive than municipalities with mayors who are member of a political party, irrespective of whether the party is in the opposition or in the government. This points to the role of connectedness and information transmission as one of the factors influencing the allocation of infrastructure transfers. In fact, a certain part of the alignment effect found earlier appears to be driven by differences in municipalities with aligned or independent mayors.

The effect of connectedness can potentially be explained by limited local administrative capacity. Most municipalities are small in terms of population and they only support small city administrations, which are responsible for a host of traditional local issues, such as zoning, immigration, or maintenance of public infrastructure. As such, the ability to lobby the state government and engage in large-scale investment projects is limited. Most importantly, the responsibility for writing grant proposals and applications for transfers rests solely with the mayor and the administration. Given the often complicated application processes and the time involved in finding a suitable state program for infrastructure transfers, it is likely that mayors may often suffer from insufficient information and do not try to obtain funding from sources which may be suitable for their municipality.[24] Südekum (2019) finds support for this argument in qualitative interviews with mayors.

Against this background, it seems likely that mayors may benefit from connections to politicians in the state capital – for instance fellow party members – who inform them about new funding programs or leftover funds which are often hastily distributed at the end of the fiscal year. These information advantages apply fairly equally to members of the opposition and the government party, as the state parliament closely monitors government activity leaving opposition lawmakers well-informed about relevant transfer programs.

[24] The number of different programs is large. As an example, the state of Lower Saxony alone had around 40 different programs for investment transfers to municipalities in 2011. Often, these programs set very detailed conditions and subsidize only specific types of local infrastructure.

2.5 Conclusion

This paper set out to investigate the role of party politics in the allocation of intergovernmental transfers in a developed economy with strong institutions. Contrary to expectations and the existing literature, I find a significant effect of alignment between the party of the mayor and the state government on the amount of infrastructure transfers given to a municipality. In the existing literature, such an alignment effect has frequently been attributed to party favoritism and pork-barrel politics. The effects in the baseline are of similar magnitude as those estimated by Brollo and Nannicini (2012) for Brazil, a surprising similarity given the differences in the institutional background between these two countries.

In terms of the mechanisms responsible for the alignment effect, this paper finds that local political conditions, not those at the state level, play an important role in explaining the alignment effect. Furthermore, the results suggest that the effect is at least somewhat driven by connectedness of local politicians to state politics, regardless of alignment. Hence, this paper agrees with the existing literature in so far as it seems doubtful that widespread party favoritism exists in a country such as Germany. However, the results call into question whether an alignment effect can be interpreted as a sign of partisanship and inefficient distribution of resources on part of the central government. Instead, it may be the result of politics and information asymmetries at the local level. As such, this paper calls for a more nuanced interpretation of alignment effects.

2.A Data Sources and Collection

The starting point for data collection in each state were the respective State Statistical Offices. In North Rhine-Westphalia, Hesse, Rhineland-Palatine, the Saarland and Thuringia, these provided fairly complete data sets on municipal elections. In the East German states of Saxony and Saxony-Anhalt, the data were scraped from online databases of the respective state government. In Lower Saxony, the information had to be digitized from handwritten records submitted by municipalities to the State Election Office.

In all states, missing information – for instance on party membership or losing candidates – was added after intensive online searches. As a last resort, missing information was obtained by calling individual municipal governments. In the end, I obtained nearly complete records of all direct mayoral elections in each state.

Due to differences in the organization of municipalities and the roles of mayors in the individual states, I had to impose certain restrictions on the data set. Since the paper focuses on the role of full-time and professional mayors, the sample is restricted only to municipalities where the mayor works full-time in this capacity. In the three East German states of Saxony, Saxony-Anhalt and Thuringia, there exist many small municipalities with honorary mayors where the responsibility for the city administration is typically delegated to an administrative union (‘Verwaltungsgemeinschaft’) of several small municipalities. The head of the administrative union is chosen by a council of representatives from each member municipality. These municipalities are dropped from the sample, as it is unclear who is responsible for writing proposals for transfers and who is going to benefit electorally from them.

In the states of Rhineland-Palatine and Lower Saxony, very small municipalities are part of municipality associations (‘Samtgemeinde’ or ‘Verbandsgemeinde’). These associations take over most of the responsibilities of normal municipal governments and are headed by a mayor who is directly elected by the entire population of the member municipalities. Hence, in these two states the level of observation is at the municipality level for those municipalities which are not part of an association and at the municipality association level for those municipalities which are part of an association.

2.B Additional Tables and Figures

Table 2.B.1: RD Results for Different Transfer Types

	Rule-based (1)	*Necessity transfers* (2)	*General-budget* (3)	*Other transfers* (4)	*Debt relief* (5)
Full Sample					
Alignment effect	0.007	0.054	0.001	0.006	0.005
	(0.043)	(0.099)	(0.105)	(0.081)	(0.019)
Observations	3029	3376	2996	2765	958
Bandwidth	19.06	20.77	17.88	16.5	5.974
Pre-Election Sample					
Alignment effect	0.010	0.139	0.154	0.101	0.0396*
	(0.111)	(0.144)	(0.122)	(0.111)	(0.021)
Observations	961	1104	1295	1112	546
Bandwidth	12.91	14.94	17.77	15.08	7.558
State/East Dummies	Yes	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes	Yes

Note: Results for nonparametric local linear RD regressions with log of different transfers per capita per year as dependent variables. The running variable is the margin of victory of the aligned candidate in the last mayoral election. Panel A contains the results for the full data set and Panel B for only those up to two years before a local election. All specifications use a uniform kernel, optimal bandwidths and standard errors clustered at the state-year level, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.2: Placebo RD Results for Municipal Characteristics

	Ln(Tax revenue p.c.) (1)	*Employment rate* (2)	*Population* (3)	*Population density* (4)
		Panel A: Full Sample		
Alignment effect	-0.019	-0.002	7334	-4.032
	(0.035)	(0.002)	(7402)	(44.7)
Observations	2029	2510	2510	2510
Bandwidth	11.99	11.99	11.99	11.99
		Panel B: Restricted Sample		
Alignment effect	0.011	0.000	22405**	49.76
	(0.055)	(0.002)	(10182)	(69.2)
Observations	940	1179	1179	1179
Bandwidth	12.49	12.49	12.49	12.49
State Dummies	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes

Note: Results for placebo RD estimations with different dependent variables,log of municipal tax revenue per capita (Column 1), share of local population in regular employment (2), population (3) and population density (4). The running variable is the margin of victory of the aligned candidate in the last mayoral election. Panel A contains the results for the full data set and Panel B for only those up to two years before a local election. All regressions use the optimal bandwidth for the nonparametric baseline regression in Table 2.2 in a local linear specification and uniform kernel. Clustered standard errors at the state-year level are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.3: Robustness Checks for Different RD Specifications

Ln(inf. transfers p.c.)	*Baseline* *(1)*	*Quadratic* *(2)*	*BW * 0.5* *(3)*	BW * 0.75 *(4)*	BW * 1.25 *(5)*	BW * 1.5 *(6)*	*Triangular* *(7)*
Panel A: Full Sample							
Alignment effect	0.215***	0.172**	0.213**	0.112*	0.119**	0.103*	0.177***
	(0.068)	(0.078)	(0.093)	(0.065)	(0.056)	(0.053)	(0.064)
Observations	2510	4086	1232	1880	3151	3729	2950
Bandwidth	11.99	19.74	5.99	8.99	14.98	17.98	14.05
Panel B: Pre-Election Sample							
Alignment effect	0.342***	0.263**	0.336**	0.284**	0.222**	0.219***	0.307***
	(0.112)	(0.119)	(0.153)	(0.116)	(0.089)	(0.081)	(0.108)
Observations	1179	2007	559	874	1471	1723	1341
Bandwidth	12.49	21.82	6.24	9.37	15.61	18.73	14.05
State Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Results for RD robustness specifications. The running variable is the margin of victory of the aligned candidate in the last mayoral election. Column (1) contains the results for the baseline from Table 2.2 for comparison. Column (2) uses a quadratic polynomial, (3)–(6) different bandwidths and (7) a triangular instead of a uniform kernel. Panel A contains the results for the full data set and Panel B for only those up to two years before a local election. All specifications use optimal bandwidths and clustered standard errors at the state-year level, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.4: Robustness Checks for Different Clustering levels

Ln(inf. transfers p.c.)	*State/Year* *(1)*	*State* *(2)*	*District* *(3)*	*Muni* *(4)*	*No* *(5)*
Panel A: Full Sample					
Alignment effect	0.215***	0.211	0.143	0.154	0.215***
	(0.068)	(0.143)	(0.122)	(0.108)	(0.053)
Observations	2510	2572	3002	2937	2510
Bandwidth	11.99	12.29	14.45	14.01	11.99
Panel B: Restricted Sample					
Alignment effect	0.342***	0.304*	0.28*	0.304**	0.268***
	(0.112)	(0.180)	(0.164)	(0.142)	(0.088)
Observations	1179	1239	1305	1239	1080
Bandwidth	12.49	12.95	13.77	12.95	11.5
State Dummies	Yes	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes	Yes

Note: Results for nonparametric local linear RD regressions with different levels of clustering. The running variable is the margin of victory of the aligned candidate in the last mayoral election. Panel A contains the results for the full data set and Panel B for only those up to two years before a local election. All specifications use a uniform kernel and optimal bandwidths, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.5: Robustness Checks for Exclusion of Observations

Ln(inf. transfers p.c.)	*Low Saxony* *(1)*	*NRW* *(2)*	*Hess* *(3)*	*RP* *(4)*	*Saar* *(4)*	*Sax* *(5)*	*Sax-Ah* *(6)*	*Thue* *(7)*	*Municipality* *(8)*	*Outliers* *(9)*
Full Sample										
Alignment effect	0.086	0.203**	0.089	0.058	0.111	0.094	0.123*	0.066	0.112	0.079
	(0.082)	(0.079)	(0.064)	(0.060)	(0.068)	(0.069)	(0.065)	(0.067)	(0.072)	(0.064)
Obs	1803	1517	1897	2557	2196	2019	2360	2166	2244	2187
Bandwidth	11.01	9.843	11.75	14.53	10.82	10.18	11.44	10.66	11.53	10.47
Pre-Election Sample										
Alignment effect	0.253**	0.411***	0.209**	0.198*	0.35***	0.302**	0.273**	0.238**	0.291***	0.239**
	(0.106)	(0.137)	(0.099)	(0.110)	(0.118)	(0.118)	(0.113)	(0.113)	(0.110)	(0.104)
Obs	1123	676	840	980	1076	1026	1074	984	1214	1065
Bandwidth	14.38	11.07	12.06	11.89	11.71	11.46	11.57	10.88	13.72	11.43
State Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Results for nonparametric RD regressions excluding different types of observations. The running variable is the margin of victory of the aligned candidate in the last mayoral election. Columns (1) – (5) contain results with individual states excluded in the order Lower Saxony (1), North Rhine-Westphalia (2), Hesse (3), Rhineland-Palatine (4), Saarland (5), Saxony (6), Saxony-Anhalt (7) and Thuringia (8). In column (9) all cities which do not belong to a district are dropped and in column (10) all observations with transfers per capita larger than 500€. Panel A contains the results for the full data set and Panel B for only those up to two years before a local election. All specifications use a uniform kernel, optimal bandwidths and standard errors clustered at the state-year level, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.6: Placebo RD Results for Other Political Characteristics

	State MP (1)	*County Exec.* (2)	*City Council* (3)	*Runoff* (4)
	Panel A: Full Sample			
Alignment effect	-0.018	-0.0691	-0.059	0.053
	(0.030)	(0.051)	(0.038)	(0.045)
Observations	2399	2259	2332	2575
Bandwidth	12.46	12.46	12.46	12.46
	Panel B: Pre-Election Sample			
Alignment effect	-0.030	-0.009	-0.115***	0.042
	(0.047)	(0.057)	(0.043)	(0.071)
Observations	1086	989	1042	1116
Bandwidth	11.9	11.9	11.9	11.9
State Dummies	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes

Note: Results for nonparametric local linear RD regressions on different political variables as outcomes. The running variable is the margin of victory of the aligned candidate in the last mayoral election. In column (1), the outcome is a dummy indicating whether the party of the directly elected local state MP is part of the state government. In column (2), it is a dummy for alignment of the party of the county executive and in column (3) a dummy for alignment between the majority of seats in the city council and the state government. In column (4) the dependent variable is a dummy indicating whether the incumbent mayor won election in a runoff. Panel A contains the results for the full data set and Panel B for only those up to two years before a local election. All specifications use a uniform kernel, state and year dummies as covariates and clustered standard errors at the the State/Year level. The bandwidths are those for the nonparametric baseline specification from Table 2.2, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 2.B.1: RD Plot for Incumbency Sample Split

(a) Incumbent running

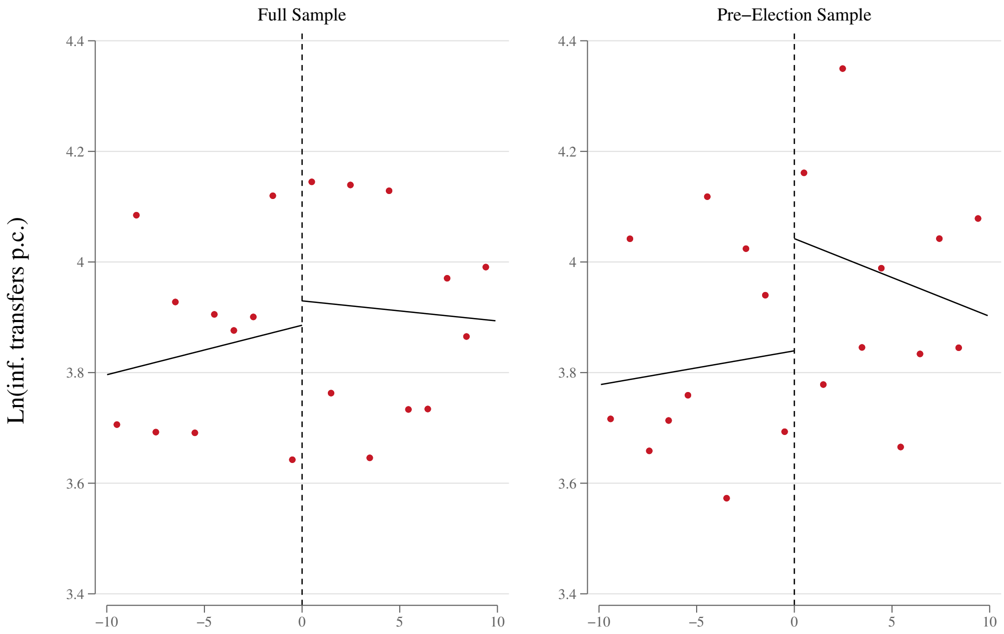

(b) Incumbent not running

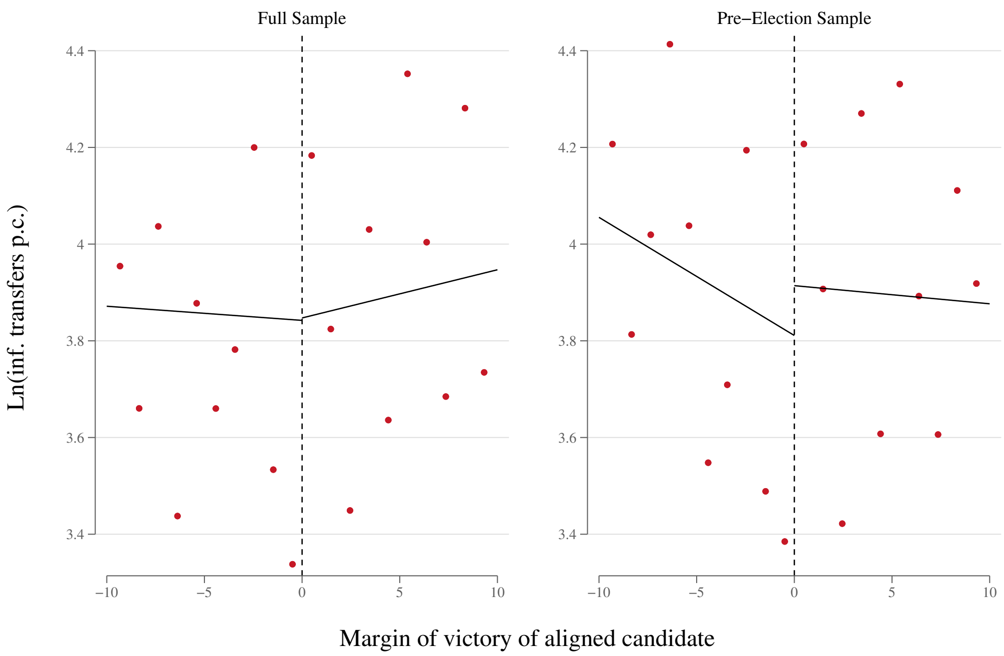

Note: RD plots for the sample split for incumbents. The x-axis is the margin of victory of the aligned candidate in the last local election and the y-axis the log of infrastructure transfers received by a municipality in a given year. Observations were binned in 10 evenly-spaced bins to the left and right of the cutoff. The line represents a linear fit of the data. The left column contains all observations in the sample, the right column only those for the two years leading up to the next local election. Panel a) shows only observations where the incumbent mayor is running for reelection, Panel b) only those where she is not. To have comparable axis ranges for all graphs, two outliers with low transfers to the left of the cutoff in the pre-election sample of Panel b) were dropped, as their inclusion significantly increased the range of the y-axis. This change did not affect the slope of the linear fits in the graph.

Chapter 3

College Openings and Local Economic Development[1]

Income and unemployment rates differ substantially across cities and regions in many countries. In the United States, for example, wages in the highest and lowest paying metropolitan areas differ by a factor of three (Moretti, 2011). Similar discrepancies in income per capita are observed between regions in the European Union (OECD, 2009). The variation in unemployment rates can be even larger. In Germany, our empirical setting, local unemployment rates across metropolitan areas vary by a factor of six (BMWi, 2013).

In response to these large regional disparities, many governments promote policies aimed at promoting regional convergence by reducing perceived inequalities. These place-based policies could be direct subsidies to firms located in or planning to move into a disadvantaged area; or, they could target workers and their families through residential subsidies. Alternatively, they can take the form of investments in infrastructure by local, state or federal governments in order to increase the economic attractiveness of a region to firms and individuals alike. One such policy

[1]This chapter is joint work with Francesco Berlingieri and Christina Gathmann. We thank Sascha Becker, Felix Bierbrauer, Anna Bindler, Michael Burda, Christian Dustmann, Bernd Fitzenberger, Matthias Heinz, Pat Kline, Hugh Macartney, Henry Overman, Pia Pinger, Hildegunn Stokke and participants at the CESifo Summer Institute, EALE, SOLE, IZA World Labor Conference, UEA, University of Cologne, Verein für Socialpolitik, DFG-SPP workshop and ZEW for very helpful comments and suggestions. We are grateful to Verena Lauber for valuable work at an early stage of the project; and Philipp Vom Berge and the staff at the FDZ-IAB for data support. We gratefully acknowledge financial support from the German Science Foundation within the Priority Programme SPP 1764 "The German Labor Market in a Globalized World: Challenges through Trade, Technology, and Demographics" and from ZEW.

is the opening of publicly funded institutions, which could be government agencies, publicly funded research institutes or institutions of higher education.

Universities and colleges in particular could be a powerful tool for regional development. First of all, universities and colleges might generate positive spillover effects to local firms through research collaborations or knowledge spillovers. Alternatively, staff and students as well as the universities and colleges themselves may stimulate the demand for local goods and services through local multiplier effects. Moreover, by improving the human capital base in the region, colleges increase the thickness of the local labor market, in particular for high-skilled workers. In labor markets with search frictions and heterogeneous firms and workers, the match quality between workers and firms would then improve when more high-skilled workers look for a job and firms offer suitable jobs in the same local labor market. Furthermore, the new high-skilled workers might generate positive spillover effects. Formal and informal interactions among individuals at work or in the neighborhood foster knowledge sharing and learning, which may result in positive production externalities across workers (see, e.g., Marshall, 1890; Lucas, 1988; Glaeser, 1999; Serafinelli, 2019). Finally, a better educated workforce and local productivity or knowledge spillovers might encourage firms to locate in the region or invest more in complementary technology and capital.

Yet, can colleges promote regional economic development independently of where they are located – even in more remote regions? If so, they could be an effective policy tool to balance agglomeration forces, which tend to concentrate economic activities and people in densely populated areas. A new college might have little impact in remote areas, however, if the new college graduates do not find adequate jobs or amenities and move to the urban centers after graduation with few gains to the local economy. Or, firms might not find it worthwhile to relocate or invest in new technology that complements high-skilled workers. Overall, the local economy might not be dynamic or advanced enough to make use of potential knowledge or productivity spillovers that colleges can create.

This paper provides answers to these important questions. An important challenge in this endeavor is that the location of colleges is not random. Many universities and colleges were established many centuries or decades ago, which makes it difficult to isolate the impact of a university from other local developments that accumulate over time. We solve this identification problem by focusing on the opening of new colleges. The college openings in our setting explicitly targeted areas outside the large urban centers to improve access to tertiary education. Yet, regions in which a new college is opened are likely to differ from other regions without a college opening in their economic base, innovative capacity or labor market dynamics. We therefore combine matching with a time-varying difference-in-differences approach. The matching approach works well in eliminating differences in observable characteristics and growth rates between event and matched control districts. Our empirical approach then compares flexibly employment and wages in regions with a college opening to employment and wages in suitable control regions without a college opening.

We have four main findings. First, the opening of a technical college results in large, persistent growth in the local student population relative to the control regions. The new colleges further successfully improve the human capital base in the region: young, high-skilled employment in the region increases by 13% within a decade after the opening. Second, we find little evidence that the new colleges raise regional employment or wages, which would indicate sustained growth or productivity gains. In particular, we find little evidence that wages of high-skilled workers change after the college opening. The absence of any wage effect on high-skilled workers stands in stark contrast to other studies on the local effects of high-skilled migration (Beerli et al., 2021) or college-induced growth in high-skilled labor (Carneiro et al., forthcoming; Fuest and Immel, 2021). Yet, our results support studies that do not find a substantial impact on high-skilled labor and the skill premium (Beaudry and Green, 2003; Blundell et al., 2022). Using a local production function approach, we show that in our setting the new colleges do generate little skill-biased technological change through changes in production (Acemoglu, 1998) or research efforts (Beaudry and Green, 2005; Caselli and Coleman, 2006). This result is important when considering the placement of colleges as a regional policy instrument. Our results suggest that gains in metropolitan areas cannot be easily translated into gains in areas with a thinner labor market and weaker economic base.

Third, we explore who actually benefits from the new colleges in the local economy. In the average region, the additional high-skilled labor is largely absorbed by high-tech firms in manufacturing, while there are no effects in the service sector. We further find few local multiplier effects and little impact on the innovative capacity of the region. Finally, we document that the state of the local labor market matters: regions with more dynamic labor markets before the opening experience a permanent growth in high-skilled workers who get employed in new firms. Less dynamic regions, in contrast, experience little sustained growth in high-skilled labor, which is mostly absorbed by incumbent firms.

Our study contributes to several strands of the literature. A number of studies have documented a strong correlation between the location of universities and patenting activity, innovation and business start-ups (Jaffe, 1989; Bania et al., 1993; Audretsch and Feldman, 1996; Cohen et al., 2002; Woodward et al., 2006). Most studies focus on the importance of academic research for the development of specific local industries, such as pharmaceuticals or electronic equipment. Closer to us are Beeson and Montgomery (1993) who study the link between the quality of a university and local employment growth; and Kantor and Whalley (2014) who use shocks to a university's financial endowment to identify wage effects outside the education sector.[2] Our study differs from this strand along several dimensions: we use the opening of new colleges as plausibly exogenous variation for identification; second, we study new colleges that are focused on teaching rather than research and start up activities. Finally, we analyze the impact on the whole local economy rather than the technology-driven spillover effects in particular industries.

Furthermore, our analysis contributes to the literature on the effects of local supply shocks. Several studies has studied the inflow of immigrants (Card, 2001;

[2] Kantor and Whalley (2019) take a long-term view tracing the role of agricultural experiments for the level of agricultural production over more than a century.

Borjas, 2003; Glitz, 2012; Manacorda et al., 2012; Ottaviano and Peri, 2012) or commuter flows (Beerli et al., 2021; Dustmann et al., 2016) on local wages and employment. Most studies using the local area approach suggest that the inflow of immigrants or commuters have a small effect on the local wages of natives. Others argue that migration induces negative employment effects locally (Glitz, 2012; Dustmann et al., 2016). As immigrants in most countries are on average less skilled than the native population, these studies focus on adjustments to an increase in low-skilled labor. And if natives and immigrants are imperfect substitutes, a large inflow of immigrants need not have a sizable impact on native wages. Hence, local adjustments to a low-skill supply shock are likely to differ from the response to a shock to high-skilled labor if there are human capital externalities or other types of knowledge spillovers. If high-skilled workers raise the productivity or innovative capacity of other workers in the same firm or other firms in the region, for instance, the benefits for the local economy might be much larger than a growth in the low-skilled workforce (Moretti, 2004; Ciccone and Peri, 2006). Our study contributes to this literature by exploiting the opening of new colleges as a plausibly exogenous shock to the high-skilled workforce, which should be perfect substitutes to other high-skilled workers in the area.

Closest to us are studies on the wage and employment effects of the growth in college graduates (see, e.g., Beaudry and Green, 2003; Blundell et al., 2022) or tertiary educational institutions (Carneiro et al., forthcoming; Fuest and Immel, 2021; Lehnert et al., 2020). Unlike the former, we focus on the *local* impact of a *local* supply shock. Unlike the latter, we analyze colleges that are more focused on teaching than research and are not located in the large metropolitan areas. We further use a different identification strategy to address concerns of differential pre-trends and endogenous location choice.

Traditional open economy models emphasize changes in the output mix produced by the local economy in response to labor supply shocks (Lewis, 2011; Dustmann and Glitz, 2015). Regions with a relative growth in low-skilled labor, for instance, experience shifts to products and sectors that make intensive use of low-skilled labor. Another reason college openings may affect the regional economy is through local multiplier effects. Both college employees and the new college graduates might create additional demand for local goods and services. Recent research suggests that local multiplier effects may be sizable (see, e.g., Moretti, 2010; Moretti and Thulin, 2013; Faggio and Overman, 2014). Our study relies on a different, plausibly exogenous source of identification to investigate changes in the output mix and local multiplier effects. We find few effects in the non-tradable sector, likely because the colleges we analyze are relatively small.

Our results provide important lessons for policy-makers. Knowing whether an increase in high-skilled labor improves the economic conditions of the regional economy and all workers; or only benefits certain sectors or firms has important implications. If there are indeed positive externalities from college openings on the local economy, this could be yet another argument for public subsidies of tertiary education. In addition, our results may also be important for the design of regional policies. National and state governments often use regional policies to support areas with high unemployment and low economic growth. Prominent examples include region-specific subsidies to firms or local governments, such as the Federal Empowerment Zones in the US (Busso et al., 2013), regional subsidy programs in France (Gobillon et al., 2012), Italy (Bronzini and de Blasio, 2012), the UK (Criscuolo et al., 2019), Germany (von Ehrlich and Seidel, 2018); or the European Structural Funds (Becker et al., 2010, 2013).[3] Our work sheds new light on the question of whether public investments like the opening of a new college can improve employment prospects and local development and thus contribute to a decline in disparities across regions differing in economic prospects.

[3]Earlier work has focused on the relationship between city (or local industry) size or density and productivity more generally (see, for example, Ciccone and Hall, 1996, for a seminal contribution and Rosenthal and Strange, 2004, for a survey).

3.1 Institutional Background

3.1.1 College Openings in West Germany

The 1960s saw a rising need for educated workers coupled with a rising demand for formal education. Faced with capacity constraints at existing universities, the federal government decided in 1968 to establish technical colleges (*Fachhochschulen*) to complement regular universities.[4] The new technical colleges played an important role in making higher education accessible to the broader population. By 2010, about one in three students pursues a degree at a technical college.

Technical colleges offer study programs that are more specialized and practice-oriented than regular universities. The degree programs are heavily concentrated in fields like engineering, law and business.[5] Degree programs at technical colleges take three to four years and are thus comparable to Bachelor programs at universities but shorter than the traditional Diploma or Masters program. Most importantly, the degree programs at technical colleges combine academic study with periods of practical training where students obtain actual work experience in companies. Unlike universities, teaching staff at technical colleges are required to have several years of work experience outside academia. Today, most technical colleges ask for a PhD. though that was not the norm in the early years. To improve the practical relevance of teaching material, technical colleges often cooperate with local companies though the colleges are mostly publicly funded with little monetary contribution from private sources.

We focus on openings of public colleges, which could be either newly founded colleges or a new campus of an existing college, which is located in a different district. We thus drop openings of private colleges because they are very small, cover only a very narrow range of fields and are of minor importance for tertiary education in Germany. We further drop college openings that were converted from existing schools of secondary education (like vocational schools, for example). These

[4]The legal basis is discussed in *Abkommen zwischen den Ländern der Bundesrepublik zur Vereinheitlichung auf dem Gebiet des Fachhochschulwesens.*

[5]In 2001, 42% of all students in technical colleges were enrolled in management or law, while roughly 30% studied engineering (Haug and Hetmeier, 2003).

openings generate only very small and smooth changes to high-skilled supply in the region (see Wissenschaftsrat, 1991; Kulicke and Stahlecker, 2004). We further exclude four cases where a new campus was opened in the same district as the parent institution. Finally, we combine college openings that occurred in the same district and year into a single event.[6] To avoid problems of changes in sample composition, we work with a balanced sample dropping openings before 1980 in order to have a least four years of data prior to the event; and drop openings after 2000 in order to follow regions for at least a decade after the event.

Our final sample consists of twenty technical colleges in West Germany that were opened between 1988 and 1998. Appendix table 3.C.1 provides a list of the college openings and the opening year of the college. The last two columns in the appendix table show that the technical colleges are small relative to public universities. Five years after the opening, the college has on average 100 employees and 900 students. Given that the event districts have on average 50,000 regular employees (in the social security system), the new college is not a quantitatively important employer for the local economy. Yet, the 900 students provide a sizable supply shock to the districts, which could boost the stock of high-skilled labor by around 50 percent.

3.1.2 Local Predictors of College Openings

While the federal government decided on the overall expansion of tertiary education, state governments chose when and where the new colleges would be located. Municipalities and local governments, in contrast, had little influence on the location decision. The costs of financing the new colleges were shared between state and federal governments. State governments aimed at a more even spatial distribution of tertiary institutions in order to facilitate access for potential students and reduce the distance to the nearest college. A second consideration was to foster structural changes in the local economy. One example is the decline of the traditional coal mining and steel industries in the Ruhr area and the Saarland (Holuscha, 2012). College openings were seen as an attractive tool to attract new companies, reduce

[6] Two colleges were opened in the same year in *Göppingen* (1988) and *Rhein-Sieg* (1995).

out-migration and avoid population decline.[7] A third consideration was that regions could not be in remote locations, provide amenities to attract enough students and have an economic base to offer employment opportunities after graduation.

The political considerations just discussed suggest that new colleges were neither opened in the large metropolitan areas, nor in predominantly rural regions. To provide more systematic evidence on the adoption process, we relate the college openings to district-level characteristics three years prior to the opening, state and year fixed effects. The results in appendix table 3.C.2 show that colleges were more likely to be located in districts without a university, which is in line with the objective of a more even spatial distribution of tertiary institutions. The college openings are also negatively correlated with population density indicating that the openings occurred outside of the big cities and urban centers.[8] Districts where a new college was opened have a good economic base as shown by employment levels, which is in line with the last objective that the local labor markets should be attractive for students and college graduates.

Interestingly, none of the other local characteristics predict a college opening: Neither the detailed industry structure, nor the age or skill structure of the existing workforce play a role. Hence, the college opening is not explained by a local lack of skilled workers, an aging workforce or an industry-specific demand for skilled workers. Even more importantly, we find no evidence that wage levels or wage growth, past population or employment growth (measured between eight and three years before the event) had any effect on the decision to found a new college. The absence of meaningful correlations between economic prospects and the college openings is reassuring and reduces concerns that politicians picked the districts with the best economic performance or economically deprived regions lobbied successfully for a new college.

[7]See Schindler et al. (1991); Landtag (1991); Wissenschaftsrat (1995) for examples of the political discussion in the individual states. Another example is the former capital of Bonn where the founding of a technical college was to compensate the city for the move of the federal government to Berlin starting in 1999 (see, e.g., Wissenschaftsrat, 1996).

[8]The proximity to an urban center has no bearing on the adoption decision (see column (2) of appendix table 3.C.2).

3.2 Data Sources and Empirical Strategy

Our empirical analysis proceeds in three steps. We first present the district-level data on labor markets and plants, on which we base our empirical analyses. We then discuss our matching approach to address the issue that districts that have a college opening differ from districts without a college opening (see the discussion in Section 3.1). The final step outlines the time-varying difference-in-differences approach we use to identify the impact of a college opening on the local economy comparing event regions to their matched control regions.

3.2.1 Data Sources

To analyze how the opening of a college affects the local economy, we draw on administrative data from Germany over several decades. Our data contain the universe of all social security records and provides detailed information on employment and wages aggregated at the district level (Stüber and Seth, 2019). It covers all employees except civil servants, military personnel and the self-employed.[9] The data start in 1975 and end in 2010, which enables us to follow districts with a new college opening several years before and up to a decade after the opening.

We focus on workers in full-time employment as of June 30 each year. Hence, we exclude apprentices, student trainees, part-time work, marginal and seasonal employment. We have detailed information on the workforce by skill and age in each district. We distinguish two skill groups based on the highest qualification obtained. Highly skilled workers are workers who have graduated from a college or university. Less skilled workers have no university or college degree but might have a vocational degree. Missing values in the education variable are imputed by exploiting the panel structure of the data (Fitzenberger et al., 2006; Kruppe et al., 2014). We further distinguish three age groups: young workers (ages 20-29), prime-aged workers (ages 30-44) and older workers (ages 45-59).[10]

[9] The social security data cover around 80% of the German labor force.

[10] Data protection rules require that a cell is set to missing if the number of employees in a cell is below twenty. This restriction affects between 1.27% to 3.38% of observations. In our analysis, we set these employment shares to zero.

Our data also contain the mean daily wage (measured on June 30th of each year). As is common in social security records, wages are right-censored at the highest level of earnings that are subject to social security contributions. Individual wages are imputed based on the procedure used in Card et al. (2013) and then aggregated to averages by age and skill group. All wages are deflated using the consumer price index with 1995 as the base year.

We supplement the district-level data with plant-level information from the German Establishment History Panel (BHP), which draws on the same population of social security records. The BHP is a 50% random sample of all establishments with at least one employee covered by the social security system in Germany (Schmucker et al., 2016). The plant-level data of the BHP allow more detailed analyses along several dimensions: first, we observe the detailed industry of a plant to investigate in which industries the new graduates find employment. Second, we can distinguish whether the local supply of highly skilled workers is absorbed by incumbent plants or through the creation of new firms. The latter are defined as plants opening their business within the past five years; here, we rely on the procedure developed by Hethey and Schmieder (2010) to identify plant openings and distinguish them from simple changes in the establishment identifier due to spin-offs or mergers, for instance. Finally, we observe the broad occupational structure of the workforce in a plant, in particular the number of engineers and natural scientists, professionals, semi-professionals and unskilled or skilled manual workers. We aggregate the plant-level data from the BHP to the district level.

The average district in West Germany has a population of around 200,000 with 51,000 full-time employees in the social security system. We compute regional employment shares by broad industry and occupation. Data on the number of college students in a district stems for German Federal Statistics Office. We add regional information on population from the European Regional Database of Cambridge Econometrics. Finally, we obtain data on population flows across districts by broad age groups from the State Statistical Offices. Every person who moves to an area has to register with the municipality of their residence.

3.2.2 Matching Procedure

Regions where a new college was opened might differ from regions that did not obtain a new college. While local politicians had little influence on where a new college opened, state governments took local conditions into account when deciding on the timing and the location of the new college (see discussion in Section 3.1). A means comparison of districts with a college opening to the average district in West Germany reveals that event regions are less densely populated and have fewer highly skilled workers than the average district (see the first columns of table 3.B.1). Moreover, event districts have a strong base in manufacturing and construction, but a smaller service sector than the average West German district. Finally, workers in event districts earn lower wages than the national average. Given that new colleges are more likely to be located in areas with a less favorable economic development than the average district, a comparison of the local development in event regions to the average local economy in West Germany would underestimate the benefits a college opening has had on the local economy.

To address the issue of regional differences, we employ a matching approach to find suitable control regions for districts with a college opening. We match on population density, the broad industry structure (ten sectors) and whether a district has another college or university. As shown in Section 3.1, these characteristics are systematically related to the location decision and might also influence regional economic performance after the opening. We use Mahalanobis distance matching, which minimizes the standardized Euclidean distance of all matching variables between treatment and control districts, to find the closest match.[11] We exclude from the pool of potential control districts those sharing a border with an event district to avoid that the development in control districts is contaminated by spillover effects from the event district.

[11]The distances between each treatment and each potential control district are normalized by the variance-covariance matrix of the pooled sample of event and possible control districts. Normalizing by the variance-covariance matrix in the control group only does not alter the results.

Figure 3.B.1 shows the geographic location of treatment districts and control districts. Most college openings during our sample period occurred in *Baden-Württemberg* and *Bavaria*, two states in Southern Germany with a strong manufacturing base. The map also confirms once more that districts with a college opening and their controls are located outside the large urban centers. The right-hand side of table 3.B.1 shows that the selected control districts are very similar to the treated districts along observable characteristics. It is important to stress that the matching procedure not only eliminates differences in the characteristics we match on; but also differences in characteristics we did not match on, such as the age structure or the skill composition of the local labor force.

3.2.3 Empirical Model

Using our matched sample of regions, we compare labor market outcomes in the event regions to those in the control regions before and after a college opening. To illustrate the evolution of our results graphically, we estimate the following event study:

$$Y_{r\tau t} = \sum_{\tau=-7}^{-2} \beta_\tau Opening_r * Period_\tau + \sum_{\tau=0}^{11} \gamma_\tau Opening_r * Period_\tau + \eta_\tau + \lambda_t + \pi_r + \varepsilon_{r\tau t} \tag{3.1}$$

where r denotes the region, t the calendar year and τ the year relative to the college opening. Note that t and τ are distinct because colleges were opened in different calendar years. We consider a period of seven years before and eleven years after the college opening (i.e., $-7 \leq \tau \leq 11$). The opening occurs in period $\tau = 0$ and the base period is the year before the college opening, i.e., $\tau = -1$.

$Y_{r\tau t}$ denotes a local labor market outcome like employment or wages in region r in a given calendar year t for the event period τ. $Opening_r$ is a binary variable equal to one if there was a college opening in region r and zero for the control regions. $Period_\tau$ is an indicator equal to one if the college opening has happened τ years ago or will happen in τ years for both treatment and control region; and zero otherwise. The coefficients of interest in equation (3.1) are the γ_τ, which trace the evolution of the outcome of interest in the event region relative to the development in the control region between τ years after the college opening and the reference period ($\tau = -1$). The empirical model in (3.1) controls for event fixed effects (η_τ) to ensure that we compare event and control districts in the same period. We further include year fixed effects (λ_t) and region fixed effects (π_r) to account for aggregate shifts and region-specific, time-invariant unobservable differences. Standard errors are clustered at the district level to account for the level of aggregation in the treatment variable.

Given our small number of treatment regions, the estimates reported in our tables aggregate event years into three broad periods: the period before the opening, a transition period shortly after the opening and the longer-run development of the new college. More specifically, we estimate variants of the following model:

$$Y_{r\tau t} = \beta^B Opening_r * Before_\tau + \gamma^T Opening_r * Transition_\tau + \gamma^P Opening_r * Post_\tau \\ + \eta^{Trans} + \eta^{Post} + \lambda_t + \pi_r + \varepsilon_{r\tau t} \quad (3.2)$$

where $Before_\tau$ denotes the period prior to the opening (from $\tau = -7$ to $\tau = -2$). The transition period $Transition_\tau$ spans the year from the opening to the graduation of the first cohort (from $\tau = 0$ to $\tau = 5$). The post period $Post_\tau$ covers the longer-run development in the event regions from six to eleven years after the opening (from $\tau = 6$ to $\tau = 11$). As before, $Opening_r$ is an indicator equal to one if a college opened in region r; and zero otherwise. The parameters of interest are γ^{Trans} and γ^{Post}, which characterize how the outcomes of interest evolve in the first few years after the college opening and in the long run, respectively. All other variables are defined as in equation (3.1) above.

The key identifying assumption is that labor market outcomes would have evolved similarly in the event and control district in the absence of a college opening conditional on our control variables. Specifically, we require that *trends* in outcomes would have involved similarly in event districts than in control districts. We show below that controlling flexibly for district-specific linear or quadratic trends to capture differential trajectories has little effect on our empirical results. The evolution of outcomes in the pre-event period provides further evidence on the plausibility of this assumption in the case of homogeneous treatment effects. We show below that the parameters β_τ in equation (3.1) resp. β^B in equation (3.2) are close to zero and statistically insignificant. We consider alternative estimators that are robust to heterogeneous treatment effects in Section 3.3.4; these do not affect our conclusions. We now turn to our main results.

3.3 Empirical Results

3.3.1 Student Population and High-Skilled Employment

We start out by analyzing the effect of a college opening on the student population in the region. A rise in the student population not only helps identifying the exact timing of the opening; but is also a prerequisite for a positive impulse on the local economy. Figure 3.B.2 traces the growth in registered college students in treated and control districts with the level normalized to one in the year before the opening ($\tau = -1$). Prior to the college openings, the evolution of college students is flat in both event and control regions.[12] We see a substantial increase in the number of students in the event region starting in the year of the college opening. Ten years after the college opening, the student population in the treatment regions has increased by a factor of 8 to about 1,600 students while we see no change in the control districts over the same period.

To quantify the impact on the student population, we estimate equation (3.2) where the dependent variables are the number of students measured either in absolute numbers or as a share of high-skilled employees in the region. The specification includes fixed effects for the district, calendar and event year. All results are cumulative estimates relative to the reference period, the year before the college opening ($\tau = -1$). Column (1) of table 3.B.2 shows that new colleges result in a large and significant increase in the student population. In the first five years after the college opening, the student population increases by 450% relative to

[12] In line with our discussion on the spatial distribution of colleges in Section 3.1, the are stark level differences, however. Event regions had only 200 students on average prior to the opening as three districts had another college or university prior to the opening; control districts in turn had almost 2,000 students on average before the event.

control districts. The effect grows to 560% over the first decade after the opening, which reflects the gradual expansion of the new college. To illustrate how much the new student population increases the human capital base of highly skilled workers in the region, we use in column (2) the student population as a share of high-skilled employees in the region. The college opening may thus increase the share of high-skilled workers in the region by more than 50% over the first decade (see column (2) of table 3.B.2). As such, the college openings imply a substantial positive shock to high-skilled labor supply in the local economy.

Yet, a local expansion in the student population does not need to translate into more high-skilled employment in the local economy. If most students leave the region after they finish their college degree to work and live in other regions, a college opening would have little permanent impact on the local skill structure or the local economy more broadly. To see whether the college opening has a lasting effect on the human capital base of the local workforce, we estimate the event study from equation (3.1) where the outcome is now full-time employment of college graduates between the ages of 20 and 29 (in logs). Figure 3.B.3 shows that the employment of young high-skilled workers moves in parallel in treatment and control districts before the opening of the new college. The relative size of young high-skilled labor starts to diverge between event and control regions three years after the college opening, when the first cohort graduates and enters the labor market. The difference widens until about seven years after the opening and then levels off. The timing of the increase in local high-skilled employment supports our identification assumption of no differential shock or growth in the demand or supply of high-skilled workers in treatment districts relative to control districts *prior* to the opening.

The employment of young high-skilled has increased by a sizable 13% in event districts six to eleven years after the opening (see column (3) of table 3.B.2).[13]

[13] The new colleges have on average 900 students five years after the opening, which would be equivalent to a 32% increase in high-skilled employment in the pre-event year. Suppose all of the 900 students have graduated eight years after the opening. In that case, more than 50% of the student population show up in local social security employment. This calculation does not include graduates who take up a job in the treatment district outside the social security system as self-employed or civil servant. It also does excludes graduates and dropouts who leave the region to find a job inside or outside the social security system in another district.

Overall, the evidence on the student population and high-skilled employment shows that the college opening results indeed in a permanent increases in the human capital base of the local economy.

3.3.2 Regional Employment and Wages

We just showed that the college opening generated a permanent growth in the local supply of high-skilled labor. This permanent growth in the human capital base might benefit the local economy through several channels. First of all, the growth creates a thicker labor market for skilled workers and thus reduces the search and hiring costs for employers and employees. This reduction in costs could raise wages and employment of high-skilled workers. In addition, highly skilled workers might have positive spillover effects on older skilled workers or less skilled workers. The availability of high-skilled labor could also encourage firms to locate in the regions or invest in capital and technology thus raising labor productivity for all workers.

We might see few effects on regional employment or wages, in turn, if young high-skilled workers simply replace older or less-skilled workers with few spillovers on the productivity of those workers. On the labor supply side, an influx of young college graduates could induce some workers to leave the region or labor market: older workers with more elastic labor supply might leave full-time employment for early retirements schemes, for instance (Dustmann et al., 2016). If the new college graduates simply replace older or less-skilled workers or these workers drop out of the labor market in response, we would observe that a college opening reduces the employment of other skill or age groups with few net effects on total employment.

To trace the average effect of a college opening on the local economy, we estimate variants of our model in equation (3.2) where the dependent variables are total employment in the local economy and employment for less- and highly-skilled workers (in logs). The first three columns in table 3.B.3 reveal that a college opening has few effects on regional employment. Neither does total employment increase (column (1) of table 3.B.3), nor is there an effect on the employment of high-skilled (columns (2)) or less-skilled workers (column (3) of table 3.B.3). While the coefficients for high-skilled employment overall are positive they do not reach statistical significance at conventional levels.

The same pattern can be seen in figure 3.B.4 where we show the year-by-year estimates based on equation (3.1). The increase in overall high-skilled employment in figure 3.B.4 mirrors the growth in young high-skilled employment in figure 3.B.3: employment rates between event and control regions start to diverge from the opening year until about seven years later and then level off. Note that the growth in overall high-skilled employment is much smaller than the growth in young college graduates, which indicates that college graduates do replace some older high-skilled workers close to retirement.[14] Also, there is little movement in the employment share of less skilled workers neither before nor after the opening of the college in the region. Hence, we do not see a simple upskilling of the workforce where employers just replace less-skilled workers with new college graduates.

Rather than employment effects, we might observe an effect on local wages. We would expect that the large increase in the supply of young, high-skilled workers at least temporarily reduces high-skilled wages. High-skilled wages might increase in the long run if firms invest in technology or capital that complement high-skilled labor. Yet, the growth in the high-skilled workforce might also increase the wages of all workers in the presence of productivity or knowledge spillovers or if attracting high-productive firms to the area.

To investigate local wage effects, we re-estimate our model in equation (3.2) where the dependent variables are now the log mean daily wages of full-time workers. As for employment, we find no effect on the average wage in the local economy compared to control regions (see column (4) of table 3.B.3). One reason for the absence of a wage effect could be that wages of high-skilled workers decline while wages of groups that are complementary in the production process increase. Yet, columns (5) and (6) of table 3.B.3 suggest no offsetting effects. The coefficients

[14] Splitting employment by age and skill together, we find no differential growth in employment between treatment and control regions. The only exception is, of course, young, high-skilled employment, for which the college opening has a large positive effect as shown in column (3) of table 3.B.2. We also find no population adjustments in the district. Hence, it is not the case that the new college graduates make older workers leave the region or the labor market.

on wages for high- and less-skilled workers are both positive in the post-period but fail to reach statistical significance.[15]

Our results imply that the new college had no effect on local wages of high-skilled or low-skilled workers. These findings stand in sharp contrast to two recent empirical studies that report a positive effect of a skilled supply shock on high-skilled wages: Beerli et al. (2021) show that high-skilled commuters increased the wages of high-skilled natives. Even closer to us, Carneiro et al. (forthcoming) find that the opening of new colleges in Norway raised high-skilled wages in the region. The authors attribute this to the fact that employers invested purposefully in skill-biased technology after the college openings. Yet, we are not the only study where a growth in high-skilled workers has had little effect on the skill premium (see Beaudry and Green (2003) for evidence from Germany and, more recently, Blundell et al. (2022) for the UK). How can we explain these very different results? Economic theory suggests that the response of the skill premium to a large shift in relative supply depends on the degree of substitutability between different types of workers, the type of technological change and potential investments in capital. We investigate these relationships in the next section.

3.3.3 Worker Substitutability, Technological Change and Capital

Despite the sizable growth in young high-skilled workers in the local economy, we observe neither an initial downward pressure on their wage, nor an increase later on.[16] If nothing else changed, we would expect a (temporary) decline in the local college premium after the new graduates enter the local labor market. One potential explanation for the absence of a short-run decline in the skill premium in response to the supply shock is that workers of different age and skill groups are

[15] An analysis of wages differentiated by skill and age groups yields a very similar pattern. We find no downward pressure on wages of young high-skilled workers though their supply expands substantially following the college opening. Similarly, wages of older high-skilled workers or young less-skilled workers who might be substitutes or complements to young high-skilled workers remain unchanged as well.

[16] We will not observe a downward pressure on nominal wages if wages are downward rigid. Yet, we should still observe a relative decline in high-skilled wages compared to the skill group, for instance.

very good substitutes. In the longer-run, firms might not respond to the growth in high-skilled labor by adjusting their production technologies, for instance. Such technology shifts could be due to profit-maximizing innovators' endogenous choice of research direction (Acemoglu, 1998) or producers' selection of an optimal production technology from a given pool of alternatives (Beaudry and Green, 2005; Blundell et al., 2022). A third explanation for the absence of a change in the skill premium could be that firms invested in physical capital instead, which raised the productivity of all workers with little impact on the skill premium (Beaudry and Green, 2003).

To investigate what our findings imply about the substitutability between workers, the nature of technological change and potential capital responses, Appendix A outlines a theoretical framework linking local skill supplies to local wages, technology and capital. We start from a regional CES production function where the local good is produced by labor of different skill and age groups as in Card and Lemieux (2001). We further allow for factor-augmenting technological change that is potentially skill-biased. For the empirical implementation, we then relate local relative wages to relative shifts in local labor supply for different age and skill groups and other control variables.

In the first step, we estimate the elasticity of substitution across age groups using high-skilled wages of young workers relative to older workers (see equation (3.A.3) in Appendix A). As OLS estimates are likely biased due to reverse causality or omitted variables, we use the college opening as an instrument for the relative supply shift of young, high-skilled workers. All the estimations use the matched sample of treatment and control districts restricted to the period four years before and eleven years after the college opening. As shown in column (3) of table 3.B.2 and figure 3.B.3, the first stage is strong in the post-event period.

Appendix table 3.C.3 reports estimates of the age premium for three alternative samples: Column (1) compares young, high-skilled workers to all prime-aged and older workers with a college degree. As the substitutability might vary across age groups, we consider the relative wages of younger to older workers with a college degree and the relative wages of younger to prime-aged workers with a college degree separately in columns (2) and (3). Appendix table 3.C.3 shows that the OLS estimates (reported in Panel A) are small and positive suggesting that high-skilled workers belonging to different age groups are actually complements rather than substitutes. The IV estimates (shown in Panel B) are negative indicating some substitutability across age groups but the estimates fail to reach statistical significance. For both estimation approaches, the estimates do not change much whether we use older or prime-aged workers as comparison. Within the CES production function framework, the elasticity of substitution across age groups is the inverse of the reported coefficients. The IV estimates thus indicates a substitution elasticity of around six, which would explain why we do not find an effect of the college opening on the wages of prime-aged or older skilled workers.

It is important to stress that we identify the response of *local* relative wages to a *local* relative supply shift, which will be different than the elasticity of substitution estimated from a national production function. A large relative shift in local supply might not affect relative wages locally if many workers leave the local labor market or drop out of the labor force, for instance. As such, we should find a smaller elasticity of substitution than based on a national production function.[17]

In the next step, we estimate the substitutability across skill groups by relating the college premium to the relative supply of skilled workers and additional controls. Specifically, we estimate a version of equation (3.A.4) where we proxy skill-biased technological change by a linear trend that is allowed to vary in the pre- and post-event period (see Appendix A for details). The results using all age groups are reported in column (4); the estimates using only the skill premium of young workers are shown in column (5) of appendix table 3.C.3.

The OLS estimates for relative supply are again positive and statistically significant irrespective of which age group we use in the estimation (Panel A). The IV estimates, in contrast, are negative but only reach statistical significance in the full sample (Panel B). The sizable difference between OLS and IV estimates indicates

[17]For Germany, OLS estimates based on a national production function range from around seven to twenty indicating a strong substitutability across age groups (Brüll and Gathmann, 2020; Fitzenberger and Kohn, 2006; Glitz and Wissmann, 2017).

some omitted variable or reverse causality issue biasing the OLS estimate upward. Based on the IV estimates, high- and less-skilled workers are (weak) substitutes with an elasticity of substitution of around 0.5. These estimates imply that a positive shock to high-skilled supply will have only a small impact on the skill premium, which is in line with the results we saw in table 3.B.3.

Interestingly, we find little support for the type of factor-augmenting technological change that favors skilled workers, which has been the focus in much of the literature on supply shifts and the skill premium. The coefficients on the linear trend variable in both the OLS and IV specifications are zero or negative and typically not statistically significant. Hence, if anything, technological change is either skill-neutral or even tends to reduce the college premium. We find a similar pattern if we restrict the specification to a common linear trend or allow for more flexible quadratic trends to capture the skill bias of technological change.

The weak substitutability across skill groups and absence of skill-biased technological change might be an artifact of the particular production function we used. The CES production function restricts technological change to be of factor-augmenting form and we have abstracted from adjustments in other inputs, especially physical capital. Yet, periods of rapid technological change might not simply raise the productivity of some production factor, but generate disruption through the arrival of new forms of production, which also require new modes of organization (see the discussion in Beaudry and Green, 2003; Blundell et al., 2022). Moreover, plants might invest in additional physical capital, which could affect the skill premium if it is not factor neutral. Beaudry and Green (2003) have long pointed out that the skill premium has risen much faster in the U.S. than in Germany, in part because employers in Germany made sizable investments in physical capital.

To assess the role of capital and broader technological advances, we extend the regional production function in two ways. First, we include physical capital as an additional input that is traded in a national market. Second, we allow for flexible shifts in technology by using a first-order linear approximation to an arbitrary production function following Blundell et al. (2022). We then obtain an augmented specification for the skill premium (see equation (3.A.6) in Appendix A). The skill premium now depends on relative skill supplies as before, the regional capital intensity and technological shifts. We estimate this augmented equation by relating the college premium to relative skill supplies, capital intensity, general TFP growth and other technological change, which we proxy by a linear time trend that may differ before and after the college openings.

The results of this more general specification are shown in column (6) for all age groups and column (7) for young workers in appendix table 3.C.3. Focusing on the IV estimates, the coefficients on relative skill supplies again suggest some substitutability between skill groups. The implied substitution elasticity in this more general specification ranges from 0.7 to 1, which is quite in a similar range than the elasticities obtained from the CES production function.

General TFP growth seems to favor high-skilled workers in the full sample (column (6)) but not among young workers (column (7)). Instead, wages of young, high-skilled workers were pushed up by investments in capital intensity indicating that employers in event regions not only hired the new college graduates but also invested in additional physical capital. The linear time trend, which proxies for factor-augmenting technological change, is again zero or negative prior to the opening but turns slightly positive in the post-event period. Hence, the shock to high-skilled labor supply did not reduce the skill premium because of some offsetting capital investments and a modest technological change favoring young college graduates. Yet, the additional capital and shifts in technology are not sizable or long-lasting enough to push up high-skilled or average wages in the even regions in the long-run.

Overall, we show find strong evidence for weak substitutability across age and skill groups at the local level – independently of how we specify the labor demand side. That implies that even large shifts in relative supply have only a small or no impact on the age or college premium, which is consistent with the absence of local wage effects – overall and across age or skill groups. We see only limited evidence for skill-biased technological change; instead, capital investments seem to play an important role for the absorption of the new college graduates in the local economy. Before we investigate the channels of the local absorption, we first demonstrate the robustness of our main results.

3.3.4 Specification Checks

The dependent variable for all robustness checks is the log of young, high-skilled employment in the region. For ease of comparison, we report our baseline estimates from table 3.B.3 again in column (1) of appendix table 3.C.4.

A first-order concern with our event study approach in equation (3.2) is that the effects may be caused by differential trends in outcomes between treated and matched control regions or by unobserved confounders. As we demonstrated in table 3.B.1, our matching approach does balance not only levels of employment and wages, but also their growth rate. Nevertheless, potential confounding factors could be other changes in regional policies like local firm subsidies; or unobserved demand and productivity shocks that are unequally distributed across regions due to differences in the underlying local industry structure. We address these concerns in several ways. First, we control for unobserved region-specific shocks by including a separate linear (in column (2)) or quadratic trend (in column (3)) for each region. The results are very similar to the baseline in column (1). Our results might also be accounted for by negative demand or supply shocks in control regions. The observed increase in young, high-skilled employment might then be explained by an employment decline in control districts rather than an expansion in treated regions. To check for confounding shifts in control regions, we run a placebo test: we match control regions to other untreated regions using the same matching approach and variables as for our main analysis (see Section 3.2.2). We then re-estimate our baseline model in equation (3.2). Column (4) of appendix table 3.C.4 shows negative coefficients – but they are never close to be statistically significant.

We next test the robustness of our results to the heterogeneity of the treatment estimates. The event study approach in equation (3.1) and the aggregate approach in equation (3.2) may not identify an average treatment effect on the treated (ATT) if effects are heterogeneous across regions or change dynamically over time. An emerging literature has demonstrated that the standard event study design with staggered adoption only identifies some weighted average of the treatment effects with weights that could be negative. To check for this possibility, we implement the re-weighting estimator of Sun and Abraham (2020) that identifies a proper ATT even in the presence of heterogeneous or dynamic treatment effects in column (5). The results are almost identical to the baseline estimates, which is perhaps less surprising if one considers that the college openings all occur within a decade and hence, a relatively short time window.

The matching approach might also give rise to concerns. Recall that we match on variables just before the college opening (in $\tau = -1$). If there are anticipation effects, the local economy might have started to adapt to the college opening, e.g., by changing its industry structure, which are part of our matching set. As an alternative, we repeat the matching using all variables measured three years before the opening ($\tau = -3$) instead. Column (6) shows statistically somewhat weaker effects for the post-period. Furthermore, our two-step estimation strategy controls for time-invariant regional confounders but does rely on a common trends assumption conditional on our matching step. An alternative approach is to match not only on regional observable characteristics, but also on pre-treatment outcomes like employment directly. The synthetic control approach matches the pre-event trend in employment for each event region, and does therefore not rely on a common trend assumption between treated and (synthetic) control regions. Yet, the approach requires pre-event confounders to be mean independent of the outcome. Column (7) in appendix table 3.C.4 shows similar, albeit noisier estimates.

Finally, college openings might not only affect the district in which the new college is located, but rather spills over to neighboring districts. Such spillover effects could be important if the college draws in young people from the surrounding districts for studying and later supplies neighboring regions with high-skilled workers. It is important to note that our matching approach does not match event regions to any neighboring district. As such, our baseline results are not contaminated by control regions experiencing positive (or negative) spillover effects. Yet, in the presence of spillovers to neighboring regions, our baseline estimates would underestimate the true benefits of the college opening. To check for spillover effects in the broader region, we define the broader region as all neighboring districts sharing a border with an event (or control) district but exclude the event resp. control district. We then re-estimate equation (3.2) where the dependent variable is now log employment of young, high-skilled workers in all neighboring districts of an event resp. control region. The results in column (8) of appendix table 3.C.4 show only a muted response beyond the event district. As such, college openings seem to have mostly local effects with few spillovers beyond district boundaries.

Overall then, all specification checks demonstrate that our results are very robust to alternative matching approaches and assumptions in the first stage as well as alternative estimators and specifications in the second stage.

3.4 Local Absorption of the New College Graduates

Our evidence so far points to few local employment and wage effects in the local economy on average, while the production function approach indicates that employers have responded to the additional skilled labor by additional capital investments. We now explore where the new high-skilled labor is absorbed in the local economy and who benefited from the college opening.

3.4.1 Employment in Manufacturing vs. Services

We first ask whether the college graduates are mainly employed in manufacturing or services. Thus, we re-estimate equation (3.2) where the dependent variables are total employment (columns (1) and (4)), high-skilled (columns (2) and (5)) and less-skilled (columns (3) and (6)) employment in the specified sector. Table 3.B.4 reveals that high-skilled labor in manufacturing increases by 16.6 percentage points, while there is little change in high-skilled employment in the service sector. We also checked whether we find any effect if we split the service sector into low- and high-skilled services. Yet, we do not find an employment effect for high-skilled services like insurance, consulting or finance sector, likely because these are typically located in the large urban centers, which are not in our sample.

Does the growth in high-skilled employment in manufacturing imply that less-skilled workers are replaced and hence, their employment declines after the college opening? Column (3) in table 3.B.4 shows a negative, albeit not significant coefficient for less-skilled employment. Hence, most college graduates find a job in local manufacturing where they replace some less-skilled manufacturing workers. As for total employment in the region, we find no effect on total employment in manufacturing (see column (1)) or services (see column (4)).[18]

To isolate the employment effect of college openings further, we split manufacturing into high-tech manufacturing – containing the chemical industry, machinery, electrical and transport equipment and some smaller manufacturing industries (see also Beerli et al., 2021) – and into other manufacturing. The results show that high-skilled employment increases mostly in high-tech manufacturing. Here, high-skilled employment grows by 28.7 percentage points within the first decade after the college opening (see column (7) in table 3.B.4). Interestingly, the long-run estimate for less-skilled workers in high-tech manufacturing is also positive (see column (8)) indicating that high-tech firms do not simply hire college graduates to replace less-skilled workers.

Given the sizable employment responses, we next turn to the development of wages in high-tech manufacturing. High-skilled wages do not grow in high-tech manufacturing after the college opening (see column (9) in table 3.B.4). The absence of a wage effect suggests that high-tech firms could satisfy their demand for skilled workers and hence, solve any labor shortages that might have existed in the region before the college opening.

The final column of 3.B.4 further shows that the additional hiring of skilled workers in high-tech manufacturing firms increased the wages of less-skilled workers

[18]Beyond these broader categories, we see few shifts in the employment composition of detailed industries. If the local supply shock is primarily absorbed through inter-regional trade, regions with a college opening would expand their skill-intensive industries and export more goods that use high-skilled workers more intensively compared to control regions. We do not see such a pattern, however.

by 5.6 percentage points within ten years after the college opening relative to high-tech firms in control regions. The wage effect across skill groups indicates that high-skilled workers raise the productivity of less-skilled workers. In turn, Section 3.A.1 showed that additional capital investments and technological change did favor high-skilled workers, but not less skilled labor.

Overall then, high-tech firms in manufacturing definitely benefited from the college opening as the new graduates helped to fill vacancies. Another group that benefits from the college opening are less-skilled workers who work in high-tech manufacturing and benefit from sizable wage growth.

3.4.2 Professionals, Managers and Engineers

We next investigate which occupations the new college graduates occupy in the local economy. To do so, we restrict the sample to workers in the age range from 20-29. We then re-estimate our model in equation (3.2) for four broad occupational groups: unskilled manual labor, skilled manual labor, semi-professionals and professionals. The results in columns (1)-(4) of table 3.B.5 reveal that new college graduates primarily enter professional occupations. In the first five years, professional employment among young workers increases by 6.2 percentage points; the cumulative effect five to ten years after the college opening has grown to 14.4 percentage points (see column (4)).

Are these additional professional jobs created in manufacturing or in the service sector? To answer this question, we study professional employment in each sector separately: columns (5) and (6) of table 3.B.5 show that professional employment grows in manufacturing by 15.4 percentage points, while we see little change in the service sector. Finally, we investigate which professional jobs the new college graduates obtain in manufacturing. In particular, we distinguish between managerial positions and engineering jobs. Columns (7) shows that a college opening does not change the number of managerial positions in the manufacturing sector. Instead, the growth of professional jobs in the manufacturing sector is concentrated among engineering jobs.

3.4.3 STEM Colleges and Innovation

As the college opening has the biggest effect in the high-tech sector with a strong growth in engineering jobs, the benefits of the college for the local economy might depend on the subjects the college offers. In particular, the effects might differ for colleges that focus more on STEM subjects and related areas and colleges that focus more on teaching business or architecture, for instance. A first reason why the type of subjects offered matters is that employers often face shortage in the supply of high-skilled occupations like engineers or other STEM subjects. Another reason could be that STEM workers create positive externalities through R&D activities, production or knowledge spillovers. To investigate the role of different study programs, we classify the colleges into those with a STEM focus and those with no STEM focus. We then re-estimate equation (3.2) by letting the coefficients on the college opening differ by college type. Panel A of in table 3.B.6 shows results for STEM colleges, Panel B for non-STEM colleges.

We find that total employment increases in regions where a STEM college was opened (see column (1)). Within the first five years, local employment increases by 11.6 percentage points in regions with a STEM college compared to the control regions, while there is no effect on local employment for non-STEM colleges. Even ten years after the college opening, the growth in local employment is still much stronger for STEM colleges (+17.5 percentage points) than for non-STEM colleges (+9.1 percentage points, which is not statistically significant). Column (2) further shows that the differential effect on total employment is not driven by a higher growth in high-skilled employment. For both STEM and non-STEM colleges, we observe little to no growth in total high-skilled employment in the region.

We next turn to the question whether STEM and non-STEM colleges also have a differential impact on local wages relative to the control regions. If graduates from STEM colleges create positive spillovers on other workers, for instance, the wages of workers whose productivity has increased through knowledge spillovers or whose demand increased because of production complementarities with STEM workers should increase. Columns (3) to (5) of table 3.B.6 show the results for the wages of less-skilled workers by age.

STEM colleges raise the wages for both young and older less-skilled workers by 1.1-1.6 percentage points relative to the control regions; we see no effect in regions that opened a college without a STEM focus. These spillover effects indicate that engineers and other graduates from STEM Colleges increase the productivity of less-skilled workers in the region.

Turning to high-skilled wages in columns (6)-(8), we find positive coefficients for all age groups in regions with STEM colleges (Panel A) but they do not reach statistical significance. For non-STEM colleges, Panel B shows that the increase in supply of young high-skilled workers reduces young, high-skilled wages initially by 2.4 percentage points (see column (6) in Panel B). That indicates that the supply of new college graduates exceeded the demand for high-skilled labor for regions with new colleges not focused on STEM subjects initially. In the long-run, the wage effect of young, skilled workers reverts to zero as employers absorb the new supply of college graduates into their workforce. We also find no statistically significant effects on high-skilled wages for prime-aged or older workers (see columns (7) and (8)).

Given the strong employment effect for engineering jobs, one might wonder whether the new engineers encourage innovation in the regions with a college opening – which could be another reason we find positive wage effects on less-skilled workers. To investigate this, we use information on patents filed in the region based on the address of the organization or person named in the patent documents as patent holder. We then use our baseline model in equation (3.2) to see whether patent activity goes up after a college opening; and whether there is any differential effect for regions in which a college with STEM focus was opened. Appendix table 3.C.6 shows the results: columns (1)-(3) report unweighted results, while columns (4)-(6) uses local employment in the pre-event period as weight. The dependent variable in all specifications is the total number of patents filed in the region. We first use the full sample (in column (1) and (4)) and then run separate analyses for STEM colleges (in columns (2) and (5)) and non-STEM colleges (in columns (3) and (6)). The main conclusion here is that we do not find any effect on patenting activity in regions with a college opening irrespective of whether the college has a STEM focus or not.

Overall then, the new graduates did not increase the innovative capacity of a region, to the extent that this can be measured by patents. The absence on any effect for patent also confirms that the new college did not themselves spur innovation through research and start-ups in the local area (as discussed in Section 3.1).

3.4.4 University Employment and Local Demand

We further investigate whether the new college itself has a direct effect on the local economy. The college creates jobs for different skill groups: on the one hand, the college needs faculty and practitioners to teach in the college. In addition, the college requires an administration and services like cleaning and maintenance. We focus first on total employment in higher education, which includes colleges and universities located in the region. As only two regions have a university prior to the college opening, employment in higher education will largely reflect additional jobs in the new colleges. Based on equation (3.2), appendix table 3.C.5 shows that there is no employment growth in higher education prior to the opening of the college in our sample. Within five years after the opening of the new colleges, total employment in higher education institutions has increased by almost 60 percentage points. In the long-run, employment in higher education has grown by 93.2 percentage points relative to control regions (see column (1)). Do colleges as employers hire more high- or rather less-skilled workers? Columns (2) and (3) show that both employment categories increase; yet, relative to the local workforce, the employment growth contributes only a small share to overall employment. Hence, it is not surprising that the impact of the college employment on the local labor market is limited.

A region might further benefit from the opening because the student body and university staff bring in additional income and raise the demand for local goods and services. The right-hand side of appendix table 3.C.5 shows that there is no effect on employment in the non-tradable sector: overall employment (in column (4)), high-skilled employment (in column (5)) and less-skilled employment (in column (6)) do not change after the college opening compared to control regions.[19] The absence of any effect on local goods and services suggests modest local multipliers in our setting, which is not surprising given that the new technical colleges are relatively small.

3.5 Are College Openings an Effective Place-Based Policy?

The evidence thus far indicates that the opening of a college improves the local human capital base. While we find no boost for the overall economy in the region, the new college graduates were quickly absorbed into the labor market, esp. by high-tech manufacturing firms to alleviate local skill shortages. The growth in high-skilled employment created additional gains for less-skilled workers in high-tech manufacturing. We also saw that STEM colleges generate larger local benefits than colleges without a STEM-focus. In contrast, job creation by the college or through local multiplier effects are modest, in part because the colleges are themselves quite small.

Yet, do these results imply that opening new institutions of higher education could be a successful place-based policy tool – a kind of a golden bullet to help the economic development of declining regions or areas facing substantial structural transformation? We saw in Section 3.1 above that the regions in which a college opened were neither the most economically advanced urban centers nor the most backward regions. Do all of the regions benefit from a college opening or just some? It could well be that college openings only generate benefits for the region in economically dynamic regions; or, it could be that less vibrant regions benefit the most from a boost to the local human capital base. Knowing the answer to this question is important to guide policy-makers: if there are no benefits to economically declining regions, it makes more sense to open colleges in regions with more favorable economic conditions, for instance.

[19] If we zoom in on the hospitality industry, we do see a transitory employment effect of about five percentage points, which is no longer statistically significant in the long-run (not reported).

3.5.1 Estimation Approach

To answer the question which regions benefit, we split our sample into regions with a more dynamic labor market and those with less dynamic labor markets. We define regions as 'dynamic' if they had above median employment growth (15 percentage points on average) in the late 1970s. We define those regions with below median employment growth (just 3 percentage points) as 'stagnant'.

Comparing the two regions along observable characteristics in appendix table 3.C.7 suggests otherwise few observable differences: Dynamic and stagnant regions are very similar in their industry structure or the age and skill structure of their workforce. Further, dynamic and stagnant regions do not differ in wage levels or wage growth prior to the college opening. Reflecting their more favorable local labor market, dynamic regions have a somewhat lower unemployment rate (by 2 percentage points) than stagnant regions (see columns (1) and (3)). Dynamic and stagnant regions also have a very similar share of high-tech manufacturing industry and are equally likely to obtain a college with a STEM focus. Based on observables, it is a-priori not clear which of the two regions might benefit the most from a college opening.

To explore the heterogeneity, we augment our baseline approach in equation (3.2) by interacting the event variable with indicators whether the region was economically dynamic or stagnant. Specifically, we estimate variants of the following model:

$$\begin{aligned} Y_{r\tau t} = \beta_B^{Dyn} Before_\tau * Opening_r^{Dyn} + \beta_T^{Dyn} Trans_\tau * Opening_r^{Dyn} \\ + \beta_P^{Dyn} Post_\tau * Opening_r^{Dyn} + \gamma_B^{Stag} Before_\tau * Opening_r^{Stag} \\ + \gamma_T^{Stag} Trans_\tau * Opening_r^{Stag} + \gamma_P^{Stag} Post_\tau * Opening_r^{Stag} \\ + \theta_\tau^{Dyn} + \theta_\tau^{Stag} + \delta_t + \alpha_r + \varepsilon_{r\tau t} \end{aligned} \tag{3.3}$$

where $Y_{r\tau t}$ are again regional employment or wages. As before, the subscripts r, t and τ denote region, calendar time and the period relative to the event, respectively. $Opening_r^{Dyn}$ is now an indicator equal to one if a college opened in a region with a dynamic labor market r. The variable $Opening_r^{Dyn}$ is zero for control districts or event districts that belong to a region with a stagnant labor market. Similarly,

$Opening_r^{Stag}$ is an indicator equal to one if a district r had a college opening but is located in a stagnant local labor market; and zero for control regions and event regions with a dynamic labor market. We focus on this pooled estimation to increase statistical power as we only have twenty event regions. As in previous sections, $Before_\tau$ denotes the period before the actual opening ($-7 \leq tau \leq -1$), $Transition_\tau$ is an indicator equal to one in the years between the opening and the graduation of the first cohort of students ($0 \leq \tau \leq 5$) and $Post_\tau$ characterizes the long-run adjustment ($6 \leq \tau \leq 11$).

The specification in equation (3.3) allows dynamic and stagnant regions to have different employment trends before and after the college opening by allowing separate event time fixed effects for dynamic (θ_τ^{Dyn}) and stagnant local labor markets (θ_τ^{Stag}). All other variables are defined as before. Compared to estimating the equation (3.2) separately for dynamic and stagnant regions, we only require calendar time to affect both types of regions similarly; yet, we do allow for differences in employment levels and trends in the pre- and post-period. The estimates β_B^{Dyn} and γ_B^{Stag} show whether dynamic (or stagnating) regions that had a college opening exhibit a differential pre-trend than their respective control regions. The main coefficients of interest are β_T^{Dyn}, β_P^{Dyn}, γ_T^{Stag} and γ_P^{Stag}, which trace whether dynamic or stagnant regions have better labor market outcomes than their respective control regions in the medium- and long-run.

3.5.2 College Openings in Dynamic and Stagnant Labor Markets

We start out with the evolution of young high-skilled employment. Figure 3.B.6 plots the coefficients for the years before and after the college opening separately for dynamic and stagnant areas. The figure shows clear differences: the employment of new college graduates increases steadily in economically dynamic regions (the red line) compared to their control regions. The picture looks different for stagnant regions: here, there seems to be only a temporary growth in the employment of young high-skilled workers (the blue line). Five years after the opening, stagnant regions see a reversal of the growth and young high-skilled employment reverts toward it pre-opening level. The corresponding estimates (reported in column (1) of table 3.B.7) indicate that the growth in local employment of young high-skilled increases more than twice as much in dynamic regions (by 16.4 percentage points, see Panel A, column (1)) relative to control regions in comparison with the growth in stagnant regions (see Panel B, column (1)). These numbers indicate that some regions seem to be able to absorb and benefit from the new supply of high-skilled workers more than others.

One reason for the observed difference is that the regions vary in the size of the supply shock because the colleges opened in economically less vibrant regions are smaller than in more dynamic areas, for instance. Columns (2) and (3) of table 3.B.7 reveal, however, that dynamic and stagnant regions do not differ in their growth of the student population. In both regions, the student population increases by a factor of five to six (in column (2)) and the potential share of high-skilled in the region by 52-54 percentage points (see column (3)). Hence, the size of the new colleges is very similar in the two regions. The differential effect is also not explained by differences in the subject mix – in vibrant regions, four colleges opened with a STEM focus, while there were five colleges with STEM focus in stagnant regions.[20]

Are regions with a stagnating economy not able to retain the high-skilled workforce trained in the region after the college opening? To investigate this, we turn to population flows between districts, which provide a better picture of actual mobility than employment flows in the social security records. Outflows in the social security records are only recorded if a person was working in a job subject to social security contributions in the district in one year and is then observed in an employment relationship subject to social security contributions in a different district in a later year, for instance. Students might never be registered in the social security records of an event district if they leave the area before or immediately after graduation and obtain their first job elsewhere. And yet, the data on population flows also have limitations. Students might not register with the municipality but

[20]The differential dynamic is also not the explained by the fact that two out of the ten dynamic regions have already had a university or college. We still find the same patterns if we drop the two districts with a university or college prior to the opening.

remain registered at their parents' residence instead; or, they might never move to the district hosting the college but commute from neighboring districts.[21] In both cases, a student entering (or leaving) the college is not registered as an inflow (or outflow). The undercounting of students in their district of study introduces measurement error in the population flows. While there is no reason the issue should be worse in dynamic than in stagnant regions, it will make our estimates less precise.

We then re-estimate equation (3.3) where the dependent variables are now population inflows and outflows of individuals between the ages of 18 and 30 years. The estimates provide suggestive support that less vibrant regions experience sizable outflows of young people from the region after the college opening (see table 3.B.7, column (5) in Panel B), which is not compensated by increased inflows (see column (4) in Panel B). We observe the opposite pattern in dynamic local economies: here, we observe a reduced outflows of young people after the college opening (see table 3.B.7, column (5) in Panel A). It seems that less vibrant regions initially attract young people to the area, but cannot retain them in the region later on. In contrast, vibrant regions are able to keep most graduates trained in the new college resulting in a sustained and permanent boost to the local human capital base.

The sustained growth of the high-skilled workforce in economically dynamic regions could also impact how the local economy benefits from the college opening. An abundant supply of high-skilled workers might attract new firms to the region, for instance. To explore this, we investigate employment in incumbent firms versus employment in new firms, i.e., that were founded within the past five years. Finally, we also look at the number of firms exiting the market. We then re-estimate equation (3.3) where the dependent variables are now total or high-skilled employment in incumbent and entering firms. Table 3.B.8 shows several interesting patterns: in stagnant regions, the additional high-skilled workers are absorbed by incumbent firms (see column (2) in Panel B). As total employment in incumbent firms does not change (see column (1) in Panel B), this implies that incumbents up-skill their workforce by replacing less-skilled workers with high-skilled workers. Dynamic

[21] Unfortunately, data on commuting flows is not available for the time period we analyze.

regions exhibit a very different pattern: here, the additional high-skilled workers are employed in new firms (see column (4) in Panel A). New firms create additional jobs and hence, increase their total employment (see column (3) in Panel A). There is no effect on firm exit.

These findings show that it matters in which economic environment a college is opened. Our evidence will disappoint policy-makers who wish to use the founding of new colleges to turn around the fate of economically declining regions. The results in this paper show that such a policy is not beneficial for stagnant regions in the longer-run. At most, existing firms will benefit from the more abundant supply of high-skilled workers. Beyond that, the benefits for economically stagnant regions are modest and temporary. The situation looks much brighter for the more dynamic regions. New colleges, even if they are relatively small, can attract new firms to the region and create additional jobs – though these do not translate into employment or wage growth at the local labor market level (see appendix table 3.C.8).[22]

Our results are consistent with insights from EU structural funds that showed net gains in income and investment per capita only in regions with a favorable human capital base and governance structure (Becker et al., 2013). While the policies implemented differ, one lesson to take away is that economically backward and struggling regions benefit little from such place-based policies. Such policies do work best in regions that have a sound economic basis and governance structure, which enables them to gain from the proposed subsidy or infrastructure investment.

3.6 Discussion and Conclusion

We exploit the opening of new technical colleges in Germany during the 1980s and early 1990s to study their impact on the local economy. Our empirical strategy combines matching with an event study approach to find suitable control regions for the event districts with a college opening. We have four main findings. First,

[22]Appendix table 3.C.8 shows that total employment increases by 3.3 percentage points (see column (1) in Panel B), while local employment in stagnating regions actually declines by 5 percentage points (see column (1) in Panel A), largely because less-skilled employment declines (see column (3) in Panel A). Yet, the standard errors also show that the estimates are noisy as none of them is statistically significant at conventional levels.

the opening of a college substantially increases the local student population. The opening further spurs a sizable growth in the share of high-skilled employment by 13% in the district where the college is located. Second, we find no effect on overall employment or wages suggesting few growth or productivity gains for the local economy on average. Third, we see that the new college graduates get absorbed by high-tech firms in manufacturing, mostly in engineering jobs, which also pushed up wages of less skilled workers in high-tech manufacturing. In contrast, we find few changes in the high- and low-skilled service sector. Finally, we document that the impact of a college opening depends on the local labor market condition: in dynamic labor markets, a college opening results in a sustained growth in the high-skilled workforce, which encourages firm entry and job creation. In less vibrant labor markets, in turn, the high-skilled share grows less and college graduates are largely absorbed by incumbent firms.

The insights from our analysis carry important lessons for regional policy. College openings are an effective strategy to increase the skill level of the regional workforce. Yet, college opening need not benefit the whole local economy by raising average employment or wages, which could provide further justification for subsidizing tertiary education. Instead, the benefits are locally concentrated in some industries and professions with larger benefits from STEM colleges. And economically backward regions are less able to reap the benefits from a college opening than regions with a more vibrant labor market.

3.A The Substitutability between Labor Inputs, Technological Change and Capital

3.A.1 Regional Production Function

Each region r produces a single output using labor. We abstract from capital and interregional trade for now. In Section 3.A.4 below, we extend our framework to allow for capital adjustments; we also analyze the potential role of inter-regional trade in Section 3.4.1. Following Card and Lemieux (2001) and Card and Lewis (2007), aggregate labor is specified as a nested CES production function using two types of labor: Less-skill ed and skilled, which we denote by L_{rt} and S_{rt} respectively.

$$Y_{rt} = \left(\theta_{lt} L_{rt}^{\psi} + \theta_{st} S_{rt}^{\psi}\right)^{\frac{1}{\psi}}, \tag{3.A.1}$$

where $-\inf < \psi \leq 1$ is a function of the elasticity of substitution between college and non-college labor ($\psi = 1 - 1/\sigma_E$). The shares of different types of labor are represented by θ_{lt} and θ_{st}, which may evolve over time due to technological change, for instance. In equation (3.A.1), labor-augmenting technical change of high-skilled workers would result in an increase in θ_{st}; and similarly for less-skilled workers (θ_{lt}). Skill-biased technical change would imply an increase in $(\theta_{st}/\theta_{lt})$ over time.

Labor in each skill group consists of a CES-aggregate of the labor of workers in $j = 3$ different age groups:

$$L_{rt} = \left[\sum_{j=1}^{3} \alpha_{lj} L_{jrt}^{\phi}\right]^{\frac{1}{\phi}}, \qquad S_{rt} = \left[\sum_{j=1}^{3} \alpha_{sj} S_{jrt}^{\phi}\right]^{\frac{1}{\phi}} \tag{3.A.2}$$

where $-\infty < \phi \leq 1$ depends on the elasticity of substitution between age groups ($\phi = 1 - 1/\sigma_A$) and the α_j's are relative efficiency parameters for age-group j, which we take as constant over time.

The specifications in equations (3.A.1) and (3.A.2) imply that the elasticity of substitution between workers of different ages are the same for all skill groups and the elasticity of substitution of different skill groups is the same for all age groups. We test these restrictions with our data below. Workers of different ages are gross substitutes when $\sigma_A > 1$, and gross complements when $\sigma_A < 1$. If different age groups within a given skill level are perfect substitutes, $\sigma_A \to \infty$.

3.A.2 Elasticity of Substitution across Age Groups

We first study the age premium among skilled workers. Assuming perfect competition and hence, that labor is paid its marginal product, this age premium is defined as:

$$ln\left(\frac{w^S_{jrt}}{w^S_{1rt}}\right) = ln\left(\frac{\alpha_{sj}}{\alpha_{s1}}\right) - \left(\frac{1}{\sigma_A}\right) ln\left(\frac{S_{jrt}}{S_{1rt}}\right) + \varepsilon_{jrt} \tag{3.A.3}$$

where $j = 2, 3$ are prime-aged and older workers and $j = 1$ is the group of young, high-skilled workers. Equation (3.A.3) shows that rising supply of young workers may increase the wages of older high-skilled workers relative to young, high-skilled workers by $\frac{1}{\sigma_A} dlog\left(\frac{S_{jrt}}{S_{1rt}}\right)$. If young and older workers are perfect substitutes in their skill group $\sigma_A \to \infty$, however, there is no effect on the age premium.

College openings provide us with an exogenous shock to the local supply of young, high-skilled labor (S_{1rt}) in the treatment regions. Using the college openings, we can then identify the elasticity of substitution across age groups (σ_A) and the ratio of efficiency parameters ($\frac{\alpha_{sj}}{\alpha_{s1}}$) from a regression of relative wages of young and older high-skilled workers on age-specific relative supplies, age group dummies, matched pair fixed effects and year fixed effects interacted with a treatment indicator. The latter allow for differential trends in relative wages between treatment and control regions over time. Matched pair fixed effects ensure that we compare each event region to its respective control. The age group dummies absorb differences in the relative efficiency parameters ($\frac{\alpha_{sj}}{\alpha_{s1}}$). By normalizing one of the α_{sj}, we obtain estimates of the efficiency parameters of the two other age groups.

It is important to stress that we identify the response of *local* relative wages to a *local* relative supply shift, which will be different than the elasticity of substitution estimated within the framework of a national production function. A large relative shift in local supply might not affect relative wages if many workers leave the local labor market or drop out of the labor force. In that case, our estimates of the elasticity of substitution should be larger in absolute terms because we do not take into account movements across local labor markets.

The results of estimation equation (3.A.3) are shown in columns (1)-(3) of appendix table 3.C.3. The top panel shows results using OLS, while the bottom panel reports instrumental variable estimates where we use the college opening as instrument for the relative supply of young, high-skilled workers. Column (1) shows the pooled estimates for all three age groups together, while columns (2) and (3) report the results if we estimate it for older and prime-aged workers (relative to young workers) separately. The OLS estimates are positive, rather than negative indicating that among high-skilled workers, different age groups might be complements rather than substitutes. The IV estimates are negative but not statistically significant. The IV estimates would imply an elasticity of from $\widehat{\sigma_A} = 3$ to $\widehat{\sigma_A} = 10$. The fact that skilled workers of different ages are not good substitutes locally provides on rationale why we find few effects on the college opening on the employment or wages of older, high-skilled workers.

3.A.3 Elasticity of Substitution across Skill Groups

We next turn to the impact on the college premium and its determinants. Using our regional production function, relative wages of skilled to less-skilled workers, in age group j are defined as:

$$ln\left(\frac{w^S_{jrt}}{w^L_{jrt}}\right) = ln\left(\frac{\alpha_{sj}}{\alpha_{lj}}\right) + ln\left(\frac{\theta_{st}}{\theta_{lt}}\right) - \left(\frac{1}{\sigma_E}\right) ln\left(\frac{S_{rt}}{L_{rt}}\right) - \left(\frac{1}{\sigma_A}\right)\left[\ln\left(\frac{S_{jrt}}{L_{jrt}}\right) - ln\left(\frac{S_{rt}}{L_{rt}}\right)\right] \tag{3.A.4}$$

Equation (3.A.4) shows how changes in relative supply may affect the college premium. Abstracting from technological change for now, the direct effect of an increase in the supply of young, skilled workers on the skill premium of young workers is given by:

$$dln\left(\frac{w^S_{1rt}}{w^L_{1rt}}\right) = -\frac{1}{\sigma_A} dln\left(\frac{S_{1rt}}{L_{1rt}}\right) + \left(\frac{1}{\sigma_A} - \frac{1}{\sigma_E}\right) dln\left(\frac{S_{rt}}{L_{rt}}\right)$$

In the absence of technological change, a college opening reduces the college premium for young workers if the supply of skilled workers increases a lot and the elasticities of substitution across age and education groups are small. Given that our estimated elasticities across age groups are large, the college premium will only decline if there is large increase in the supply of skilled labor and a low elasticity of substitution across skill groups. The indirect effect of an increase in the supply of young, skilled workers on the skill premium of older workers is in turn:

$$dln\left(\frac{w^S_{jrt}}{w^L_{jrt}}\right) = \left(\frac{1}{\sigma_A} - \frac{1}{\sigma_E}\right) dln\left(\frac{S_{rt}}{L_{rt}}\right)$$

In the absence of technological change, the skill premium among older workers can only increase if $|\sigma_E| > |\sigma_A|$. Also note that if young and old workers are perfect substitutes within skill groups (i.e., $\sigma_A \rightarrow \infty$), then the effect on the skill premium for young workers is identical to the effect for older workers with its size determined by σ_E.If there is skill-biased technological change, i.e., an increase in $\frac{\theta_{st}}{\theta_{lt}}$ in event regions over time, the skill premium could rise for all age groups in response to the positive supply shock of skilled workers.[23]

We cannot estimate the parameters of equation (3.A.4) directly because of the unobservable aggregate skill supplies (S_{rt} and L_{rt}), which depend on the age-specific skill supplies and parameters of the model (σ_A, α_{sj} and α_{lj}). However, we can use our estimates from equation (3.A.3) above to calculate aggregate skill supplies. If the efficiency parameters are the same across skill groups ($\alpha_{sj} = \alpha_{lj}$ for all j), estimation of equation (3.A.3) is sufficient to identify the efficiency parameters for all age and skill groups. The evidence indicates that the restricted version of uniform efficiency parameters is valid. Using our estimates of $\widehat{\sigma_A}$ and $\widehat{\alpha_{.j}}$, we can calculate the aggregate skill supplies (S_{rt} and L_{rt}) from equation (3.A.2).

It remains to make an assumption how the rate of relative skill-biased technological change ($\theta_l t$ and $\theta_s t$) evolves over time. The college opening will affect the relative skill supplies of young, skilled workers; moreover, it may raise the rate of skill-biased technological change in the event regions relative to control regions. Therefore, the college opening cannot serve as an instrument to separate the two channels. We start out with the assumption the skill-biased technological change follows a linear trend t according to $\frac{\theta_{st}}{\theta_{lt}} = d_0 + d_1 * t + d_2 * t^{After} + \vartheta_{rt}$ (Katz and Murphy, 1992). This assumption implies that we can capture the impact

[23]The college premium might also increase with relative changes in the age-specific efficiency units ($\frac{\alpha_{sj}}{\alpha_{lj}}$). We follow the literature here and abstract from that possibility here.

of skill-biased technological change with linear time trends where the trends are allowed to differ after the college opening.

Columns (4) and (5) in appendix table 3.C.3 report the estimates for the skill premium allowing for factor-augmenting technological change. Column (4) uses the whole sample, while column (5) uses the skill premium for young workers only. The OLS estimates in Panel A are again positive, while the IV estimates in Panel B are negative. Based on the IV estimates, workers belonging to different skill groups are substitutes with an elasticity of substitution of 0.5. The linear trend is negative or zero suggesting little change in technology in response to the supply shock. Based on these estimates, we would expect that the relative supply shock reduces the skill premium a little.

3.A.4 Allowing for Adjustments in Capital and Flexible Technological Change

An alternative explanation why a sizable growth in skilled workers has little impact on relative wages is because of adjustments in capital (Beaudry and Green, 2003; Beaudry et al., 2010), which we have abstracted from for now. If capital and skilled labor are complements (or become more complementary due to technological change), a decline in the skill premium induces firms to invest more in capital. In addition, the specification of the production function in equation (3.A.1) above is restrictive as it only allows for factor-augmenting technological change but not for disruptive shifts in technology and modes of organization as seen during the IT revolution, for instance (Blundell et al., 2022).

To investigate the role of capital and flexible technological change, we extend the regional production function according to $Y_{rt} = F(\theta_{st}S_{rt}, \theta_{lt}L_{rt}, K_t)$ where K_t denotes physical capital, which we assume to be traded in a national market. As before, S_{rt} and L_{rt} are the regional supplies of skilled and less-skilled workers, which combines the labor of different age groups according to the CES production function in equation (3.A.2) above. θ_{st} and θ_{lt} again denote the skilled and less-skilled labor-enhancing technological change parameters. Unlike in equation 3.A.1, we do not specify a specific functional form for the production function in order to nest alternative models of technological change (Blundell et al., 2022). Rather, we only require the production function F(.,.,.) to be constant returns to scale.

Assuming competitive labor markets where each skill group is paid their marginal product, the first-order condition for skilled workers in age group j is:

$$w^S_{jrt} = \theta_{st} F_1 \left(\frac{\theta_s S_{rt}}{K_t}, \frac{\theta_l L_{rt}}{K_t}, 1\right) \alpha_{sj} \left(-\frac{1}{\sigma_A}\right) \left(\frac{S_{jrt}}{S_{rt}}\right)$$

where F_1 denotes the first derivative with respect to skilled labor. Similarly, for less-skilled workers we obtain:

$$w^L_{rt} = \theta_{lt} F_2 \left(\frac{\theta_{st} S_{rt}}{K_t}, \frac{\theta_{lt} L_{rt}}{K_t}, 1\right) \alpha_{lj} \left(-\frac{1}{\sigma_A}\right) \left(\frac{L_{jrt}}{L_{rt}}\right)$$

where F_2 denotes the first derivative with respect to less-skilled labor. Using a first-order linear approximation, we can write the wages of skilled workers belonging to age group j as:

$$lnw^S_{jrt} \approx ln\alpha_{sj} - \frac{1}{\sigma_A} ln\left(\frac{S_{jrt}}{S_{rt}}\right) + ln\theta_{st} + \delta_1 ln\left(\frac{\theta_{st} S_{rt}}{\theta_{lt} L_{rt}}\right) + \delta_2 ln\left(\frac{K_t}{\theta_{st} S_{rt}}\right)$$

and likewise for less-skilled workers:

$$lnw^L_{jrt} \approx ln\alpha_{lj} - \frac{1}{\sigma_A} ln\left(\frac{L_{jrt}}{L_{rt}}\right) + ln\theta_{lt} + \gamma_1 ln\left(\frac{\theta_{st} S_{rt}}{\theta_{lt} L_{rt}}\right) + \gamma_2 ln\left(\frac{K_t}{\theta_{lt} L_t}\right)$$

Concavity of the production function implies $\delta_1 - \delta_2 \leq 0$ and $\gamma_1 + \gamma_2 \geq 0$. The skill premium can then be written as:

$$\begin{aligned} ln\left(\frac{w^S_{jrt}}{w^L_{jrt}}\right) = ln\left(\frac{\alpha_{sj}}{\alpha_{lj}}\right) + (\delta_1 - \delta_2 - \gamma_1) ln\left(\frac{S_{rt}}{L_{rt}}\right) - \frac{1}{\sigma_A}\left[ln\left(\frac{S_{jrt}}{L_{jrt}}\right) - ln\left(\frac{S_{rt}}{L_{rt}}\right)\right] \\ +(\gamma_2 - \delta_2) ln\theta_{lt} + (1 + \delta_1 - \delta_2 - \gamma_1) ln\left(\frac{\theta_{st}}{\theta_{lt}}\right) + (\delta_2 - \gamma_2)\left(\frac{K_t}{L_{rt}}\right) \end{aligned} \tag{3.A.5}$$

The first term denotes the evolution of the age-specific efficiency parameters of the skill group, which we will capture by age group fixed effects. The second term captures how the relative supply of skilled workers affects the skill premium, which under standard assumptions about technological change is governed by the elasticity of substitution across skill groups. The third term characterizes how the difference between the age-specific relative skill supply from the overall relative skill supply affects the skill premium, which is governed by the elasticity of substitution across age groups. The last three terms represent the impact of general productivity increases, additional skill-biased technological change and adjustments in capital intensity, respectively.

To estimate equation (3.A.5), we again need to make some assumption about the underlying productivity parameters θ_{st} and θ_{lt} as they are unobserved. We follow Beaudry and Green (2005) and Blundell et al. (2022) by assuming that general productivity increases are captured by TFP growth and include a linear time trend to capture any exogenous skill-biased productivity shifts.[24]

The specification for the skill premium, which extends equation (3.A.4), is now defined as:

$$\begin{aligned} ln\left(\frac{w^S_{jrt}}{w^L_{jrt}}\right) &= d_0 ln\left(\frac{\alpha_{sj}}{\alpha_{lj}}\right) + d_1 ln\left(\frac{S_{rt}}{L_{rt}}\right) + d_2\left[ln\left(\frac{S_{jrt}}{L_{jrt}}\right) - ln\left(\frac{S_{rt}}{L_{rt}}\right)\right] \\ &+ d_3\left(\frac{lnTFP_t}{sh_{lt}+sh_{st}}\right) + d_4 ln\left(\frac{K_t}{L_{rt}}\right) + d_5 * t + d_6 * t^{After} + \varepsilon_{jrt} \end{aligned} \tag{3.A.6}$$

where sh_{lt} and sh_{st} denote skill shares. The fact that TFP growth and capital have no regional subscript reflects the assumptions that technologies are available in all regions and that the capital market is national.

The coefficient d_1 in equation (3.A.6) reflects the elasticity of substitution between skilled and less-skilled labor just like in the CES production function framework, i.e., $d_1 = -\frac{1}{\sigma_E}$. Skill-biased technological change might show up as $d_3 > 0$ (if general TFP growth favors skilled workers) or as $d_5 > 0$ (if there is a positive time trend in the skill premium just as in the CES production function framework). If the estimated elasticity of substitution between skill groups is large when controlling for technological change then the effect of a shift in relative skill supply on the relative wage should be large and positive. The results of estimating equation (3.A.6) are shown in columns (6) and (7) of appendix table 3.C.3.

[24] Following the productivity literature, TFP growth can be approximated as $TFP_t \approx sh_{ht} ln\theta_{st} + sh_{lt} ln\theta_{lt}$.

963.A. The Substitutability between Labor Inputs, Technological Change and Capital

The IV results again indicate a moderate degree of substitutability between high- and less-skilled workers. The elasticity of substitution ranges between 0.7 and 1 is therefore close to the elasticity based on the CES production function. As before, the linear trend is negative but turns slightly positive after the college opening – indicating that technological change exhibits some skill bias. General productivity gains favor high-skilled workers in the full sample (see column (6)) though the effect is not significant for the sample of young workers (see column (7)). Finally, young high-skilled workers benefit from additional capital investments. Overall then, the college opening would have decreased the skill premium but was offset by both capital investments and a moderate skill biased technological change.

3.B Main Tables and Figures

Table 3.B.1: Comparison of Treatment and Control Regions

	Treated Districts	Other West German Districts	Control Districts	Difference Treated vs. All Districts		Difference Treated vs. Control Districts	
				Coeff.	S.E.	Coeff.	S.E.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Matched Characteristics							
University or College in Region	0.100	0.354	0.100	-0.254***	0.073	0.000	0.098
Population per square km	349.19	1090.58	357.92	-741.390***	157.60	-8.73	93.05
Industry structure:							
Share in Agriculture and Fishing	0.010	0.008	0.008	0.001	0.001	0.001	0.001
Energy and Mining	0.026	0.018	0.012	0.007	0.015	0.014	0.016
Manufacturing	0.437	0.355	0.489	0.083**	0.032	-0.052	0.043
Construction	0.079	0.067	0.067	0.013***	0.004	0.012*	0.007
Trade	0.128	0.142	0.121	-0.014**	0.006	0.007	0.008
Transport and Communication	0.036	0.051	0.034	-0.015***	0.004	0.002	0.004
Financial services	0.026	0.043	0.026	-0.017***	0.004	0.000	0.002
Other Services	0.179	0.228	0.169	-0.049***	0.012	0.010	0.016
Non-Profit Organizations	0.019	0.025	0.018	-0.006*	0.003	0.002	0.003
Public Administration	0.060	0.063	0.056	-0.003	0.004	0.004	0.006
Panel B: Characteristics Not Matched							
Age structure:							
Ages 20-29	0.328	0.305	0.326	0.022**	0.009	0.001	0.011
Ages 30-44	0.406	0.406	0.392	0.000	0.009	0.014	0.012
Ages 45-59	0.267	0.289	0.282	-0.022**	0.009	-0.015	0.011
Education:							
High-skilled Share	0.056	0.092	0.062	-0.036***	0.006	-0.006	0.007
Less-skilled Share	0.944	0.908	0.938	0.036***	0.006	0.006	0.007
Other Regional characteristics:							
Unemployment Rate	8.201	8.558	7.298	-0.357	0.778	0.903	1.044
Employment	49,518	51,373	44,670	-1,855	6,797	4,848	9,686
Population (in Thousands)	282.52	396.34	252.35	-113.82	83.84	30.17	68.41
Employment Growth (past 5 years)	0.076	0.045	0.090	0.031	0.020	-0.015	0.030
Average Daily Wage	118.38	125.95	122.77	-7.572***	2.319	-4.39	4.03
Wage Growth (past 5 years)	0.050	0.052	0.071	-0.001	0.012	-0.020	0.023

Notes: The table compares characteristics between treatment regions (column (1)), the average region in West Germany (column (2)) and the matched control region (column (3)) in the pre-event period (t=-1). Columns (4)-(5) show the difference (and standard errors) between the treatment and average West German districts in observable characteristics, while columns (7) and (8) show the differences in observables between treatment and matched control regions. With the exception of employment and the indicator for the presence of another university in the district, all observations are weighted by district employment in the year just before the event ($\tau = -1$). The matched control regions are identified through Mahalanobis matching using the variables shown in Panel A in the pre-event period (t=-1). Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.2: College Openings, the Student Population and the Supply of Young High-skilled Workers

	Number of Students (1)	Student to High-skilled Workers (2)	Young, High-skilled Employment (3)
Opening*Before	-0.173	-0.017	0.018
	(0.196)	(0.038)	(0.034)
Opening*Transition	4.500***	0.211***	0.055
	(0.523)	(0.059)	(0.034)
Opening*Post	5.603***	0.526***	0.128**
	(0.644)	(0.117)	(0.062)
Event Time Fixed Effects	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes
Observations	760	760	760
R^2	0.917	0.924	0.979

Notes: The table shows estimates from the event study in equation (3.2). The dependent variables are the log number of students in the district (in column (1)), the ratio of students to the number of full-time workers with a university degree (in column (2)), and the log of full-time employees between the ages of 20 and 29 with a university degree (in column (3)). The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.3: Effects of a College Opening on Regional Employment and Wages

	Employment			Wages		
	Total	High-skilled	Less-skilled	Total	High-skilled	Less-skilled
	(1)	(2)	(3)	(4)	(5)	(6)
Opening*Before	0.002	-0.006	0.001	0.001	0.002	0.001
	(0.014)	(0.016)	(0.014)	(0.005)	(0.010)	(0.004)
Opening*Transition	-0.001	0.017	-0.002	0.006	-0.001	0.005
	(0.013)	(0.018)	(0.014)	(0.005)	(0.006)	(0.005)
Opening*Post	-0.008	0.027	-0.009	0.016	0.011	0.014
	(0.030)	(0.042)	(0.031)	(0.011)	(0.013)	(0.009)
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	760	760	760	760	760	760
R^2	0.992	0.991	0.991	0.961	0.951	0.960

Notes: The table reports estimates of the event study in equation (3.2). The dependent variables are the log of total employment (column (1)), employment of workers with a university or college degree (in column (2)) and employment of workers without a college or university degree (in column (3)). Column (4)-(6) shows estimates for the mean log daily wage overall (column (4)), high-skilled workers (column (5)) and less-skilled workers (column (6)) as dependent variables. The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.4: College Openings and the Industry Composition of Employment

	Manufacturing			Services			High-Tech Manufacturing			
	Total Employment (1)	High-skilled Employment (2)	Less-skilled Employment (3)	Total Employment (4)	High-skilled Employment (5)	Less-skilled Employment (6)	High-skilled Employment (7)	Less-skilled Employment (8)	High-skilled Wages (9)	Less-skilled Wages (10)
Opening*Before	0.016	0.031	-0.002	-0.013	0.001	-0.019	0.050	0.020	-0.014	-0.015
	(0.021)	(0.034)	(0.022)	(0.017)	(0.034)	(0.017)	(0.077)	(0.037)	(0.020)	(0.009)
Opening*Transition	0.012	0.068	0.015	0.005	-0.014	0.007	0.091	0.002	-0.008	0.003
	(0.021)	(0.041)	(0.020)	(0.015)	(0.037)	(0.014)	(0.057)	(0.042)	(0.017)	(0.009)
Opening*Post	-0.026	0.166*	-0.023	0.011	-0.031	0.021	0.287*	0.043	0.037	0.056***
	(0.059)	(0.095)	(0.060)	(0.035)	(0.069)	(0.034)	(0.147)	(0.101)	(0.033)	(0.019)
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	760	760	760	760	760	760	760	760	760	760
R-squared	0.985	0.969	0.982	0.988	0.978	0.989	0.946	0.973	0.843	0.843

Notes: The table reports estimates of equation (3.2) where the dependent variable is the logarithm of employment (total, high-skilled and less-skilled) in manufacturing and services. High-tech manufacturing contains the chemical industry, machinery, electrical and transport equipment and some smaller manufacturing industries (see also Beerli et al., 2021). The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.5: College Openings and Employment by Occupation

	Employment of Young Workers				**Professional Employment by Industry**		**Professional Occupations in Manufacturing**	
	Unskilled Manual (1)	Skilled Manual (2)	Semi-Professional (3)	Profess-ional (4)	Manufac-turing (5)	Services (6)	Managers (7)	Engineers (8)
Opening*Before	-0.117	0.003	0.011	0.026	0.000	-0.011	-0.027	0.043
	(0.083)	(0.016)	(0.024)	(0.032)	(0.034)	(0.025)	(0.031)	(0.063)
Opening*Transition	0.064	0.016	0.027	0.062*	0.068*	0.012	0.000	0.082
	(0.065)	(0.015)	(0.024)	(0.034)	(0.036)	(0.030)	(0.038)	(0.060)
Opening*Post	0.059	0.057	0.050	0.144***	0.154*	0.034	-0.012	0.252**
	(0.114)	(0.044)	(0.050)	(0.053)	(0.082)	(0.055)	(0.072)	(0.124)
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	760	760	760	760	760	760	760	760
R-squared	0.892	0.985	0.984	0.979	0.977	0.981	0.980	0.961

Notes: The table reports the estimates in equation (3.2). The dependent variable is the logarithm of employment in different occupations. Column (1) shows results for employees in unskilled manual occupations (such as agricultural workers or unskilled industrial workers); column (2) for skilled manual workers (such as mechanics or hairdressers); column (3) for semi-professional occupations (such as nurses or technicians); and column (4) for professionals (such as doctors, engineers or managers). The dependent variables are the logarithm of employment of professionals in manufacturing firms (column (5)) and in all other industries (column (6)). The last two columns separate professionals in manufacturing into employment of managers (column (7)) and employment of engineers (column (8)). The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.6: Employment and Wage Effects in STEM and Non-STEM Colleges

	High-Skilled Employment		Less-skilled Wages			High-skilled Wages		
	Age 20-29	Total	Age 20-29	Age 30-44	Age 45-60	Age 20-29	Age 30-44	Age 45-59
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Colleges with STEM Focus								
Opening*Before	0.008	0.012	-0.001	-0.007	0.001	0.019	-0.013	-0.014
	(0.055)	(0.021)	(0.006)	(0.007)	(0.006)	(0.018)	(0.013)	(0.023)
Opening*Transition	0.116***	0.014	0.011*	0.003	0.007	0.015	0.010	-0.008
	(0.033)	(0.020)	(0.005)	(0.005)	(0.004)	(0.011)	(0.008)	(0.017)
Opening*Post	0.175**	0.017	0.018	0.007	0.016*	0.027	0.013	0.014
	(0.062)	(0.046)	(0.013)	(0.012)	(0.008)	(0.017)	(0.017)	(0.035)
Observations	342	342	342	342	342	342	342	342
R-squared	0.985	0.992	0.975	0.966	0.983	0.912	0.969	0.952
Panel B: Colleges without STEM Focus								
Opening*Before	0.026	-0.005	-0.000	-0.002	-0.001	0.013	-0.011	0.021
	(0.043)	(0.020)	(0.006)	(0.004)	(0.005)	(0.012)	(0.009)	(0.019)
Opening*Transition	0.006	-0.013	0.002	0.002	0.001	-0.024**	-0.008	-0.001
	(0.050)	(0.018)	(0.008)	(0.007)	(0.009)	(0.010)	(0.009)	(0.017)
Opening*Post	0.091	-0.029	0.007	0.008	0.012	-0.006	-0.010	0.013
	(0.101)	(0.041)	(0.013)	(0.009)	(0.014)	(0.020)	(0.017)	(0.031)
Observations	418	418	418	418	418	418	418	418
R-squared	0.973	0.992	0.922	0.965	0.961	0.857	0.941	0.921
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports the estimates of equation (3.2) when the sample is split between colleges with a STEM-focus and those without. The dependent variables are the log of employment of workers with a college or university degree aged 20 to 29, log employment of all workers with a college or university degree. Column (3)-(8) use log wages of workers by qualification and age group as dependent variables. The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.7: Students, High-Skilled Employment and Population in Dynamic and Stagnant Regions

	Number of Students (1)	Students to High-Skilled Workers (2)	Young, High-skilled Employment (3)	Population 18-30 Inflows p.c. (4)	Population 18-30 Outflows p.c. (5)
			Panel A: Dynamic Regions		
Opening*Before	-0.298	-0.072	0.027	-0.105	-0.076
	(0.275)	(0.060)	(0.042)	(0.066)	(0.087)
Opening*Transition	4.191***	0.223***	0.081**	-0.057	-0.070
	(0.775)	(0.059)	(0.038)	(0.074)	(0.069)
Opening*Post	5.379***	0.540***	0.169**	-0.082	-0.125
	(0.798)	(0.143)	(0.081)	(0.085)	(0.084)
			Panel B: Stagnant Regions		
Opening*Before	-0.049	0.038	0.009	0.000	-0.015
	(0.271)	(0.045)	(0.051)	(0.061)	(0.044)
Opening*Transition	4.811***	0.200*	0.028	0.032	0.056
	(0.696)	(0.103)	(0.047)	(0.047)	(0.045)
Opening*Post	5.832***	0.515***	0.084	0.035	0.088
	(1.002)	(0.186)	(0.077)	(0.086)	(0.077)
Observations	760	760	760	751	751
R-squared	0.918	0.925	0.981	0.939	0.913
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes
Quadratic Time Trend	No	No	No	Yes	Yes

Notes: The table reports the estimates of heterogeneous treatment effects on employment and the student population according to equation (3.3). Column (1) shows estimates for the effect of a college opening on the log number of students, column (2) for the ratio of students to workers with a college or university degree, and column (3) for workers with a college or university degree aged 20 to 29. The last two columns use the log number of yearly in- and outmigration of individuals aged 18-30 per capita as dependent variable. Panel A contains the estimates for dynamic regions, panel B for stagnant regions where regions are split by the median employment growth between 1975 and 1980. The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.8: Employment in Incumbent and New Firms in Dynamic and Stagnant Regions

	Employment Incumbent Firms		**Employment New Firms**		**Exiting**
	Total	High-skilled	Total	High-skilled	**Firms**
	(1)	(2)	(3)	(4)	(5)
			Panel A: Dynamic Regions		
Opening*Before	-0.034	0.027	0.096	0.184	-0.050
	(0.026)	(0.037)	(0.107)	(0.238)	(0.067)
Opening*Transition	-0.017	-0.067*	-0.000	0.119	-0.055
	(0.024)	(0.039)	(0.114)	(0.263)	(0.081)
Opening*Post	-0.007	-0.118	0.311*	0.345	-0.014
	(0.056)	(0.092)	(0.181)	(0.344)	(0.093)
			Panel B: Stagnant Regions		
Opening*Before	0.028	0.001	-0.067	-0.175	0.000
	(0.021)	(0.032)	(0.080)	(0.188)	(0.049)
Opening*Transition	0.014	0.055*	0.010	-0.100	0.041
	(0.022)	(0.030)	(0.081)	(0.193)	(0.064)
Opening*Post	-0.000	0.131*	-0.156	-0.345	0.047
	(0.048)	(0.075)	(0.142)	(0.266)	(0.067)
Observations	760	760	760	760	760
R-squared	0.992	0.986	0.939	0.879	0.968
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes
Quadratic Time Trend	No	No	No	Yes	Yes

Notes: The table reports the estimates of heterogeneous treatment effects on employment in incumbent firms and new firms as well as the number of exiting firms according to equation (3.3). Entering firms are defined as new firms entering within a five-year window (from t to t-4). Exiting firms are defined as the number of firms closing down in a given calendar year. The dependent variables are full-time employment in incumbent firms (in logs) in column (1) and full-time high-skilled employment (in logs) in incumbent firms in column (2). The dependent variables are full-time employment (in logs) in entering firms in column (3) and high-skilled employment (in logs) in entering firms. In column (5), the dependent variable is the number of firms exiting the region. Panel A contains the estimates for dynamic regions, panel B those for stagnant regions where regions are split by the median employment growth between 1975 and 1980. The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 3.B.1: Geographic Location of Treatment and Control Districts

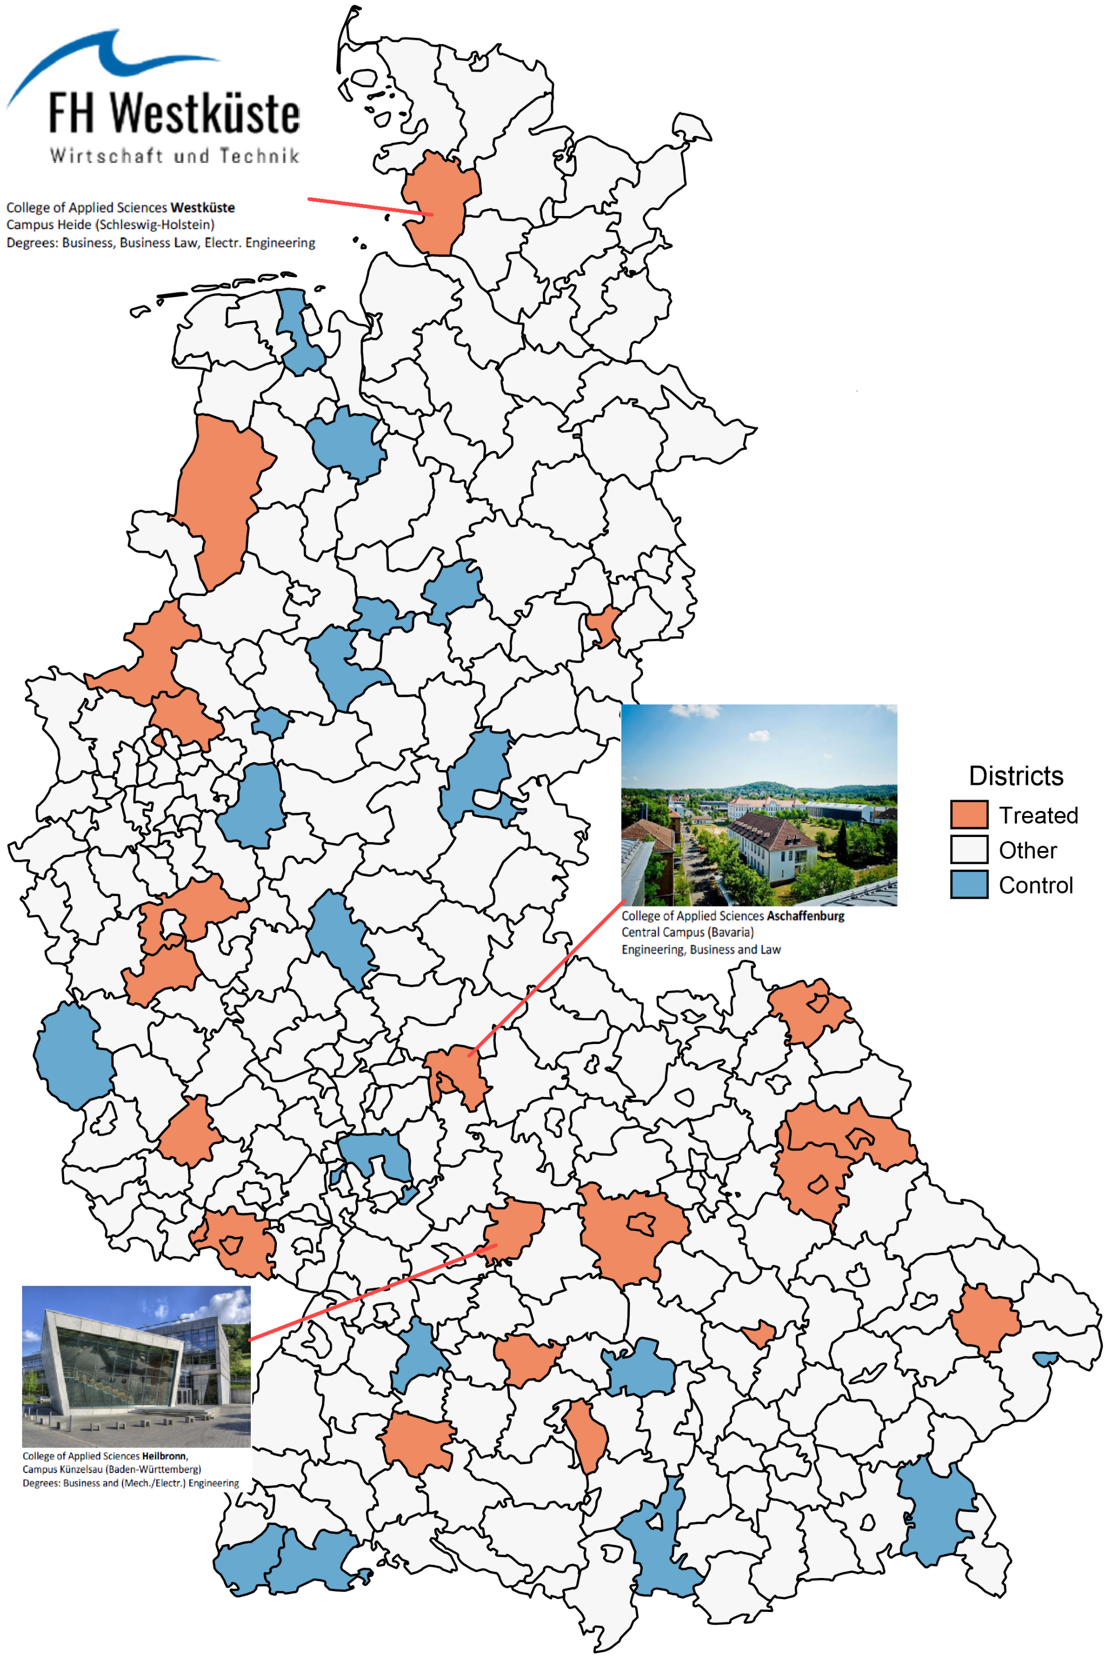

Figure 3.B.2: Number of Students in Treatment and Control Districts

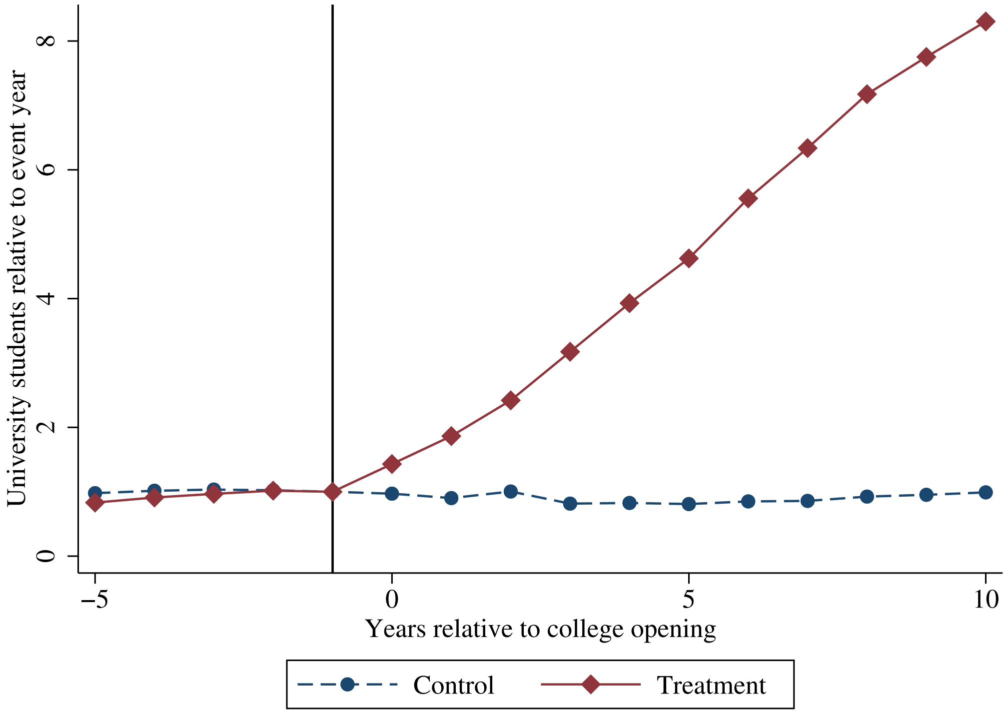

Notes: Average number of full-time students enrolled in universities of treatment and control districts. The year before the college opening is normalized to one.

Figure 3.B.3: College Openings and the Employment of Young High-Skilled Labor

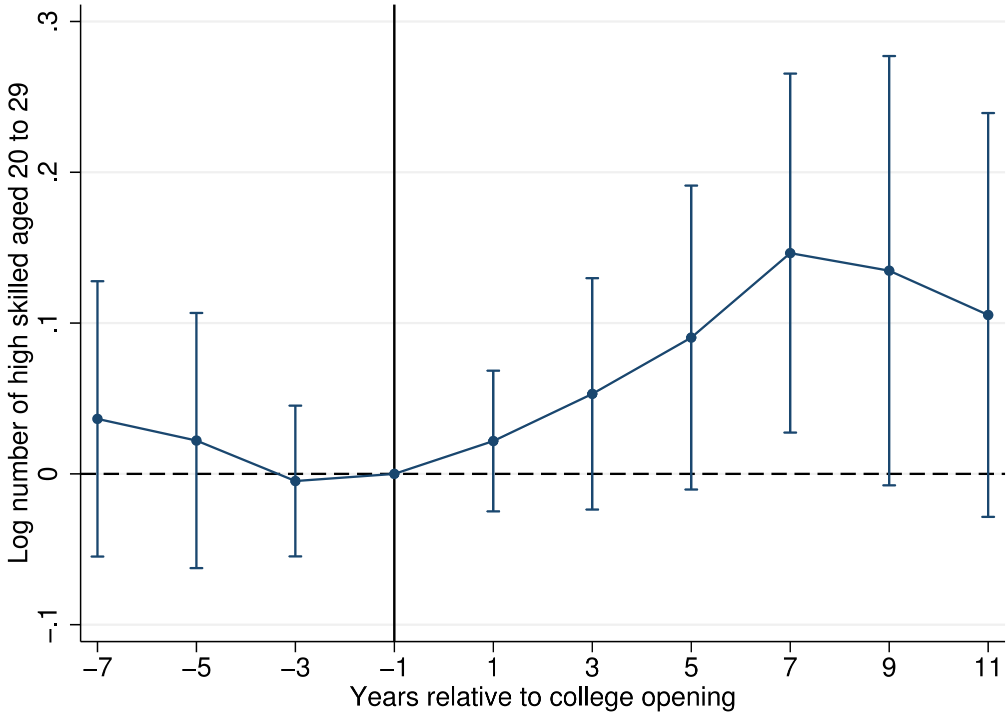

Notes: The figure reports the point estimates and 95% confidence intervals of the regression described in equation (3.1). The dependent variable is the logarithm of full-time employees between the ages of 20 and 29 with a college or university degree. The unit of observation is district-year. Standard errors are clustered at the district level.

Figure 3.B.4: College Openings and Employment of Less- and High-Skilled Workers

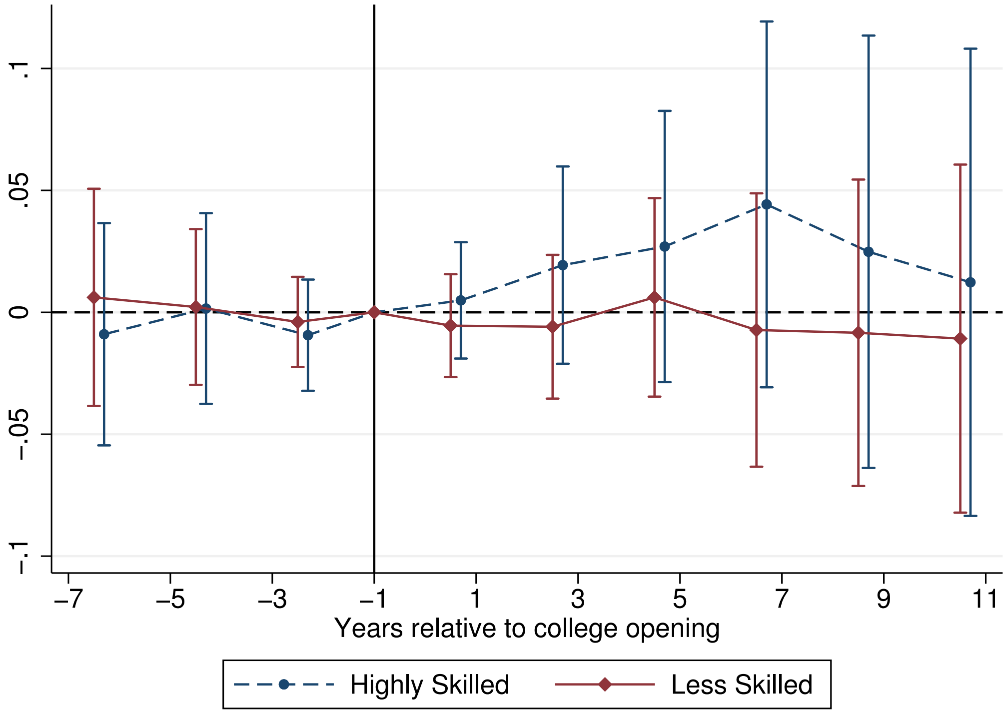

Notes: The figure reports the point estimates and 95% confidence intervals of the regression described in equation (3.1). The dependent variable is the logarithm of full-time employees by skill level. The unit of observation is district-year. Standard errors are clustered at the district level.

Figure 3.B.5: College Openings and High-Skilled Employment by Sector

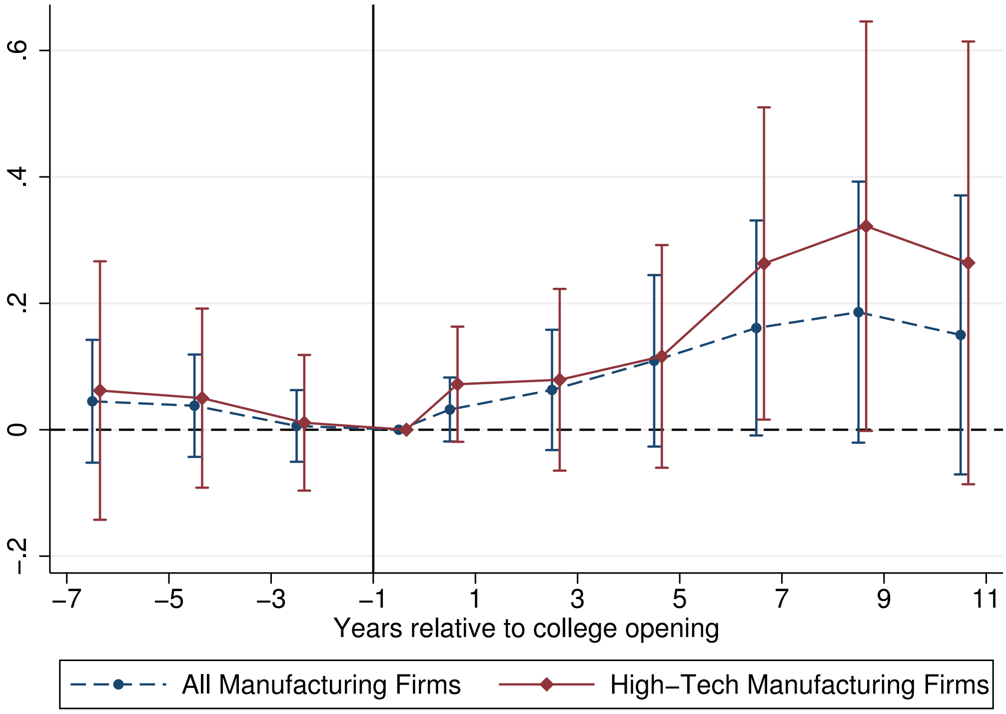

Notes: The figure reports the point estimates and 95% confidence intervals of the regression described in equation (3.1). The dependent variable is the logarithm of full-time employees with a college or university degree in manufacturing firms or in high tech manufacturing firms. The unit of observation is district-year. Standard errors are clustered at the district level.

Figure 3.B.6: College Openings and Young High-Skilled Employment in Dynamic and Stagnant Regions

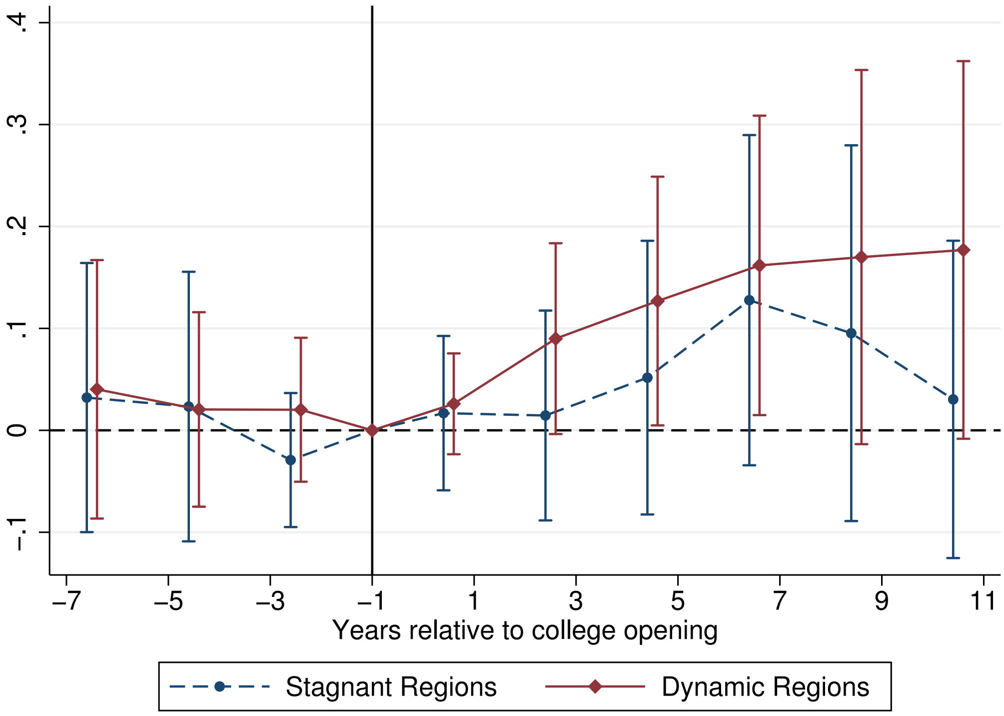

Notes: The figure reports the point estimates and 95% confidence intervals of the regression described in equation (3.3) separately for dynamic and stagnant regions. Regions are split by the median employment growth between 1975 and 1980. The dependent variable is the logarithm of full-time employees aged 20-29 with a college or university degree. The unit of observation is district-year. Standard errors are clustered at the district level.

3.C Additional Tables and Figures

Table 3.C.1: List of Treatment Colleges

College	City	Opening	5 Years after Opening	
		Year	Students	Employees
Hochschule Esslingen, Hochschule Nürtingen-Geislingen	Göppingen, Geislingen	1988	798	80
Hochschule Heilbronn	Künzelsau	1988	438	62
Hochschule Albstadt-Sigmaringen	Albstadt-Ebingen	1988	885	118
Westfälische Hochschule Gelsenkirchen	Bocholt	1992	626	62
FH Westküste	Heide	1993	601	92
FH Braunschweig-Wolfenbüttel	Salzgitter	1993	671	72
Hochschule Kaiserslautern	Zweibrücken	1994	1186	123
Technische Hochschule Ingolstadt	Ingolstadt	1994	863	94
Technische Hochschule Deggendorf	Deggendorf	1994	1121	132
Hochschule für angewandte Wissenschaften Hof	Hof	1994	792	106
FH Neu-Ulm	Neu-Ulm	1994	855	110
Hochschule Osnabrück	Lingen	1995	319	50
Hochschule Bonn-Rhein-Sieg	Sankt Augustin, Rheinbach	1995	1762	138
Westfälische Hochschule	Recklinghausen	1995	797	71
Ostbayerische Technische Hochschule Amberg-Weiden	Amberg	1995	476	108
Ostbayerische Technische Hochschule Amberg-Weiden	Weiden	1995	524	59
FH Aschaffenburg	Aschaffenburg	1995	762	112
Hochschule Trier	Birkenfeld	1996	1148	149
Hochschule für angewandte Wissenschaften Ansbach	Ansbach	1996	915	133
Hochschule Koblenz	Remagen	1998	1650	156

Table 3.C.2: Regional Determinants of College Adoption

	(1)	(2)	(3)	(4)
Tertiary institution in region	-0.004**		-0.004***	-0.004**
	(0.001)		(0.001)	(0.001)
Population (in 1000s) per sq km	-0.003***		-0.003*	-0.003**
	(0.001)		(0.001)	(0.001)
Share in Agriculture and Fishing	-0.192		-0.183	-0.166
	(0.109)		(0.110)	(0.102)
Energy and Mining	-0.014		-0.019	-0.008
	(0.046)		(0.047)	(0.044)
Manufacturing	-0.027		-0.025	-0.016
	(0.042)		(0.041)	(0.038)
Construction	-0.091		-0.089	-0.066
	(0.069)		(0.062)	(0.063)
Trade	-0.034		-0.028	-0.020
	(0.032)		(0.031)	(0.028)
Transport and Communication	-0.008		-0.005	0.007
	(0.046)		(0.034)	(0.039)
Financial services	-0.055		-0.056	-0.044
	(0.057)		(0.058)	(0.059)
Other services	-0.057		-0.055	-0.045
	(0.050)		(0.050)	(0.046)
Non-Profit Organizations	0.010		0.006	0.029
	(0.119)		(0.109)	(0.117)
Highly skilled share		-0.053**	0.006	
		(0.022)	(0.017)	
Share aged 20-29		-0.075	-0.055	
		(0.058)	(0.047)	
Share aged 45-59		-0.049	-0.052	
		(0.041)	(0.043)	
Urban region		-0.001	-0.001	
		(0.004)	(0.003)	
Urban neighboring region		-0.001	-0.001	
		(0.002)	(0.002)	
Population				-0.021
				(0.028)
Average daily wage				-0.029
				(0.040)
Highly skilled employment				0.014***
				(0.004)
Less skilled employment				-0.031
				(0.028)
Year fixed effects	Yes	Yes	Yes	Yes
Federal State FE	Yes	Yes	Yes	Yes
Observations	5088	5088	5088	5088
Adjusted R^2	0.006	0.005	0.005	0.006

Notes: The table reports the estimates of a linear probability model with college opening as the dependent variable for the years 1985-2000. The unit of observation is district-year. All explanatory variables are lagged by one year. Growth variables refer to the growth in characteristics between τ-6 and τ-1. Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.C.3: Estimates of Production Functions and Elasticities

	Age Premium High-Skilled				Skill Premium		
	Pooled	Old/Young	Prime/Young	HS/LS All	HS/LS Young	HS/LS All	HS/LS Young
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: OLS Estimates							
Relative Supply	0.062***	0.062***	0.041***	0.445***	0.811***	0.817***	0.572***
	(0.010)	(0.016)	(0.009)	(0.063)	(0.096)	(0.067)	(0.070)
Time Trend				-0.004*	-0.001	-0.012***	-0.004***
				(0.002)	(0.002)	(0.001)	(0.001)
Time Trend x Post				-0.002	-0.000	0.009***	0.003***
				(0.002)	(0.002)	(0.001)	(0.001)
Total Factor Productivity						0.249***	0.102
						(0.078)	(0.082)
Capital/LS Labor						-0.008***	0.010***
						(0.003)	(0.003)
Observations	1120	560	560	1680	560	1638	546
R Squared	0.854	0.873	0.725	0.883	0.481	0.885	0.459
Panel B: IV Estimates							
Relative Supply	-0.161	-0.339	-0.105	-2.194***	-2.412	-1.454***	-1.073*
	(0.144)	(0.393)	(0.069)	(0.593)	(1.597)	(0.525)	(0.619)
Time Trend				0.000	0.004	-0.005**	0.001
				(0.000)	(0.004)	(0.002)	(0.002)
Time Trend x Post				-0.001	0.003	0.009***	0.003**
				(0.003)	(0.004)	(0.001)	(0.001)
Total Factor Productivity						0.261***	0.107
						(0.101)	(0.115)
Capital/LS Labor						-0.002	0.013***
						(0.004)	(0.005)
Observations	1120	560	560	1680	560	1638	546
F-Statistic (1st Stage)	7.401	1.678	11.580	37.905	5.498	44.961	13.443
Match Pair FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	No	No
Age Group FE	Yes	No	No	Yes	No	Yes	No

Notes: The table reports the estimates of the production function linking relative wages across age or skill groups to relative supplies and controls for technology and capital. The unit of observation is district x year. The results in columns (1)-(3) are from estimating the regression in equation (3.A.3) where the dependent variable is the wage of older high-skilled workers relative to younger high-skilled workers (pooled estimates for older and prime-aged workers relative to young workers in column (1), separately for older and prime-aged relative to young workers in columns (2) and (3), respectively). The results in columns (4)-(5) are based on equation (3.A.4) where the dependent variable is the relative wage of high-skilled to less-skilled workers for all age groups in column (4) and for young workers only in column (5). The results in columns (6) and (7) are based on equation (3.A.6) where the dependent variable is again the skill premium for all workers in column (6) and for young workers in column (7). Panel A reports OLS estimates, Panel B instrumental variable estimates where the college opening is used as an instrument for the relative supply of young, high-skilled workers. All specifications include calendar year and matched pair fixed effects. Pooled estimates in columns (1), (4) and (6) also control for age group fixed effects. Columns (4)-(7) include a linear time trend that is allowed to differ between pre- and post-opening period. Columns (6)-(7) further control for TFP growth and capital intensity at the national level. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.C.4: Specification Checks

	Baseline Estimates (1)	Linear trend (2)	Quadratic trend (3)	Placebo Estimates (4)	Heterogeneous effects (5)	Matching in t=-3 (6)	Synthetic Control (7)	Spillovers Broader Region (8)
Opening*Before	0.018	0.012	0.012	-0.015	0.019	0.034	-0.008	-0.015
	(0.034)	(0.034)	(0.034)	(0.037)	(0.032)	(0.041)	(0.033)	(0.016)
Opening*Transition	0.055	0.061*	0.061*	-0.027	0,055*	0.029	0.032	0.026
	(0.034)	(0.031)	(0.031)	(0.037)	(0.032)	(0.040)	(0.032)	(0.020)
Opening*Post	0.128**	0.145***	0.145***	-0.057	0,128**	0.055	0.083	0.045
	(0.062)	(0.052)	(0.052)	(0.075)	(0.059)	(0.076)	(0.069)	(0.035)
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	760	760	760	722	760	760	760	760
R-squared	0.979	0.991	0.991	0.973	0.980	0.974	0.976	0.996

Notes: The table reports several robustness checks for log young, high-skilled employment. The baseline results are shown in column (1). Columns (2) and (3) contain estimates when a linear or quadratic regional trend are added to equation (3.2), respectively. Column (4) measures the placebo treatment effect by matching all control districts to other untreated districts. Column (5) provides estimates using the heterogeneous treatment effects estimator of Sun and Abraham (2020). The estimates in column (6) are based on a matched sample using the same variables but in $\tau = -3$. Column (7) presents estimates for a synthetic control estimator. Column (8) estimates the spillover effect of treatment into neighboring districts. The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and *Post* the period from $\tau = 6$ to $\tau = 11$. Unless state otherwise, the estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.C.5: Employment Effects in Higher Education and the Non-tradable Sector

	Employment in Universities			**Employment in Non-tradable sector**		
	Total (1)	High Skilled (2)	Less Skilled (3)	Total (4)	High Skilled (5)	Less Skilled (6)
Opening*Before	-0.116	0.009	-0.095	-0.007	0.002	-0.007
	(0.077)	(0.068)	(0.064)	(0.016)	(0.028)	(0.016)
Opening*Transition	0.593***	0.360**	0.450***	-0.001	-0.026	0.001
	(0.198)	(0.149)	(0.154)	(0.017)	(0.035)	(0.017)
Opening*Post	0.932***	0.692*	0.738***	-0.008	-0.047	-0.002
	(0.339)	(0.297)	(0.251)	(0.036)	(0.061)	(0.036)
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	760	760	760	760	760	760
R-squared	0.832	0.739	0.860	0.987	0.981	0.987

Notes: The table reports the estimates of equation (3.2) for employment in higher education and the non-tradable sector (i.e., public administration and all services excluding financial services). The dependent variables are the log of total employment (column 1), log employment of workers with a college or university degree (2) and workers without a degree (3) in higher education. Columns (4)-(6) present estimates equivalent to the first three columns for employment in the non-tradable sector. The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.C.6: College Opening and Patent Activity

	Patents (Unweighted)			Patents (Weighted)		
	Total (1)	STEM (2)	Non-Stem (3)	Total (4)	STEM (5)	Non-Stem (6)
Opening*Before	-0.204	-0.142	-0.286	-0.245	-0.239	-0.259
	(0.237)	(0.355)	(0.304)	(0.193)	(0.276)	(0.275)
Opening*Transition	-0.107	-0.229	0.038	-0.098	-0.198	0.020
	(0.169)	(0.255)	(0.209)	(0.141)	(0.192)	(0.203)
Opening*Post	0.027	-0.049	0.116	0.045	-0.024	0.125
	(0.172)	(0.215)	(0.256)	(0.150)	(0.176)	(0.233)
Observations	760	418	342	760	418	342
R-squared	0.853	0.811	0.924	0.869	0.840	0.921
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports the estimates of equation (3.2) on patent activity. The dependent variables are the log number of granted patents in a district (columns 1-3), and the log number of patents weighted by the number of people from a district on the patent application. Column (1) and (4) present estimates for all district, the other columns for a sample split based on college focus in STEM or non-STEM subjects. The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.C.7: Comparison between Treatment and Control Districts in Dynamic and Stagnant Regions

	Stagnant		Dynamic		Stagnant vs Dynamic Treatment Districts		Stagnant Treated vs. Control		Dynamic Treated vs. Control	
	Treated Districts	Control Districts	Treated Districts	Control Districts	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Matched Characteristics										
Institutions of higher education:										
Tertiary institution in region	0.000	0.000	0.200	0.200	0.200	0.133	0.000	0.000	0.000	0.189
Population per square km	389.39	384.05	306.17	328.079	-83.226	156.824	5.343	155.251	-21.912	94.596
Industry structure:										
Share in Agriculture and Fishing	0.008	0.008	0.011	0.009	0.003	0.002	0.001	0.002	0.002	0.002
Energy and Mining	0.043	0.014	0.008	0.009	-0.035	0.027	0.028	0.028	-0.002	0.002
Manufacturing	0.437	0.499	0.437	0.477	0.001	0.062	-0.062	0.072	-0.04	0.049
Construction	0.076	0.065	0.083	0.070	0.007	0.008	0.012	0.009	0.013	0.011
Trade	0.121	0.118	0.135	0.125	0.014	0.011	0.004	0.01	0.01	0.012
Transport and Communication	0.034	0.035	0.039	0.033	0.005	0.004	-0.001	0.007	0.006	0.004
Financial services	0.025	0.026	0.027	0.026	0.002	0.002	-0.001	0.002	0.001	0.002
Other services	0.179	0.162	0.179	0.177	0	0.023	0.016	0.025	0.002	0.019
Non-Profit Organizations	0.015	0.018	0.024	0.017	0.008	0.006	-0.003	0.003	0.006	0.006
Public Administration	0.062	0.056	0.058	0.055	-0.004	0.009	0.006	0.010	0.003	0.007
Panel B: Characteristics Not Matched										
Age structure:										
Age: 20-29	0.317	0.326	0.339	0.327	0.022	0.018	-0.009	0.012	0.012	0.018
Age: 30-44	0.411	0.392	0.400	0.392	-0.01	0.018	0.019	0.017	0.008	0.015
Age: 45-59	0.272	0.282	0.261	0.281	-0.011	0.019	-0.010	0.019	-0.020	0.013
Further regional characteristics:										
Unemployment rate	9.362	6.774	6.959	7.896	-2.403*	1.286	2.588	1.612	-0.938	0.795
Employment	51,195	47,634	47,841	41,706	-3,354	12,029	3,561	14,710	6,135	12,660
Population (in Thousands)	306.22	243.58	257.15	262.36	-49.07	116.85	62.64	103.97	-5.21	89.53
Education:										
Highly skilled share	0.053	0.068	0.059	0.054	0.006	0.008	-0.015	0.010	0.005	0.008
Less skilled share	0.947	0.932	0.941	0.946	-0.006	0.008	0.015	0.010	-0.005	0.008
Employment Growth (past 5 years)	0.049	0.096	0.104	0.085	0.055	0.037	-0.047	0.046	0.020	0.033
Average Daily Wage	116.01	126.49	120.92	118.518	4.918	3.85	-10.481	6.504	2.406	3.082
Wage Growth (past 5 years)	0.051	0.087	0.049	0.05	-0.002	0.025	-0.035	0.038	-0.003	0.019

Notes:The table compares mean characteristics between dynamic treatment regions (column (1)), their matched controls (column (2)) as well as stagnant treatment regions (column (3)) and their matched controls (column (4)) in the pre-event period (t=-1). Columns (5)-(6) show differences (and standard errors) between dynamic and stagnant treatment regions along observable characteristics. Columns (7)-(10) show differences in observables between dynamic or stagnant treatment regions and their respective matched control regions. With the exception of employment and the indicator for the presence of another university in the district, all observations are weighted by district employment in the year just before the event ($\tau = -1$). The matched control regions are identified through Mahalanobis matching using the variables shown in Panel A in the pre-event period ($\tau = -1$). Regions are divided into dynamic and stagnant based on a median split of their employment growth between 1975 and 1980. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.C.8: Employment and Wages in Dynamic and Stagnant Regions

	Employment			**Average Wages**		
	Total	High-skilled	Less-skilled	Total	High-skilled	Less-skilled
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: Dynamic Regions			
Opening*Before	-0.004	0.008	-0.004	0.001	0.002	0.001
	(0.017)	(0.017)	(0.018)	(0.007)	(0.010)	(0.007)
Opening*Transition	0.012	0.022	0.010	0.012	0.004	0.011
	(0.015)	(0.022)	(0.016)	(0.008)	(0.010)	(0.009)
Opening*Post	0.033	0.038	0.030	0.019	0.004	0.017
	(0.028)	(0.046)	(0.028)	(0.018)	(0.020)	(0.016)
			Panel B: Stagnant Regions			
Opening*Before	0.008	-0.019	0.007	0.001	0.003	0.000
	(0.021)	(0.025)	(0.022)	(0.006)	(0.017)	(0.005)
Opening*Transition	-0.014	0.011	-0.014	-0.001	-0.006	-0.000
	(0.018)	(0.026)	(0.019)	(0.005)	(0.007)	(0.005)
Opening*Post	-0.049	0.014	-0.047	0.012	0.017	0.011
	(0.038)	(0.056)	(0.039)	(0.010)	(0.014)	(0.010)
Observations	760	760	760	760	760	760
R-squared	0.994	0.993	0.993	0.963	0.956	0.963
Event Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Quadratic Time Trend	No	No	No	Yes	Yes	Yes

Notes: The table reports the estimates of heterogeneous treatment effects on employment and wages according to equation 3.3. Columns (1)-(3) present estimates for the effect of a college opening on log full-time employment (in column (1)), log of workers with college or university degree (in column (2)) and the log of workers without a tertiary degree (in column (3)). Columns (4)-(6) use the log of the average daily wage for the respective group as dependent variable. Panel A contains the estimates for dynamic regions, panel B those for stagnant regions according to employment growth in the decade prior to the college opening. The unit of observation is district x year. All regressions are weighted by employment in the year prior to the event ($\tau = -1$). $Before$ denotes the period from $\tau = -7$ to $\tau = -2$, $Transition$ the period from $\tau = 0$ to $\tau = 5$, and $Post$ the period from $\tau = 6$ to $\tau = 11$. The estimates are based on a matched sample comparing districts with a college opening to suitable control regions using Mahalanobis matching (see details in text). Standard errors are clustered at the district level. Significance level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Chapter 4

The Local Consequences of War: Evidence from Germany and World War I[1]

Economists have long studied the effects of negative population shocks on local economies and labor markets. Typical examples of such shocks are outmigration (Hornbeck and Naidu, 2014; Andersson et al., forthcoming), pandemics (Brainerd and Siegler, 2003; Voigtländer and Voth, 2013a; Karlsson et al., 2014), natural disasters (Vigdor, 2008; Boustan et al., 2012) or wars (Voigtländer and Voth, 2013b). These shocks typically affect local populations and firms both through the human and physical capital stock.

Identifying the causal effect of a labor supply shock is challenging. For instance, natural disasters and wars often have a direct effect on the number of workers and infrastructure. In the case of migration, individuals' decision to migrate is endogenous to local economic conditions in the home country. These conditions are likely to also determine the economic response to the population shock. The closest setting to an exogenous negative population shock that leaves other economic inputs unchanged are pandemics. The Black Death killed an estimated 40% of the local population in the Middle Ages (Voigtländer and Voth, 2013a), but it

[1]This chapter is joint work with Tobias Korn and Lukas Wellner. We thank seminar audiences at the 2019 BBQ seminar and CES Dissertation Workshop at Harvard University for helpful comments and suggestions. We are very grateful to Dr. Jesper Zedlitz of the Verein für Computergenealogie for providing valuable information on the digitized loss lists and sharing data. Tobias Korn and Lukas Wellner acknowledge generous support by the German Research Foundation (DFG) as part of the Research Training Group 1723.

is unlikely that the economic effect of such a shock on a pre-industrial economy translates to more modern settings. In recent times, pandemics take a smaller toll on local populations. For instance only about 2% of the population died during the Spanish Flu in the early 20th century (Karlsson et al., 2014) and mortality due to Covid19 was significantly below 0.5% of the population in most countries (Johns Hopkins University, 2022). It is doubtful whether such small population shocks have a large impact on local labor markets.

In this paper, we investigate the effect of a large, negative population shock, namely German casualties in World War I. An estimated two million German soldiers died during the war, about 15% of the male population between 20 and 45, and many more were wounded or incapacitated. Figure 4.1 provides descriptive evidence of the size of the shock by plotting the ratio of women to men between the age of 20 to 45 in the German population by census years.

Figure 4.1: Gender Ratio of German Population over Time

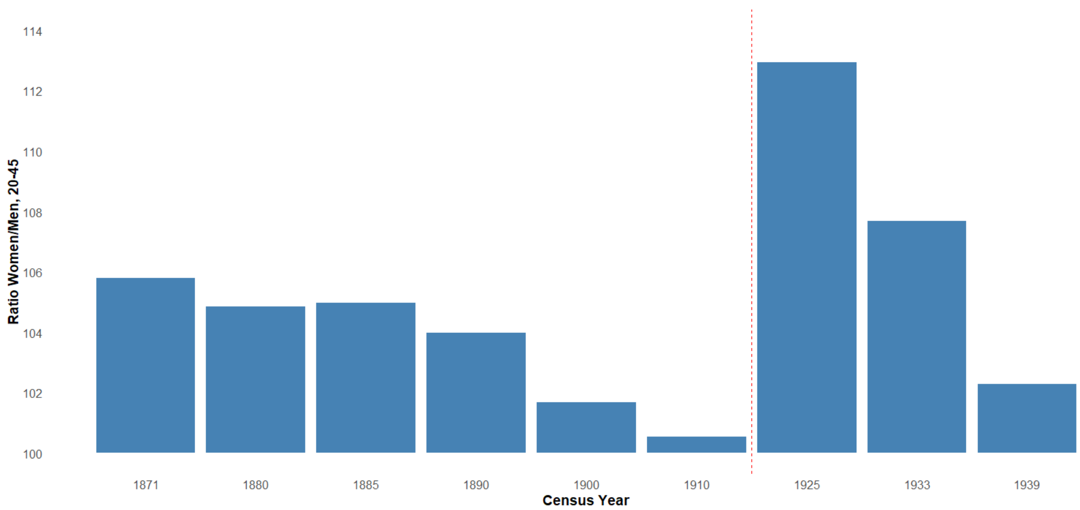

Note: This graph reports the gender ratio of women to men aged 20-45 in the German Empire over time, as reported in the respective census wave. Dashed lines indicate the period of World War I from July 1914 to November 1918.

The Figure points to a sizable increase in the ratio of women to men from near parity before to more than 115 women per 100 men right after the war. Notable in our setting is the fact that - different from many other settings - the shock did not affect the capital stock of the economy through war-related destruction. Except for a short period at the beginning of the war, all combat took place outside of German borders and military technology was not advanced enough for large-scale destruction of firms and infrastructure behind German lines. Also, while the war led to territorial changes with Germany losing a significant part of land, these territorial changes were not accompanied by large-scale migration from former German territories which might have offset the effects of the shock. Lastly, the war occurred at a critical time in economic development, when a host of labor-saving technologies, such as trucks, tractors or harvesting engines, became more available and women started entering the regular workforce in large numbers. Hence, our study may provide an interesting case study which also offers insights for more recent examples of negative population shocks, for instance outmigration from developing countries.

Economic theory suggests that an isolated shock to the population has two main channels to affect the local economy: First, increased scarcity of labor may lead to more participation of previously non-working groups, for instance women, in the labor market (Goldin, 1991; Acemoglu et al., 2004; Goldin and Olivetti, 2013) and increased adoption of technologies if they are substitutes for labor (Habakkuk, 1962; Acemoglu, 2010; Hornbeck and Naidu, 2014). If this labor supply effect dominates in our setting, it would manifest itself in higher wages, both for males and females, because labor is more valuable, and increased use of labor-saving technologies. Second, a negative population shock may manifest itself on the demand side as a smaller consumer base for local goods and services which in turn decreases firm profits and demand for labor. In this case, we would expect lower wages and less economic activity in regions hit by the shock. Which of these two channels dominates is highly relevant for economic policy: If the shock mostly works through the labor supply side, it may have positive overall effects for economic development. For instance, Young (2005) argues that labor scarcity resulting from the AIDS epidemic will lead to higher wages and more economic development in African countries. Voigtländer and Voth (2013b) attribute medieval urbanization and economic development in Europe in large part to deaths from war and plagues. However, if the labor demand effect dominates, the effect on economic development will be negative. For instance, Karlsson et al. (2014) shows that Swedish regions more affected by the Spanish Flu had higher rates of population in poor houses and lower returns to capital.

In the context of wars or natural disasters, such negative economic effects of a labor supply shock have been found to be short-lived. However, in many cases it is difficult to separate the effects of the shock on population from those on the physical capital stock (Brakman et al., 2004; Miguel and Roland, 2011; Feigenbaum et al., 2018).

In this paper, we investigate whether the negative population shock to the German economy in World War I manifested itself mostly on the demand or supply side. We exploit quasi-random regional variation in the number of casualties in a county to estimate the causal effect of the casualty shock. To this end, we geocoded the birthplaces of 8.5 million German soldiers who were mentioned in the official casualty lists published by the German army during the war. We find large regional variation in the casualty rate, i.e., the number of list entries divided by pre-war male population. We link these data to detailed census records - especially newly digitized wage data for men and women - for the pre- and post-war period which enables us to estimate the effect of a higher casualty rate in a difference-in-differences setting. Our estimation strategy allows us to account for institutional and economic factors which may have influenced the casualty rate.

Our results suggest that the casualty shock had a more prominent demand side effect. Counties with a higher casualty rate exhibited lower wage growth, both for men and women. The gender wage gap increases with the casualty rate, suggesting that a scarcity of male labor did not improve the relative labor market outcomes for women. Furthermore, a higher casualty rate leads to lower population growth, again both for male and female population. For women, this is due to migration away from more affected counties. In the case of men, the decline in population growth can be attributed to deaths from the war as well as outmigration. We interpret these findings as a sign that deteriorating economic conditions led to an outflow of the population, in line with an effect on the demand side. To further investigate the validity of this explanation, we use post-war data on tax revenues and find that the casualty rate had a negative effect on per-capita tax revenues for corporate, turnover and payroll taxes. This also points to a decrease in overall economic activity due to the labor supply shock.

To parse out the mechanism behind the demand side explanation we use data on industry and occupation structure. A higher casualty rate is associated with lower employment shares in trade and industry but a higher share in agriculture, pointing to a deindustrializing effect of the labor supply shock which likely explains our findings for lower wage growth. Furthermore, the share of self-employed workers increased while the share of salaried employees and wage workers in counties with higher casualty rates decreased. We use detailed data from agricultural censuses to explore the within-industry effects of the labor supply shock. Higher casualties led to a consolidation of agricultural landholdings in medium-sized farms at the expense of small and subsistence-size farms. However, employment in medium-sized farms does not increase, suggesting that agricultural workers from smaller farms were displaced. The supply side effects are negligible in our setting: We do not find increased use of machines and agricultural technology associated with the casualty rate. Lastly, we show that the negative economic effects also translated into the political sphere: Counties with higher casualty rates have higher vote shares for extreme right wing parties, especially the Nazi party in the final years of the Weimar Republic. Parties of the radical left and democratic center performed significantly worse in more exposed counties.

Our findings contribute to three strands of literature. First, we contribute to the aforementioned literature on the effects of negative labor supply shocks. In terms of technology adoption, Hornbeck and Naidu (2014) and Andersson et al. (forthcoming) find that areas more affected by outmigration have higher rates of technology adoption in agriculture. For individual-level outcomes, Goldin (1991) and Goldin and Olivetti (2013) find that the labor supply of women increased due to male mobilization in World War II. Acemoglu et al. (2004) show that this increase in female labor force participation lowered wages because female labor became a better substitute for male labor after the war. In a closely related setting to our study, Boehnke and Gay (forthcoming) show that French casualties in World War I led to higher female labor force participation after the war. This increase in women entering the labor market was driven by deteriorating conditions on the marriage market and less by demand side effects such as employers substituting female for male labor. By contrast, we find evidence that casualties acted as a negative shock to local demand which in turn decreased labor market opportunities for women and lowered wage growth and increased the gender wage gap.

Second, we contribute to a growing literature which examines the impact of World War I on voting behavior in the interwar period. Acemoglu et al. (forthcoming) use Italian casualties in World War I to argue that the death toll in the war was a major cause for the rise of communism in Italy, which in turn caused the counter-movement and rise of the right-wing government led by Mussolini in the mid 1920s. The impact of local exposure to the war on the rise of the Nazi party in Germany is examined in a similar fashion by Koenig (2020) and De Juan et al. (2021). Both papers establish a direct connection between World War I casualties or the return of veterans to their homes and increased vote shares of extreme right and the Nazi Party in the Weimar Republic. They attribute these effects to the organizational skill of right-wing veterans' organizations and the common experience of suffering and human loss associated with casualties. Our findings suggest that counties more affected by the war through casualties underwent major economic transitions and had lower economic activity as well as higher vote shares for extreme right wing and the Nazi Party. Increased support for right-wing parties may be a direct result of the War as well as an intermediate result of economic hardship resulting from a negative shock to the local economy. There already exists some evidence that voters in the Weimar Republic voted for extreme right-wing parties due to economic considerations (King et al., 2008; Voth et al., 2020; Galofré-Vilà et al., 2021). Our paper establishes a link between the local effects of World War I, economic outcomes, and voting behavior. It therefore contributes to a larger literature investigating the rise of fascism in Europe in the interwar period.

Third, several recent studies rely on the raw data set of German casualty lists in World War I from Computergenealogie (2019a) that we also use in our analysis. For instance, De Juan et al. (2021) investigate the effect of casualties in a county on voting behavior in the post-war era. Ciccone (2021) uses casualty list entries as a negative, exogenous local population shock for the state of Württemberg. Kersting and Wolf (2021) use the names and birthplaces of soldiers to identify patterns in first names as signs of nation-building after German unification in 1871. All these studies rely on the original geocoding provided by Computergenealogie (2019a). We improve these data set by providing a new and more accurate geocoding for more than 84% of observations in the data set. Hence, we hope that our improved data can also serve as a resource for other scholars working with this data set.

4.1 Historical Background

World War I was the first fully modern and global war, lasting from July 1914 to November 1918. During the conflict between the "Central Powers" of Germany, Austria-Hungary and the Ottoman Empire on the one side and the Allied countries of Great Britain, France, Russia and later the United States on the other, over 60 million soldiers were mobilized and an estimated 10 million of them perished. After initial German advances in the West, offensives on both sides quickly stalled due to technological advances and outdated military tactics. As a result, the war on the Western Front devolved into deadly trench warfare with trenches stretching from Flanders in Belgium to the Franco-Swiss border in the South. On the Eastern Front, an initial Russian advance into East Prussia was successfully repelled and after the Battle of Tannenberg in August 1914, the entirety of the war was fought on Russian soil. The conflict in the East ended with the Russian surrender and the treaty of Brest-Litovsk in March 1918.

After the entry of the United States into the war in April 1917, the situation of the Central Powers quickly deteriorated and a ceasefire was signed in November

1918. The losing Central Powers faced harsh repercussions: The Ottoman Empire and Austria-Hungary were broken up into several states. Germany was stripped of its colonial possessions and lost significant territory to France, Belgium, newly-formed Poland and Czechoslovakia.

The German war effort involved a nearly universal mobilization of men aged 17-45. However, the need for manpower was accompanied with a need for large amounts of military supplies such as guns and ammunition. The German economy was ill-equipped to handle the sudden rise in demand. Especially in the Fall of 1914, the German Army used a far greater number of shells and ammunition than was produced, spurring fears among military leadership that German forces may run out of ammunition (Afflerbach, 2018). In response to this development, existing firms were forced to shift from civilian to military production and German High Command ordered the construction of additional plants.[2] However, none of these shifts created long-lasting changes to German industrial structure: After the war, most firms returned to their pre-war civilian production.

The war had a significant impact on the composition of the workforce. Due to a lack of male labor, many women started taking up employment in previously male-dominated industries. For instance, the female employment share rose from 7% in 1914 to 23% in 1916 in the chemical industry or from 24% to 55% in the electrical industry in the same time period (Helfferich, 1925). Prisoners of war or deported civilians from occupied territories were forced to support the German war effort as well. After the war, returning soldiers reentered the workforce and took back their previous positions and female labor force participation fell again.

Lastly, in the aftermath of the war, Germany was obliged to pay high reparations to the Allied countries. These consisted of direct monetary payments as well as deliveries of coal, steel and other industrial goods. While these reparations amounted to a significant share of German industrial production, Germany's industrial base was left largely untouched. The Allies decided not to dismantle German factories

[2] These centralized programs, for instance the so-called "Hindenburg Program" were largely unsuccessful in increasing military production. For instance, of over 40 furnaces built as part of the program, none were used during and most were demolished after the war (Helfferich, 1925).

or relocate them to Allied territory. Furthermore, the territorial loss, especially in former Eastern Prussia did not lead to large-scale migration from those territories.

The war had negative overall economic consequences for Germany. Political turmoil and reparation payments contributed to hyperinflation in 1923 which delayed the post-war economic recovery. Most of the outcomes we use in our analysis for the post-war period are therefore for 1925 and later - a relatively stable time for the German economy in the interwar period up to the Great Depression starting in 1929.

4.2 Data and Identification Strategy

4.2.1 Casualty Data

We use the universe of all German entries in the so called loss lists (‘Verlustlisten’) which were digitized in a crowd sourced effort by Verein für Computergenealogie (2019a). In total, these lists contain 8.5m entries on the fate of German soldiers during WWI. Each entry includes information on the soldier’s name, birthplace, rank, regiment and the type of casualty reported, i.e., whether he was captured, wounded or killed.[3] The data also includes the soldier’s place of birth. This information has been used by Computergenealogie (2019a) to geocode the soldiers’ birthplaces. However, there exist serious concerns whether this geocoding is accurate, either because the underlying information in the lists is incorrect or the geocoding was imprecise. For instance, 955 list entries with the birthplace “Zürich” are assigned to a small hamlet with a few dozen inhabitants of the same name in the county of Steinfurt in Westphalia, not the large city in Switzerland. In other cases, parts of the birthplace information was discarded, despite of being important, or available information was wrongly interpreted. For instance, 3,286 list entries containing the spelling “Straßfurt i. E.” are assigned to a farm with less than ten inhabitants in Northern Germany. Instead, these entries likely referred to soldiers born in the city of Strasbourg in Alsace.

[3]While some information is simply not reported in the original casualty lists, some information that is available has not been digitized yet. Casualty status and regiment information for instance has only been collected for about 700,000 entries.

Some of these mistakes may lead to systematic errors in the number of casualties when aggregating to higher geographic units.[4] For instance, the original coding will likely be biased to higher number of casualties in urban areas or large cities with common names (for instance, 'Neustadt'). Therefore, we recode the entire universe of German WWI casualties on the county level. To this end, we use a comprehensive list of German municipality names in 1910 from Schubert (2020)[5] and exploit the hierarchical structure of entries that typically consist of the name of the municipality of birth followed by the county or administrative district. We are able to assign about 7.2 million entries from the casualty lists to a county, a coverage of about 84%. We describe our geocoding process and differences between our coding and the original one in Appendix 4.A.

Our main explanatory variable is the casualty rate, i.e., the number of casualty list entries assigned to a county divided by pre-war male population. Figure 4.2 shows the geographic distribution of the casualty rate. It is higher in the more rural areas of East and West Prussia, Pomerania, Posen, and Silesia and lower in more urban areas such as Berlin, Danzig and the Ruhr area and Southern Germany.

We assume that the casualty rate is an accurate proxy of the size of the local casualty shock. To test this assumption, we measure whether the casualty rate accurately reflects the number of dead or incapacitated soldiers from a county. The casualty lists contain information on whether a soldier was killed, wounded, or captured for each entry, but as of now this information is digitized for only 1.1m entries, about 15% of the sample. Hence, the number of list entries would not be an accurate measure of the size of the shock if the number of deaths were not highly correlated with the number of list entries. We use the partly-digitized information on the casualty status from the lists and find that the correlation between the number of entries and the number of deaths is 0.91.[6] This suggests,

[4]One should note that the original geocoding was never meant to be used for this purpose in the first place.

[5]We thank Uli Schubert for graciously sharing his data with us.

[6]To show that this finding is not related to the order in which entries were digitized, we calculate both the share of entries with digitized information on the casualty status as well as the share of digitized entries with death as status of all entries in a county. Figure 4.B.2 in the Appendix shows scatter plots of the casualty rate, the share of digitized entries and the share of

Figure 4.2: Spatial Distribution of Casualty Rate

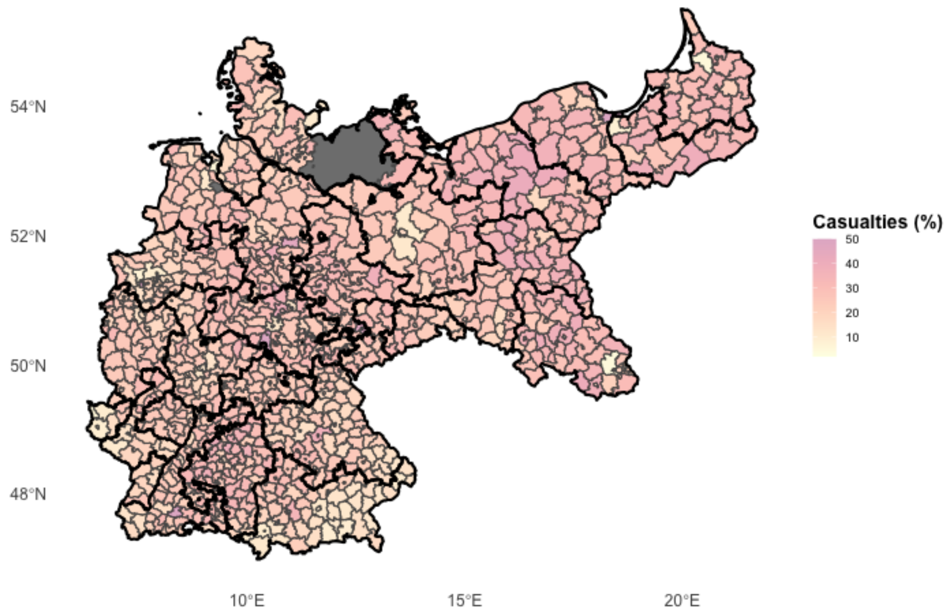

Note: Map of the number of casualty list entries divided by male county population in 1910. Thick black lines depict army district borders. Thin gray lines depict county borders. Grey areas denote missing data. The shapefile is based on MPIDR (2011).

that the casualty rate accurately reflects the share of dead soldiers and the size of the shock to a county. As a robustness check we later also use the death rate, i.e., the number of entries coded as dead divided by male population in 1910 in a county, as explanatory variable, which leaves our results unchanged.

Soldiers' assignment to specific military units within the German Army relied on their place of residence.[7] At the beginning of the war, Germany was divided into so-called army districts and each military unit, for instance a regiment, was based within this district with most of its recruits hailing from the surrounding area. Figure 4.B.1 in the Appendix shows a contemporary map of the army districts at the beginning of the war. This structure affected the distribution of casualties across army districts. As an example, Figure 4.3 shows regiment-specific casualty

entries with death status. There is no systematic relationship between the casualty share and the share of deaths among all entries, suggesting that our main variable is on average a good proxy for the number of deaths. Furthermore, there appears to be no systematic relationship between the share of digitized entries and the casualty share.

[7]As the number of casualties increased over time, this rule was relaxed significantly (Stachelbeck, 2010).

rates for two military units, the 49th Infantry Regiment based in the city of Gnesen in the province of Posen and the 57th Infantry Regiment based in Wesel, Westphalia. The map shows that the vast majority of casualties for both regiments are from soldiers whose birthplaces lie within the borders of the respective army districts.

There are two conclusions to be drawn from these graphs: First, differences in deployment of military units to more or less deadly theaters of the war may drive regional variation in casualty rates. While we found no evidence of such systematic differences across army districts in contemporary sources, we nonetheless account for the army district structure in our identification strategy. Second, the graphs show that assigning casualties based on birthplaces accurately reflects where soldiers lived before the war. In the presence of high internal migration in the pre-war period, our casualty rate may be inaccurate because we do not observe the last place of residence before the war. There existed some internal migration in Germany in the pre-war period (Grant, 2005), especially from rural towards urban regions. However, it is unclear how much this applied to soldiers who were young and therefore had limited time to migrate before being drafted in the army. In order to approximate the extent of internal migration we use the two regiments from Figure 4.3 and calculate the share of soldiers whose place of birth lies within the borders of the army district assigned to their regiment. The shares are 55% for the 49th infantry regiment in Gnesen and 61% for the 57th in Wesel. These shares increase to 66% for the 49th and stays stable at 58% for the 57th regiment when we only look at list entries at the beginning of the war in 1914, when the assignment to regiments was more closely related to the place of residence than at the later stages of the war. This strongly suggests that most soldiers were still living in or very close to their birthplace. Hence, we are confident that our casualty rate gives an accurate reflection of the number of casualties sustained in a county.[8]

[8]To further check the extent to which internal migration might influence the casualty rate, we use an additional set of lists of missing soldiers which were collected after the war and were digitized by Computergenealogie (2019b). The lists contain soldiers' birthplaces and the place of residence of their closest relative, usually their parents or spouse. We match birthplaces and the residence of the closest relative and find that for about 60% of entries, the address of the closest relative was identical or very similar to the place of birth of the missing soldiers. This only includes cases in which birthplaces and places of residence were virtually identical except for

Figure 4.3: Spatial Distribution of Regiment-Specific Casualty Rates

(a) 49$^{\text{th}}$ Infantry Regiment

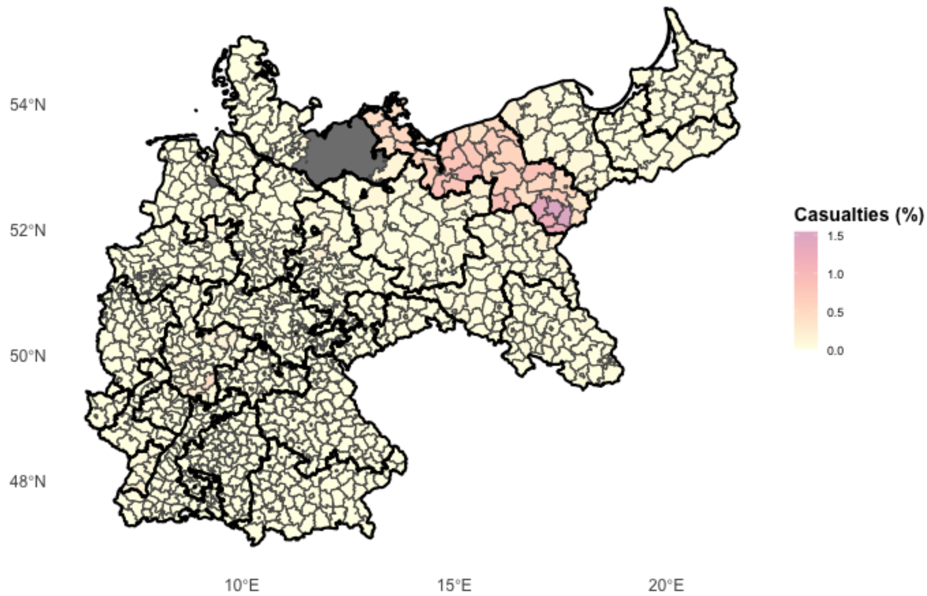

(b) 57$^{\text{th}}$ Infantry Regiment

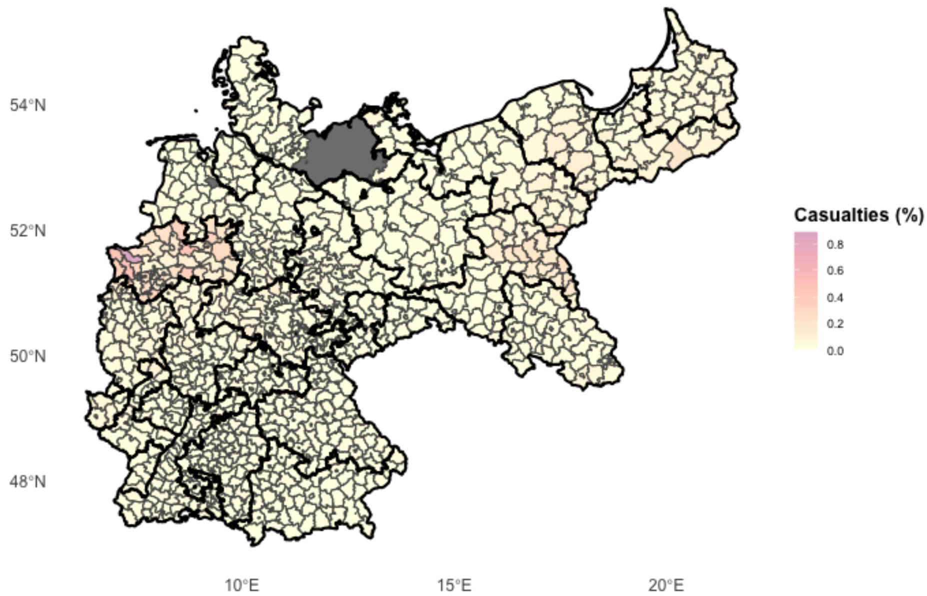

Note: Map of the number of casualty list entries divided by male county population in 1910 for the respective regiment. The upper map shows the casualty rate for casualties reported from the 49$^{\text{th}}$ Infantry Regiment and the lower map shows the casualty rate for casualties reported from the 57$^{\text{th}}$ Infantry Regiment. Thick black lines depict army district borders. Thin gray lines depict county borders. Grey areas denote missing data. Shapefiles are based on MPIDR (2011).

In order to measure the population changes associated with the casualty shock, we use data from the population censuses in 1910, 1916 and 1925. With these data we are able to calculate the gender ratios before, during and after the war. Furthermore, we are able to calculate population growth by gender and the intensity of migration from a county as the difference between the overall population change and the number of births and deaths recorded in a county between 1910 and 1925. Using data from the 1916 census during the war, we can also analyze differences in mobilization rates across counties. These differences can be the result of several factors, for instance decisions of local draft boards or the number of workers in war-related industries, such as miners or steelworkers, who were exempt from the draft. As a proxy for mobilization rates we calculate the share of women over men in the population in 1916.[9] A higher share implies that more men were away in the war and the county experienced higher mobilization rates.

The first graph of Figure 4.4 shows the geographic distribution of the female population share in 1916. As is evident, it varied greatly, from near parity to about 65% females in the local population. The female share is lower in industrial centers such as the Ruhr area and higher in rural areas. The second graph in Figure 4.4 shows that a higher female share in 1916, i.e., higher mobilization rates, is associated with higher casualty shares. Hence, our identification strategy needs to account for differences in mobilization rates since it is likely related to pre-war economic outcomes which may also influence the reaction to the casualty shock.

In order to achieve a consistent size of all counties over time, we account for aggregations and county splits between 1907 and 1925. For our regression results, we drop all counties that experienced municipality-level border changes or which were not listed in any of the relevant censuses. Our final data set consists of 780 counties that we observe before and after the war. Based on the discussion above,

differences in spelling. It does not include cases where soldiers and their relatives lived in different towns in the same county or very close to each other.

[9] The census provides information for the male population in three categories: civilian males, male military personnel and foreign prisoners of war. We exclude prisoners of war because they were clearly not part of the local male population and we exclude soldiers because according to the information provided in the census they were either on leave or in local training camps at the time of the census.

Figure 4.4: Mobilization and Casualty Rates

(a) Share of Female Population in 1916

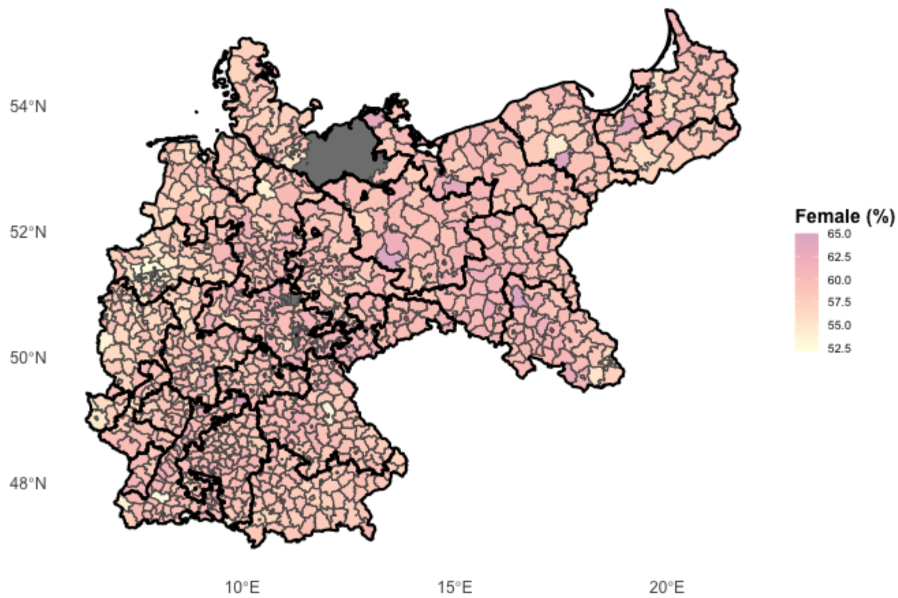

(b) Share of Female Population in 1916 and Casualty Rate

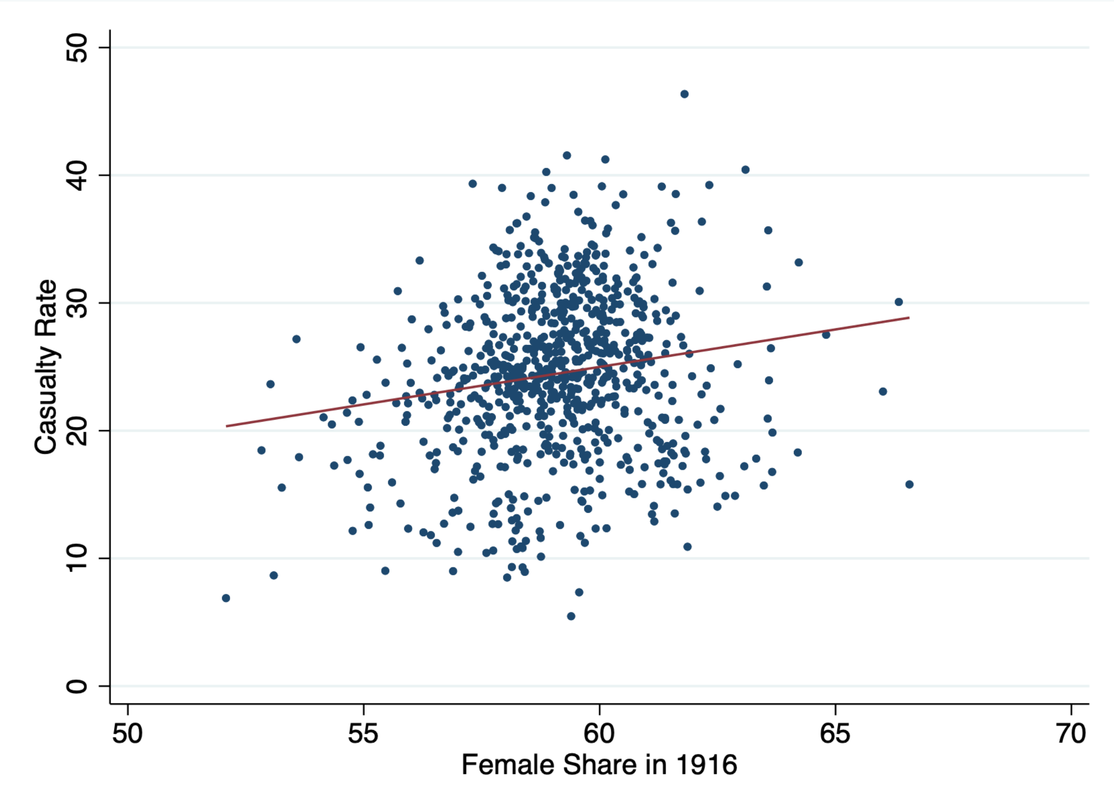

Note: Graph a) shows the share of female population in a county in 1916. Graph b) is a scatter plot of the female population share in 1916 and our main explanatory variable, the casualty rate, i.e., the number of casualty list entries divided by male population in 1910. The red line represents a linear fit of the data. The shapefile is based on MPIDR (2011).

we run several regressions to estimate the determining factors of the casualty rate using our final data set. Table 4.1 presents the results of these regressions.

Table 4.1: Pre-War Determinants of Casualty Rate

	Casualty Rate					
	(1)	(2)	(3)	(4)	(5)	(6)
Share Female 1916	0.587***	0.224	0.731***	0.816***	0.816***	0.760***
	(0.137)	(0.143)	(0.133)	(0.131)	(0.130)	(0.131)
Urban=1			-6.921***	-4.591***	-4.612***	-5.101***
			(0.609)	(0.863)	(0.867)	(0.865)
Pop Density 1910				-0.001***	-0.001***	-0.001***
				(0.000)	(0.000)	(0.000)
Distance to Coal					-0.009**	-0.012*
					(0.004)	(0.007)
Army District FE	No	Yes	Yes	Yes	Yes	Yes
Admin. District FE	No	No	No	No	No	Yes
Observations	778	778	778	778	778	778
R^2	0.027	0.341	0.450	0.466	0.470	0.545

Notes: OLS Regressions using the casualty rate, i.e., the number of entries to the casualty lists divided by male population in 1910 as dependent variable. $ShareFemale1916$ is the share of female population in 1916 and a proxy for draft intensity. $Urban$ is a dummy indicating whether a county is urban. $PopDensity1910$ is population density in 1910 and $DistancetoCoal$ is the distance of a county centroid from the nearest coal deposit. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The coefficient for the female population share in 1916 has a positive sign in all models suggesting that higher mobilization rates are in fact associated with a higher casualty rate. Furthermore, simply controlling for the female share in 1916 and the army district structure explains about 34 percent of the variation in the casualty rate. The casualty rate is higher in rural areas, as evidenced by the consistently negative coefficient for the urban dummy[10] and population density in 1910. Distance to the nearest coal deposit as a proxy for the presence of war-related industries has a small, negative coefficient. This is surprising but likely due to the fact that we are already controlling more directly for such industries both through the mobilization rate and the urban dummy. Adding administrative district

[10]The dummy take the value 1 for urban counties (*Stadtkreise*) and immediate cities (*unmittelbare Städte*). A complete list of urban counties in our sample can be found in Table 4.B.1 in the Appendix.

('Regierungsbezirk') fixed effects in the last column does not significantly change any of the coefficients but increases the R^2.

4.2.2 Wage Data

For our main outcomes for local economic conditions, we use data on pre- and post-war social security reference wages ('Ortsübliche Tagelöhne') which determined the amount benefits awarded to injured workers based on the public accident insurance scheme. According to contemporary accounts, these reference wages reflected regional variation in wage levels and the cost of living. They were published at different intervals in the Proceedings of the Ministry of the Interior ('Zentralblatt des deutschen Reiches'). For our analysis, we digitized the last pre-war revision of the reference wages from January 1914 and the first major post-war revision for January 1922. The reference wages were provided for three different age groups (older than 21, 16 to 21 and under 16) as well as by gender.

In order to test whether these reference wages reflect actual wage levels, we compare them with payroll tax data by Brockmann et al. (2022) for 1926, the earliest year available from their data.[11] While our post-war wage data are from 1922 and reflect the wages of unskilled laborers, the payroll tax data reflect the wage level of high wage earners, as lower incomes were mostly exempt from the payroll tax.[12] The correlation between the daily wage of male workers aged over 21 in 1922 and the amount of payroll tax per capita in a county in 1926 is about 0.5, suggesting that while there are some differences, our social security wage data provides an accurate reflection of local wage levels.

4.2.3 Identification Strategy

We want to estimate the effect of the casualty shock on the local economy in German counties. A naive approach would simply treat our main variable, the casualty rate, as an exogenous shock to the local population. Contemporaneous accounts suggest that individual survival chances were largely determined by chance. For instance,

[11]We thank Brockmann et al. for sharing their data with us.

[12]Brockmann et al. (2022) estimate the share of exempt wage earners to be around 50%.

Jünger (1920) writes in his account of the war, "Behind every parapet was fate, waiting to randomly grab a new victim." Or as Remarque (1928) puts it: "Above all is chance. When a bullet is coming, all I can do is duck. I can neither know nor steer where it lands."[13] However as we already discussed above, institutional factors such as the military district structure, differences in mobilization rates and pre-war economic characteristics may be correlated with the casualty rate and the economic outcomes we want to analyze. In this case, a naive OLS regression using the casualty rate as explanatory variable would yield biased coefficients.

Hence, we aim to control for these factors in our identification strategy. We exploit the fact that we observe most of our main outcomes both in the pre- and post-war period. We can therefore specify a difference-in-differences model according to the following equation:

$$\begin{aligned} Outcome_{i,t} = &\beta Casualty_i \times Post_t + \gamma Post_t + \delta_i \\ &+ \theta Army_d \times Post_t + \zeta FemShare\_1916_i \times Post_t + \epsilon_{i,t}, \end{aligned} \tag{4.1}$$

where i is the county, t the time period, either pre- or post-war. The outcome in question needs to be observed in both time periods. Our main explanatory variable, the casualty rate $Casualty_i$ varies by county but is fixed over time. Hence, the coefficient of interest is the interaction of the casualty rate with the dummy indicating the Post period, $Post_t$. County-level FE, δ_i absorb all time-invariant county characteristics, such as assignment to an army district or the draft intensity. However, in order to also account for any differential effects of the district structure or mobilization rate on the outcome we also include interactions of $Post_t$ and dummies indicating the 22 army districts, $Army_d$, and our measure of the mobilization rates, the share of female population in 1916, $FemShare\_1916_i$.

The time-invariant effect of the army district fixed effects and female share is already captured by the county fixed effects. By interacting them with the dummy for the post period, we allow for differential effects in the post period. For instance,

[13] These individual assessments are also backed up by the data provided by the German army in its official casualty reports. The most common cause of death for soldiers was artillery fire, for about 50% of casualties (Reichswehrministerium, 1934).

controlling for differential effects of the female share in 1916 rules out that our wage effects are biased due to changes in the substitutability of types of labor due to differences in mobilization rates in the spirit of Acemoglu et al. (2004). Our identifying assumption is that conditional on county fixed effects, differential effects of the military district structure and mobilization rates over time, the remaining variation in casualty rates is orthogonal to the error term. In this case, the effect captured by β stems from the remaining quasi-random variation in the casualty rate and allows a causal interpretation.

Potential violations of our identifying assumption are county-specific, time-varying shocks which are correlated with the casualty rate and affect the outcome in question. For instance, the estimate of β may be biased if urban and rural counties exhibit different economic trends in the pre-war period or if their ability to absorb the negative population shock differs. We address this concern in our robustness checks by dropping all urban counties and show that our main estimates are robust to the exclusion of these observations. A further concern are other war-related shocks, such as destruction of the capital stock or expulsions of the local population which may be correlated with the casualty rate, even though the war was fought mostly outside of German borders. Nonetheless, we control for such a possibility in our robustness checks and show that our estimates are robust to excluding all counties which were directly affected by combat. In an additional check, we also exclude all counties which were affected by the post-war occupation by Belgian and French troops to show that our estimates are not a result of potential expulsions of the local population or forced de-industrialization.

4.3 Results

For our main results, we use wage levels in 1914 and 1922 as dependent variables in our difference-in-differences setup. This will shed some first light on whether the casualty shock operated more through the labor supply or demand channel. If wages rise, this points to a labor supply effect as workers earn higher wages due to labor scarcity. If wages fall with higher casualty rates, this could be interpreted as a sign of demand side effects: Demand for goods and services as well as firm profits decrease and labor demand falls as a result. In such a scenario, we expect lower wages in response to the casualty shock. Figure 4.5 presents the point estimates of our main regressions for the daily wage in 1914 and 1922 for three age groups, over 21 years, between 16 and 21 years, and under 16 years, by gender.

Figure 4.5: Effect of Casualty Rate on Daily Wages

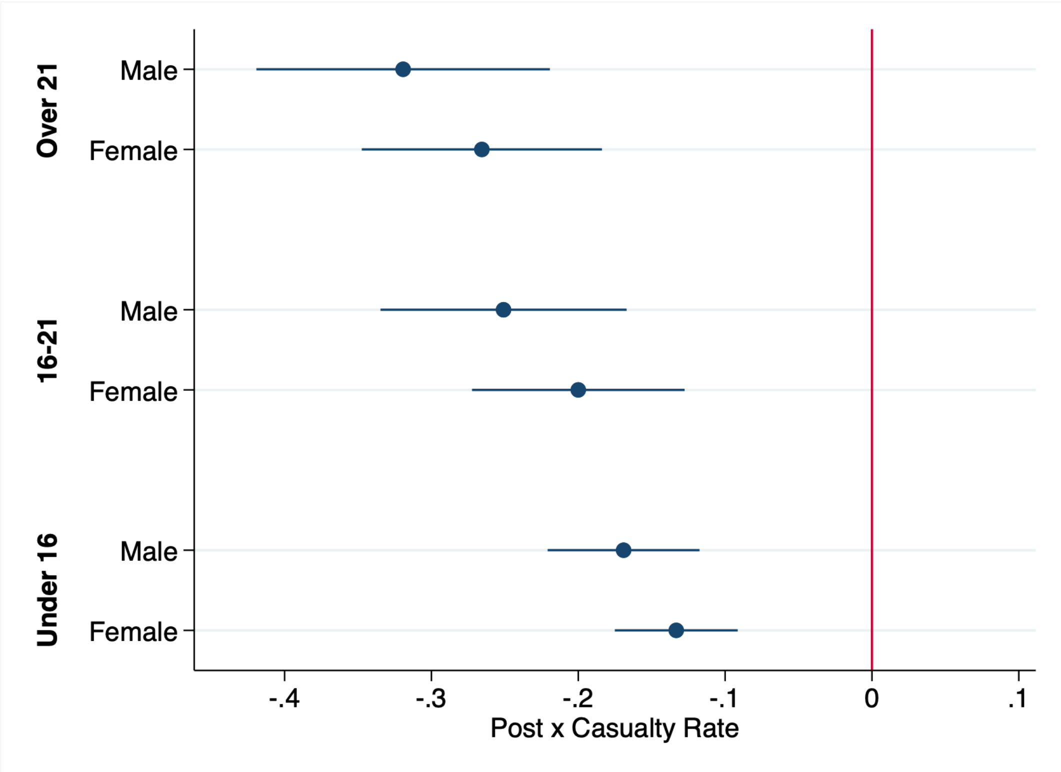

Note: This graph reports the coefficients and 95% Confidence Intervals from regressing daily wages in 1914 and 1922 on the casualty rate in difference-in-differences estimations based on equation (4.1).

We find that the point estimate for our coefficient of interest is negative and significant in each specification. A one standard deviation increase in the casualty rate (6.5) leads to an about 2.1 Reichsmark (RM) lower wage for male workers over 21. For female workers in the same age group, the wage growth was about 1.5 RM lower for a one standard deviation increase in the casualty rate. Note that our estimations include year fixed effects, which control for price changes across periods, which allows us to interpret these wage growth differences as corresponding differences in real wage changes.[14] These differences are qualitatively similar across age groups, though the effect is a little smaller among younger workers.

Our wage results are surprising in so far as they point to significant negative consequences of the casualty shock for workers. Clearly, the casualty shock did not increase wages by increasing the value of labor, especially for women. We therefore investigate how the shock affected the relative standing of women in the labor market, the composition of the local population and whether the local population reacted to the shock, for instance through migration. Table 4.2 presents results using the gender wage gap, defined as the share of female daily wages relative to male daily wages in 1914 and 1922, the gender ratio, defined as the share of women in the population in 1910 and 1925, as dependent variables in our difference-in-differences setup. Furthermore, we use net migration between 1910 and 1925 divided by population in 1910 as dependent variable. For the migration variables, we drop the two-period setup, since we only observe these variables once for the entire period and run the following regression:

$$Outcome_i = \beta Casualty_i + \gamma Controls_i + \theta Army_d + \zeta FemShare\_1916_i + \epsilon_i. \quad (4.2)$$

Again, our coefficient of interest is β, the effect of the casualty rate on the outcome in question. As previously, we also control for the army district structure and the mobilization rate through the female share in 1916. We further add county-specific pre-war controls, $Controls_i$. In our main specification these are a dummy indicating whether a county is urban, population density in 1910, distance of the county centroid to the nearest coal deposit in kilometers and administrative district ('Regierungsbezirk') fixed effects.

We find that among workers above the age of 16, the gender wage gap increased significantly in locations that were harder hit by casualties. Both for workers over 21 and workers between 16 and 21, the ratio of female to male wages decreased by around 1.4 percentage points for a one standard deviation increase in the casualty rate. Hence, the wage difference between genders became more pronounced in

[14] This is important, as our post-war wage data for 1922 already show signs of the beginning inflation period in 1923.

Table 4.2: Effect of Casualty Rate on Gender Wage Gaps and Population Movement

	Gender Wage Gap			Share	Net Migration	
	Over 21	16 - 21	Under 16	Females	Female	Male
	(1)	(2)	(3)	(4)	(5)	(6)
Casualty Rate					-0.466***	-0.266***
					(0.071)	(0.071)
Post x Casualty Rate	-0.223***	-0.216***	-0.019	-0.069***		
	(0.050)	(0.054)	(0.056)	(0.009)		
R^2	0.791	0.766	0.669	0.878	0.331	0.238
Post x Army District	Yes	Yes	Yes	Yes	No	No
Post x Female Share 1916	Yes	Yes	Yes	Yes	No	No
County FE	Yes	Yes	Yes	Yes	No	No
Army District FE	No	No	No	No	Yes	Yes
Female Share 1916	No	No	No	No	Yes	Yes
Admin. District FE	No	No	No	No	Yes	Yes
Further Controls	No	No	No	No	Yes	Yes
Observations	1552	1552	1552	1556	724	725

Notes: Results of regressions described in equation (4.1) for the two-period setup and 4.2 for the one-period setup. Outcomes are the ratio of female to male wages in different age groups (Columns (1)-(3)) in 1914 and 1922, the share of female population in 1910 and 1925 (Column (4)) and the share of net migration between 1910 and 1925 by gender (Columns (5) and (6). Further controls are population density in 1910, a dummy indicating whether a county is urban, and the distance to the nearest coal deposit in kilometers. Note that for columns (1)-(4), army district FE, the female share in 1916 and administrative district FE are included through the county FE. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

locations that incurred higher casualties during the war. This trend occurs amid a national decrease in real wages. Again, these results suggest that the casualty shock manifested itself more as a demand side shock in our setting. Even in the face of a higher scarcity of male labor, women do not start to earn higher wages. Instead, relative to men, their wages decrease.

We provide one explanation for this trend in the gender wage gap among counties with higher casualty rates in the remaining columns of Table 4.2. In Column (4), we again run a difference-in-differences estimation, but with the share of female inhabitants in a county as the dependent variable. We find that this share follows a significantly lower trend in counties that experienced a larger casualty shock. On average, a one standard deviation increase in the casualty rate leads to a 0.45 percentage point lower share of women in the population. This effect is relatively sizable compared to an overall increase in the share of female inhabitants of around

10 percentage points nationally. We provide an explanation for this lower increase in the female population share among harder hit counties in Columns (5) and (6). Here, we depict results from OLS cross-section estimates based on equation (4.2) with net migration between 1910 and 1925 as the dependent variable. We calculated this share as the population change net of births and deaths between 1910 and 1925, relative to pre-war population in 1910. We find evidence that net migration was negative in harder hit counties, especially for women. Hence, a stronger net outmigration of women between 1910 and 1925 led to a change of the gender composition in these counties even though the casualty shock predominantly affected the male population. This suggests that women escaped deteriorating economic conditions, especially lower wages, by moving towards counties that were less affected by the casualty shock.

4.4 Mechanisms

4.4.1 Local Economic Activity

We now want to investigate the demand side effect directly by looking at the intensity of local economic activity and employment shares. Since we lack accurate data on output, we use per capita tax revenues on payroll ('Lohnsteuer'), corporate profits ('Köperschaftssteuer'), and turnover ('Umsatzsteuer') from Brockmann et al. (2022) as proxies for overall output and economic activity in a county. We use employment shares in manufacturing, trade and agriculture from Falter and Hänisch (1990) as dependent variables to measure the effect of the casualty shock on sectoral composition. Unfortunately, all these data are only available for the post-war period. Hence, we can only estimate cross-section OLS regressions based on equation (4.2) with these data as the dependent variables. Even though we are confident that our control variables capture a large part of the endogeneity of the exposure to war casualties as discussed above, these correlational results should still be interpreted with caution. We report the results of these estimates in Table 4.3.

Across all types of taxes, we find that per capita revenues are significantly lower in counties with a higher casualty rate. The effects are sizable: A one standard

Table 4.3: Effect of Casualty Rate on Tax Revenues and Industry Shares

	Tax Revenues			Industry Employment Shares		
	Payroll	Corporate	Turnover	Manufacturing	Trade	Agriculture
	(1)	(2)	(3)	(4)	(5)	(6)
Casualty Rate	-0.040***	-0.036***	-0.014***	-0.747***	-0.344***	1.343***
	(0.004)	(0.006)	(0.002)	(0.087)	(0.040)	(0.109)
R^2	0.577	0.357	0.496	0.620	0.693	0.719
Army District FE	Yes	Yes	Yes	Yes	Yes	Yes
Female Share 1916	Yes	Yes	Yes	Yes	Yes	Yes
Admin. District FE	Yes	Yes	Yes	Yes	Yes	Yes
Further Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	772	772	772	778	778	778

Notes: Results of regressions described in equation (4.2) for the one-period setup. Outcomes are the log revenue of payroll (Column (1)), corporate (Column (2)) and turnover (Column (3)) taxes per capita in 1926. Columns (4)-(6) use employment shares in different industries in 1925 as outcomes. Column (4) shows the share of population working in manufacturing (0-100), column (5) the share of population working in trade, and column (6) the share of population working in agriculture. Further controls are population density in 1910, a dummy indicating whether a county is urban, and the distance to the nearest coal deposit in kilometers. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

deviation increase in the casualty rate is associated with about 20% lower corporate tax revenues per capita in a county. While we do not want to claim causality for these results, this points towards an overall lower economic activity in more-affected counties and to a demand effect. The casualty shock negatively affected the local consumer base which led to lower spending, consumption and firm profits, which in turn depressed wages and led to outmigration of the local population, especially women. These results are backed up by our findings on the sectoral composition: In more affected counties, employment in the manufacturing and trading sectors is significantly lower than in less affected counties. At the same time, the employment share in the agricultural sector is significantly higher; a one standard deviation increase in the WWI casualty share can be associated with a more than 8.7 percentage points higher employment share in agriculture. This hints at a partial de-industrialization in more affected counties and may explain the negative effects on wages and migration we found earlier. We present further evidence of this de-industrializing effect in Table 4.B.2 in the Appendix by looking at the occupational structure in manufacturing and agriculture in counties hit by the casualty shock. We find that the casualty rate is negatively associated with the share of wage workers and salaried employees in manufacturing in a county but with a higher share of self-employed individuals, which likely reflects worse labor market outcomes. The occupational structure in agriculture seems to be less affected by the casualty shock.

4.4.2 Agricultural Sector

We now turn to examine the effect of the casualty shock on the agricultural sector in more detail, as it seems to be less affected by the casualty shock. Furthermore, the results for sectoral composition from Table 4.3 suggest that the agricultural sector may have absorbed some of the individuals previously working in trade or manufacturing. We use information from the comprehensive agricultural censuses conducted in 1907 and 1925. The censuses contain rich data on the total number of farms by county, the total amount of plotted land in hectare and number of farm workers. All information is provided for five different farm sizes, very small (below 2 ha), small (2 - 5 ha), medium sized (5 - 20 ha), large (20 - 100 ha), and very large (above 100 ha) farms. With these data, we are able to analyze the effect of casualties both on the overall importance of the agricultural sector in the local economy as well as on distributional responses within the sector. In order to assess the effect of casualties on technology adoption, we use further county-level data on the number of machines owned by farms in 1925 in Prussia,[15] as well as district-level ('Regierungsbezirk') data on machine use in agriculture in 1907 and 1925.

Table 4.4 presents the results of difference-in-differences estimates using the number of individuals working on farms, the land size and the number of farms by farm size group as the dependent variable.

The results point to a consolidation of farms and land within the agricultural sector due to the casualty shock. The number of farms, employees and land cultivated by very small farms decreases with a higher casualty rate. The land is absorbed by medium-sized farms and the number of farms in this size group increases, but not the number of workers. Overall, the total number of farms and

[15] Unfortunately, these data are not available for the pre-war period.

Table 4.4: Effect of Casualty Rate on Agricultural Sector

	Total	Very Small <2ha	Small <5ha	Medium <20ha	Large <100ha	Very large >100ha
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: Log number of workers			
Post x Casualty Rate	-0.015***	-0.028***	-0.007***	0.001	-0.002	-0.014
	(0.003)	(0.004)	(0.002)	(0.002)	(0.003)	(0.009)
R^2	0.986	0.963	0.991	0.992	0.979	0.955
			Panel B: Log land size			
Post x Casualty Rate	-0.000	-0.018***	-0.003	0.006***	0.003	-0.004
	(0.001)	(0.002)	(0.002)	(0.002)	(0.004)	(0.012)
R^2	0.997	0.987	0.995	0.995	0.980	0.935
			Panel C: Log number of farms			
Post x Casualty Rate	-0.018***	-0.025***	-0.003*	0.005***	0.002	-0.002
	(0.003)	(0.004)	(0.002)	(0.002)	(0.002)	(0.003)
R^2	0.973	0.962	0.995	0.995	0.993	0.984
Post x Army District	Yes	Yes	Yes	Yes	Yes	Yes
Post x Female Share	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1556	1556	1556	1556	1556	1556

Notes: Results of regressions described in equation (4.1) for the two-period setup. Outcomes are the log number of agricultural workers (Panel A), the log number of hectares cultivated by a farms (B), and the log number of farms (C) in a county. Column (1) shows estimates for all farms in a county, (2) for very small farms with less than 2 hectares of land, (3) for small farms with between 2 and 5 hectares, (4) for medium farms with between 5 and 20 hectares, (5) for large farms with between 20 and 100 hectares and (6) for very large farms with more than 100 hectares. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

agricultural workers decreases. We can interpret these findings in two ways. First, the consolidation towards medium-sized farms could indicate an increase in the value of labor, hence leading to higher opportunity costs of owning and working on smaller farms. However, as we discuss above, we find that real wages of low-skilled labor actually decrease, especially in counties that were more exposed to the casualty shock. Furthermore, we would expect consolidation to benefit mostly large and very large farms which are able to substitute for labor through investments in technology. However, the two largest farm groups appear relatively unaffected by the casualty shock. A second possible explanation is that former workers and owners of small farms leave their farms and seek employment elsewhere through migration in response to the casualty shock and deteriorating economic conditions. The remaining farms absorb the newly-unused land but this is not accompanied by higher production intensity through increase in labor or labor-saving technologies. This in line with results in Table 4.B.3 in the Appendix which show that the number of workers per hectare decreased, especially in medium-sized farms which absorbed land from smaller farms. Furthermore, using district-level data on agricultural machines, such as mechanized ploughs, we do not find any increase both in the absolute number or the number of machines per hectare across all farm size groups. In fact, the use of such machines decreases with a higher casualty rate. Hence, we conclude that the casualty shock did not increase the adoption of labor-saving technologies or lead to more efficient production in agriculture which would have been in line with a labor supply effect of the casualty shock as found by Hornbeck and Naidu (2014) and Andersson et al. (forthcoming).

4.5 Robustness Checks

In this section, we briefly investigate the robustness of our main estimates. There are two main concerns: First, our estimates may be driven by specific groups of observations or capture war-related effects on the local capital stock if - contrary to contemporary accounts - the post-war occupation of Western German areas or the Russian advance into Eastern Prussia at the beginning of the war had lasting effects on the local capital stock and was correlated with the casualty rate. Furthermore, we want to ensure that our results are not due to outliers in the casualty rate and that the casualty rate accurately measures the size of the shock to the local population. Table 4.B.4 in the Appendix shows the results of robustness checks to address these issues for two of our main outcomes, the gender wage gap for workers over 21 and the female share in the local population. Column (1) in Table 4.B.4 repeats our main estimates and serves as a comparison. In column (2), we exclude urban counties from our sample to only consider rural areas to ensure that our results are not driven by either rural or urban counties. While this sample adjustment leads to slightly lower point estimates, our findings remain qualitatively unchanged. To exclude destruction of the local capital stock as an explanation for our results we drop the counties that were affected by the early Russian advance in 1914 in column (3) and all Western German counties that were occupied by French and Belgian forces in the post-war period in column (4). This leaves our estimates basically unchanged. Finally, we conduct two alterations to our estimation specification. In column (5), we code our continuous casualty rate into a dummy variable which distinguishes between above and below median values in the casualty rates variable. With this variable, we then estimate a standard difference-in-differences estimation with a binary treatment variable. This confirms our findings for the continuous treatment variable, as the results remain qualitatively similar and the coefficient is negative and highly significant. On average, above-median affected counties experienced an almost two percent increase in the gender wage gap, as well as a roughly 0.5 percent lower share of female inhabitants as compared to counties with a below-median exposure to WWI casualties. In Column (6), we use the death rate, i.e., the number of digitized entries indicating that a soldier died divided by pre-war male population as explanatory variable. This yields qualitatively similar results as using the casualty rate as explanatory variable. Overall, we conclude that our results are robust, both to potential destruction of the capital stock and to different specifications.

A specific concern inherent to our migration results is the fact that we observe this variable only in the post-war period and are therefore constrained to estimating cross-sectional OLS regressions only. To investigate the likelihood that unobserved differences that cannot be accounted for by the county fixed effects that we use in our difference-in-differences estimations bias our results, we provide a number of alternative estimations in Table 4.B.5. Here, we fade in our control variables one by one. Then, the variance of the point estimates of our variable of interest, the casualty share, tells us how sensible our estimates are towards the inclusion of additional controls. For both the female and male migration shares, we observe that the point estimates are remarkably stable across specifications. The male migration share is virtually unaffected by including or excluding any of the covariates. The female migration share is similarly robust, even though the point estimate decreases in absolute size by about 0.2 points when we control for urban/rural counties. Apparently, urban counties experienced a significantly higher in-migration which, once accounted for, decreases our point estimates slightly. Still, the point estimates are not significantly different from each other across specifications, which encourages our interpretation that the remaining bias in the cross-sectional estimates is at most modest. To further check the validity of our estimates, we also employ the procedure by Oster (2019). Following her best-practice recommendations, we find that the selection on unobservables would need to be at least 2.8 or 1.2 times larger than the selection on observables to explain our results ($\delta \geq 2.8$ for the male migration share, $\delta \geq 1.2$ for the female migration share).

4.6 Political Consequences

As a last part of our analysis, we examine how the casualty shock affected the political climate in counties with a higher casualty rate. A large literature argues that the rise of fascism in Germany was directly linked to the experience of loss in the First World War, such as increased acceptance of violence in the political arena or narratives such as the stab-in-the-back legend which scapegoated Jews and the political left for Germany's defeat (see, for instance, Ziemann, 2018; Hett, 2018). Koenig (2020) and De Juan et al. (2021) provide first micro-level evidence of this phenomenon by using either the return of veterans from the war or the share of dead soldiers among all World War I casualties as an exogenous shock to local exposure to the war. Both papers find that more exposed counties exhibit higher vote shares for extreme right parties, especially the Nazi party. At the same time, there are various studies that link economic deterioration to increased voter polarization and especially support for extreme right parties in the Weimar Republic (King et al., 2008; Voth et al., 2020; Galofré-Vilà et al., 2021).

We want to investigate whether our results allow us to establish a link from the economic deterioration caused by the casualty shock to support of right-wing parties in the post-war period. We again use the data by Falter and Hänisch (1990) which contain the results of all elections in the Weimar Republic from 1920 to 1933 at the county level. We roughly divide the parties into four groups across the political spectrum, extreme left, democratic, extreme right and the Nazi party. We run regressions with the county vote share for these party groups in the eight elections of the Weimar Republic as the dependent variable. We present the results in Figure 4.6.

Figure 4.6: Effect of Casualty Rate on Election Outcomes

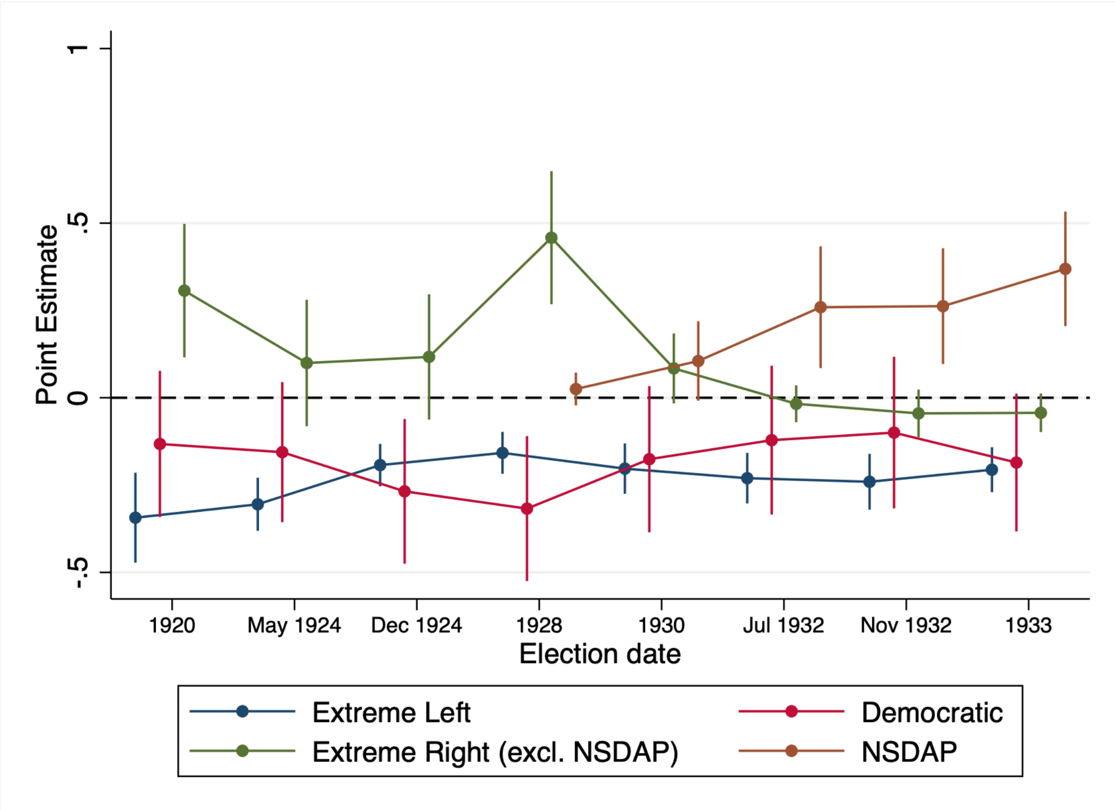

Note: This graph reports the coefficients and 95% Confidence Intervals from regressing the vote shares for different parties on the casualty share and further controls based on equation (4.2). Extreme left wing parties are the German Communist Party ('DKP') and the Independent Social Democratic Party ('USPD'). Democratic Parties are the Social Democratic Party ('SPD'), Center ('Zentrum'), German Democratic Party ('DDP') later known as German State Party ('DstP'). Extreme Right Wing Parties are the German National People's Party ('DNVP') and German People's Party ('DVP'). We control for army and administrative district fixed effects and the female share in 1916. Further controls are population density in 1910, a dummy indicating whether a county is urban, and the distance to the nearest coal deposit in kilometers.

Our findings confirm that extreme-right parties enjoyed higher support in counties with more WWI casualties. In all elections but the ones in 1924, the point estimate for extreme right wing parties is positive. For example in the 1928 elections, extreme right wing parties gained around three percentage points more votes for a one SD increase in the casualty share. Note that this was the last election when the NSDAP was not yet an important electoral force, collecting only 2.6% of the overall vote. Once the NSDAP emerged as a serious voting alternative in 1930, they alone gained around three percentage points in votes for a one SD increase in war casualties and the effect for other extreme right wing parties disappeared, suggesting that disaffected voters that previously supported the other extreme right wing parties switched their votes to the Nazi party, once it became a viable alternative.

We find the opposing effect for radical left wing parties. Across all elections, left wing parties collected fewer votes in counties that were more exposed to WWI casualties. Similarly, less-polarized parties which we term 'Democratic' and consist of parties like the Center ('Zentrum'), Social Democrats ('SPD') and German Democratic Party ('DDP') received significantly fewer votes across most interwar elections. Hence, the economic hardship caused by WWI casualties only benefited right wing parties, and drew votes away - especially from radical left wing parties. Hence, economic deterioration caused by the casualty shock may have been an important factor in the increased support of right wing parties in the Weimar Republic. Therefore, our results present an interesting addition to the results of Koenig (2020) and De Juan et al. (2021) who argue that this increase in support was a direct result of higher exposure to the war, not local economic conditions.

4.7 Conclusion

This paper examines how the loss of life during World War I affected the local German economy as well as its political environment. The German context of WWI provides an exceptional case study to investigate the effect of a severe population shock in isolation of other economic destruction that usually accompanies such shocks. Due to the fact that fighting mostly took place outside of German territory, the infrastructure and capital stock were left largely untouched. Hence, we can identify the sole effect of the shock on economic and political outcomes.

For our analysis, we collected, combined, and harmonized various unique data sets. First of all, we geocoded more than 8.5 million casualty list entries. We digitized several censuses and archival data on local economic outcomes, and harmonized them across the German territory to account for various reforms to Germany's subnational borders. This allowed us to not only investigate national trends in different economic outcomes, but to also estimate how counties that suffered more under WWI casualties deviated from this national trend in a difference-in-differences setting. We hope that these data sets will also serve as a resource for other scholars working on these topics.

Our results demonstrate that the labor supply shock induced by casualties had severe negative consequences for more affected counties. Wage growth decreased and gender wage disparities increased due to higher casualties, prompting outmigration by women and lower population growth. These changes appear to be associated with lower economic activity as evidenced by lower tax revenues and employment shares in industry and trade. Interestingly, counties with higher casualties experienced a temporary halt to industrialization, as workers moved more towards the agricultural sector. We find that this transition benefited medium-sized farms, mainly at the expense of small farms. We interpret this as evidence that workers moved towards the agricultural sector due to lack of employment opportunities in other sectors of the economy. We explain this finding by a severe demand shock resulting from decreased population due to casualties, which overshadowed the wage-promoting effect of labor scarcity which has been found to dominate in the previous literature on negative population shocks (Young, 2005; Voigtländer and Voth, 2013b; Hornbeck and Naidu, 2014). Finally, we find that these negative economic effects benefited the extreme right political parties and gave a significant boost to the electoral chances of the Nazi Party. As such, our paper also contributes to a larger literature on the political effects of wars (De Juan et al., 2021; Acemoglu et al., forthcoming).

4.A Geocoding of Casualties

We improved significantly on the existing geocoding of casualty entries. The original data set by Verein für Computergenealogie (2019a) was automatically geocoded via the Historic Gazetteer as part of a Master Thesis by Sen (2016). This automatic coding contained a series of coding errors. Our geocoding process limits these errors as we rely on a string-matching algorithm, which compares the manually digitized location names with a list of communities based on Germany's administrative structure of 1910, i.e., four years before the start of the war by Schubert (2020).

The original entries in the casualty lists included information on the birth place of a respective soldier. This information however follows no specific structure. The majority of entries consists of only one word, e.g., "Berlin," or a combination of two words separated by a comma, e.g., "Spandau, Berlin." In some instances, this information contains up to five separate words, all separated by a comma. Unfortunately, these comma-separated location descriptions do not follow a specified structure. In many instances, the first entry identifies a town or village, while the later entries identify higher administrative entities. In several instances however, this order may be reversed, e.g., to "Berlin, Spandau." Information on a higher local level does not always identify an official administrative entity, but local information on the region or the next bigger town. Examples for these are entries of the form "Calw, Schwarzwald" or "Heiligenstadt, Mühlhausen." Such cases complicate a direct matching of entries and our list of communities.

We proceeded in the following steps. First, we separated all entries by the comma to attach all distinct location information for each entry. Next, we assigned each list-entry to an entry in our community list, starting from the highest order of information. Our community list contains, for each community, information at five administrative levels, i.e., the "Land," "Provinz," "Bezirk," "Kreis," and "Gemeinde." For each casualty list entry, we started with the last information of the entry's string, for example with "Baden" in the entry "Neuenheim, Heidelberg, Baden." For this information, we started by looking for matches in the highest administrative level, i.e., the "Land," and then moving down to lower levels until a direct match was found. In the example here, we would directly find a match for "Baden" as a "Land." We would then proceed with the second part of information, looking at lower administrative levels in our community list, though subset to the region we already identified.

This direct pattern matching involved a number of difficulties. First, several entries were either misspelled or incorrectly typed during the digitization process. Therefore, we corrected numerous entries where we could not identify a direct match. Second, several entries would not result in an exact match. For example, the entry "Mühlhausen" could either refer to the town Mühlhausen near Erfurt in what is today the state of Thuringia, or to the district Mulhouse, which is part of the state or Alsace which today is in France. We worked through all instances of inexact matches manually, sorting the entries to the respective administrative area based on each entry's lower level information.

In the end, we were able to match around 7.5 million of the total 9 million entries to an entry at the fourth ("Kreis") or fifth ("Gemeinde") administrative level. Of the remaining entries, many do either not identify a location in Germany (e.g., "Philadelphia, Vereinigte Staaten") or only contain information on a higher administrative level (e.g., "Bayern") that does not allow matching it to the county level.

Our coding deviates in several ways from the original coding. In fact, the correlation between casualties per county according to our coding and the original coding is only 0.47. Figure 4.A.1 illustrates these coding disparities by plotting the casualties per county according to our coding over the casualties per county according to the original coding. Here, we first of all see three obvious cases where our coding leads to significantly more casualties than the old coding. Most obviously, Berlin sticks out with much more casualties according to our coding. We suspect that several Berlin-entries ended up in the neighboring district Niederbarnim (a significant outlier to the right), as the coordinates assigned to Berlin in the original coding were not accurately matching Berlin's administrative boundaries. Second, Frankfurt am Main and Mülheim an der Ruhr also deviate upwards. Here, we trace the difference to the original coding mistakenly sorting entries like "Frankfurt" to the (much less populated) "Frankfurt Oder" in Eastern Germany. Similarly, several cases for "Mülheim" were mistakenly allocated to Mühlheim am Rhein.

Figure 4.A.1: Differences in Geocoding

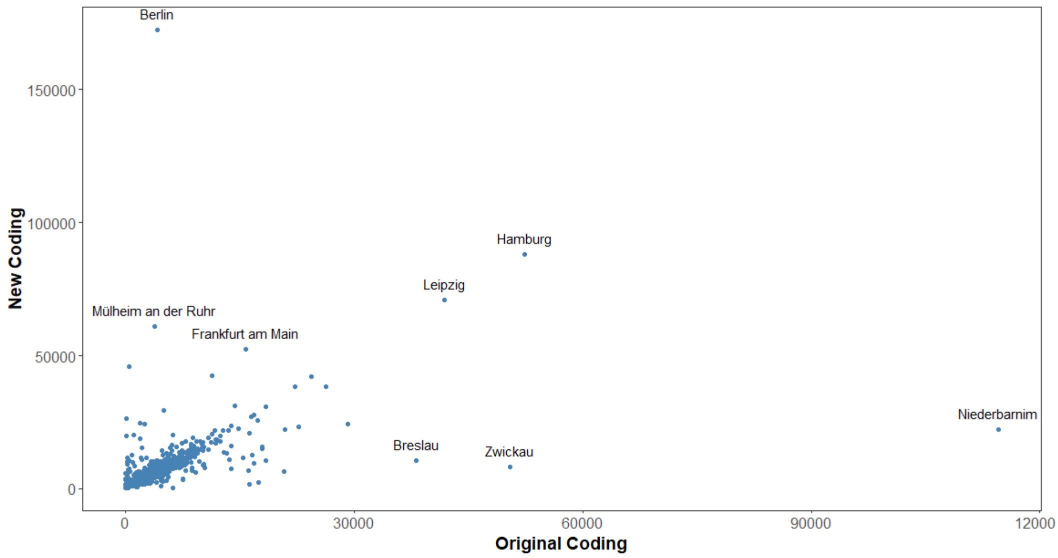

Note: This graph illustrates the differences in geocoded casualties between the original coding and our coding. On the horizontal axis, we plot the casualties per county according to the original coding, while the vertical axis displays the casualties per county according to our new coding.

There are also several deviations towards the right side, i.e., where the original coding suggests higher casualty counts as our coding. These deviations are better illustrated by the map in Figure 4.A.2. Here, dark-red to violet colors indicate counties where the ratio of casualties according to the original coding vs. our coding is between 2 and 15. In addition, we plot the centroids of higher administrative units, i.e., districts, provinces, and states. This shows that extraordinarily higher casualties in the original geocoding tend to overlap with administrative centroids. For example, Niederbarnim, the large outlier in the graph in the North-East, close to Berlin, contains the centroids of both the province Brandenburg and the state of Prussia. The many entries which only held the information "Preußen" or "Brandenburg" were sorted into Niederbarnim. Similar cases are, e.g., Breslau, Düsseldorf, Ebersberg, and Landau an der Isar.

Figure 4.A.2: Geographic Differences in Geocoding

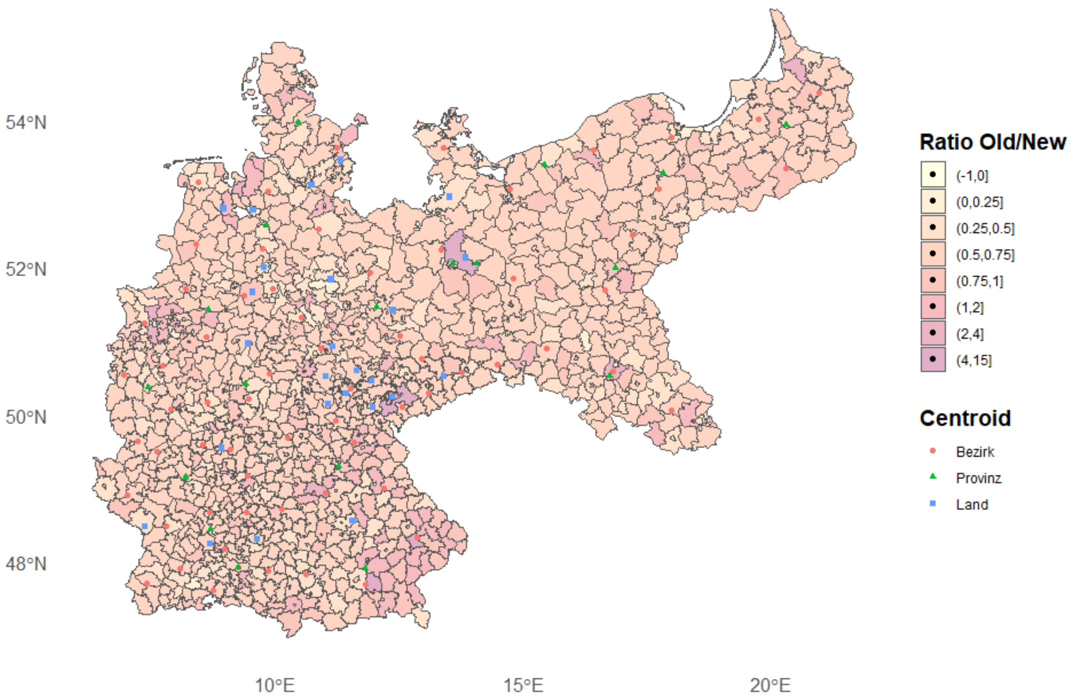

Note: This map illustrates the geographic variation in the differences in geocoded casualties. The color palette illustrates the ratio of a county's casualties in the original coding vs. our coding. The points depict the geographic centroids of higher administrative units, i.e., districts ("Bezirk"), provinces ("Provinz""), and states ("Land").

4.B Additional Figures and Tables

Figure 4.B.1: German Empire Army Districts in 1910

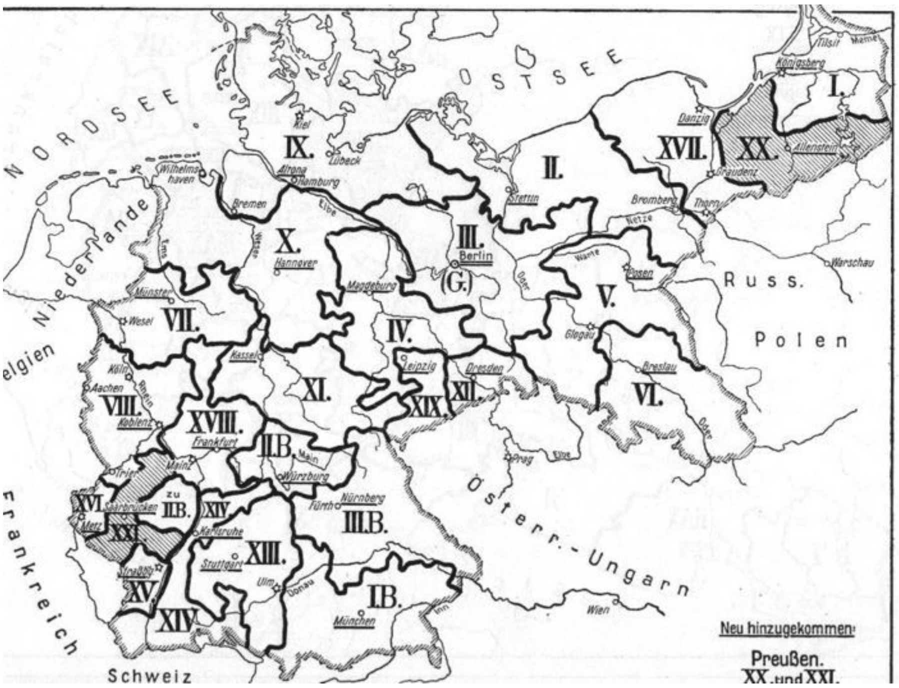

Notes: This figures displays the distribution of the 21 army districts of the German Empire at its boundaries of 1914.

Figure 4.B.2: Share of Dead and Share of Digitized Entries

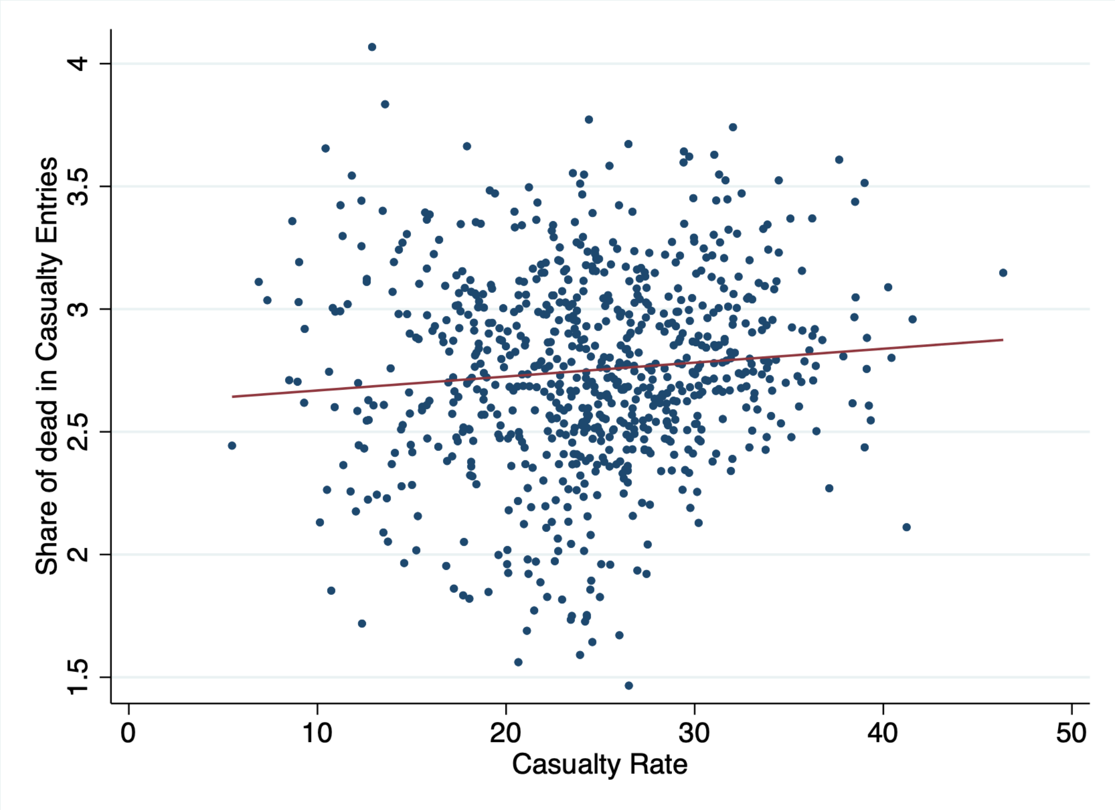

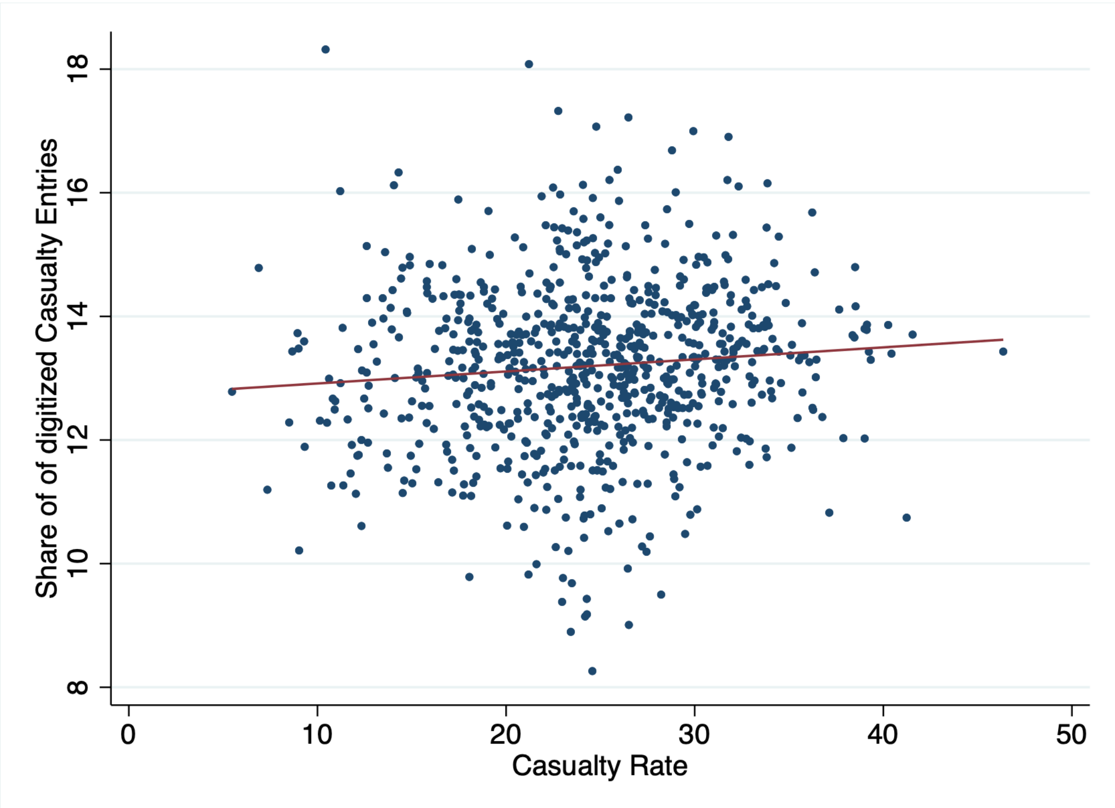

Note: Scatter plots of the share of dead entries (upper graph) and share of digitized entries (lower graph) on the y-axis and the share of casualty list entries on the x-axis. The red line represents a linear fit of the data.

Table 4.B.1: List of Urban Counties

County Type	Name	County Type	Name
Stadtkreis	Königsberg	unmittelbare Stadt	Amberg
Stadtkreis	Insterburg	unmittelbare Stadt	Neumarkt i. Oberpf
Stadtkreis	Potsdam	unmittelbare Stadt	Bamberg
Stadtkreis	Brandeburg	unmittelbare Stadt	Bayreuth
Stadtkreis	Landsberg (Warthe)	unmittelbare Stadt	Forchheim
Stadtkreis	Frankfurt an der Oder	unmittelbare Stadt	Hof
Stadtkreis	Guben	unmittelbare Stadt	Kulmbach
Stadtkreis	Cottbus	unmittelbare Stadt	Ansbach
Stadtkreis	Forst (Lausitz)	unmittelbare Stadt	Dinkelsbühl
Stadtkreis	Stettin	unmittelbare Stadt	Eichstätt
Stadtkreis	Stargard	unmittelbare Stadt	Fürth
Stadtkreis	Stolp in Pommern	unmittelbare Stadt	Rothenburg ob der Tauber
Stadtkreis	Stralsund	unmittelbare Stadt	Schwabach
Stadtkreis	Breslau	unmittelbare Stadt	Weißenburg i. Bayern
Stadtkreis	Brieg	unmittelbare Stadt	Aschaffenburg
Stadtkreis	Schweidnitz	unmittelbare Stadt	Kissingen
Stadtkreis	Liegnitz	unmittelbare Stadt	Kitzingen
Stadtkreis	Görlitz	unmittelbare Stadt	Schweinfurt
Stadtkreis	Oppeln	unmittelbare Stadt	Würzburg
Stadtkreis	Gleiwitz	unmittelbare Stadt	Freising
Stadtkreis	Ratibor	unmittelbare Stadt	Ingolstadt
Stadtkreis	Magdeburg	unmittelbare Stadt	Landsberg
Stadtkreis	Aschersleben	unmittelbare Stadt	München
Stadtkreis	Halberstadt	unmittelbare Stadt	Rosenheim
Stadtkreis	Halle an der Saale	unmittelbare Stadt	Traunstein
Stadtkreis	Weißenfels	unmittelbare Stadt	Deggendorf
Stadtkreis	Zeitz	unmittelbare Stadt	Landshut
Stadtkreis	Nordhausen	unmittelbare Stadt	Passau
Stadtkreis	Mühlhausen in Thüringen	unmittelbare Stadt	Straubing
Stadtkreis	Flensburg	unmittelbare Stadt	Dillingen
Stadtkreis	Neumünster	unmittelbare Stadt	Donauwörth
Stadtkreis	Wandsbek	unmittelbare Stadt	Günzburg
Stadtkreis	Altona	unmittelbare Stadt	Kaufbeuren
Stadtkreis	Göttingen	unmittelbare Stadt	Kempten
Stadtkreis	Celle	unmittelbare Stadt	Memmingen
Stadtkreis	Lüneburg	unmittelbare Stadt	Neuburg an der Donau
Stadtkreis	Harburg	unmittelbare Stadt	Neu-Ulm
Stadtkreis	Emden	unmittelbare Stadt	Nördlingen
Stadtkreis	Münster	Bezirksamt u. Stadt	Germersheim
Stadtkreis	Bielefeld	Stadt	Dresden
Stadtkreis	Hamm in Westfalen	Stadt	Plauen
Stadtkreis	Hagen	Stadt	Zwickau
Stadtkreis	Iserlohn	Stadt	Chemnitz
Stadtkreis	Lüdenscheid	Stadt	Stuttgart
Stadtkreis	Kassel	Stadt	Bremerhaven
Stadtkreis	Hanau	Stadt	Vegesack
Stadtkreis	Wiesbaden		
Stadtkreis	Koblenz		
Stadtkreis	Krefeld		
Stadtkreis	Duisburg		
Stadtkreis	Düsseldorf		
Stadtkreis	Elberfeld		
Stadtkreis	Remscheid		
Stadtkreis	Solingen		
Stadtkreis	Rhendt		
Stadtkreis	Bonn		
Stadtkreis	Trier		

Note: This table lists all urban counties in our final data set. The left side contains the names for all Prussian urban counties, the right side for all non-Prussian urban counties.

Table 4.B.2: Effect of Casualty Rate on Occupation Structure

	Agriculture			Manufacturing		
	Workers	Salaried	Self-Employed	Workers	Salaried	Self-Employed
	(1)	(2)	(3)	(4)	(5)	(6)
Casualty Rate	-0.053	-0.021	0.084*	-0.332***	-0.208***	0.520***
	(0.078)	(0.018)	(0.046)	(0.057)	(0.024)	(0.062)
R^2	0.752	0.543	0.724	0.411	0.620	0.489
Army District FE	Yes	Yes	Yes	Yes	Yes	Yes
Female Share 1916	Yes	Yes	Yes	Yes	Yes	Yes
Admin. District FE	Yes	Yes	Yes	Yes	Yes	Yes
Further Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	778	778	778	778	778	778

Notes: Results of regressions described in equation (4.2) for the one-period setup. Outcomes are the share of workers (‘Arbeiter’), salaried employees (‘Angestellte’) and self-employed (‘Selbstständige’) among all individuals working in Agriculture (columns (1)-(3)) and in manufacturing (columns (4)-(6)) in 1925. Further controls are population density in 1910, a dummy indicating whether a county is urban, and the distance to the nearest coal deposit in kilometers. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 4.B.3: Developments in the Agricultural Sector

(a) Number of Employees

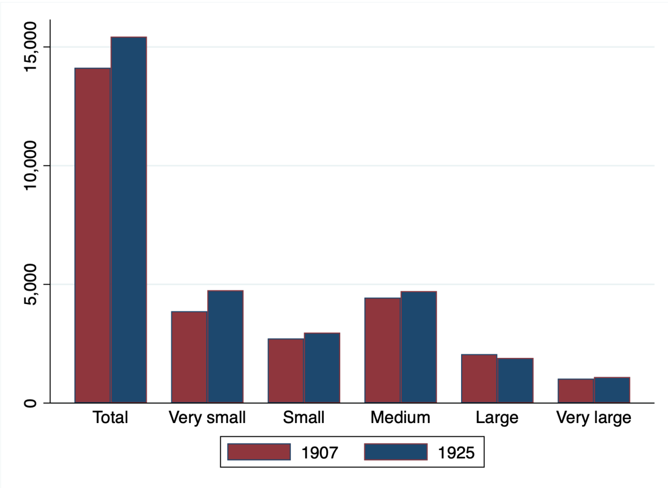

(b) Number of Farms

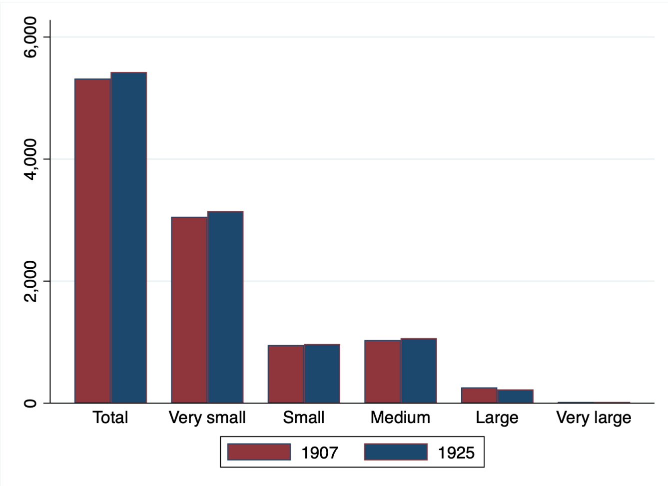

(c) Total area (in hectares

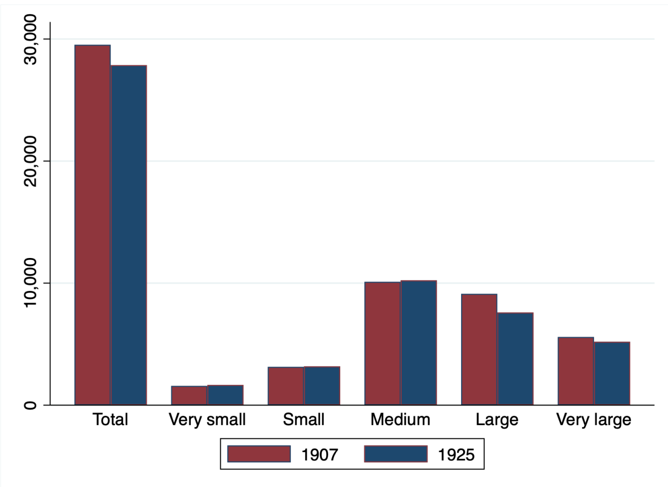

Notes: The graphs report the change in agricultural employees (A), farms (B), and cultivated area (C) between 1907 and 1925. We report these numbers separately by farm size, where "very small" farms constitute farms with less than 2ha area, small farms are between 2–5ha, medium farms between 5–20ha, large farms between 20–100ha, and very large farms are above 100ha.

Table 4.B.3: Effect of Casualty Rate on Agricultural Labor and Technology Use

	Total	Very Small <2ha	Small <5ha	Medium <20ha	Large <100ha	Very large >100ha
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Number of workers per hectare					
Post x Casualty Rate	-0.008***	-0.008***	-0.002***	-0.002***	-0.001*	-0.003*
	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)
Observations	1556	1556	1556	1556	1524	1146
R^2	0.932	0.809	0.828	0.890	0.776	0.704
	Panel B: Log number of mechanized ploughs					
Post x Casualty Rate	-0.000	-0.032**	-0.023	-0.036	-0.003	0.043
	(0.028)	(0.015)	(0.021)	(0.031)	(0.033)	(0.029)
R^2	0.940	0.607	0.582	0.783	0.887	0.917
	Panel C: Number of mechanized ploughs per hectare					
Post x Casualty Rate	0.005	-0.003**	-0.002	-0.002	-0.008	-0.024**
	(0.004)	(0.001)	(0.001)	(0.002)	(0.007)	(0.011)
Observations	138	138	138	138	138	136
R^2	0.810	0.584	0.573	0.687	0.733	0.895
Post x Army District	Yes	Yes	Yes	Yes	Yes	Yes
Post x Female Share	Yes	Yes	Yes	Yes	Yes	Yes
County / District FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Results of regressions described in equation (4.1) for the two-period setup. Outcomes are the number of agricultural workers per hectare of land (Panel A), the log number of mechanized ploughs used (B), and the number of mechanized ploughs per hectare (C) in 1907 and 1925 as dependent variable. The data are at the county-level in Panel A and at the district-level in Panel B and C. Column (1) shows estimates for all farms in a county, (2) for very small farms with less than 2 hectares of land, (3) for small farms with between 2 and 5 hectares, (4) for medium farms with between 5 and 20 hectares, (5) for large farms with between 20 and 100 hectares and (6) for very large farms with more than 100 hectares. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4.B.4: Robustness Checks for DiD-Estimates

		Dropping Observations			Alternative Specifications	
	Baseline	Urban Counties	Russian Advance	Ruhr Occupation	Median Split	Dead Rate
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Gender Wage Gap over 21					
Post x Casualty Rate	-0.223***	-0.180***	-0.227***	-0.176***		
	(0.050)	(0.055)	(0.051)	(0.050)		
Post x Above Median Casualty Rate=1					-1.977***	
					(0.590)	
Post x Dead Rate						-7.307***
						(1.621)
Observations	1552	1348	1492	1376	1556	1552
R^2	0.791	0.802	0.789	0.812	0.790	0.791
	Panel B: Female Share					
Post x Casualty Rate	-0.069***	-0.051***	-0.068***	-0.068***		
	(0.009)	(0.009)	(0.010)	(0.010)		
Post x Above Median Casualty Rate=1					-0.496***	
					(0.080)	
Post x Dead Rate						-1.907***
						(0.269)
Observations	1556	1352	1496	1378	1558	1556
R^2	0.878	0.896	0.880	0.871	0.870	0.875
Post x Army District	Yes	Yes	Yes	Yes	Yes	Yes
Post x Female Share	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Robustness checks for regressions described in equation (4.1) for the two-period setup. Column (1) shows estimates from Table 4.2 for the ratio of female to male wages for workers aged over 21 (Panel A) in 1914 and 1922 and the share of female population in 1910 and 1925 (Panel B). Column (2) shows estimates for the same outcomes without any urban counties (for a list see Table 4.B.1 in the Appendix). In column (3), all counties affected by the initial Russian offensive into Eastern Prussia in 1914 are dropped and in Column (4) all counties affected by the post-war occupation by Allied forces are dropped. Column (5) uses a dummy for a median split of the casualty rate as explanatory variable and column (6) the number of casualty list entries classified as dead divided by pre-war male population as explanatory variable. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4.B.5: Robustness Checks for OLS-Estimates

	Baseline (1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Female net migration share					
Casualty Rate	-0.466***	-0.480***	-0.640***	-0.666***	-0.452***	-0.571***
	(0.071)	(0.047)	(0.063)	(0.075)	(0.075)	(0.062)
Share Female 1916	0.278		0.924***	0.857**	0.200	0.695*
	(0.343)		(0.312)	(0.352)	(0.303)	(0.358)
Pop Density 1910	-0.001					0.001**
	(0.001)					(0.000)
Urban=1	9.493***				8.069***	
	(1.939)				(1.351)	
Distance to Coal	0.004					0.010
	(0.010)					(0.007)
R^2	0.331	0.116	0.208	0.272	0.328	0.218
	Panel B: Male net migration share					
Casualty Rate	-0.266***	-0.231***	-0.373***	-0.380***	-0.267***	-0.324***
	(0.071)	(0.049)	(0.062)	(0.072)	(0.073)	(0.061)
Share Female 1916	-0.125		0.229	0.215	-0.125	0.059
	(0.327)		(0.282)	(0.307)	(0.301)	(0.322)
Pop Density 1910	-0.000					0.001**
	(0.001)					(0.000)
Urban=1	4.327**				4.223***	
	(1.998)				(1.389)	
Distance to Coal	0.006					0.004
	(0.011)					(0.007)
R^2	0.238	0.034	0.144	0.218	0.237	0.150
Observations	724	724	724	724	724	724
Army District FE	Yes	No	Yes	Yes	Yes	Yes
Admin. District FE	Yes	No	No	Yes	Yes	No

Notes: Specification checks for regressions described in equation (4.2) for the one-period setup. Column (1) shows estimates from Table 4.2 for the share of net migration between 1910 and 1925 for women (Panel A) and men (Panel B). $ShareFemale1916$ is the share of female population in 1916 and a proxy for draft intensity. $Urban$ is a dummy indicating whether a county is urban. $PopDensity1910$ is population density in 1910 and $DistancetoCoal$ is the distance of a county centroid from the nearest coal deposit. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Bibliography

Abadie, Alberto, Susan Athey, Guido Imbens, and Jeffrey Wooldridge (2017) "When should you adjust Standard Errors for Clustering?," *NBER Working Paper*, Vol. 24003.

Acemoglu, Daron (1998) "Why do new Technologies complement Skills? Directed Technical Change and Wage Inequality," *Quarterly Journal of Economics*, Vol. 113, pp. 1055–1089.

——— (2010) "When does Labor Scarcity encourage Innovation?," *Journal of Political Economy*, Vol. 118, pp. 1037–1078.

Acemoglu, Daron, David H Autor, and David Lyle (2004) "Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury," *Journal of Political Economy*, Vol. 112, pp. 497–551.

Acemoglu, Daron, Giuseppe De Feo, Giacomo De Luca, and Gianluca Russo (forthcoming) "War, Socialism, and the Rise of Fascism: An Empirical Exploration," *Quarterly Journal of Economics*.

Afflerbach, Holger (2018) *Auf Messers Schneide - Wie das Deutsche Reich den Ersten Weltkrieg verlor*, München: C.H. Beck.

Akhmedov, Akhmed and Ekaterina Zhuravskaya (2004) "Opportunistic political Cycles: Test in a young Democracy Setting," *Quarterly Journal of Economics*, Vol. 179, pp. 1301–1338.

Albouy, David (2013) "Partisan Representation in Congress and the geographic Distribution of Federal Funds," *Review of Economics and Statistics*, Vol. 95, pp. 127–141.

Andersson, David, Mounir Karadja, and Erik Prawitz (forthcoming) "Mass Migration and Technological Change," *Journal of the European Economic Association*.

Asher, Sam and Paul Novosad (2017) "Politics and local economic Growth: Evidence from India," *American Economic Journal: Applied Economics*, Vol. 9, pp. 229–273.

Audretsch, David B and Maryann P Feldman (1996) "R&D Spillovers and the Geography of Innovation and Production," *American Economic Review*, Vol. 86, pp. 630–640.

Autor, David, David Dorn, Gordon Hanson, and Kaveh Majlesi (2020) "Importing political Polarization? The electoral Consequences of rising Trade Exposure," *American Economic Review*, Vol. 110, pp. 3139–83.

Azulai, Michel (2017) "Public Good Allocation and the Welfare Costs of political Connections: Evidence from Brazilian Matching Grants," *Working Paper*.

Bania, Neil, Randall W. Eberts, Michael S. Fogarty et al. (1993) "Universities and the Startup of New Companies: Can we generalize from Route 128 and Silicon Valley?" *Review of Economics and Statistics*, Vol. 75, pp. 761–766.

Baskaran, Thushyanthan and Zohal Hessami (2017) "Political Alignment and Intergovernmental Transfers in Parliamentary Systems: Evidence from Germany," *Public Choice*, Vol. 171, pp. 75–98.

Beaudry, Paul, Mark Doms, and Ethan Lewis (2010) "Should the Personal Computer be considered a Technological Revolution? Evidence from US Metropolitan Areas," *Journal of Political Economy*, Vol. 118, pp. 988–1036.

Beaudry, Paul and David A Green (2003) "Wages and Employment in the United States and Germany: What explains the Differences?" *American Economic Review*, Vol. 93, pp. 573–602.

——— (2005) "Changes in US Wages, 1976–2000: Ongoing Skill Bias or major Technological Change?" *Journal of Labor Economics*, Vol. 23, pp. 609–648.

Becker, S., P.H. Egger, and M. Von Ehrlich (2013) "Absorptive Capacity and the Growth Effects of Regional Transfers: A Regression Discontinuity Design with Heterogeneous Treatment Effects," *American Economic Journal: Economic Policy*, Vol. 5(4), pp. 29–77.

Becker, Sascha O., Peter H. Egger, and Maximilian von Ehrlich (2010) "Going NUTS: The effect of EU Structural Funds on Regional Performance.," *Journal of Public Economics*, Vol. 94, pp. 578–590.

Beerli, Andreas, Jan Ruffner, Michael Siegenthaler, and Giovanni Peri (2021) "The Abolition of Immigration Restrictions and the Performance of Firms and Workers: Evidence from Switzerland," *American Economic Review*, Vol. 111, pp. 976–1012.

Beeson, Patricia and Edward B. Montgomery (1993) "The Effects of Colleges and Universities on Local Labor Markets," *Review of Economics and Statistics*, Vol. 75, pp. 753–61.

Bertelsmann Stiftung (2008) *Beruf Bürgermeister/in - Eine Bestandsaufnahme für Deutschland*, Gütersloh: Bertelsmann Stiftung.

Besley, Timothy and Anne Case (1995) "Does electoral Accountability affect economic Policy Choices? Evidence from gubernatorial Term Limits," *Quarterly Journal of Economics*, Vol. 110, pp. 769–798.

Blundell, R., D.A. Green, and W. Jin (2022) "The U.K. as a Technological Follower: Higher Education Expansion and the College Wage Premium," *Review of Economic Studies*, Vol. 89, pp. 142–180.

BMWi (2013) *Wirtschaftsdaten Neue Bundesländer*: Bundesministerium für Wirtschaft und Technologie.

BmWK (2022) *Boosting the regional Economy*, Berlin: Federal Ministry for Economic Affairs and Climate Action, URL: `https://www.bmwk.de/Redaktion/EN/Dossier/regional-policy.html`.

Boehnke, Jörn and Victor Gay (forthcoming) "The Missing Men: World War I and Female Labor Force Participation," *Journal of Human Resources*.

Borjas, George J (1989) "Economic Theory and International Migration," *International Migration Review*, Vol. 23, pp. 457–485.

——— (2003) "The Labor Demand Curve is downward sloping: Reexamining the Impact of Immigration on the Labor Market," *Quarterly Journal of Economics*, Vol. 118, pp. 1335–1374.

Boustan, Leah Platt, Matthew E Kahn, and Paul W Rhode (2012) "Moving to higher Ground: Migration Response to natural Disasters in the early Twentieth Century," *American Economic Review*, Vol. 102, pp. 238–44.

Bracco, Emanuele, Ben Lockwood, Francesco Porcelli, and Michela Redoano (2015) "Intergovernmental Grants as Signals and the Alignment Effect: Theory and Evidence," *Journal of Public Economics*, Vol. 123, pp. 78–91.

Brainerd, Elizabeth and Mark V Siegler (2003) "The Economic Effects of the 1918 Influenza Epidemic," *CEPR Discussion Papers*, Vol. 3791.

Brakman, Steven, Harry Garretsen, and Marc Schramm (2004) "The Strategic Bombing of German Cities during World War II and its Impact on City Growth," *Journal of Economic Geography*, Vol. 4, pp. 201–218.

Brender, Adi and Allan Drazen (2005) "Political Budget Cycles in New versus Established Democracies," *Journal of Monetary Economics*, Vol. 52, pp. 1271–1295.

Brockmann, Paul, Christoph Halbmeier, and Eva Sierminska (2022) "Geocoded Tax Data for the German Interwar Period: A Novel Database for Regional Analyses," *Working Paper*.

Brollo, Fernanda and Tommaso Nannicini (2012) "Tying your Enemy's Hands in close Races: The Politics of Federal Transfers in Brazil," *American Political Science Review*, Vol. 106, pp. 742–761.

Bronzini, R. and G. de Blasio (2012) "Evaluating the Impact of Investment Incentives: The Case of Italy's Law 488/1992," *Journal of Urban Economics*, Vol. 60, pp. 327–349.

Brüll, Eduard and Christina Gathmann (2020) "Evolution of the East German Wage Structure," *ZEW Discussion Paper*, Vol. 20-081.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel (2015) "The Value of Democracy: Evidence from Road Building in Kenya," *American Economic Review*, Vol. 105, pp. 1817–1851.

Busso, M., J. Gregory, and P. Kline (2013) "Assessing the Incidence and Efficiency of a prominent Place-based Policy," *American Economic Review*, Vol. 103, pp. 897–947.

Calonico, Sebastian, Matias D Cattaneo, Max H Farrell, and Rocío Titiunik (2019) "Regression Discontinuity Designs using Covariates," *Review of Economics and Statistics*, Vol. 101, pp. 442–451.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik (2014) "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, Vol. 82, pp. 2295–2326.

Card, David (2001) "Immigrant Inflows, Native Outflows, and the local Labor Market Impacts of higher Immigration," *Journal of Labor Economics*, Vol. 19, pp. 22–64.

Card, David, Jörg Heining, and Patrick Kline (2013) "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, Vol. 128, pp. 967–1015.

Card, David and Thomas Lemieux (2001) "Can falling Supply explain the rising Return to College for Younger Men? A Cohort-Based Analysis," *Quarterly Journal of Economics*, Vol. 116, pp. 705–746.

Card, David and Ethan Lewis (2007) "The Diffusion of Mexican Immigrants during the 1990s: Explanations and Impacts," in George Borjas ed. *Mexican Immigration to the United States*, Cambridge, MA: National Bureau of Economic Research Conference Report.

Carneiro, Pedro, Kai Liu, and Kjell G Salvanes (forthcoming) "The Supply of Skills and endogenous Technical Change: Evidence from a College Expansion Reform," *Journal of the European Economic Association*.

Case, Anne and Angus Deaton (2020) *Deaths of Despair and the Future of Capitalism*, Princeton: Princeton University Press.

Caselli, Francesco and Wilbur John Coleman (2006) "The World Technology Frontier," *American Economic Review*, Vol. 96, pp. 499–522.

Cattaneo, Matias D, Michael Jansson, and Xinwei Ma (2018) "Manipulation Testing based on Density Discontinuity," *Stata Journal*, Vol. 18, pp. 234–261.

Ciccone, Antonio (2021) "Gibrat's Law for Cities: Evidence from World War I Casualties," *CESifo Working Paper*, Vol. 9006.

Ciccone, Antonio and Robert E Hall (1996) "Productivity and the Density of Economic Activity," *American Economic Review*, Vol. 86, pp. 54–70.

Ciccone, Antonio and Giovanni Peri (2006) "Identifying Human Capital Externalities: Theory with Applications," *Review of Economic Studies*, Vol. 73, pp. 381–412.

Cohen, Wesley M, Richard R Nelson, John P Walsh et al. (2002) "Links and Impacts: The Influence of Public Research on Industrial R&D," *Management Science*, Vol. 48, pp. 1–23.

Criscuolo, Chiara, Ralf Martin, Henry Overman, and John Van Reenen (2019) "Some causal Effects of an industrial Policy," *American Economic Review*, Vol. 109, pp. 48–85.

Curto-Grau, Marta, Albert Solé-Ollé, and Pilar Sorribas-Navarro (2018) "Does Electoral Competition curb Party Favoritism?," *American Economic Journal: Applied Economics*, Vol. 10, pp. 378–407.

De Juan, Alexander, Felix Haass, Carlo Koos, Sascha Riaz, and Thomas Tichelbaecker (2021) "War and Nationalism: Evidence from World War I and the Rise of the Nazi Party," *Working Paper*.

Desmet, Klaus and J Vernon Henderson (2015) "The Geography of Development within Countries," in *Handbook of Regional and Urban Economics, vol. 5*, Vol. 5: Elsevier B.V. pp. 1457–1517.

Dustmann, Christian and Albrecht Glitz (2015) "How do Industries and Firms respond to Changes in Local Labor Supply?" *Journal of Labor Economics*, Vol. 33, pp. 711–750.

Dustmann, Christian, Uta Schönberg, and Jan Stuhler (2016) "Labor Supply Shocks, Native Wages, and the Adjustment of Local Employment," *Quarterly Journal of Economics*, Vol. 132, pp. 435–483.

Eggers, Andrew C, Anthony Fowler, Jens Hainmueller, Andrew B Hall, and James M Snyder (2015) "On the Validity of the Regression Discontinuity Design for estimating electoral Effects: New Evidence from over 40,000 close Races," *American Journal of Political Science*, Vol. 59, pp. 259–274.

von Ehrlich, M. and T. Seidel (2018) "The persistent Effects of Place-based Policy? Evidence from the West-German Zonenrandgebiet," *American Economic Journal: Economic Policy*, Vol. 10, pp. 344–74.

Engelmaier, Florian and Till Stowasser (2017) "Electoral Cycles in Savings Bank Lending," *Journal of the European Economic Association*, Vol. 15, pp. 296–354.

European Commission (2021) "Cohesion Policy Regulation 2021-2027," *Official Journal of the European Union*, Vol. L231.

Eurostat (2022) *Unemployment Statistics at Regional Level*, Brussels: European Commission, URL: `https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Unemployment_statistics_at_regional_level`.

Faggio, Giulia and Henry Overman (2014) "The effect of public Sector Employment on local Labour Markets," *Journal of Urban Economics*, Vol. 79, pp. 91–107.

Falter, Jürgen W and Dirk Hänisch (1990) "Wahl- und Sozialdaten der Kreise und Gemeinden des Deutschen Reiches von 1920 bis 1933," GESIS Data base, Data set ZA8013, Köln.

Federal Republic of Germany (1949) *Basic Law for the Federal Republic of Germany (as amended September 2020)*, Berlin: Federal Ministry of Justice.

Feigenbaum, James J, James Lee, and Filippo Mezzanotti (2018) "Capital Destruction and Economic Growth: The Effects of Sherman's March, 1850-1920," *NBER Working Paper*, Vol. 25392.

Ferraz, Claudio and Frederico Finan (2011) "Electoral Accountability and Corruption: Evidence from the Audits of Local Governments," *American Economic Review*, Vol. 101, pp. 1274–1311.

Fetzer, Thiemo (2019) "Did Austerity cause Brexit?," *American Economic Review*, Vol. 109, pp. 3849–86.

Fitzenberger, Bernd and Karsten Kohn (2006) "Skill Wage Premia, Employment, and Cohort Effects: Are Workers in Germany all of the same Type?," *ZEW Discussion Papers*, Vol. 06-044.

Fitzenberger, Bernd, Aderonke Osikominu, Robert Völter et al. (2006) "Imputation Rules to improve the Education Variable in the IAB Employment Subsample," *Schmollers Jahrbuch: Journal of Applied Social Science Studies/Zeitschrift für Wirtschafts-und Sozialwissenschaften*, Vol. 126, pp. 405–436.

Freier, Ronny (2015) "The Mayor's Advantage: Causal Evidence on Incumbency Effects in German Mayoral Elections," *European Journal of Political Economy*, Vol. 40, pp. 16–30.

Fuest, Clemens and Lea Immel (2021) "University Openings and their Long-term Impact on Regional Wages: Evidence from West Germany," *ifo Dresden berichtet*, Vol. 28, pp. 24–28.

Galofré-Vilà, Gregori, Christopher M Meissner, Martin McKee, and David Stuckler (2021) "Austerity and the Rise of the Nazi Party," *The Journal of Economic History*, Vol. 81, pp. 81–113.

Gehne, David H (2012) *Bürgermeister - Führungskraft zwischen Bürgerschaft, Rat und Verwaltung*, Stuttgart: Richard Boorberg Verlag.

Gehne, David H and Lars Holtkamp (2002) "Bürgermeisterkandidaten zwischen Verwaltungsprofis, Parteisoldaten und Schützenkönigen," in *Im Westen was Neues*, Wiesbaden: VS Verlag für Sozialwissenschaften, pp. 55–88.

Gehring, Kai and Stephan A Schneider (2018) "Toward the Greater Good? EU Commissioners' Nationality and Budget Allocation in the European Union," *American Economic Journal: Economic Policy*, Vol. 10, pp. 214–239.

Gelman, Andrew and Guido Imbens (2019) "Why high-order Polynomials should not be used in Regression Discontinuity Designs," *Journal of Business & Economic Statistics*, Vol. 37, pp. 447–456.

Glaeser, Edward L. (1999) "Learning in Cities," *Journal of Urban Economics*, Vol. 46, pp. 254–277.

Glitz, Albrecht (2012) "The Labor Market Impact of Immigration: A Quasi-Experiment exploiting Immigrant Location Rules in Germany," *Journal of Labor Economics*, Vol. 30, pp. 175–213.

Glitz, Albrecht and Daniel Wissmann (2017) "Skill Premiums and the Supply of young Workers in Germany," *CESifo Working Paper*, Vol. 6576.

Gobillon, Laurent, Thierry Magnac, and Harris Selod (2012) "Do unemployed Workers benefit from Enterprise Zones: The French Experience," *Journal of Public Economics*, Vol. 96, pp. 881–892.

Golden, Miriam and Brian Min (2013) "Distributive Politics around the World," *Annual Review of Political Science*, Vol. 16, pp. 73–99.

Goldin, Claudia D (1991) "The Role of World War II in the Rise of Women's Employment," *American Economic Review*, pp. 741–756.

Goldin, Claudia and Claudia Olivetti (2013) "Shocking Labor Supply: A Reassessment of the Role of World War II on Women's Labor Supply," *American Economic Review*, Vol. 103, pp. 257–62.

Grant, Oliver (2005) *Migration and Inequality in Germany 1870-1913*, Oxford: Oxford University Press.

Grossman, Philip J (1994) "A Political Theory of Intergovernmental Grants," *Public Choice*, Vol. 78, pp. 295–303.

Habakkuk, Hrothgar John (1962) *American and British Technology in the Nineteenth Century: The Search for Labour Saving Inventions*, Cambridge: Cambridge University Press.

Haug, Heike-Franziska and Heinz-Werner Hetmeier (2003) *Bericht zur finanziellen Lage der Hochschulen*, Wiesbaden: Statistisches Bundesamt.

Helfferich, Karl (1925) *Der Weltkrieg*, Berlin: Ullstein & Co.

Hethey, Tanja and Johannes F. Schmieder (2010) "Using Worker Flows in the Analysis of Establishment Turnover: Evidence from German administrative Data," FDZ-Methodenreport 06/2010, Institut für Arbeitsmarkt- und Berufsforschung (IAB), Nürnberg.

Hett, Benjamin Carter (2018) *The Death of Democracy: Hitler's rise to Power and the Downfall of the Weimar Republic*, New York: Henry Holt and Company.

Hodler, Roland and Paul A Raschky (2014) "Regional Favoritism," *Quarterly Journal of Economics*, Vol. 129, pp. 995–1033.

Holtmann, Everhard, Christian Rademacher, and Marion Reiser (2017) *Kommunalpolitik - Eine Einführung*, Wiesbaden: Springer.

Holuscha, Elisabeth (2012) "Das Prinzip Fachhochschule: Erfolg oder Scheitern? Eine Fallstudie am Beispiel Nordrhein-Westfalen," Ph.D. dissertation, Marburg University, Germany.

Hornbeck, Richard and Suresh Naidu (2014) "When the Levee Breaks: Black Mgration and Economic Development in the American South," *American Economic Review*, Vol. 104, pp. 963–90.

Jaffe, Adam B (1989) "Real Effects of Academic Research," *American Economic Review*, Vol. 79, pp. 957–970.

Johns Hopkins University of Medicine (2022) "Covid19 Mortality Analyses," Website: https://coronavirus.jhu.edu/data/mortality.

Jünger, Ernst (1920) *In Stahlgewittern*, Stuttgart: Klett-Cotta.

Kantor, Shawn and Alexander Whalley (2014) "Knowledge Spillovers from Research Universities: Evidence from Endowment Value Shocks," *Review of Economics and Statistics*, Vol. 96, pp. 171–188.

——— (2019) "Research Proximity and Productivity: Long-term Evidence from Agriculture," *Journal of Political Economy*, Vol. 127, pp. 819–854.

Karlsson, Martin, Therese Nilsson, and Stefan Pichler (2014) "The Impact of the 1918 Spanish Flu Epidemic on economic Performance in Sweden: An Investigation into the Consequences of an extraordinary Mortality Shock," *Journal of Health Economics*, Vol. 36, pp. 1–19.

Katz, L. F. and K. M. Murphy (1992) "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, Vol. 107, pp. 35–78.

Kersting, Felix and Nikolaus Wolf (2021) "On the Origins of National Identity. German Nation-Building after Napoleon," *CEPR Discussion Paper*, Vol. 16314.

King, Gary, Ori Rosen, Martin Tanner, and Alexander F Wagner (2008) "Ordinary Economic Voting Behavior in the Extraordinary Election of Adolf Hitler," *The Journal of Economic History*, Vol. 68, pp. 951–996.

Koenig, Christoph (2020) "Loose Cannons: War Veterans and the Erosion of Democracy in Weimar Germany," *CAGE Working Paper*, Vol. 497.

Kruppe, Thomas, Britta Matthes, and Stefanie Unger (2014) "Effectiveness of Data Correction Rules in process-produced Data. The Case of educational Attainment."

Kulicke, Marianne and Thomas Stahlecker (2004) *Forschungslandkarte Fachhochschulen Potenzialstudie*, Bonn, Berlin: Bundesministerium für Bildung und Forschung (BMBF).

Landtag, Schleswig-Holsteinischer (1991) "Landeshochschulplan für Schleswig-Holstein," Drucksache 12/1632, Kiel.

Lee, David S and Thomas Lemieux (2010) "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, Vol. 48, pp. 281–355.

Lehnert, Patrick, Curdin Pfister, and Uschi Backes-Gellner (2020) "Employment of R&D Personnel after an educational Supply Shock: Effects of the Introduction of Universities of Applied Sciences in Switzerland," *Labour Economics*, Vol. 66, pp. 1018–83.

Lewis, Ethan (2011) "Immigration, Skill Mix, and Capital Skill Complementarity," *Quarterly Journal of Economics*, Vol. 126, pp. 1029–1069.

Lucas, Robert Jr. (1988) "On the Mechanics of Economic Development," *Journal of Monetary Economics*, Vol. 22, pp. 3–42.

Manacorda, Marco, Alan Manning, and Jonathan Wadsworth (2012) "The Impact of Immigration on the Structure of Wages: Theory and Evidence from Britain," *Journal of the European Economic Association*, Vol. 10, pp. 120–151.

Marshall, Alfred (1890) *Principles of Economics*, New York: Macmillan.

Miguel, Edward and Gerard Roland (2011) "The Long-run Impact of Bombing Vietnam," *Journal of Development Economics*, Vol. 96, pp. 1–15.

Moretti, E. and P. Thulin (2013) "Local Multipliers and Human Capital in the United States and Sweden," *Industrial and Corporate Change*, Vol. 22, pp. 339–362.

Moretti, Enrico (2004) "Workers' Education, Spillovers, and Productivity: Evidence from Plant-level Production Functions," *American Economic Review*, Vol. 94, pp. 656–690.

——— (2010) "Local Multipliers," *American Economic Review, Papers & Proceedings*, Vol. 100, pp. 373–377, URL: `http://www.jstor.org/stable/27805022`.

——— (2011) "Local Labor Markets," in David Card and Orley Ashenfelter eds. *Handbook of Labor Economics*, Vol. 4, Amsterdam: Elsevier, Chap. 14, pp. 1237–1313.

MPIDR - Max Planck Institute for Demographic Research and Chair for Geodesy and Geoinformatics University of Rostock (2011) "Population History GIS Collection," Rostock.

Neumark, David and Helen Simpson (2015) "Place-Based Policies," in *Handbook of Regional and Urban Economics, vol. 5*, Vol. 5: Elsevier B.V. pp. 1197–1287.

Nordhaus, William D (1975) "The Political Business Cycle," *Review of Economic Studies*, Vol. 42, pp. 169–190.

OECD (2009) *How Regions Grow: Trends and Analysis*: OECD Publishing, Paris.

Oster, Emily (2019) "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Journal of Business & Economic Statistics*, Vol. 37, pp. 187–204.

Ottaviano, Gianmarco IP and Giovanni Peri (2012) "Rethinking the Effect of Immigration on Wages," *Journal of the European Economic Association*, Vol. 10, pp. 152–197.

Reichsamt des Innern (1914) *Zentralblatt für das Deutsche Reich*, Berlin: Carl Heymanns Verlag.

Reichsministerium des Innern (1922) *Zentralblatt für das Deutsche Reich*, Berlin: Carl Heymanns Verlag.

Reichswehrministerium (1934) *Sanitätsbericht über das Deutsche Heer im Weltkriege 1914/1918.*, Berlin: Mittler und Sohn.

Remarque, Erich Maria (1928) *Im Westen nichts Neues*, Köln: Kiepenheuer und Witsch.

Rosenthal, Stuart S. and William C. Strange (2004) "Evidence on the Nature and Sources of Agglomeration Economies.," in J. Vernon Henderson and Jacques-Francois Thisse eds. *Handbook of Regional and Urban Economics. Volume 4. Cities and Geography*, Amsterdam: Elsevier, pp. 2119 – 2171.

Schindler, Götz, Louis van Harnier, Ricarda Länge-Soppa, and Bernhard Schindler (1991) *Neue Fachhochschulstandorte in Bayern*, München: Bayerisches Staatsinstitut für Hochschuleforschung und Hochschulplanung (IHF).

Schmucker, Alexandra, Stefan Seth, Johannes Ludsteck, Johanna Eberle, Andreas Ganzer et al. (2016) "Establishment History Panel 1975-2014," FDZ-Methodenreport 03/2016, Institut für Arbeitsmarkt- und Berufsforschung (IAB), Nürnberg.

Schneider, Christina (2010) "Fighting with one Hand tied behind the Back: Political Budget Cycles in the West German States," *Public Choice*, Vol. 142, pp. 125–150.

Schubert, Uli (2020) "Gemeindeverzeichnis des Deutschen Reiches am 1.12.1910," URL: `www.gemeindeverzeichnis.de`.

Sen, Dennis (2016) "Toponym Resolution on Historical Serial Sources," *Master Thesis, University of Kiel.*

Serafinelli, Michel (2019) "Good Firms, Worker Flows and Firm Productivity," *Journal of Labor Economics*, Vol. 37, pp. 747–792.

Shi, Min and Jakob Svensson (2006) "Political Budget Cycles: Do they differ across Countries and why?," *Journal of Public Economics*, Vol. 90, pp. 1367–1389.

Stachelbeck, Christian (2010) *Militärische Effektivität im Ersten Weltkrieg: die 11. Bayerische Infanteriedivision 1915 bis 1918*, Paderborn: Schöningh.

Stratmann, Thomas and Martin Baur (2002) "Plurality Rule, proportional Representation, and the German Bundestag: How Incentives to pork-barrel differ across electoral Systems," *American Journal of Political Science*, Vol. 46, pp. 506–514.

Stüber, Heiko and Stefan Seth (2019) "The Public Release Data of the Administrative Wage and Labor Market Flow Panel," *Journal of Economics and Statistics*, Vol. 239, pp. 333–344.

Südekum, Jens (2019) “Europas Regionen besser fördern,” Friedrich-Ebert-Stiftung.

Südekum, Jens (2021) “Place-based Policies – How to do them and why,” *DISCE Discussion Paper*, Vol. 367.

Sun, Liyang and Sarah Abraham (2020) “Estimating dynamic Treatment Effects in Event Studies with heterogeneous Treatment Effects,” *Journal of Econometrics*.

Tavits, Margit (2009) “Geographically targeted Spending: Exploring the electoral Strategies of Incumbent Governments,” *European Political Science Review*, Vol. 1, pp. 103–123.

Verein für Computergenealogie (2019a) “Datenbank deutsche Verlustlisten im 1. Weltkrieg.”

——— (2019b) “Datenbank deutsche Vermisstenliste im 1. Weltkrieg.”

Vigdor, Jacob (2008) “The economic Aftermath of Hurricane Katrina,” *Journal of Economic Perspectives*, Vol. 22, pp. 135–54.

Voigtländer, Nico and Hans-Joachim Voth (2013a) “How the West ‘invented’ Fertility Restriction,” *American Economic Review*, Vol. 103, pp. 2227–64.

——— (2013b) “The Three Horsemen of Riches: Plague, War, and Urbanization in Early Modern Europe,” *Review of Economic Studies*, Vol. 80, pp. 774–811.

Voth, Hans-Joachim, Sebastian Doerr, Stefan Gissler, and José Luis Peydró (2020) “From Finance to Fascism: The real Effect of Germany’s 1931 Banking Crisis,” *CEPR Discussion Papers*.

Wissenschaftsrat (1991) *Empfehlungen zur Entwicklung der Fachhochschulen in den 90er Jahren*, Köln: Wissenschaftsrat.

——— (1995) “Stellungnahme zu den Planungen des Landes Rheinland-Pfalz zur Gründung eines neuen Standortes der Fachhochschule Rheinland-Pfalz in Birkenfeld,” Drucksache 2042/95, Berlin.

——— (1996) "Stellungnahme zu den Planungen des Landes Rheinland-Pfalz zur Gründung eines neuen Standortes der Fachhochschule Rheinland-Pfalz in Remagen," Drucksache 2356/96, Berlin.

Woodward, Douglas, Octavio Figueiredo, and Paulo Guimaraes (2006) "Beyond the Silicon Valley: University R&D and High-Technology Location," *Journal of Urban Economics*, Vol. 60, pp. 15–32.

Young, Alwyn (2005) "The Gift of the Dying: The Tragedy of AIDS and the Welfare of future African Generations," *Quarterly Journal of Economics*, Vol. 120, pp. 423–466.

Ziemann, Benjamin (2018) "The First World War and National Socialism," in *A Companion to Nazi Germany*, Hoboken: John Wiley & Sons, pp. 47–61.