The London School of Economics and Political Science

Essays in Applied Economics

Stephan Ernst Maurer

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy. London, May 2017

Meinen Eltern

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 50,000 words.

Statement of inclusion of previous work

I confirm that Chapter 3 of this thesis is a heavily revised version of the EC599 research paper I submitted for my MRes degree, awarded by the LSE in July 2013.

Statement of conjoint work

I certify that chapters 1 and 2 of this thesis were co-authored. Chapter 1 was co-authored with Jörn-Steffen Pischke and Ferdinand Rauch, while chapter 2 was co-authored with Andrei Potlogea. I contributed 33% of the work in chapter 1, and 50% of the work in chapter 2.

Statement of use of third party for editorial help

I confirm that chapter 3 of this thesis was copy edited for conventions of language, spelling and grammar by Michael Beaney, LSE Language Centre.

Acknowledgements

This thesis would not have been possible without the help and support of many people.

First of all, I would like to thank my supervisor Steve Pischke and my advisor Guy Michaels. Throughout my PhD, they have been extremely helpful and approachable, offering valuable guidance and thoughtful advice, but also encouraging me and cheering me up when needed. I have learned a lot about Economics, Econometrics, and the research process from them. Maybe even more importantly, through their own example, they have taught me to always question my own results and to be my own best sceptic.

I am also grateful to Esteban Aucejo, Jeremiah Dittmar, and Alan Manning for many helpful comments and discussions. Ferdinand Rauch helped me a lot during the job market, and as co-author of the first chapter of this thesis. I also would like to thank him for having invited me to present my work in Oxford.

As I learned over the course of the past six years, doing research can be frustrating at times. It involves a great deal of failed projects, ideas that do not materialize, and work that does not reap rewards. I was lucky to have many wonderful friends in London that shared these toils and cheered me up when needed. Andrei Potlogea and his wife Susannah have been great friends since our joint days at UPF. Andrei also is the co-author of the second chapter of this thesis, and I could not have found a better one for it. To Florian Blum, Kilian Huber, and Francesco Sannino, I am grateful for their humour and company through a lot of PhD life ups and downs. Another big thank you is due to Diego Battiston, Soledad Giardili, Nicola Limodio, Maddalena Liverani, Federico Rossi, and Francesco Sannino for countless mock presentations, reading groups, and Friday night dinners. I also would like to thank María Sánchez Vidal for her patience, humour, and support during the last year and the job market frenzy: moltes gràcies.

LSE and CEP have provided a stimulating environment for research during the last six years, but also a source of great colleagues and friends. Marta De Philippis, Georg Graetz, and Yu-Hsiang Lei have been the best officemates one could ask for, and I further thank Michael Amior, Michel Azulai, Andrés Barrios, Shiyu Bo, Rui Costa, Réka Juhász, Felix König, Monica Langella, Pily López Uribe, Maria Molina-Domene, Tom O'Keeffe, Jonathan Pinder, Frank Pisch, Roberto Sormani, Claudia Steinwender, Lisa Windsteiger, Junichi Yamasaki, and many others. Mark Wilbor and John Curtis have helped a lot with the administrative hurdles that sometimes had to be overcome. Financial support from the LSE and the Studienstiftung des Deutschen Volkes is also most gratefully acknowledge.

It has been a long educational journey that led to this thesis, and I would like to express my gratitude to those that have sparked and fostered my curiosity along its way, in particular Klaus Breinlinger, Werner Geigle, Claudia Sielaff, Christina Felfe, Alex Lefter, Davide Cantoni, and Kurt Schmidheiny. During all those years, my friends at home, notably Pascal Albrecht, have also been a great source of support. The same goes for my family, especially Edith, Andreas, Petra, Jojo, and Tim.

Last but not least, I would like to thank my parents Karin and Wilfried. When they drove me to St.Gallen nearly 10 years ago, they probably would have never thought that I would embark on a journey that would last so long and bring me to four different countries and as many universities. Without their unfaltering love and support, I could never have done it. It is to them that I dedicate this thesis, as well as these words by Hebel:

"Es ist für mich wahr und bleibt für mich wahr, der Himmel ist nirgends so blau, und die Luft nirgends so rein, und alles so lieblich und so heimlich als zwischen den Bergen von Hausen."

Abstract

This thesis consists of three papers that belong to the broad realm of Applied Economics.

The first chapter studies the causal connection between trade and development, using one of the earliest massive trade expansions in prehistory: the first systematic crossing of open seas in the Mediterranean during the time of the Phoenicians. For each point on the coast, we construct the ease with which other points can be reached by crossing open water. We show that an association between better connected locations and archaeological sites emerges during the Iron Age when sailors routinely crossed open water. We corroborate these findings at the world scale.

In the second chapter, we use oil discoveries in the US South between 1900 and 1940 to analyse whether male-biased demand shocks reduce women's labour force participation. We find that oil wealth has a zero net effect on female labour force participation due to two opposing channels. Oil discoveries raise male wages, which leads to an increased marriage rate of young women and thus could have depressed female labour supply. But oil wealth also increases demand for women in services, which counterbalances the marriage effect. Our findings demonstrate that when the nontradable sector is open to women, male-biased demand shocks in the tradable sector need not reduce female labour force participation.

The third chapter analyses whether the German National Socialists used economic policies to reward their voters after coming to power in 1933. Using newly-collected data on public employment from the German censuses in 1925, 1933, and 1939 and addressing the potential endogeneity of the NSDAP vote share in 1933 by way of an instrumental variables strategy based on a similar party in Imperial Germany, I find that cities with higher NSDAP vote shares experienced a relative increase in public employment: for every additional percentage point in the vote share, the number of public employment jobs increased by around 2.5%.

Contents

List of Figures

List of Tables

Introduction

This thesis consists of three papers that belong to the broad realm of Applied Economics, with a particular focus on how natural advantages and politics can shape the geographic distribution of economic activities. From a methodological point of view, all three chapters also share the use of variation provided by history in order to shed lights on economic questions.

In the first chapter, together with Jörn-Steffen Pischke and Ferdinand Rauch, we study the causal effect of trade opportunities on economic growth. Providing clean evidence for this effect is very difficult due to reverse causality, path dependence, the endogenous location of economic activities and omitted variables such as institutions. Our solution is to combine fine geographic variation with one of the earliest major trade expansions in human history: the first systematic crossing of the Mediterranean during the Iron Age. In particular, for each point on a 10km $\times$10km grid of the Mediterranean coast, we calculate the number of other coastal points that can be reached over sea within given distances (e.g. 500km). Using data on the locations of archaeological sites from various sources, we then show that coastal points from which more other coastal points can be reached over sea have a greater local density of human activity. However, this relationship only emerges during the Iron Age, when open sea voyages were common. Before, when sailors were mostly following the coast, our results are weaker and less consistent. In essence, we use geographical connectedness as an instrumental variable for trade, but due to the absence of trade data for the Iron Age, we estimate a reduced form effect of connectedness on economic growth. Our results indicate a positive effect of trade on growth, where trade should be understood in a broad sense that also includes the spread of ideas or people. In addition, our study also provides further evidence for the persistent role of locational fundamentals in the spatial distribution of cities.

Thus, as chapter 1 shows, natural advantages can have important broad implications for growth. In addition, they can also have substantial distributional effects, for example by favouring some workers more than others. In this vein, the second chapter (joint with Andrei Potlogea) uses another geographic advantage - natural resources - to study the effect of male-biased labour demand shocks on gender labour market differences. Between 1900 and 1940, many large oil fields were discovered in the US South. We provide a theoretical framework that models oil discoveries as male-biased demand shocks in a small open economy. If women are allowed to work in services, they are subject to two opposing labour demand effects: A demand reduction in the tradable sector, and a demand increase in the nontradable sector. The more productive the oil industry, the more likely it is that the latter effect will dominate. We then use data on the location and discovery years of oil fields in the US South, coupled with individual-level census data, to compare the evolution of female labour market outcomes in oil-rich counties relative to control counties. We find that oil wealth has a zero net effect on female labour force participation: Oil discoveries increase demand for male labour in oil mining and manufacturing and consequentially raise male wages. This leads to an increased marriage rate of young women, which could have depressed female labour force participation. But consistent with our model, oil wealth also increases demand for women in services, which counterbalances the marriage effect and leaves women's overall labour force participation rate unchanged. Our findings demonstrate that when the nontradable sector is open to women, male-biased demand shocks in the tradable sector need not reduce female labour force participation.

In the third chapter, finally, I look at how politics can induce changes in the spatial distribution of economic activity, with a particular focus on public employment. A sizeable literature has shown how firms (e.g. Fisman 2001) or regions (e.g. Anderson and Tollison 1991) can benefit from being politically connected. I add to this literature by analysing whether cities in Nazi Germany benefited from having voted for the National Socialists before they came to power, using newly-connected data on public employment from the German censuses in 1925, 1933, and 1939 as outcome measure. The vote share for the National Socialists is likely to be endogenous for a variety of reasons such as reverse causality (Voigtländer and Voth 2016) and differential impacts of the post-1929 economic crisis (King et al 2008). To address this, I employ an instrumental variables strategy based on the persistence of local political attitudes: I instrument the 1933 vote share of the National Socialist party (NSDAP) by way of the 1912 vote share of a smaller party in Imperial Germany that catered to similar voters as the NSDAP later. Doing so, I find that cities with higher NSDAP vote shares experienced a relative increase in public employment: for every additional percentage point in the vote share, the number of public employment jobs increased by around 2.5 percent. The findings are robust to including or excluding cities that underwent substantial changes in their population and territory during the period of observation, to using the 1930 or 1932 elections instead of the 1933 one as explanatory variable, and to employing different definitions of "public employment" as outcome variables.

Chapter 1

Of Mice and Merchants: Trade and Growth in the Iron Age[1]

1.1 Introduction

We investigate to what degree trading opportunities affected economic development at an early juncture of human history. In addition to factor accumulation and technical change, Smithian growth due to exchange and specialization is one of the fundamental sources of growth. An emerging literature on the topic is beginning to provide compelling empirical evidence for a causal link from trade to growth. We contribute to this literature and focus on one of the earliest trade expansions in pre-history: the systematic crossing of open seas in the Mediterranean at the time of the Phoenicians from about 900 BC. We relate trading opportunities, which we capture through the connectedness of points along the coast, to early development as measured by the presence of archaeological sites. We find that locational advantages for sea trade matter for the foundation of Iron Age cities and settlements, and thus helped shape the development of the Mediterranean region, and the world.

A location with more potential trading partners should have an advantage if trade is important for development. The particular shape of a coast has little influence over how many neighbouring points can be reached from a starting location within a certain distance as long as ships sail mainly close to the coast. However, once sailors begin to cross open seas, coastal geography becomes more important: Some coastal points are in the reach of many neighbours while other can reach only few. The general shape of the coast and the location of is-

[1]We thank Jan Bakker and Juan Pradera for excellent research assistance, and David Abulafia, Tim Besley, Andrew Bevan, Francesco Caselli, Jeremiah Dittmar, Avner Greif, Carl Knappett, Andrea Matranga, Guy Michaels, Dennis Novy, Luigi Pascali, Dominic Rathbone, Tanner Regan, Corinna Riva, Susan Sherratt, Pedro CL Souza, Peter Temin, John van Reenen, Ruth Whitehouse, and participants at various seminars and conferences for their helpful comments and suggestions.

lands matters for this. We capture these geographic differences by dividing the Mediterranean coast into grid cells, and calculating how many other cells can be reached within a certain distance. Parts of the Mediterranean are highly advantaged by their geography, e.g. the island-dotted Aegean and the "waist of the Mediterranean" at southern Italy, Sicily, and modern Tunisia. Other areas are less well connected, like most of the North African coast, parts of Iberia and southern France, and the Levantine coast.

We relate our measure of connectivity to the number of archaeological sites found near any particular coastal grid point. This is our proxy for economic development. It is based on the assumption that more human economic activity leads to more settlements and particularly towns and cities. While these expand and multiply, there are more traces in the archaeological record. We find a pronounced relationship between connectivity and development in our data set for the Iron Age around 750 BC, when the Phoenicians had begun to systematically traverse the open sea, using various different data sources for sites. We find a weaker and less consistent relationship between connectivity and sites for earlier periods. This is consistent with the idea that earlier voyages occurred, maybe at intermediate distances, at some frequency already during the Bronze Age. Our interpretation of the results suggests that the relationship between coastal geography and settlement density, once established in the Iron Age, persists through the classical period. This is consistent with a large literature in economic geography on the persistence of city locations. While our main results pertain to the Mediterranean, where we have good information on archaeological sites, we also corroborate our findings at a world scale using population data for 1 AD from McEvedy and Jones (1978) as outcome.

Humans have obtained goods from far away locations for many millennia. While some of the early trade involved materials useful for tools (like the obsidian trade studied by Dixon, Cann, and Renfrew 1968), as soon as societies became more differentiated a large part of this early trade involved luxury goods doubtlessly consumed by the elites. Such trade might have raised the utility of the beneficiaries but it is much less clear whether it affected productivity as well. Although we are unable to measure trade directly, our work sheds some light on this question. Since trade seems to have affected the growth of settlements even at an early juncture this suggests that it was productivity enhancing. The view that trade played an important role in early development has recently been gaining ground among economic historians; see e.g. Temin (2006) for the Iron Age Mediterranean, Algaze (2008) for Mesopotamia, and Temin (2013) for Ancient Rome.

Our approach avoids issues of reverse causality and many confounders by using a geography based instrument for trade. In fact, we do not observe trade itself but effectively estimate a reduced form relationship, relating opportunities for trade directly to economic development. This means that we do not necessarily isolate the effect of the exchange of goods per se. Our results could be driven by migration or the spread of ideas as well, and when we talk about "trade" we interpret it in this broad sense. We do believe that coastal connectivity captures effects due to maritime connections. It is difficult to imagine any other channel why geography would matter in this particular manner, and we show that our results are not driven by a variety of other geographic conditions.

Since we do not use any trade data we avoid many of the measurement issues related to trade. We measure trading opportunities and development at a fine geographic scale, hence avoiding issues of aggregation to a coarse country level. Both our measure of connectedness and our outcome variable are doubtlessly extremely crude proxies of both trading opportunities and of economic development. This will likely bias us against finding any relationship and hence makes our results only more remarkable.

The periods we study, the Bronze and Iron Ages, were characterized by the rise and decline of many cultures and local concentrations of economic activity. Many settlements and cities rose during this period, only to often disappear again. This means that there were ample opportunities for new locations to rise to prominence while path dependence and hysteresis may have played less role compared to later ages. The political organization of the Mediterranean world prior to the Romans was mostly local. The Egyptian Kingdoms are the main exception to this rule but Egypt was mostly focused on the Nile and less engaged in the Mediterranean. As a result, institutional factors were less important during the period we study.

There is a large literature on trade and growth. Canonical studies are the investigations by Frankel and Romer (1999) and Redding and Venables (2004). These papers use distance from markets and connectivity as measured by gravity relationships to capture the ease with which potential trading partners can be reached. However, these measures do not rely purely on geography but conflate economic outcomes like population and output, which are themselves affected by the development process.

The more recent literature has circumvented this by analysing exogenous events related to changes in trade. Most similar to our study are a series of papers which also exploit new trade relationships arising from discoveries, the opening of new trade routes, and technological change. Acemoglu, Johnson, and Robinson (2005) link Atlantic trade starting around 1,500 AD to the ensuing shift in the focus of economic activity in Europe from the south and centre of the continent to the Atlantic periphery. Redding and Sturm (2008) focus on the natural experiment created by the division and reunification in Germany, which changed the access to other markets sharply for some locations but not others. Various papers exploit the availability of new transport technologies; Feyrer (2009) uses air transport, Donaldson (forthcoming) and Donaldson and Hornbeck (2016) use railroads, and Pascali (forthcoming) steam ships. These papers generally find that regions whose trading opportunities improved disproportionately saw larger income growth. That we find similar results for a much earlier trade expansion suggests that the productivity benefits of trade have been pervasive throughout history.

Our paper also relates to a literature on how changes in locational fundamentals shape the location of cities (Davis and Weinstein 2002, Bleakley and Lin 2012, Bosker and Buringh 2017, Michaels and Rauch forthcoming). Our contribution to this literature is to give evidence on one of the most important locational fundamentals, market access. In a world with multiple modes of transport for the transportation of different goods, it is typically hard to measure market access and changes of market access of a city. Our measure relates to a world where much long distance trade took place on boats, which makes it easier to isolate a measure of market access.

Also closely related is the paper by Ashraf and Galor (2011a). They relate population density in various periods to the relative geographic isolation of a particular area. Their interest is in the impact of cultural diversity on the development process, and they view geographic isolation effectively as an instrument for cultural homogeneity. Similar to our measure, their geographic isolation measure is a measure of connectivity of various points around the world. They find that better connected (i.e. less isolated) countries have lower population densities for every period from 1 to 1,500 AD, which is the opposite of our result. Our approach differs from Ashraf and Galor (2011a) in that we only look at coasts and not inland locations. They control for distance to waterways in their regressions, a variable that is strongly positively correlated with population density. Hence, our results are not in conflict with theirs.

Our paper is also related to a number of studies on pre-historic Mediterranean connectivity and seafaring. McEvedy (1967) creates a measure of "littoral zones" using coastal shapes. He produces a map which closely resembles the one we obtain from our connectivity measure but does not relate geography directly to seafaring. This is done by Broodbank (2006), who overlays the connectivity map with archaeological evidence of the earliest sea-crossings up to the end of the last Ice Age. He interprets the connections as nursery conditions for the early development of nautical skills, rather than as market access, as we do for the later Bronze and Iron Ages. Also related is a literature in archaeology using network models connecting archaeological sites; Knappett, Evans, and Rivers (2008) is an excellent example for the Bronze Age Aegean. None of these papers relate to the changes arising from open sea-crossings, which is the focus of our analysis. Temin (2006) discusses the Iron Age Mediterranean through the lens of comparative advantage trade but offers no quantitative evidence as we do.

1.2 Brief history of ancient seafaring in the Mediterranean

The Mediterranean is a unique geographic space. The large inland sea is protected from the open oceans by the Strait of Gibraltar. The tectonics of the area, the African plate descending under the Eurasian one, have created a rugged northern coast in Europe and a much straighter one in North Africa. Volcanic activity and the more than 3,000 islands also tend to be concentrated towards the north. The climatic conditions in the Mediterranean are generally relatively favourable to agriculture, particularly in the north. The Mediterranean is the only large inland sea with such a climate (Broodbank 2013). Its east-west orientation facilitated the spread of agriculture from the Levant (Diamond 1997). Despite these common features, the size of the Mediterranean and an uneven distribution of natural resources also implies great diversity. Modern writers on the Mediterranean, most notably Horden and Purcell (2000) have stressed that the area consists of many micro-regions. Geography and climate make the Mediterranean prone to many risks, such as forest fires, earthquakes, plagues of locusts, droughts, floods, and landslides. As a consequence, trade networks that allow to moderate shocks are of great mutual interest in the region, and trade has played a central role since its early history.[2]

Clear evidence of the first maritime activity of humans in the Mediterranean is elusive. Crossings to islands close to the mainland were apparently undertaken as far back as 30,000 BC (Fontana Nuova in Sicily). In a careful review of the evidence, Broodbank (2006) dates more active seafaring to around 10,000 BC based on the distribution of obsidian (a volcanic rock) at sites separated by water (see Dixon, Cann, and Renfrew 1965, 1968). This points to the existence of active sea-faring of hunter-gatherer societies, and suggests that boats must have travelled distances of 20-35 kilometres around that time. We have no evidence on the first boats but they were likely made from skin and frame or dugout canoes.

The beginning of agriculture around the Mediterranean happened in the Levant between 9,500 BC and 8,000 BC. From there it spread initially to Anatolia and the Aegean. Signs of a fairly uniform Neolithic package of crops and domesticated animals can be found throughout the Mediterranean. The distribution of the earliest evidence of agriculture, which includes islands before reaching more

[2]The following discussion mainly draws on Abulafia (2011) and Broodbank (2013).

peripheral parts of the mainland, suggests a maritime transmission channel.
The Neolithic revolution did not reach Iberia until around 5,500 BC. By that time, many islands in the Aegean had been settled, there is evidence for grain storage, and metal working began in the Balkans. Because of the uneven distribution of ores, metals soon became part of long range transport. Uncertainty must have been a reason for the formation of networks both for insurance and exchange. The first archaeological evidence of a boat also stems from this period: a dugout canoe, about 10m long, at La Marmotta north of Rome. A replica proved seaworthy and allowed travel of 20 - 25km per day in a laden boat.
The Levant, which was home to the first cities, remained a technological leader in the region, yet there is little evidence of sea-faring even during the Copper Age. This changed with the rise of large scale societies in Mesopotamia and Egypt. Inequality in these first states led to rich elites, who soon wished to trade with each other. Being at the cross-roads between these two societies, the Levant quickly became a key intermediary.
Two important new transport technologies arrived in the Mediterranean around 3,000 BC: the donkey and the sail. The donkey was uniquely suited to the climatic conditions and rugged terrain around the Mediterranean (better than camels or horses). Donkeys are comparable in speed to canoes. Sailboats of that period could be around 5-10 times faster in favourable conditions, ushering in a cost advantage of water transport that would remain intact for many millennia to come. The land route out of Egypt to the Levant ("The Way of Horus") was soon superseded by sea routes leading up the Levantine coast to new settlements like Byblos, with Levantine traders facilitating much of Egypt's Mediterranean trade. Coastal communities began to emerge all the way from the Levant via Anatolia to the Aegean and Greece.
There is no evidence of the sail spreading west of Greece at this time. Canoes, though likely improved into high performance water craft, remained inferior to sail boats but kept facilitating maritime transport in the central and western Mediterranean. The major islands there were all settled by the early Bronze Age. While not rivalling the maritime activity in the eastern Mediterranean, regional trade networks arose also in the west. One example is the Beaker network of the 3rd Millennium BC; most intense from southern France to Iberia, with fewer beakers found in the western Maghreb, northern Italy, and Sardinia but also stretching all the way into central Europe, the Baltic, and Britain. Land routes probably dominated but sea trade must have played a role. The Cetina culture of the late 3rd Millennium BC in the Adriatic is another example. Occasional sea-crossings up to 250km were undertaken.
A drying spell around 2,200 BC and decline in Egypt disrupted the active mar-

itime network in the eastern Mediterranean and the population it supported. The oldest known shipwreck in the Mediterranean at the island of Dokos in southern Greece dates from this period. The 15 metres long boat could carry a maximum weight of 20 tons. The wreck contained largely pottery, which was likely the cargo rather than carrying liquids, and also carried lead ingots. The ship probably was engaged in local trade.

Decline in the eastern Mediterranean soon gave rise to new societies during the 2nd millennium BC: palace cultures sprang up all over the eastern Mediterranean. Minoan Crete and Mycenae in Greece were notable examples but similar cities existed along the Anatolian coast and in the Levant. The palaces did not simply hold political power, but were centres of religious, ceremonial, and economic activity. At least initially, craftsmen and traders most likely worked for the palace rather than as independent agents. Sail boats still constituted an advanced technology, and only the concentration of resources in the hands of a rich elite made their construction and operation possible. The political reach of the palaces at coastal sites was local; larger polities remained confined to inland areas as in the case of Egypt, Babylon, or the Hittite Empire.

An active trade network arose again in the eastern Mediterranean stretching from Egypt to Greece during the Palace period. The Anatolian land route was replaced by sea trade. Some areas began to specialize in cash crops like olives and wine. A typical ship was still the 15 metre, 20 ton, one masted vessel as evidenced by the Uluburn wreck found at Kas in Turkey, dating from 1,450 BC. Such vessels carried diverse cargoes including people (migrants, messengers, and slaves), though the main goods were likely metals, textiles, wine, and olive oil. Evidence for some of these was found on the Uluburun wreck; other evidence comes from archives and inscriptions akin to bills of lading. Broodbank (2013) suggests that the value of cargo of the Uluburun ship was such that it was sufficient to feed a city the size of Ugarit for a year. Ugarit was the largest trading city in the Levant at the time with a population of about 6,000 - 8,000. This highlights that sea trade still largely consisted of high value luxury goods. The Ugarit archives also reveal that merchants operating on their own account had become commonplace by the mid 2nd millennium. Levantine rulers relied more on taxation than central planning of economic activities. Trade was both risky and profitable; the most successful traders became among the richest members of their societies.

Around the same time, the Mycenaeans traded as far as Italy. Sicily and the Tyrrhenian got drawn into the network. While 60 - 70km crossings to Cyprus or Crete and across the Otrano Strait (from Greece to the heel of Italy) were commonplace, coast hugging still prevailed among sailors during the 2nd millennium BC. After crossing the Otrano Strait, Greek sailors would continue along

the coast of the Bay of Taranto, the instep of Italy's boot, as is suggested by the distribution of Greek pottery at coastal sites. Indigenous sea-farers from the central Mediterranean now joined these routes, and the sail finally entered the central Mediterranean around 1,200 BC. While there were no big breakthroughs, naval technology also improved in the late 2nd millennium. Better caulking and keels added to sea-worthiness (Abulafia 2011), while brail rigging and double prows improved manoeuvrability. Most notably, latitude sailing was developed and allowed sailors to steer a straight east-westerly course. "This was a leap in the scope of connections, a permanent shift in Mediterranean history and a crucial stage in tying together the basin's inhabitants across the soon-to-be shrinking sea," observes Broodbank (2013, p. 431) before warning that "we should not exaggerate, nor anticipate, the importance of such connections at this early juncture. Not until the Iron Age did relations become close enough to fundamentally reshape the culture and economies of outlying regions." (p. 441)
A new period of decline around 1,200 BC reduced the power of Egypt, wiped out cities like Ugarit, and ended the reign of the last palace societies in the eastern Mediterranean. In the more integrated world that the eastern Mediterranean had become, troubles spread quickly from one site to others. The Bronze Age came to an end with iron coming on the scene. Rather than being technologically all that much superior to bronze, iron ore was far more abundant and widespread than copper and hence much more difficult to monopolize. As was the case many times before, decline and change opened up spaces for smaller players and more peripheral regions. Cyprus flourished. Many Levantine cities recovered quickly. Traders from the central Mediterranean also expanded. Traditionally, decline during the Bronze Age collapse was often blamed on the anonymous "Sea Peoples." Modern scholarship seems to challenge whether these foreigners were simply just raiders and pirates, as the Egyptians surely saw them, rather than also entrepreneurial traders who saw opportunities for themselves to fill the void left by the disappearance of imperial connections and networks. Some of these new interlopers settled in the Levant (Broodbank 2013).

While there is much academic debate about the origin of the Phoenicians, there is little doubt that the Levantine city states which had taken in these migrants were the origin of a newly emerging trade network. Starting to connect the old Bronze Age triangle formed by the Levantine coast and Cyprus, they began to expand throughout the entire Mediterranean after 900 BC. The Phoenician city states were much more governed by economic logic than was the case for royal Egypt. One aspect of their expansion was the formation of enclaves, often at nodes of the network. Carthage and Gadir (Cadiz) are prime examples but many others existed. At least initially these were not colonies; the Phoenicians did not try to dominate local populations. Instead, locals and other settlers were invited to pursue their own enterprise and contribute to the trading network. The core of the network consisted of the traditional sea-faring regions, the Aegean and the Tyrrhenian. The expanding trade network of the early 1st millennium BC did not start from scratch but encompassed various regional populations. Tyrrhenian metal workers and Sardinian sailors had opened up connections with Iberia at the close of the 2nd millennium. But the newly expanding network not only stitched these routes together, it also created its own, new, long-haul routes.

These new routes began to take Phoenician and other sailors over long stretches of open sea. While this had long been conjectured by earlier writers like Braudel (2001, writing in the late 1960s) and Sherratt and Sherrat (1993), contemporary scholars are more confident. Cunliffe (2008) writes about the course of a Phoenician sailor: "Beyond Cyprus, for a ship's master to make rapid headway west there was much to be said for open-sea sailing. From ... the western end of Cyprus he could have sailed along the latitude to the south coast of Crete ... where excavation has exposed a shrine built in Phoenician fashion. Travelling the same distance again ..., once more following the latitude, would have brought him to Malta" (p. 275-276), a route which became known as the "Route of the Isles." Abulafia (2011) describes their seafaring similarly: "The best way to trace the trading empire of the early Phoenicians is to take a tour of the Mediterranean sometime around 800 BC. ... Their jump across the Ionian Sea took them out of the sight of land, as did their trajectory from Sardinia to the Balearics; the Mycenaeans had tended to crawl round the edges of the Ionian Sea past Ithaka to the heel of Italy, leaving pottery behind as clues, but the lack of Levantine pottery in southern Italy provides silent evidence of the confidence of Phoenician navigators." (p. 71).

This involved crossing 300 - 500km of open sea. One piece of evidence for sailing away from the coast are two deep sea wrecks found 65km off the coast of Ashkelon (Ballard et al. 2002). Of Phoenician origin and dating from about 750 BC, the ships were 14 metres long, and each carried about 400 amphorae filled with fine wine. These amphorae were highly standardized in size and shape. This highlights the change in the scale and organization of trade compared to the Uluburun wreck with its diverse cargo. It also suggests an early form of industrial production supporting this trade.

An unlikely traveller offers a unique lens on the expansion of trade and the density of connections which were forged during this period. The house mouse populated a small area in the Levant until the Neolithic revolution. By 6,000 BC, it had spread into southern Anatolia before populating parts of north eastern Africa and the Aegean in the ensuing millennia (there were some travellers on the Uluburun ship). There were no house mice west of Greece by 1,000 BC.

Then, within a few centuries, the little creature turned up on islands and on the mainland throughout the central and western Mediterranean (Cucchi, Vigne, and Auffray 2005).

The Phoenicians might have been at the forefront of spreading mice, ideas, technology, and goods all over the Mediterranean but others were part of these activities. At the eve of Classical Antiquity, the Mediterranean was constantly criss-crossed by Greek, Etruscan, and Phoenician vessels as well as smaller ethnic groups. Our question here is whether this massive expansion in scale led to locational advantages for certain points along the coast compared to others, and whether these advantages translated into the human activity which is preserved in the archaeological record. A brief, rough time line for the period we investigate is given in figure 1.1.

1.3 Data and key variables

For our Mediterranean dataset we compute a regular grid of 10×10 kilometres that spans the area of the Mediterranean and the Black Sea using a cylindrical equal area projection. This projection ensures that horizontal neighbours of grid points are on the same latitude at 10km distance from each other, and that each cell has an equal area over the surface of the earth.[3] We define a grid-cell as water if falls completely into water, using a coastline map of the earth from Bjorn Sandvik's public domain map on world borders[4]. We define it as coastal if it is intersected by a coastline. We classify grid cells that are neither water nor coastal as land. Our estimation dataset consists of coastal cells only, and each cell is an observation. There are 3,646 cells in the data set.

We compute the distance between coastal point i and coastal point j moving only over water d_{ij} using the cost distance command in ArcGIS. Our key variable in this study, called c_{di}, measures the number of other coastal cells which can be reached within distance d from cell i. Destinations may include islands but we exclude islands which are smaller than $20km^2$. We also create separate measures, one capturing only connectedness to islands, and a second measuring connectedness to other points on the mainland coast. While we use straight line or shortest distances, we realize that these would have rarely corresponded to actual shipping routes. Sailors exploited wind patterns and currents, and often used circular routes on their travels (Arnaud 2007). Our measure is not supposed to mimic sailing routes directly but simply capture opportunities.

Figure 1.2 displays the measure c_{500} for a distance of 500km; darker points in-

[3]As the Mediterranean is close enough to the equator distortions from using another projection are small in this area of interest.

[4]We use version 3, available from `http://thematicmapping.org/downloads/world_borders.php`.

dicate better connected locations. Measures for other distances are strongly positively correlated and maps look roughly similar. The highest connectedness appears around Greece and Turkey partly due to the islands, but also western Sicily and the area around Tunis. The figure also highlights substantial variation of the connectedness measure within countries. The grid of our analysis allows for spatial variation at a fine scale. Since our measure of connectedness has no natural scale that is easy to interpret, we normalize each c_d to have mean 0 and standard deviation 1. Figure 1.3 shows a histogram of our normalized connectedness measure for a distance of 500km. Its distribution is somewhat bimodal, with a large spike at around -0.5, and a second, smaller one around 2. Basically all values above 1 are associated with locations in and around the Aegean, by far the best connected area according to our measure.

We interpret the measure c_d as capturing connectivity. Of course, coastal shape could proxy for other amenities. For example, a convex coastal shape forms a bay, which may serve as a natural harbour. Notice that our or 10×10 kilometre grid is coarse enough to smooth out many local geographic details. We will capture bays 50 kilometres across but not those 5 kilometres across. It is these more local features which are likely more relevant for locational advantages like natural harbours. Our grid size also smooths out other local geographic features, like changes in the coastline which have taken place over the past millennia, due, for example, to sedimentation. The broader coastal shapes we capture have been roughly constant for the period since 3,000 BC, which we study (Agouridis, 1997).

Another issue with our measure of connectivity is whether it only captures better potential for trade or also more exposure to external threats like military raids. Overall, it was probably easier to defend against coastal attacks than land-based ones (e.g. Cunliffe, 2008, p. 447) so this may not be a huge concern. But at some level it is obvious that openness involves opportunities as well as risks. In this respect we measure the net effect of better connectivity.

We also compute a global dataset based on a global grid. We increase the cell size to 50×50 kilometres. This is for computational convenience, but also our outcome variable at the global level varies only at the country level and thus spatial precision is less relevant than in the Mediterranean data set. We focus on the part of the world between -60 degrees and 60 degrees latitude, as units outside that range are unlikely candidates for early urbanization for climatic reasons. In the Southern Hemisphere there is hardly any landmass apart from the Antarctic below 60 degrees, while in the Northern Hemisphere 60 degrees is close to Helsinki, Aberdeen, and Anchorage, well north of climatic conditions particularly favourable to early settlement.[5] We again compute the distance

[5]In defining our connectedess measure, we restrict attention to the area between 78 and -54 degrees latitude, which roughly mark the northernmost and southernmost cities of the world.

from each coastal grid point to each other coastal grid point by moving only over water. Figure 1.4 shows the global connectedness measure c_{500}. The most connected coastal points are located again near Greece, but also in Southeast Asia, Chile, Britain, and Northern Canada, while Western Africa and Eastern South America have few well connected coastal points.

One limitation of the proposed connectivity measure c_d is that it gives all coastal points that can be reached within distance d equal weight. We would however expect a connection with a well connected coastal point to be more beneficial than a connection with a remote coastal cell. To address this limitation we could weight destination cells by their own c_d, and recompute a weighted version called $c2_d$. After normalization we could compute an additional measure $c3_d$, where we use $c2_d$ as the weight. Repeating this infinitely often, the measure converges to a variable called "network centrality." This is a standard measure in various disciplines to capture the importance of nodes in a network. To compute the centrality measure, we create a symmetric matrix A for all binary connections, with entries that consist of binary variables indicating distances smaller than d. We set the diagonal of the matrix to zero. We solve equation $Ax = \lambda x$ for the largest possible eigenvalue λ of matrix A. The corresponding eigenvector x gives the centrality measure.

Our main source of data on settlements in pre-history is the Pleiades dataset (Bagnall et al. 2014) at the University of North Carolina, the *Stoa Consortium*, and the *Institute for the Study of the Ancient World* at New York University maintained jointly by the *Ancient World Mapping Center*.[6] The Pleiades dataset is a gazetteer for ancient history. It draws on multiple sources to provide a comprehensive summary of the current knowledge on geography in the ancient world. The starting point for the database is the *Barrington Atlas of the Greek and Roman World* (Talbert 2000); but it is an open source project and material from multiple other scholarly sources has been added.

The Pleiades data are available in three different formats of which we use the "pleiades-places" dataset. It offers a categorization as well as an estimate of the start and end date for each site. We only keep units that have a defined start and end date, and limit the data set to units that have a start date before 500 AD. We use two versions of these data, one more restricted (which we refer to as "narrow") and the other more inclusive ("wide"). In the narrow one we only keep units that contain the word "urban" or "settlement" in the categorization. These words can appear alongside other categorizations of minor constructions, such as bridge, cemetery, lighthouse, temple, villa, and many others. One prob-

This excludes the Southern tip of Tierra del Fuego in South America as coastal destination, but should be fairly inconsequential.

[6]Available at `pleiades.stoa.org`. We use a version of the dataset downloaded in June 2014.

lem with the narrow version is that the majority of Pleiades sites do not have a known category. So that we do not lose these sites we include all sites irrespective of their category; including both those classified as “unknown” or with any other known classification in the wide version of the data.

Some of the entries in the Pleiades dataset are located more precisely than others. The dataset offers a confidence assessment consisting of the classifications precise, rough, and unlocated. We only keep units with a precisely measured location.[7] For both datasets, as we merge the Pleiades data onto our grid we round locations to the nearest 10×10 kilometres and are thus robust to some minor noise.

Since the Pleiades data is originally based on the *Barrington Atlas* it covers sites from the classical Greek and Roman period well and adequate coverage seems to extend back to about 750 BC. Coverage of older sites seems much more limited as the number of sites with earlier start dates drops precipitously. For example, our wide data set has 1,491 sites in 750 BC and 5,649 in 1 AD but only 63 in 1,500 BC. While economic activity and populations were surely lower in the Bronze Age, there are likely many earlier sites missing in the data. As a consequence, our estimation results with the Pleiades data for earlier periods may be rather unreliable.

We therefore created an additional data set of sites from the *Archaeological Atlas of the World* (Whitehouse and Whitehouse 1975). The advantage of the *Whitehouse Atlas* is that it focuses heavily on the pre-historic period, and therefore complements the Pleiades data well. A disadvantage is that it is 40 years old. Although there has been much additional excavation in the intervening period, there is little reason to believe that it is unrepresentative for the broad coverage of sites and locations. The interpretation of the archaeological evidence may well have changed but this is of little consequence for our exercise. Another drawback of the *Whitehouse Atlas* is that the maps are much smaller than in the *Barrington Atlas*. As a result, there may have been a tendency by the authors to choose the number of sites so as to fill each map without overcrowding it, leading to a distribution of sites which is too uniform (something that would bias our results against finding any relationship with our connectivity measure). This, however, is offset by the tendency to include maps for smaller areas in locations with many sites. For example, there are separate maps for each of Malta, Crete, and Cyprus but only three maps for all of Iberia.

[7] Pleiades contains some sites that have the same identifier, but different locations. This could reflect, among others, sites from different eras in the same location or different potential locations for the same site. We deal with this by dropping all sites that have the same Pleiades identifier and whose coordinates differ by more than 0.1 degree latitude or longitude in the maximum. The latter restrictions affects around one percent of the Pleiades data. The remaining identifiers with several sites are dealt with by counting them as one unit, averaging their coordinates. For overlapping time spans, we use the minimum of such spans as the start date, the respective maximum as end date.

We geo-referenced all entries near the coasts on 28 maps covering the Mediterranean in the *Whitehouse Atlas* ourselves. Using the information in the map titles and accompanying text, we classified each map as belonging to one of three periods: the Neolithic, the Bronze Age, or the Iron Age and later. Some maps contain sites from multiple periods but give a classification of sites, which we use. Other maps straddle periods without more detailed timing information. In this case, we classified sites into the three broad periods ourselves using resources on the internet. In a few cases, it is not possible to classify sites clearly as either Neolithic or Bronze Age in which case we classified them as both (see the appendix for details).

To measure the urbanization rate near each coastal grid point for time t we count the number of sites from either Pleiades or Whitehouse that exist at time t within 50 kilometres of that coastal point on the same landmass. We also count the number of land cells that are within 50 kilometres of that coastal grid point on the same landmass. We normalize the number of sites by the number of land cells within this radius. We prefer this density measure to the non-normalized count of sites in order to avoid that coastal shape (which enters our connectivity measure) mechanically influences the chance of having an archaeological site nearby. On the other hand, we want to classify a small trading islands as highly urbanized. We normalize our measure of urban density to have mean 0 and standard deviation 1 for each period to facilitate comparison over time when the number of settlements changes.

1.4 Specification and results

We run regressions of the following type:

$$u_{it} = X_i \gamma_t + c_{di} \beta_{dt} + \epsilon_{it}, \tag{1.1}$$

where u_{it} is the urbanization measure for grid point i, X_i are grid point control variables and c_{di} is a connectivity measure for distance d. We only measure connectivity of a location, not actual trade. Hence, when we refer to trade this may refer to the exchange of goods but could also encompass migration and the spread of ideas. u_{it} measures the density of settlements, which we view as proxy for the GDP of an area. Growth manifests itself both in terms of larger populations as well as richer elites in a Malthusian world. We would expect that the archaeological record captures exactly these two dimensions.

We use latitude, longitude, and distance to the Fertile Crescent, which all do not vary over time, as control variables. We explore dropping the Aegean, to address concerns that our results may be driven exclusively by developments around the Greek islands, by far the best connected area in the Mediterranean. We also show results dropping North Africa to address concerns that there may be fewer archaeological sites in North Africa due to a relative lack of exploration. This may spuriously correlate with the fact that the coast is comparatively straight. We cluster standard errors at the level of a grid of 2×2 degree following Bester, Conley and Hanson (2011). We normalize u_{it} and c_{di} to mean 0 and standard deviation 1 to make our estimates comparable across years with different numbers of cities, and different magnitudes of connectedness measures.

Our measure of connectedness depends only on coastal and maritime geography and therefore is plausibly exogenous. However, it might be spuriously correlated with other factors that affect early growth, such as agricultural productivity, topographic conditions, or rivers, which provide inland connections. Those factors are hard to measure precisely. Hence, instead of including them on the right-hand side of our regression equation as control variables, we follow the suggestion of Pei, Pischke and Schwandt (2017) and show that they are not systematically related to our measure of coastal connectivity. The results of these balancing regressions are shown in table 1.1.

In the first row, we relate connectedness to agricultural productivity, which we construct using data from the FAO-GAEZ database and following the methodology of Galor and Özak (2016): We convert agroclimatic yields of 48 crops in $5' \times 5'$ degree cells under rain-fed irrigation and low levels of input into caloric yields and assign the maximal caloric yield to 10×10km cells. For each coast cell, we then calculate the average on the same landmass within 50km from the coast cell. In the second row, we use Nunn and Puga (2012)'s measure of ruggedness, again averaged over 50km radii around our coast cells. Finally, the third row looks at distance to the nearest river mouth. For this, we used Wikipedia to create a list of all rivers longer than 200km, geocoded their mouths and mapped them to our coast cells (Nile and Danube have large deltas that map to multiple cells). We then calculate the distance of each coastal cell to the nearest river mouth. All three measures are standardized to have mean 0 and standard deviation 1. As a result, the sizes of coefficients are directly comparable to those in our connectedness regressions.

Columns (1) starts by showing the results of balancing regressions just controlling for latitude and longitude. Column (2) also adds a control for distance to the Fertile Crescent. This may be important because agriculture spread from the Fertile Crescent throughout the Mediterranean Basin, and various authors have linked the timing of the Neolithic Revolution to later development (Diamond 1997; Hibbs and Olsson 2004; Comin, Easterly, and Gong 2010). Conditional on the full set of controls we use in our analysis, neither agricultural productivity, ruggedness, nor distance to the nearest river mouth seem to have a large association with our measure of connectedness. Columns (3) and (4) show that dropping the Aegean from the sample leads to bigger associations but also impairs precision. Outside of North Africa, a slight negative association between connectedness and both ruggedness and agricultural productivity arises, but only the latter is statistically significant. Overall, our measure of connectedness does not appear to be systematically related to the three variables examined in table 1.1, especially once we control for distance to the Fertile Crescent. As a result, we will use all of latitude, longitude, and distance to the Fertile Crescent as controls in the analyses that follow.

1.4.1 Basic results

In table 1.2 then, we start by showing results for connections within 500km and the settlement densities in 750 BC from our different data sets. At this time, we expect sailors to make extensive use of direct sea connections, and hence the coefficients β_{dt} from equation 1.1 should be positive. This is indeed the case for a wide variety of specifications. We find the strongest results in the Pleiades data with the wide definition of sites, and the association is highly significant. The coefficient is slightly lower for the narrow site definition, and for Iron Age sites from the Whitehouse Atlas. Dropping the Aegean in column (2) leads to a loss of precision, with standard errors going up noticeably for all three outcome variables. Coefficients are similar in magnitude or increase, indicating that the Aegean was not driving the results in column (1). Dropping North Africa in column (3) makes little difference compared to the original results.

The effects of better connectedness seem sizable: a one standard deviation increase in connectedness increases settlement density by 20 to 50 percent of a standard deviation. While the parameterization with variables in standard deviation units should aid the interpretation, we offer an alternative view on the size of the coefficients in table 1.3. Here we presents results when we replace our previous density measure by a coarser binary variable that simply codes whether a coastal cell has at least one archaeological site within a 50km radius. The effects are basically positive but somewhat more sensitive to the particular specification and data set. Coefficients range from a high of 0.32 in the narrow Pleiades data excluding the Aegean to zero in the Whitehouse data excluding North Africa.

While this specification is much coarser than the previous one, it facilitates a discussion of the magnitude of our results. Recall that there are two modes around -0.5 and 2 in the distribution of the connectedness variable in figure 1.3. Going from -0.5 to 2 roughly corresponds to moving from a point on the coast of Southern France (say Toulon) to western Turkey (say Izmir). Using the coefficient for the wide Pleiades data in column (1) of 0.12, this move would increase the probability of having a site within 50km by 30 percentage points. Such an increase is sizable—the unconditional probability of having any site nearby in the wide Pleiades data is 61% and it is 49% among cells with connectivity below zero. Of course, a 2.5 standard deviations increase in connectedness is also a substantial increase. Most of our estimates of the effects of connectedness are far from trivial but they also leave lots of room for other determinants of growth.

We now return to our original site definition as in table 1.2. A potential concern with our results might be that we are not capturing growth and urbanization, but simply the location of harbours. To address this, table 1.4 repeats the analysis of table 1.2, but omitting coastal cells themselves from the calculation of settlement density. Here we are investigating whether a better connected coast gives rise to more settlements further inland. The results are similar to those from the previous table, indicating that the effects we observe are not driven by coastal locations but also manifest themselves in the immediate hinterland of the coast. This bolsters the case that we are seeing real growth effects of better connections. The number of observations in table 1.4 is slightly lower than before since we omit coastal cells that have only other coastal cells within a 50km radius (e.g. the north-eastern tip of Cyprus).

Table 1.5 shows some further robustness checks of our results for different subsamples. Column (1) repeats our baseline results from table 1.2. Columns (2) to (4) use only continental cells as starting points, dropping island locations. In column (2), we keep both continent and island locations as potential destinations. Results are similar or, in the case of the Whitehouse data, stronger. Columns (3) and (4) explore whether it is coastal shape or the locations of islands which drive our results. Here, we calculate connectedness using either only island cells as destinations (in column 3) or only continental cells (in column 4). Both matter, but islands are more important for our story. Coefficients for island connections in column (3) are about twice the size of those in column (4). Finally, column (5) replaces our simple connectedness measure with the eigenvalue measure of centrality; results are again very similar. These results suggest that the relationships we find are not driven only by a particular subsample or connection measure.

Our previous results are for connections within a 500km radius. Figure 1.5 displays coefficients for connectivities at different distances, using the basic specification with the wide Pleiades set of sites. It demonstrates that coefficients are fairly similar when we calculate our connectivity measure for other distances. This is likely due to the fact that these measures correlate pretty closely across the various distances. There is a small hump with a peak around 500km, probably distances which were important during the Iron Age when sailors started to make direct connections between Cyprus and Crete or Crete and Sicily. But we don't want to make too much of that.

Figure 1.6 shows results from the wide Pleiades data over time. The figure has various features. Coefficients are small and mostly insignificant until 1,000 BC but increase sharply in 750 BC, consistent with the Iron Age expansion of open sea routes. There are smaller and less significant effects of connectivity during the late Bronze Age in 2,000 and 1,500 BC. From 500 BC, the effects of connectivity decline and no correlation between sites and connectivity is left by the end of the Roman Empire. In table 1.2, we have demonstrated that the large association between connectedness and the presence of sites is replicated across various data sets and specifications for the year 750 BC, so we are fairly confident in that result. Figure 1.6 therefore raises two questions: Is the upturn in coefficients between 1,000 BC and 750 BC real or an artefact of the data? And does the association between sites and connectedness vanish during the course of the Roman Empire? On both counts there are reasons to be suspicious of the Pleiades data. Coverage of sites from before 750 BC is poor in the data while coverage during the Roman period may be too extensive.

We use the Whitehouse data to probe the findings for the earlier period. In figure 1.7 we plot coefficients again against distances varying from 100km to 1,000km. The red line, which refers to sites in the Iron Age and later, shows sizable and significant coefficients in the range between 0.2 and 0.3, as we had already seen in table 1.2, and these coefficients vary little by distance. The blue line shows coefficients for Bronze Age sites. Coefficients are small and insignificant for small distances and for very large ones but similar to the Iron Age coefficients for intermediate ones. The Bronze Age results are unfortunately a bit noisier and the pattern by distance unfortunately does not completely resolve the issue about the emergence of the sites - connections correlation either.

1.4.2 Persistence

Once geographical conditions have played a role in a site location, do we expect this relationship to be stable into the future? There are two reasons why the answer would be affirmative. Connections should have continued to play a role during the period of the Roman Empire when trade in the Mediterranean reached yet a more substantial level. Even if the relative role of maritime connectivity declined—maybe because sailors got better and distance played less of a role, or other modes of transport, e.g. on Roman roads, also became cheaper—human agglomerations created during the Phoenician period may have persisted. A large literature in urban economics and economic geography has addressed this question and largely found substantial persistence of city locations, sometimes across periods of major historical disruption (Davis and Weinstein 2002, Bleakley and Lin 2012, Bosker and Buringh 2017, Michaels and Rauch forthcoming among others). Either explanation is at odds with the declining coefficients over time

in figure 1.6 after 750 BC.
We suspect that the declining coefficients in the Pleiades data stems from the fact that the site density is becoming too high during the Roman period. In 750 BC there are 1,491 sites the data set and this number increases to 5,640 in 1 AD at the height of the Roman Empire. There are only 3,464 coastal gridpoints in our data set. As a result, our coastal grid is quickly becoming saturated with sites after the start of the Iron Age. We suspect that this simply eliminates a lot of useful variation within our data set: By the height of the Roman Empire most coastal grid points are located near some sites. Moreover, existing sites may be concentrated in well-connected locations already and maybe these sites grow further. New settlements after 750 BC, on the other hand, might arise in unoccupied locations, which are less well connected.
In order to investigate this, we split the sites in the Pleiades data into those which existed already in 750 BC but remained in the data in subsequent periods and those which first entered at some date after 750 BC. Figure 1.8 shows results for the period 500 BC to 500 AD. The blue, solid line shows the original coefficients for all sites. The black, broken line shows coefficients for sites present in 750 BC which remained in the data while the red, dashed line refers to sites that have newly entered since 750 BC. The coefficients for remaining sites are very stable, while the relationship between connectedness and the location of entering sites becomes weaker and even turns negative over time. Because the new entrants make up an increasing share of the total over time, the total coefficients (solid line) are being dragged down by selective site entry during the Roman era. This is consistent with the results of Bosker and Buringh (2017) for a later period, who find that having a previously existing city close by decreases a location's chance of becoming a city seed itself.

1.4.3 Results for a world scale

Finally, we corroborate our findings for the Mediterranean at a world scale. We have only a single early outcome measure: population in 1 AD from McEvedy and Jones (1978). This is the same data as used by Ashraf and Galor (2011b) for a similar purpose. Population density is measured at the level of modern countries, and the sample includes 123 countries. Recall that we compute connectivity for coastal cells on a 50 x 50km grid points for this exercise.
Mimicking our estimates for the Mediterranean, we start by running regressions at the level of the coastal cell, of which we have 7,441. Both connectivity and population density are normalized again. We obtain an estimate of 0.20 with a standard error, clustered at the country level, of 0.16. Here we control for the absolute value of latitude (distance from the equator) only but this control

matters little.[8]

Alternatively, we aggregate the world data to the level of countries, which is the unit at which the dependent variable is measured anyway. We normalize variables after aggregating. Figure 1.9 is a scatter plot of c_{500} against mean population density at the country level. The weights in this figure correspond to the number of coastal grid points in each country. The line in the figure comes from a standard bivariate regression and has a slope of 0.44 (0.33). These estimates are in the same ballpark of the ones for the Mediterranean in table 1.2. Note that many Mediterranean countries can be found in the upper right quadrant of this plot, highlighting how connectivity in the basin may have contributed to the early development of this region.

1.5 Conclusion

We argue that connectedness matters for human development. Some geographic locations are advantaged because it is easier to reach a larger number of neighbours. We exploit this idea to study the relationship between connectedness and early development around the Mediterranean. We argue that this association should emerge most potently when sailors first started crossing open seas systematically. This happened during the time when Phoenician, Greek, and Etruscan sailors and settlers expanded throughout the Mediterranean between 800 and 500 BC. Barry Cunliffe (2008) calls this period at the eve of Classical Antiquity "The Three Hundred Years That Changed the World" (p. 270).

This is not to say that sea trade and maritime networks were unimportant earlier. While we find clear evidence of a significant association between connectedness and the presence of archaeological sites for 750 BC, our results are more mixed as to whether this relationship began to emerge at that period because the data on earlier sites are more shaky. On the other hand, we find that once these locational advantages emerged the favoured locations retain their urban developments over the ensuing centuries. This is in line with a large literature on urban persistence.

While our paper speaks to the nexus between trade and growth we are unable to link connectedness directly to trade in goods or other channels of sea based transport like migrations or the spread of ideas. Some of the issues we hope to explore in future work are the interactions between maritime connections and other locational advantages, like access to minerals. Finally, we hope to probe the persistence of these effects more by linking the results to data on more modern city locations.

[8] Neither east-west orientation nor distance from the Fertile Crescent seems to make as much sense on a world scale. Unlike for the Mediterranean, there were various centres of early development around the world.

1.6 Tables and Figures

Figure 1.1: Timeline

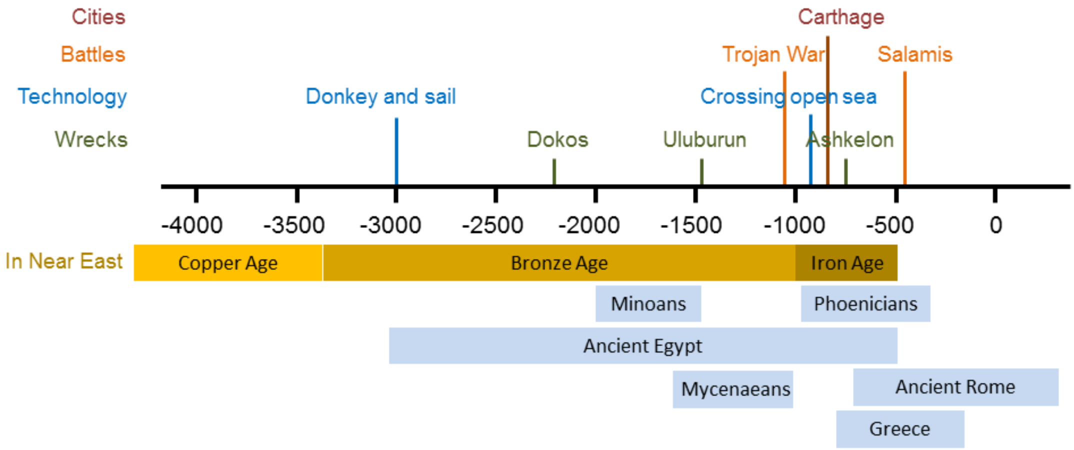

Figure 1.2: Connectedness in the Mediterranean for a 500km distance

Figure 1.3: Distribution of our connectedness variable at 500km distance

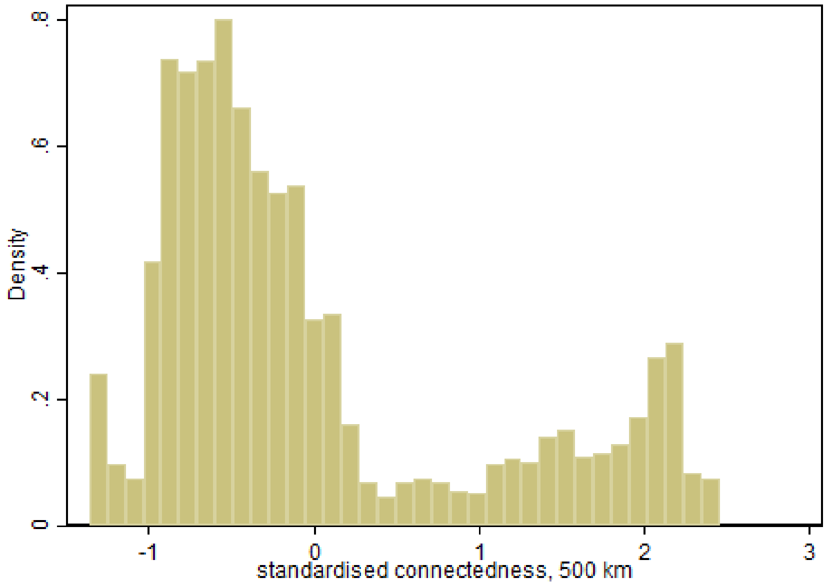

Figure 1.4: Connectedness in the world for a 500km distance

Figure 1.5: Coefficients for wide Pleiades sites by distance

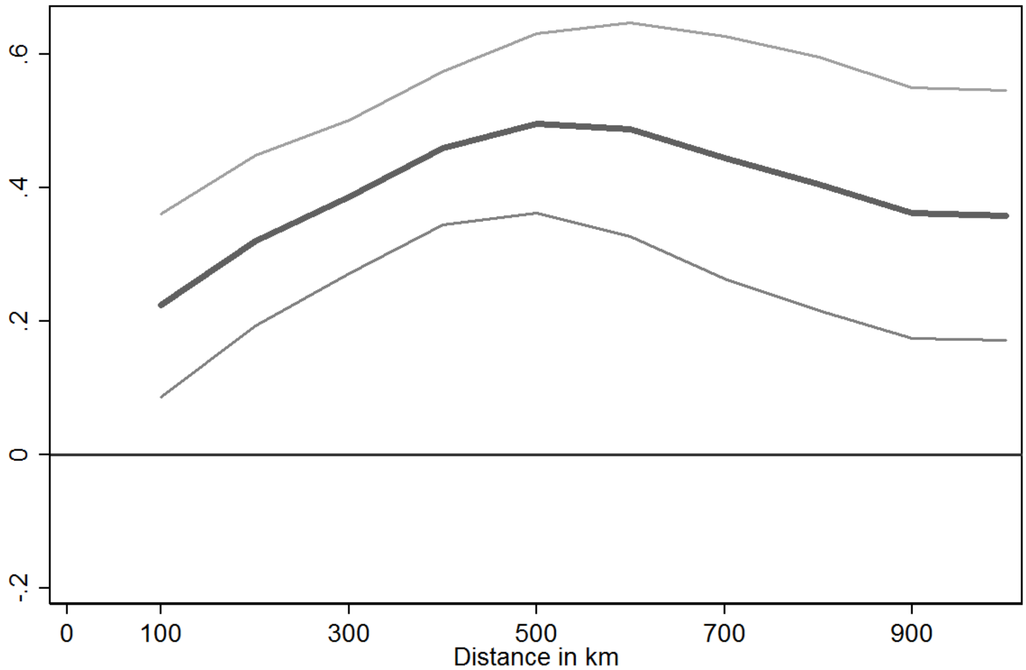

Figure 1.6: Coefficients for wide Pleiades sites over time

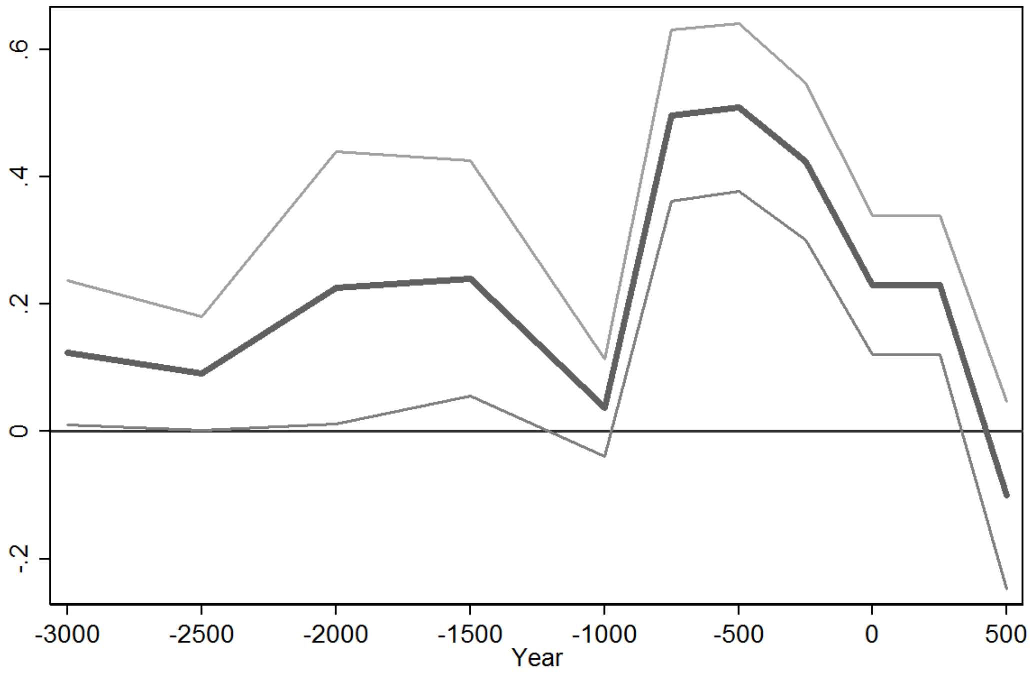

Figure 1.7: Coefficients for Whitehouse sites for different periods

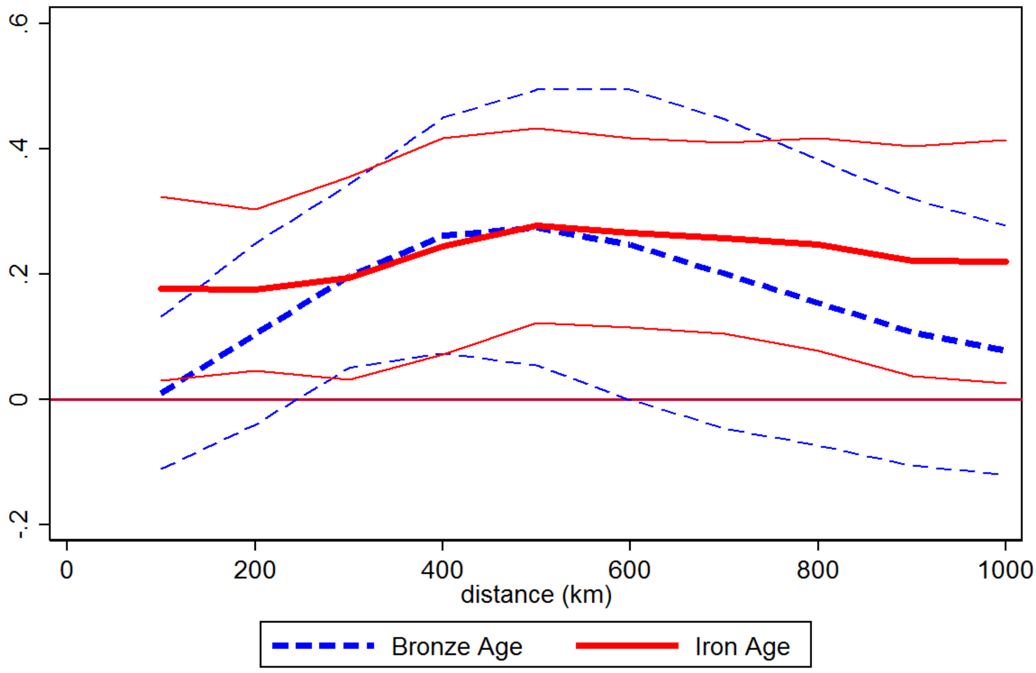

Figure 1.8: Coefficients for Wide Pleiades sites: Entry, Existing, Total

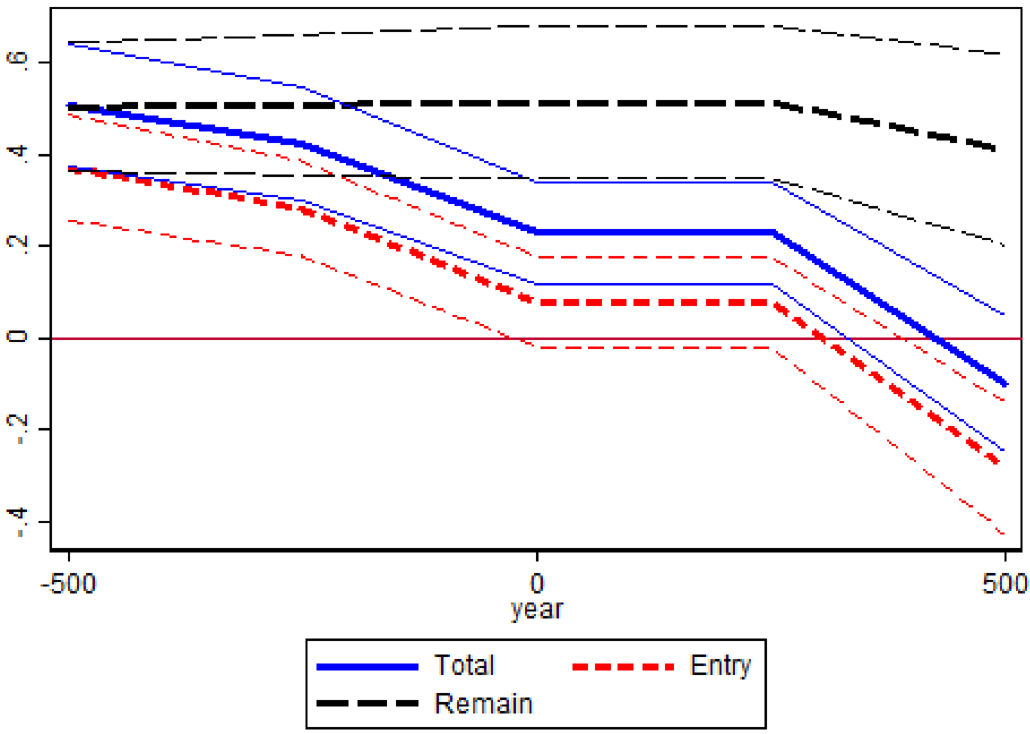

Figure 1.9: Global correlation between connectedness and population density around 1AD

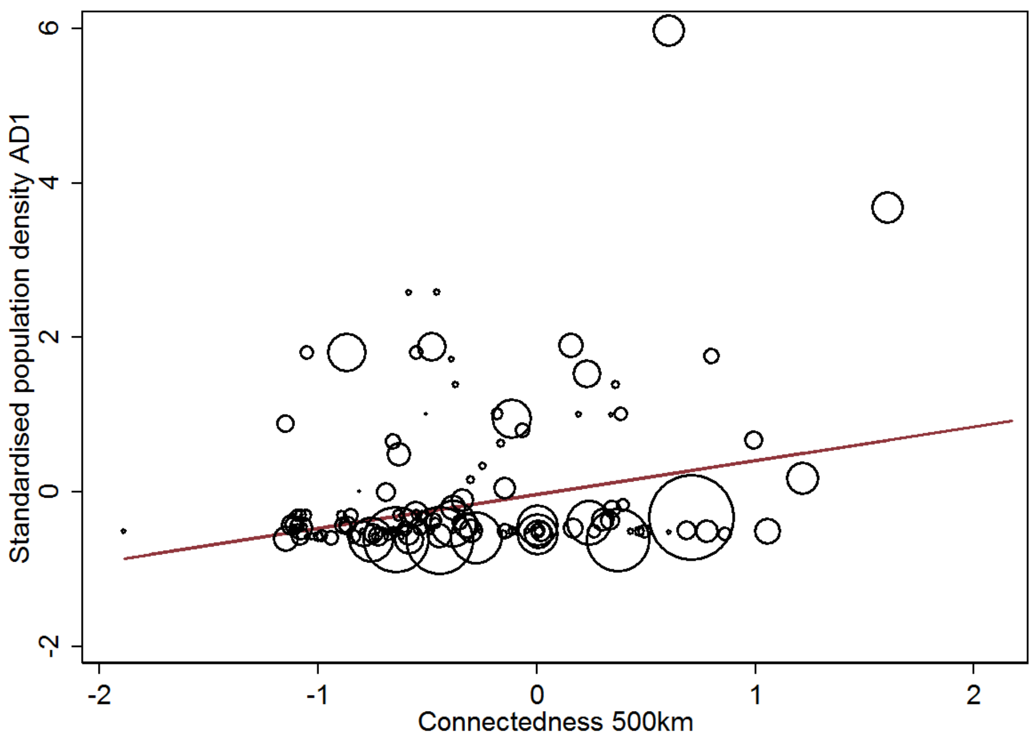

Weights reflect length of coasts of countries.

Table 1.1: Balancing checks

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
Agricultural productivity	0.11	-0.02	0.24	0.08	-0.06	-0.19
(following Galor and Ozak (2016))	(0.05)	(0.06)	(0.14)	(0.13)	(0.05)	(0.05)
Ruggedness	0.10	0.10	-0.23	-0.26	-0.16	-0.16
(following Nunn and Puga (2012))	(0.08)	(0.10)	(0.30)	(0.28)	(0.09)	(0.10)
Distance to the nearest	-0.15	-0.12	-0.26	-0.30	-0.06	0.00
river mouth	(0.09)	(0.10)	(0.27)	(0.28)	(0.08)	(0.09)
Observations	3646	3646	2750	2750	3044	3044
Controls:						
Longitude and Latitude	X	X	X	X	X	X
Distance to the Fertile Crescent		X		X		X
Dropping Aegean			X	X		
Dropping North Africa					X	X

Coefficients from a regression of various dependent variables on 500km connectedness. Standard errors clustered at the level of 2x2 degree cells, in parentheses.

Table 1.2: Basic results

Dependent variable	(1)	(2)	(3)
Pleiades Wide 750 BC	0.50	0.42	0.45
	(0.07)	(0.14)	(0.08)
Pleiades Narrow 750 BC	0.25	0.48	0.22
	(0.08)	(0.16)	(0.09)
Whitehouse Atlas Iron Age	0.28	0.50	0.20
	(0.08)	(0.18)	(0.09)
Observations	3646	2750	3044
Controls:			
Longitude and Latitude	X	X	X
Distance to the Fertile Crescent	X	X	X
Dropping Aegean		X	
Dropping North Africa			X

Coefficients from regressions on 500km connectedness. Standard errors clustered at the level of 2x2 degree cells, in parentheses.

Table 1.3: Results with a binary outcome variable

Dependent variable	(1)	(2)	(3)
Pleiades Wide 750 BC	0.12	0.17	0.11
	(0.04)	(0.14)	(0.04)
Pleiades Narrow 750 BC	0.07	0.32	0.06
	(0.05)	(0.11)	(0.06)
Whitehouse Atlas Iron Age	0.04	0.03	-0.01
	(0.04)	(0.10)	(0.04)
Observations	3646	2750	3044
Controls:			
Longitude/Latitude	X	X	X
Distance to the Fertile Crescent	X	X	X
Dropping Aegean		X	
Dropping North Africa			X

Coefficients from regressions on 500km connectedness. Standard errors clustered at the level of 2x2 degree cells, in parentheses.

Table 1.4: Results excluding coastal cells from outcome definition

Dependent variable	(1)	(2)	(3)
Pleiades Wide 750 BC	0.49	0.34	0.46
	(0.09)	(0.13)	(0.09)
Pleiades Narrow 750 BC	0.16	0.38	0.16
	(0.15)	(0.19)	(0.15)
Whitehouse Atlas Iron Age	0.31	0.53	0.30
	(0.07)	(0.27)	(0.08)
Observations	3234	2539	2647
Controls:			
Longitude and Latitude	X	X	X
Distance to the Fertile Crescent	X	X	X
Dropping Aegean		X	
Dropping North Africa			X

Coefficients from regressions on 500km connectedness. Standard errors clustered at the level of 2x2 degree cells, in parentheses. Coastal cells and their sites are omitted from the outcome definition.

Table 1.5: Results for different measure of connectedness

	Standard 500km connectedness				Centrality
	(1)	(2)	(3)	(4)	(5)
Pleiades Wide 750 BC	0.50	0.43	0.47	0.20	0.46
	(0.07)	(0.15)	(0.14)	(0.12)	(0.09)
Pleiades Narrow 750 BC	0.25	0.42	0.46	0.20	0.19
	(0.08)	(0.13)	(0.11)	(0.12)	(0.09)
Whitehouse Iron Age	0.28	0.49	0.47	0.29	0.23
	(0.08)	(0.07)	(0.05)	(0.08)	(0.09)
Observations	3646	2658	2658	2658	3646
From	All	Continent	Continent	Continent	All
To	All	All	Island	Continent	All

Coefficients from a regression of density measures from different sources on measures of 500km connectedness or eigenvalue centrality. Robust standard errors, clustered at the level of 2x2 degree cells, in parentheses. All regressions control for longitude, latitude, and distance to the Fertile Crescent.

1.7 Appendix A: Coding of Whitehouse sites

We classified the maps contained in the *Whitehouse Atlas* into three broad time periods: Neolithic, Bronze Age, and Iron Age or later based on the map title, accompanying texts, and labels for individual sites. Table 1.6 provides details of our classification of the maps. The maps on pages 72, 76, 90, and 96 straddle both the Neolithic and Bronze Age period, while the map on page 102 could refer to either the Bronze or Iron Age. For these maps, we narrowed down the dating of sites based on resources we could find on the Internet about the respective site. Table 1.7 provides details of our dating.

Table 1.6: Classification of maps in the *Whitehouse Atlas*

Pages	Map title/details	Time period
72f.	Neolithic to Bronze Age sites in Anatolia	Bronze Age or earlier
74f.	Hittites and their successors	Bronze Age
76f.	Late prehistoric and proto-historic sites in Near East	Bronze Age or earlier
90f.	Neolithic to Bronze Age sites in Western Anatolia and the Cyclades	Bronze Age or earlier
92f.	Neolithic sites in Greece	Neolithic
94f.	Cyprus	various
96f.	Crete	Bronze Age or earlier
98f.	Mycenaean and other Bronze Age sites in Greece	Bronze Age
100f.	The Mycenaeans abroad	Bronze Age
102f.	The Phoenicians at home	Bronze Age or Iron Age
104f.	The Phoenicians abroad	Iron Age or later
106f.	Archaic and Classical Greece	Iron Age or later
108f.	The Greeks overseas	Iron Age or later
110f.	Neolithic sites in the central Mediterranean	Neolithic
112f.	Copper and Bronze Age sites in Italy	Bronze Age
114f.	Copper and Bronze Age sites in Sicily and the Aeolian Islands	Bronze Age
116f.	Copper and Bronze Age sites in Corsica and Sardinia	Bronze Age
118f.	Early Iron Age sites in the central Mediterranean	Iron Age or later
120f.	The central Mediterranean: Carthaginians, Greeks and Etruscans	Iron Age or later
122	Malta	Bronze Age or earlier
123ff.	Neolithic sites in Iberia	Neolithic
126ff.	Copper and Bronze Age sites in Iberia	Bronze Age
129ff.	Early Iron Age sites in Iberia	Iron Age or later
140f.	Neolithic and Copper age sites in France and Switzerland	Neolithic
164f.	Bronze Age sites in France and Belgium	Bronze Age
172f.	The spread of Urnfield Cultures in Europe	Iron Age or later
174f.	The Hallstatt and La Tene Iron Ages	Iron Age or later
176f.	Iron Age sites in Europe	Iron Age or later

Table 1.7: Classification of specific sites in the *Whitehouse Atlas*

Map page	Site name	Neolithic	Bronze Age	Iron Age	Source
72	Dundartepe	1	1	0	see notes
72	Fikirtepe	1	1	0	Whitehouse
72	Gedikli	1	1	1	TAY Project
72	Karatas	0	1	1	Wikipedia
72	Kayislar	1	1	0	TAY Project
72	Kizilkaya	0	1	1	Wikipedia (Kizilkaya/Burdur)
72	Kumtepe	1	0	0	Wikipedia
72	Maltepe	1	1	1	TAY Project
72	Mentese	1	0	0	TAY Project
72	Mersin	1	1	1	Wikipedia
72	Silifke	0	1	1	Wikipedia
72	Tarsus	1	1	1	Wikipedia
72	Tilmen Huyuk	1	1	1	TAY Project
72	Troy	0	1	1	Wikipedia
76	Amrit/Marathus	0	1	0	Wikipedia
76	Amuq	1	1	0	Whitehouse
76	Aradus	0	1	1	Wikipedia (Arwad)
76	Atchana/Alalakh	0	1	0	Wikipedia
76	Beisamoun	1	0	0	see notes
76	Byblos	1	1	1	Wikipedia
76	Gaza	0	1	1	Wikipedia
76	Gezer	0	1	1	Wikipedia
76	Hazorea	1	1	0	Whitehouse
76	Kadesh	1	1	0	Wikipedia (Kadesh (Syria))
76	Megiddo	1	1	1	Wikipedia
76	Mersin	1	1	1	Wikipedia
76	Samaria	1	1	1	New World Encyclopedia
76	Sidon	1	1	1	Wikipedia
76	Tainat	1	1	0	Whitehouse
76	Tell Beit Mirsim	0	1	1	see notes
76	Tyre	0	1	1	Wikipedia
76	Ugarit/Ras Shamra	1	1	0	Wikipedia
90	Akrotiraki	1	1	0	see notes
90	Chalandriani	0	0	0	Wikipedia
90	Dhaskalio	0	1	0	Wikipedia
90	Dokathismata	0	1	1	Wikipedia (see notes)
90	Emborio	1	1	0	see notes
90	Fikirtepe	1	1	0	Whitehouse
90	Glykoperama	1	1	0	Whitehouse
90	Grotta	0	1	0	see notes
90	Heraion	1	1	0	Whitehouse
90	Kephala	1	1	0	Whitehouse
90	Kumtepe	1	0	0	Wikipedia
90	Mavrispilia	1	1	0	Whitehouse
90	Paroikia	1	1	0	Whitehouse
90	Pelos	1	1	0	Whitehouse
90	Phylakopi	0	1	0	Wikipedia
90	Poliochni	1	1	0	Wikipedia (see notes)
90	Protesilaos	1	1	0	Whitehouse
90	Pyrgos	1	1	0	Whitehouse

Table 1.7: Classification of specific sites in the *Whitehouse Atlas*, continued

Map page	Site name	Neolithic	Bronze Age	Iron Age	Source
90	Saliagos	1	0	0	Wikipedia
90	Spedos	0	1	0	Wikipedia
90	Thermi	0	1	0	Wikipedia (Lesbos)
90	Tigani	1	1	0	Whitehouse
90	Troy	0	1	1	Wikipedia
90	Vathy	1	1	0	Whitehouse
90	Vryokastro	0	1	0	see notes
94	Alambra	0	1	0	Whitehouse
94	Amathous	0	0	1	Whitehouse
94	Anoyira	0	1	0	Whitehouse
94	Arpera	0	1	0	Whitehouse
94	Athienou/Golgoi	0	0	1	Whitehouse
94	Ayia Irini	0	1	0	Whitehouse
94	Ayios Iakovos	0	1	0	Whitehouse
94	Ayios Sozomenos	0	1	0	Whitehouse
94	Dhenia	0	1	0	Whitehouse
94	Enkomi	0	1	0	Whitehouse
94	Erimi	1	0	0	Whitehouse
94	Idalion	1	1	0	Whitehouse
94	Kalavassos	1	0	0	Whitehouse
94	Kalopsidha	0	1	0	Whitehouse
94	Karmi	0	1	0	Whitehouse
94	Karpasia	0	0	1	Whitehouse
94	Kato Paphos	1	1	0	Whitehouse
94	Khirokitia	1	0	0	Whitehouse
94	Kition	0	0	1	Whitehouse
94	Kouklia/ Old Paphos	0	1	0	Whitehouse
94	Kourion	1	1	1	Whitehouse
94	Krini	0	1	0	Whitehouse
94	Ktima	0	0	1	Whitehouse
94	Kyrenia	0	0	1	Whitehouse
94	Kythrea	1	0	0	Whitehouse
94	Lapithos	1	0	0	Whitehouse
94	Myrtou	0	1	0	Whitehouse
94	Nikosia	0	1	1	Whitehouse
94	Nitovikla	0	1	0	Whitehouse
94	Palaiokastro	0	1	0	Whitehouse
94	Palaioskoutella	0	1	0	Whitehouse
94	Petra tou Limniti	1	0	0	Whitehouse
94	Philia	0	1	0	Whitehouse
94	Pyla-Kokkinokremmos	0	1	0	Whitehouse
94	Salamis	0	1	1	Whitehouse
94	Sinda	0	1	0	Whitehouse
94	Soli/Ambelikou	1	0	0	Whitehouse
94	Sotira	1	0	0	Whitehouse
94	Troulli	1	0	0	Whitehouse
94	Vasilia	0	1	0	Whitehouse
94	Vouni	1	1	0	Whitehouse
94	Vounous	0	1	0	Whitehouse

Table 1.7: Classification of specific sites in the *Whitehouse Atlas*, continued

Map page	Site name	Neolithic	Bronze Age	Iron Age	Source
96	Amnisos	0	1	0	Wikipedia
96	Apesokari	1	1	0	Wikipedia
96	Apodhoulou	1	1	0	Whitehouse
96	Arkhanes	0	1	0	Wikipedia
96	Armenoi	1	1	0	Minoan Crete
96	Ayia Triadha	0	1	1	Wikipedia (Hagia Triadna)
96	Diktaean Cave	1	1	0	Wikipedia (Psychro Cave)
96	Erganos	1	1	0	Whitehouse
96	Fournou Korifi	0	1	0	Minoan Crete
96	Gournes	1	1	0	Whitehouse
96	Gournia	0	1	0	Minoan Crete
96	Idaean Cave	1	1	0	Wikipedia
96	Kamares Cave	1	1	0	Wikipedia
96	Karfi	0	1	0	Wikipedia
96	Katsamba	1	1	0	Whitehouse
96	Khania	1	1	1	Wikipedia
96	Knossos	1	1	1	see notes
96	Krasi	1	1	0	Wikipedia (Malia, Crete)
96	Mallia	0	1	0	see notes
96	Mirsini	1	1	0	Whitehouse
96	Mirtos	1	1	0	Minoan Crete
96	Mitropolis	1	1	0	Whitehouse
96	Mochlos	0	1	0	Minoan Crete
96	Monastiraki	0	1	0	Wikipedia
96	Mouliana	1	1	0	see notes
96	Palaikastro	0	1	0	Minoan Crete
96	Petras	0	1	0	Wikipedia
96	Phaistos	1	1	1	Wikipedia
96	Pirgos (Nirou Khani)	0	1	0	Wikipedia
96	Platanos	1	1	0	Whitehouse
96	Plati	1	1	0	Whitehouse
96	Praisos	1	1	1	Wikipedia
96	Pseira	1	1	0	Wikipedia
96	Rousses	1	1	0	Whitehouse
96	Sklavokampos	0	1	0	Wikipedia
96	Stavromenos	0	1	0	see notes
96	Tylissos	0	1	0	Wikipedia
96	Vasiliki	0	1	0	Wikipedia
96	Vathypetro	0	1	0	Minoan Crete
96	Zakro	0	1	0	Wikipedia
96	Zou	1	1	0	Minoan Crete

Table 1.7: Classification of specific sites in the *Whitehouse Atlas*, continued

Map page	Site name	Neolithic	Bronze Age	Iron Age	Source
102	Adana (Ataniya)	1	1	1	Wikipedia
102	Al Mina	0	0	1	Wikipedia
102	Amrit/Marathus	0	1	0	Wikipedia
102	Antioch	0	0	1	Wikipedia
102	Aradus	0	1	1	Wikipedia
102	Askalon	1	1	1	Wikipedia
102	Atchana/Alalakh	0	1	0	Wikipedia
102	Atlit	0	1	1	Wikipedia
102	Beersheba	1	1	1	Wikipedia
102	Berytus	0	0	1	Wikipedia
102	Byblos	1	1	1	Wikipedia
102	Enkomi	0	1	0	Wikipedia
102	Gaza	0	1	1	Wikipedia
102	Hazor	0	1	1	Wikipedia
102	Jaffa	1	1	1	Wikipedia
102	Kadesh	1	1	0	Wikipedia
102	Kourion	1	1	1	Wikipedia
102	Megiddo	1	1	1	Wikipedia
102	Minet el-Beida	0	1	1	see notes
102	Nikosia	0	1	1	Wikipedia
102	Salamis	0	1	1	Wikipedia
102	Samaria	1	1	1	New World Encyclopedia
102	Sarepta	0	1	1	Wikipedia
102	Shechem	1	1	1	Wikipedia
102	Sidon	1	1	1	Wikipedia
102	Simyra	0	1	1	Wikipedia
102	Tarsus	1	1	1	Wikipedia
102	Tripolis	0	0	1	Wikipedia
102	Tyre	0	1	1	Wikipedia
102	Ugarit/Ras Shamra	1	1	0	Wikipedia
122	Bahrija	0	1	0	Whitehouse
122	Borg in Nadur	0	1	0	Whitehouse
122	Ghar Dalam	1	1	0	Whitehouse
122	Skorba	1	0	0	Whitehouse
122	Tarxien	1	1	0	Whitehouse

Sources and notes for site classification

Dundartepe: The Cambridge Ancient History, 3rd ed. Vol. 1, Part 2, Early History of the Middle East, eds. I. E. S. Edwards, C. J. Gadd, N. G. L. Hammond, 1971, p. 400 and Ancient West and East, Vol 1, Number 2, 2002, ed. Gocha R. Tsetskhladze, p.245

TAY Project: `http://www.tayproject.org/veritabeng.html` under the site name

Wikipedia: `https://en.wikipedia.org` under the site name

Beisamoun: Israel Antiquities Authority, Beisamoun (Mallaha), `http://www.hadashot-esi.org.il/report_detail_eng.aspx?id=809`

New World Encyclopedia: `http://www.newworldencyclopedia.org` under the site name

Tell Beit Mirsim: Biblewalks, `http://www.biblewalks.com/Sites/BeitMirsim.html`

Akrotiraki: `http://www.aegeanislands.gr/discover-aigaio/archaeology-aigiao/archaeology-aigaio.html`

Dokathismata: Entry under Amnorgos, end date unclear but clearly settled during the Classical period

Emborio: `www.archaeology.wiki/blog/2016/03/07/history-chios-seen-exhibits-archaeological-museum/`

Grotta: `http://www.naxos.gr/en/naxos/sights-and-sightseeing/archaeological-sites/article/?aid=19`

Poliochni: End date is unclear

Vryokastro: `http://www.tinosecret.gr/tour/museums/512-vryokastro.htm`

Minoan Crete: `http://www.minoancrete.comusingpull-downmenus`

Knossos: Wikipedia lists Knossos as abandoned around 1100 BC but the Whitehouse Atlas has it appear again on Iron Age map 106

Mallia: `http://www.perseus.tufts.edu/hopper/artifact?name=Mallia&object=Site`

Mouliana: `https://moulianaproject.org`

Stavromenos:
`https://greece.terrabook.com/rethymno/page/archaelogical-site-of-stavromenos`

Minet el-Beida: Wikipedia. No independent dating info for Minet el-Beida. It is routinely referred to as the harbour of Ugarit. Hence dating the same as Ugarit

1.8 Appendix B: Additional specifications

As the historical record in section 1.2 has shown, seafaring at short distances was already common in the Mediterranean before the Iron Age. However, with the advent of the Phoenicians, regular long-distance travel and trade across the open sea emerged. Our measure of 500km connectedness might conflate short-distance connectedness with long-run one. In table 1.8, we therefore create a different version of connectedness that is based on points that can be reached at a distance of 100 to 500km, i.e. excluding points within 100km. The results are very similar to the basic ones in table 1.2.

Not only connections over sea mattered but land trade was important as well. Rivers constituted important entry points for land based trade. As a result, cells close to river mouths should be relatively advantaged. The inland connections afforded by river mouths might also interact with the maritime connections we have analysed so far. Sea and inland trade maybe either complements or substitutes. Table 1.9 investigates these possibilities by entering a regressor for distance to the nearest river mouth in our regressions, and interacting it with our sea based connectivity measure. The main effect of the connectedness variable is affected little by these additions. The main effect for distance to a river mouth is positive, indicating that locations close to a river mouth have fewer settlements, not more. But the coefficients are small and not significant. The interaction terms are similarly small, except in the specification without the Aegean where we find a positive (but at most marginally significant) interaction effect. This would suggest that coastal and inland connections are substitutes rather than complements.

Table 1.8: Results when excluding short-distance connectedness

Dependent variable	(1)	(2)	(3)
Pleiades Wide 750 BC	0.50	0.43	0.45
	(0.07)	(0.13)	(0.08)
Pleiades Narrow 750 BC	0.26	0.51	0.22
	(0.08)	(0.15)	(0.09)
Whitehouse Atlas Iron Age	0.27	0.52	0.19
	(0.08)	(0.17)	(0.09)
Observations	3646	2750	3044
Controls:			
Longitude and Latitude	X	X	X
Distance to the Fertile Crescent	X	X	X
Dropping Aegean		X	
Dropping North Africa			X

Coefficients from regressions on connectedness between 100 and 500km. Standard errors clustered at the level of 2x2 degree cells, in parentheses.

Table 1.9: Results with interactions

Regressor	(1)	(2)	(3)
Connectedness	0.49	0.43	0.45
	(0.07)	(0.14)	(0.08)
Distance to river mouth	0.08	0.17	0.07
	(0.06)	(0.13)	(0.09)
Interaction	0.09	0.19	0.02
	(0.05)	(0.11)	(0.07)
Observations	3646	2750	3044
Controls:			
Longitude/Latitude	X	X	X
Distance to the Fertile Crescent	X	X	X
Dropping Aegean		X	
Dropping North Africa			X

Pleiades wide data. Coefficients from regressions on 500km connectedness. Standard errors clustered at the level of 2x2 degree cells, in parentheses.

Chapter 2

Male-biased Demand Shocks and Women's Labour Force Participation: Evidence from Large Oil Field Discoveries[1]

2.1 Introduction

Across the globe and over time, differences in female labour force participation are substantial (Olivetti 2013). Across sectors, sex ratios vary considerably, even within countries. For example, in 2008, women accounted for only 9 percent of all the hours worked in construction in the United States, but for 76 percent of all the hours worked in health services (Ngai and Petrongolo forthcoming). Because of such differences in the occupational structure, sectoral shocks can be gender-biased and affect men and women very differently.

One such example is oil wealth. The discovery of oil or natural gas resources is a male-biased labour demand shock that leads to an increase in male wages. However, as argued by Ross (2008), oil wealth may not only be beneficial to men, but also detrimental to women's labour force participation: Dutch Disease effects can harm other tradable sectors that employ women and reduce

[1]We thank Guy Michaels, Steve Pischke, Giacomo Ponzetto, Esteban Aucejo, Ghazala Azmat, Michel Azulai, Nitika Bagaria, Dan Black, Florian Blum, Zelda Brutti, Francesco Caselli, Vasco Carvalho, Alejandro Cuñat, Marta De Philippis, Jeremiah Dittmar, Richard Hornbeck, Kilian Huber, Yu-Hsiang Lei, Alan Manning, Andrea Matranga, Rachel Ngai, Raphaël Parchet, Frank Pisch, Roland Rathelot, Ferdinand Rauch, Pedro Souza, Jaume Ventura, Alwyn Young, several anonymous referees, and seminar participants at the London School of Economics, University of Oxford, Institut d'Economia de Barcelona, Halle Institute for Economic Research, Ludwig Maximilian University of Munich, University of Konstanz, the 2014 EDP Jamboree in Paris, the 15th Trans-Atlantic Doctoral Conference at LBS and the 2016 SAEe Simposio in Bilbao for their helpful comments and suggestions. To Guy Michaels, we are particularly indebted for also sharing his dataset on oil discoveries with us. All remaining errors are our own. Potlogea acknowledges financial support from the Spanish Ministry of Economy and Competitiveness.

demand for female labour. Moreover, oil discoveries increase male wages, which, in household models of labour supply, may reduce married women's labour supply (Ashenfelter and Heckman 1974, Becker 1981). Thus, both supply of and demand for female labour decrease, and as a result female employment decreases as well.

In this paper, we use the discoveries of large oil fields in the Southern United States between 1900 and 1940 as a series of large and exogenous male-biased demand shocks to study the implications of such shocks for female labour force participation and for the marriage market. We use data on the location and discovery of large oil fields, coupled with individual-level census data, to compare the evolution of female labour market outcomes in oil-rich counties relative to baseline counties. Using this difference-in-differences strategy, we find that oil discoveries have a zero net effect on female labour force participation. While oil shocks are initially male-biased and lead to increased male wages and male employment in oil mining and manufacturing, the service sector also expands and employs more women. If anything, we even find a slight increase in labour force participation for single women. At the same time, increased male wages increase the supply of "marriageable" men in the county and thus increase young women's marriage rates. Since married women were less likely to work, this has a depressing effect on female labour market involvement.

Our study suggests that the effect of oil wealth on women is not necessarily negative, but depends on social norms and institutions that establish whether other sectors are open to women and what the socially acceptable role of married women in the labour market is. In a broader context, our paper shows how the effects of a gender-biased shock to the tradable sector can be offset by income effects that lead to growth in the nontradable sector. We formalize this insight in a model of structural transformation in a small open economy. If women are allowed to work in services, our model shows that oil discoveries give rise to two opposing labour demand effects for females: A demand reduction in the tradable sector, and a demand increase in the nontradable sector. The more productive the oil industry, the more likely it is that the latter effect will dominate.

Our paper belongs to a sizeable literature that has analysed the implications of resource shocks and other demand shocks for local economies and local labour markets (Pratt 1980, Corden and Neary 1982, Carrington 1996, Sachs and Warner 1999, 2001, Black, McKinnish and Sanders 2005, Papyrakis and Gerlagh 2007, Marchand 2012, Furchtgott-Roth and Gray 2013). We confirm the findings of a series of a recent papers that have found positive effects of resource wealth on local economies (e.g. Michaels 2011, Aragon and Rud 2013, Fetzer 2014, Allcott and Kenniston 2014). Like them, we find that oil discoveries lead to sustained and substantial population growth and to a shift of the economy from agriculture into oil mining, manufacturing, and services. However, the focus of this literature has typically been on the more general implications of resource wealth in terms of employment and welfare, and gender differences have not been analysed in much detail.

At the same time, the vast changes in female labour force participation over the last one hundred years have sparked considerable interest among labour economists. Research has highlighted the importance of both supply-side and demand-side factors as well that of changing social norms. Reasons for supply shifts include medical advances (Goldin and Katz 2002, Albanesi and Olivetti 2016), labour-saving technological progress in the production of household goods (Greenwood, Seshadri and Yorukoglu 2005, Jones, Manuelli and McGrattan 2015), declining child care costs (Attanasio, Low and Sanchez-Marcos 2008) and the persistent effects of male wartime mobilization (Acemoglu, Autor and Lyle 2004). On the demand-side, the literature has documented the importance of structural transformation (Akbulut 2011, Voigtländer and Voth 2013, Ngai and Petrongolo forthcoming) and trade integration (Do, Levchenko and Raddatz 2016, Gaddis and Pieters 2017) as shifters of gender-specific demands. Finally, Fernandez (2011, 2013) and Alesina, Giuliano and Nunn (2013) present evidence for the importance of social norms. Klasen and Pieters (2015) analyse the determinants of continued low female labour force participation in India, and find evidence of adverse factors operating through both demand-side and supply-side channels, as well as through conservative social norms.

To the best of our knowledge, Michael Ross (2008, 2012) was the first to connect resource wealth to female labour force participation. Several papers have followed since then (Kang 2009, Norris 2010, Aragon, Rud and Toews 2016, Kotsadam and Tolonen 2016). However, a sizeable part of this literature has so far focused on cross-country comparisons. Such comparisons suffer from the potential problem of not being able to control satisfactorily for differences in institutions, culture and other unobservables at the country level. Our setting is able to address this issue.

We rely on within-country variation in oil wealth that is plausibly exogenous. This allows us to avoid many of the shortcomings of cross-country studies such as the difficulty to adequately control for differences in institutions and culture. Using information about the timing of oil field discoveries, we can employ a clean difference-in-differences research design, comparing the evolution of oil-rich counties before and after discoveries to that of non-oil counties over the same time period. Oil discoveries as a source of variation in labour demand are also not subject to endogeneity concerns that arise in connection with other measures such as contemporaneous oil production. Closest to our paper in this respect are two contemporaneous studies by Kotsadam and Tolonen (2016) and

Aragon, Rud and Toews (2016). The former uses mine openings and closures in Africa and finds a negative effect on female labour force participation. The key difference to our paper is that these mines are more short-lived than the large oil fields we consider- with an average life span of 10 years, they lead to local boom and bust cycles. The oil wealth in our sample, on the other hand, had very broad and long-lasting growth implications. In this respect, our paper is more similar to Aragon, Rud and Toews (2016), who however look at a persistent demand shock that is biased against males, namely the collapse of the UK coal industry. Our setting is therefore arguably closer to the theoretical mechanisms described by Ross (2008).

Finally, the case of the Southern United States offers us the advantages of analysing consistent, high-quality data from a region that at that time still relied heavily on agriculture. This is important because much of the discussion surrounding the link between resource abundance and gender issues has concentrated on developing, often heavily agrarian economies.

The rest of this paper is organized as follows: Section 2.2 discusses the channels through which oil discoveries could influence women's labour market outcomes and presents a simple model of sectoral reallocation in the wake of oil discoveries. The data and empirical strategy of our paper are then discussed in section 2.3. Section 2.4 presents the results and several robustness checks, section 2.5 concludes.

2.2 Theoretical Mechanisms

Before we proceed to the main body of our analysis, this section provides a brief discussion of the mechanisms through which oil wealth could affect the labour market outcomes of women. Existing work (see for example Ross 2008 and 2012), has identified both demand side and supply side channels that could lead to mineral wealth having an adverse impact on women's involvement in the labour market.

2.2.1 A Model of Structural Transformation and Labour Demand

On the labour demand side, oil discoveries are likely to lead to substantial structural transformation in the local economy that can change demand for male and female labour. This has already been noted by Ross (2008) who describes how oil wealth can lead to a "Dutch Disease" that reduces demand for female labour by crowding out tradable goods industries that employ women. In what follows, we present a simple model of structural transformation in the wake of oil discoveries. It will not only guide our empirical analysis, but also allow us to speak about the implications of our findings beyond our empirical setting.

Set-up

Consider a local economy (county) endowed with exogenous populations of men and women (denoted L_M and L_W respectively). This local economy can participate in the production of two tradable goods, a consumption good Y_c (we may think of this as an agricultural good), and oil (Y_o). It also produces a nontradable good called local services (Y_s).

The consumption good is produced by combining male labour, female labour, and sector specific capital (we can think of this as land) via the Cobb-Douglas technology:

$$Y_c = A_c (L_M^c)^{\alpha} (L_W^c)^{\beta} (N_c)^{1-\alpha-\beta} \tag{2.1}$$

where A_c denotes the county's productivity in the delivery of the consumption good, L_M^c and L_W^c denote male and female labour employed in the production of the consumption good, while N_c represents the quantity of capital specific to the consumption good employed in production. We assume that the county has an exogenous endowment of such capital (land) given by $\bar{N}_c$.

Oil, on the other hand can be produced using only male labour via the linear technology:

$$Y_o = A_o L_M^o \tag{2.2}$$

where A_o reflects the county's productivity in the oil sector, while L_M^o denotes the number of male workers employed in oil production.

Services, on the other hand, are produced using female labour via the linear technology:

$$Y_s = A_s L_W^s \tag{2.3}$$

where A_s reflects the county's productivity in the services sector, while L_W^s denotes the number of women employed in services.

We consider counties to be small open economies, and as a result the prices of tradable goods are exogenously determined: p_c for the consumption good, p_o for oil. The price of (nontradable) services, on the other hand, is determined locally and is endogenous. We denote it as p_s. We assume all good and factor markets are perfectly competitive.

Preferences of all agents in the economy (men, women and capital owners) are given by the Cobb-Douglas utility function:

$$U(c, s) = c^{\gamma} s^{1-\gamma} \tag{2.4}$$

where c denotes the quantity consumed of the consumption good while s denotes the quantity of services consumed.

Aside from working in the tradable sector or the services sector, women have access to an outside option given by domestic activities. This outside option yields an exogenous level of utility $\bar{u}$. Men don't have an outside option, so they will always be working, either in the tradable sector or in oil.[2]

In our empirical exercise we analyse a historical time period featuring relatively low female labour force participation. To make the model relevant to this set-up we assume that the outside option provided by the domestic activity is sufficiently attractive such that some women don't work, at least before oil is discovered (i.e. when oil productivity is zero or very low). We thus impose the restriction

$$\bar{u} > \left\{ \frac{\beta\gamma^{\beta} A_c (L_M)^{\alpha} (\bar{N}_c)^{1-\alpha-\beta}}{\left[\frac{\beta L_W}{1-\gamma+\gamma\beta}\right]^{1-\beta}} \right\}^{\gamma} [(1-\gamma)A_s]^{1-\gamma} \tag{2.5}$$

With the set-up above, we model oil discoveries as an exogenous increase in the productivity of the oil sector. In what follows, we explore the effects of rising productivity in the oil sector on the labour market outcomes of women. In particular, we are interested in the impact of oil discoveries (modelled as increases in A_o) on women's labour force participation. The main prediction of the model is presented in Proposition 1 below. The proof of this result is presented in Appendix A.

Main Result

Proposition 1. *There exist three thresholds of productivity in the oil sector, A_1^*, A_2^*, A_3^*, with $A_1^* < A_2^* < A_3^*$ such that:*

1. *If $0 \leq A_o < A_1^*$ no oil is produced in the county in equilibrium. Along this range, any increase in A_o has no effect on women's labour force participation (or wages). Increases in A_o along this range also do not affect male wages.*

2. *if $A_1^* < A_o < A_2^*$ oil is produced in the county in equilibrium. Along this range, any increase in A_o leads to a reduction in women's labour force participation but keeps women's wages unchanged. Increases in A_o along this range lead to an increase in male wages and to the proportion of men working in oil.*

[2]The reason for this assumption is twofold. Firstly, we are primarily interested in modelling female labour force participation. Secondly, the average male labour force participation rate in our sample is around 90%, making it also an empirically less relevant case to study in our setting.

3. *if* $A_2^* < A_o < A_3^*$ *oil is produced in the county in equilibrium. Along this range, any increase in* A_o *leads to an increase in women's labour force participation but keeps women's wages unchanged. Increases in* A_o *along this range lead to an increase in male wages and to the proportion of men working in oil.*

4. *if* $A_o > A_3^*$ *oil is produced in the county in equilibrium. Along this range, women's domestic outside option is no longer binding, and all women work in equilibrium. As a result, further increases in* A_o *are no longer associated with increases in women's labour force participation (which is now* 100%*), but to increases in female wages. Increases in* A_o *along this range lead to an increase in male wages and to the proportion of men working in oil.*

Corollary 1. *In the absence of the service sector (i.e. if* $\gamma = 1$*), any increase in* A_o *beyond the threshold at which the production of oil is economical in the county is associated with declines in women's labour force participation.*

Proof: See Appendix A.
For small values of A_o (i.e. before oil discoveries), the production of oil is uneconomical in the county. Moreover, in this range of low oil productivities, small increases in oil productivity have no effect, as oil remains too costly to produce in the county. As oil productivity crosses the threshold A_1^*, the county becomes productive enough in oil to begin production. This puts upward pressure on male wages, as the oil sector contributes to the demand for male labour. Rising male wages in turn affect the demand for female labour and implicitly women's labour force participation. The link between male wages and female labour demand operates through two channels, one via the tradable sector and one via the services sector.

Within the tradable sector, an increase in male wages has two opposite effects on the demand for female labour. First, it tends to crowd out the tradable sector, as the costs of one of its inputs go up. We call this a scale effect, which leads to the loss of female employment in this sector. Second, the increase in male wages tends to encourage the substitution of male labour with female labour, thus increasing demand for the latter. The strength of this substitution effect depends on how easy it is to replace male with female labour (i.e. the greater the substitutability of male and female labour, the stronger this effect is). In the case of our Cobb-Douglas technology it can be shown that the scale effect is guaranteed to dominate the substitution effect, such that rising male wages are associated with declines in labour demand coming from the tradable sector[3].

The increase in male wages caused by growth in oil related activities also

[3]For a more general specification in which male and female labour are combined into a labour aggregate with an elasticity of substitution given by ϵ it can be shown that the scale effect dominates the substitution effect as long as $\epsilon < \frac{1}{\alpha+\beta}$.

raises the demand for local services. This income effect tends to increase female employment in the service sector, thus potentially counterbalancing the depressing effect of oil on demand for female labour in the tradable sector. The strength of this effect crucially depends on the expenditure share (i.e. size) of the service sector. The larger the share of income spent on local services, the more likely it is that growth in this sector compensates for loss of female employment in the tradable sector.

For an initial range of oil productivities beyond A_1^* (and up to A_2^*), the overall negative effect of growth in the oil sector on tradable sector female labour demand dominates growing female employment in the services sector. As a result, increases in A_o lead to decreases in female labour force participation. Beyond A_2^*, the relationship between these two forces is reversed, with the service sector employment growth dominating female job losses in the tradable sector. As a result, beyond A_2^* increases in A_o are associated with increases in female labour force participation. Finally, beyond A_3^*, male wages become sufficiently high such that high demand in the services sector draws the entire female population into the labour force (the outside option of domestic activity no longer binds). For $A_o > A_3^*$ further increases in A_o are no longer reflected in higher female labour force participation (which is now 100%), but in higher female wages.

Given the analysis above, it should be clear that in the absence of the service sector, any increase in A_o beyond the threshold that makes oil production economical in the county is associated with declines in female labour force participation. This is because the negative effect of rising male wages on the tradable sector is still present, while the compensating force of increasing female labour demand due to growth in services is no longer operational.

Finally, care has to be taken when taking the predictions of this model to the data. While we have modelled oil discoveries as a continuous increase in A_o, the setting we study is perhaps better described by a discrete jump in A_o. In this context, the predictions of the model regarding the impact of oil discoveries on female labour force participation are ambiguous. As the relationship between (the continuous) A_o and female labour force participation is U-shaped, a discrete jump in A_o may be associated with a decline, an increase or no change in female labour force participation, depending on which portion of the U-shape the local economy jumps to.

Our model offers several important insights. Firstly, if the local service sector is absent or if women are precluded from working in it, oil discoveries reduce demand for female labour in the tradable sector and therefore drive them out of the labour market. On the other hand, if the service sector is present and is open to women, oil discoveries continue to harm demand for female labour in the tradable sector, but increase demand for female labour in the nontradable sector. The more productive the oil sector is, the more the latter effect is likely to dominate, and female labour force participation may even increase.

2.2.2 Reducing Labour Supply

In addition to these demand side effects, oil can also depress women's incentives to participate in the workplace, leading to an inward shift of their labour supply curve. Resource mining and its associated activities tend to display heavily male labour forces (something that will also be apparent in our setting) and thus constitute a male-biased labour demand shock. Such a shock translates into substantially higher male wages, which constitutes an increase in non-labour income for the spouses of these males. Moreover, in a Becker (1981) model of household labour supply, an increase in male wages would increase men's comparative advantage in market work. Both arguments give rise to a negative elasticity of married women's labour supply with respect to their husband's wages, which is also consistent with a sizeable empirical literature in labour economics (e.g. Ashenfelter and Heckman 1974). In addition to boosting the non-labour income of the wives of oil workers, oil could also increase government revenue and thus government transfers. (Ross, 2008). However, recent studies have not found that oil wealth translates into substantial increases in public goods provision (Caselli and Michaels 2013, Martinez 2016). The more important effect would thus be on married women whose husbands experience wage increases. For this group, we would expect a decrease in labour supply.

Taken together, our theoretical discussion indicates two effects of oil on female labour supply: As shown in the previous section, the labour demand effect is ambiguous and depends on the size of the oil discovery shock. In addition, there is a negative labour supply effect for women due to their husbands' wages increasing. The overall effect is therefore ambiguous, and we proceed to assess it empirically.

2.3 Data and Empirical Strategy

In order to identify major oil field discoveries, we draw on the dataset compiled by Guy Michaels (2011), which lists all counties in the Southern and Southwestern United States that are situated above an oil field of 100 million barrel or more before any oil was extracted. We will refer to these counties simply as oil-rich counties. In addition, the dataset also contains information about the earliest major discovery in each county. We treat a county as having discovered its oil wealth after the first major discovery has taken place. Figures 2.1 to 2.5 give an overview over the geography of oil discoveries between 1900 and 1940 based on county shapefiles from NHGIS (Minnesota Population Center, 2016). There are only relatively few oil discoveries in the early years, with most discoveries happening in the 1920s and 1930s. In terms of geographic scope, we follow Michaels (2011) and include all counties within 200 miles of the oil-rich counties in Texas, Louisiana, and Oklahoma. This gives us a sample of 774 counties in total.[4]

For our outcome variables, we use the individual-level US census data available from the Integrated Public Use Microdata Series (IPUMS) for 1900 to 1940 (Ruggles et al 2015). This includes three novel, recently made available full count (100%) samples for 1920, 1930, and 1940. For 1900 and 1910, we use the largest samples available- a 5% sample for 1900, and a 1.4% sample that over-samples minorities for 1910. We generally focus on the part of the population which is of working age, defined to be between 17 and 65.

Most of our variables of interest are available for our whole period of analysis. This includes labour force participation, the sector in which an individual works, sex, race, and marriage status. Employment status conditional on labour force participation was not recorded in the census of 1900 and 1920, but is available for the remaining three census decades. For earnings, we observe annual wage income in 1940 only, but we also have an "occupational earnings score" variable that is available for 1900-1940. This variable assigns each individual the percentile rank of its occupation in terms of median earnings in 1950.

The staggered nature of oil discoveries across space and over time lends itself quite naturally to a difference-in-differences research design. The basic regression we run uses the IPUMS individual level data and is of the form

$$y_{ct} = \alpha_c + \tau_t + \beta DiscoveredOilField_{ct} + X_c'\gamma_t + u_{ct} \tag{2.6}$$

where y_{ct} denotes outcome y (e.g. the labour force participation rate) in county c and year t. $DiscoveredOilField$ is a dummy that equals 1 if county c is oil-rich and if at least one of its major oil fields has already been discovered, and 0 otherwise. τ and α are year- and county fixed effects. X is a vector of control variables that vary at the county level. In line with Michaels (2011), we control for several geographic features that might be spuriously correlated with oil wealth: Longitude, latitude, aridity, average annual rainfall, distance to the closest navigable river, and distance to the ocean. All of these variables vary at the county level only, but we allow for them to have time-varying effects.

[4]One potential issue is that especially in the Western part of the sample, population was growing and counties therefore often changed, were split up, or newly created. In order to address this, we compare the area of each county in each census decade to its area in 1940, and drop all observations from a given county-year cell if the absolute difference of the county's area in that year compared to 1940 exceeds 5%. In addition, we drop all counties from Oklahoma in 1900, as this state at this point was still largely unorganized and divided into the Oklahoma and Indian Territory.

Since our key variation lies at the county x year level, we aggregate the individual census data to this level.[5] We then run our regressions at this level, weighting by the respective cell sizes and clustering standard errors at the county level. This approach produces identical point estimates and standard errors as running all our regressions at the individual level (Angrist and Pischke 2009). In all regressions, we cluster standard errors at the county level. When calculating standard errors, we use Young (2016)'s procedure that calculates an effective degrees of freedom correction for clustered and robust standard errors. These effective degrees of freedom are also used when performing hypothesis tests.

For wages, we only have data from 1940, requiring a cross-sectional specification with stronger identifying assumptions:

$$y_c = \beta DiscoveredOilField_c + X_c'\gamma + s_s + u_c \tag{2.7}$$

In this case, $DiscoveredOilField_c$ codes whether county c has an oil field which was already discovered by 1940. X is the same vector of control variables as above, and s are state fixed effects. The more demanding identification assumption behind this specification is that counties without oil wealth and counties that had not yet discovered their oil wealth by 1940 are a valid counterfactual for those counties with already discovered oil fields in 1940. In Appendix B, we show several robustness checks, the results of which support this assumption.

The setting of our analysis is the American Southwest five to nine decades after the American Civil War. After the war, the US South remained a predominantly agricultural region that lagged considerably behind the rest of the country and caught up only very slowly (Wright 1986, Caselli and Coleman 2001). This is also borne out by the county-level summary statistics displayed in table 2.1. We show the labour force shares of 4 major sectors (agriculture, manufacturing, services, oil mining) for the beginning and end of our sample. As can be seen, agriculture is by far the most important sector in 1900, but loses importance over the next 40 years. Manufacturing and services, on the other hand, are slowly gaining ground. Over the whole time period, our area of observation is a very rural region, with an urban population share of only 15%. In addition, the summary statistics also show a very low level of female involvement in the labour force. Female labour force participation stood at 16.6% in 1900 and increased by 3 percentage points by 1940. This is no Southern peculiarity: Across the whole United States, female labour force participation rose from 20.6% in 1900 to 25.8% in 1940 (Goldin 1990, Chapter 2). However, there was considerable heterogeneity along two dimensions: Black women were much more likely to work than white women (Goldin 1977), and single women had much higher

[5]When aggregating, we weight individual observations by the person weights provided by IPUMS in order to improve the representativity of our sample.

labour force participation than married ones (Goldin 1990).

All in all, our basic setting is thus a still heavily rural and agricultural economy that is slowly moving towards more "modern" sectors, and which over the whole period of analysis displays relatively low female labour force participation rates.

2.4 Results

2.4.1 Oil and gender labour market differences

In 1900, the economy of the Southern United States was thus still predominantly agricultural and nearly no major oil field had yet been discovered, as can be seen from figure 2.1. To set the stage for our main analysis relating resource shocks to gender outcomes, we first check whether oil discoveries indeed represent important developments for the discovering counties. Table 2.2 shows difference-in-differences results for the log of the working age population, female population share, and the share of the labour force employed in, respectively, oil mining, manufacturing, services, and agriculture. We do not find any evidence for a resource curse, but rather detect substantial growth: Population increases by around 42%, and the labour force shifts from agriculture into oil mining and, to a lesser extent, manufacturing and services. These findings largely confirm the results of Michaels (2011) for the same region, but are based on weaker identifying assumptions. They show that oil discoveries constitute substantial shocks to a county's economy that lead to growth and structural transformation. Interestingly, the population growth seems gender-balanced, with the female population share staying unchanged. In tables 2.3 and 2.4, we then analyse whether the economy's structural transformation affects men and women differentially. The results show that this is clearly the case: Oil discoveries lead to a decrease in the share of the male labour force employed in agriculture by 8.8 percentage points, while increasing the male labour force share employed in oil mining and manufacturing by 6.5 and 1.2 percentage points, respectively. As there is virtually no oil mining prior to oil discoveries, this means that all of a sudden, 6.5% of the whole male labour force work in the oil extraction industry, a sizeable change to the local economy's structure. Women, on the other hand, are nearly not affected at all by the oil industry. However, they also experience a decrease in the importance of agriculture of 6.3 percentage points, whereas the service sector's labour force share grows by 5.2 percentage points. Thus, both genders leave agriculture in relative terms, but while men go increasingly into the oil and manufacturing industries, women flock to the service sector. In terms of our model, it seems that the service sector in our economy is clearly open towards women.

In this case, our model has an ambiguous prediction in terms of female labour supply depending on the strength of the oil shock. Table 2.5 shows difference-in-differences estimates of the effect of oil wealth on labour force participation by gender. In columns 1 and 3, we show results for male and female labour force participation, while columns 2 and 4 show results for the employment rate conditional on being in the labour force, a measure which we do not have for 1900 and 1920. Interestingly, we do not find large effects on either gender: While the male labour force participation seems to increase at least slightly in oil-rich counties, we do not observe any changes in female labour force participation. Similarly, we fail to identify any effect of oil discoveries on either the male or female employment rates. The respective point estimates are close to zero across all specifications and statistically insignificant. Thus, in contrast to the cross-country results of Ross (2008), we do not find evidence for a negative relationship between oil wealth and female labour market participation. The most likely explanation for this is the growth of the service sector. As discussed in section 2.2, even though the initial demand shock of oil is male-biased (and columns 1 of tables 2.3 and 2.4 suggest it indeed is), oil only reduces women's labour market involvement if there is no other sector that benefits from oil and is open to women. In our case, as table 2.4 shows, the service sector is playing this role: As the economy grows, demand for services increases, and more women go to work in such jobs.[6]

So far, we have focused on quantities transacted in the labour market. To analyse the effect of oil discoveries on wages, we look at two outcome variables. The first is the most direct measure of wages, annual wage income. However, this variable is only available for 1940, which requires a cross-sectional specification with stronger identifying assumptions. To at least partially remedy this, we also look at the occupational earnings score, available from IPUMS. The occupational earnings score is a percentile rank of occupations based on their 1950 median earnings. It thus does not capture within-occupation changes in wages, but it does capture movements of people into jobs that pay relatively more or less. The advantage of this variable is that it is available for the whole time period of 1900-1940 so that we can use our difference-in-differences specification.

In panel A of table 2.6, we report the results. In columns 1 and 2 we report the averages of log annual wage earnings for men and women separately. Unlike our previous results on labour force participation and employment, we identify substantial effects of oil wealth on local wage rates that differ by gender. While male wages increase by 26 log points on average, female wages stay nearly unchanged. In columns 3 and 4, we convert annual wage income to percentile

[6]In Appendix B, we show that nearly the whole growth in the importance of the service sector as employer of female labour is due to the personal service and entertainment sector. This suggests that income effects rather than demand linkages from other industries are the driver of this growth.

ranks of the 1940 wage distribution to facilitate comparison with the occupational score variable.[7] We see the same relative movement: Men move up 7.5 percentiles, whereas women increase their earnings by at most one percentile. Columns 5 and 6 display cross-sectional estimates for 1940 for the occupational score variables. They show the same relative picture: Men gain 5 percentage ranks, while women in this specification even lose ground. In the panel difference-in-differences specification of columns 7 and 8, the gap between the estimates is a bit smaller, but still sizeable and in line with previous findings. Clearly, the structural transformation that counties experience in the wake of oil discoveries is not gender-neutral: While men obtain high-paying jobs in oil mining and manufacturing, the service sector jobs for women tend to be relatively low-paid occupations in the personal service sector. The gender gap therefore increases. In addition, comparing columns 3 and 4 to 5-8, we see that a large part of the wage growth is due to occupational changes, but there is also evidence for within-occupational wage increases, as the wage percentile coefficients are larger than then occupational score ones. One concern with these findings is that the male wage growth could be driven by the high wages in oil mining and thus only accrue to the small group of oil miners. In panel B we therefore drop all individuals working in the oil extraction sector. As can be seen, while the wage growth decreases a bit, it stays economically and statistically significant: Outside of the oil industry, men gain 3 percentile ranks in the occupational earnings distribution, 6 percentiles in the wage distribution and 22.5 log points in wages. For women, on the other hand, the coefficients are nearly unchanged. Given the low number of women in the oil mining industry, this should not come as a surprise.

All in all, so far, our results display a much more nuanced effect of the initial male-biased demand shock on women: While men greatly benefit from new jobs in oil mining and manufacturing and can substantially increase their wages, female wages remain more or less unchanged. Moreover, oil does not appear to drive women out of the labour market, instead female labour force participation remains unchanged. The likely explanation is an increase in service sector jobs that absorbs any "excess" female labour. However, the service sector increase comes predominantly in personal service sector jobs that are not as well paid, such that the gender pay gap increases.

[7] In creating percentile ranks, we follow the IPUMS procedure behind occupational earnings scores very closely: We first use the 1940 national full count to create a z-score of wage income and then transform this into a percentile rank.

2.4.2 Heterogeneity along marriage status and race

One potential issue with our findings so far, in particular with the absence of a negative effect on female labour force participation, could be that most women were not part of the labour force anyways. As Goldin (1990, 2006) has shown, female labour force participation was still very low and only slowly rising during the first third of the twentieth century. This is also visible in our sample, where the female labour force participation rate over the whole time period is just below 24%. The vast majority of women thus was most likely already in a corner solution with no labour force participation, which could explain why we do not observe any changes in labour supply. While this argument has merit, it overlooks two important sources of heterogeneity in female labour market involvement: Marriage status and race. We therefore analyse whether we observe heterogeneous responses to oil discoveries along these dimensions, and whether these potential heterogeneous responses can explain the absence of an overall change in female labour market involvement.

Marriage status was a crucial determinant of labour force participation in the early 20th century United States. Until the 1920s, it was typical for women to exit the labour force upon marriage (Goldin 2006), and in many places, employers had explicit "marriage bars" that prevented the hiring or led to the firing of married women. Not surprisingly then, in 1900, across the whole country, only 5.6% of all married women were working, as opposed to 43.5% of all single women, and the former share had only increased to 13.8% by 1940. (Goldin 1990) Our failure to find a significant labour force participation result could thus be due to us pooling irresponsive married women with more responsive singles. If this is the case, we would expect to find a reaction for singles at least. Table 2.7 therefore shows results separately for single and married women. As can be seen, while we find a small negative point estimate for married women, if anything, our point estimate for single women is positive and amounts to 1.3 percentage points. While not very precisely estimated, this would mean that far from crowding singles out of the labour market, oil if anything *increased* their propensity to work. For the more responsive group of singles, our conclusion that oil does not drive them out of the labour market thus holds a fortiori.

However, in addition to changing female labour force participation directly, oil might also have other demographic effects. As we have shown, male wages increase substantially after oil discoveries. This could increase the local supply of "marriageable" men and thereby increase marriage rates (Buckley 2003, Jelnov 2016). Increased marriage rates in turn have a dampening effect on female labour force participation due to married women being less likely to work.

One challenge in assessing the plausibility of this channel is the problem that there is a large "stock" of already married women who are not affected by oil discoveries at all. We therefore need to focus on "marginal" women whose decision when and whether to marry can still be affected by oil wealth- these will typically be younger women, among which the marriage rate and thus the stock of married women is lower. In table 2.8, we show the effects of oil discoveries on a woman's probability of being married. Column 1 shows results for all women and illustrates our stock problem: Among all women, the marriage rate is close to 70% (and many of those not married are widowed already), and we do not find an effect of oil on marriage rates. For women aged 30 or below, or 25 or below, where more women are in the process of making their marriage decision, we find that oil increases the probability of being married by a statistically significant 1.5 percentage points.

Overall, separating the analysis for single and married women supports our previous conclusion that oil does not reduce female labour force participation. If anything, single women even experience a slight increase in labour market involvement. However, this effect is counterbalanced by the fact that young women also are more likely to get married, and married women during this period rarely work. Social norms about the appropriate role of married women thus seem to keep women from benefiting more from oil than could have been possible.

In addition to marriage status, race was another important determinant of female labour force participation in our setting. Black women were considerably more likely to work. In 1920, the labour force participation of white females was 21.6%, while for non-white women, it stood at 43.1% (Goldin 1990). Reasons for this difference include economic factors such as black husbands' lower average incomes, but also cultural differences and the legacy of slavery. (Goldin 1977, Boustan and Collins 2013)

Our overall unchanged female labour force participation rate thus might mask racial heterogeneity. Black women with a generally greater labour force involvement should be more responsive to the shocks brought about by oil discoveries. In addition, they were particularly present in the low-skilled personal service sector jobs that experienced particular employment increases. (Goldin 1977) To assess whether black and white women react differentially to oil discoveries, we split our sample by race. Table 2.9 shows the results for labour force participation, average earnings scores and the share of the labour force working in services and agriculture- in columns 1-4 for whites, in columns 5-8 for blacks. The results show that black women indeed react much more to oil discoveries: The importance of the agricultural sector as employer decreases by nearly 9 percentage points for them, while the services experience a corresponding increase. This movement from agriculture to services also appears financially beneficial, as black women gain more than a percentile rank in the occupational earnings score. For white women, on the other hand, all these responses are qualitatively

of the same sign, but much more muted. The likely explanation for this heterogeneity is that white women start from considerably better paying jobs- they are on average in the 32nd percentile of the median occupational earnings distribution, whereas blacks are in the 7th, and few white women work in agriculture. For whites, who already have relatively well-paid service sector jobs, the new personal service jobs are thus not very attractive. Black women, on the other hand, can considerably improve over their existing agricultural jobs. Finally, however, overall labour force participation stays unchanged for both races.
All in all, analysing the effect of oil along the dimensions of race and marriage status has uncovered several important additional results. First off, oil discoveries do increase the earnings of women, but only of black women who on average had the worst-paying jobs and for whom the new personal service sector jobs are an improvement. White women, the majority of which is already in the service sector, do not gain and also do not react a lot in general. Secondly, if anything, single women even experience an increase in labour market opportunities, but this effect is counterbalanced by an increased propensity to marry and a social environment where married women are strongly discouraged from working. Again, as in the case of the service sector's openness, this highlights the role of social norms. Had the labour market been more open to married women, we might even have seen a slight increase in female labour force participation. Instead, women seem to benefit less than they could have. Finally, our general results seem robust to splitting the sample of women both by race and marriage status. In Appendix B, we perform further robustness checks and show that our results are also robust to using a leads and lags specification instead of a simple dummy for oil discoveries, to including county-specific linear time trends, to dropping all counties without large oil fields, and to excluding all counties who do not have a discovered oil field themselves, but are adjacent to one. Using the methodology of Pei, Pischke and Schwandt (2017), we also show that oil discoveries are not systematically associated with the presence of the boll weevil, another important local economic shock during our period of analysis (Lange, Olmstead and Rhode 2009, Ager, Brueckner and Herz forthcoming).

2.4.3 Migration

As we have shown in table 2.2, counties that discover oil experience great population growth. This immediately raises the question of migration. Migration is not a problem for the identification strategy of our paper, but it poses potential challenges for the interpretation of our results: Do the changes and potential benefits of oil accrue only to migrants, or does the incumbent population also share the benefits? The ideal test of this hypothesis would be to split the sample by migrant status. Unfortunately, however, we do not observe migrant status.

What we do know is each individual's state of birth. We can thus at least identify a subsample of migrants, namely those who migrated across state borders. These are also the migrants who on average travelled greater distances. If there is a concern of selective migration and spatial sorting, this would therefore be the group where we would suspect these factors to operate in the strongest way. As one admittedly imperfect way of addressing the questions posed by migration, we therefore split the sample by birth- one group consists of people that were born in a different state and thus clearly migrated to their new residence at some point in their life (albeit not necessarily after oil discoveries). The second group consists of people that live in their state of birth. This group is made up of short-distance migrants and individuals who have not moved at all. In tables 2.10 and 2.11, we show our key results separately for these two groups. Table 2.10 shows results for the overall sex ratio and for the male-related outcomes: labour force participation rate, employment rate, average earnings scores and the sectoral allocation of the labour force. For men, we can indeed see some evidence for selective migration: For men born out of state, the occupational earnings score results and the sectoral reallocation from agriculture to oil and manufacturing are considerably stronger than for men born in their state of residence. However, neither group displays a huge change in labour force participation rate. Table 2.11 shows the results for our female-related outcome: Labour force participation and employment rates, marriage rates of women aged 25 or below, average occupational earnings scores and sectoral composition of the labour force. Here, the case for selective migration seems considerably weaker: The marriage rate increase is driven almost entirely by women born in state, and the sectoral reallocation from agriculture to services is very similar for both groups. In addition, again no group displays a change in labour force participation rates or earnings scores. Overall, for men, selective migration is clearly a concern and at least part of our findings are due to migrants being systematically different from the existing population. However, for women, this case seems much weaker.

2.5 Conclusion

In this paper, we have analysed the implications of a male-biased labour demand shock in the tradable sector on labour market outcomes by gender. Our particular focus has been on testing the hypothesis that oil discoveries adversely affect female labour market participation.

We have provided a theoretical framework in which we model the oil sector as male-dominated and oil discoveries as male-biased labour demand shocks. In our model, oil discoveries lead to two opposing changes in female labour demand: Demand for female labour in the tradable sector is reduced, whereas demand for

women in the nontradable sector increases. If the service sector does not exist or is not open to women, female labour force participation unambiguously declines. On the other hand, if the service sector is open to women, the nontradable sector might have enough pull to actually increase female labour force participation.
We then use oil discoveries in the US Southwest between 1900 and 1940 to empirically estimate the effect of oil wealth on gender labour market differences. Contrary to previous studies (e.g. Ross 2008), we do not find a negative relationship between oil wealth and female labour force participation, which is instead found to be unchanged. Our explanation for this is a process of sectoral reallocation by gender that is consistent with our model: The initial oil shock indeed brings about male-biased employment increases in oil mining and manufacturing with associated wage gains for men. However, the service sector also expands and offers ample employment opportunities for women, who therefore do not leave the labour force. If anything, we even find that the labour market prospects of single women slightly improve. However, at the same time, there is an increased supply of marriageable men in the county, and the marriage rate of young women increases. In our period of analysis, married women were strongly discouraged from market work, so that this increased marriage rate depresses any potential gains women could have made from oil discoveries.
Our findings point to the importance of social norms and institutions in transmitting gender-biased demand shocks. A first such norm is the openness of the service sector to women. In our setting, the service sector is open to women and expands considerably. Our model indicates that if the latter was not the case, we would indeed expect to find a negative employment effect for women. This could also explain the negative cross-country result in Ross (2008), which, as noted by himself, seems to be absent in countries where women are common in the service sector.
A second relevant social norm concerns the "appropriate" role of married women. In our setting, married women are socially discouraged (and sometimes outright barred) from working, and as a consequence, they potentially benefit less from oil discoveries than they could have under different social norms.
The broader implication of our paper is that an initial gender-biased shock to a tradable sector does not necessarily lead to aggregate gender-biased outcomes. Through income effects (or potentially backward or forward linkages), the initial gender-biased implications can easily be attenuated or even fully netted out.

2.6 Tables and Figures

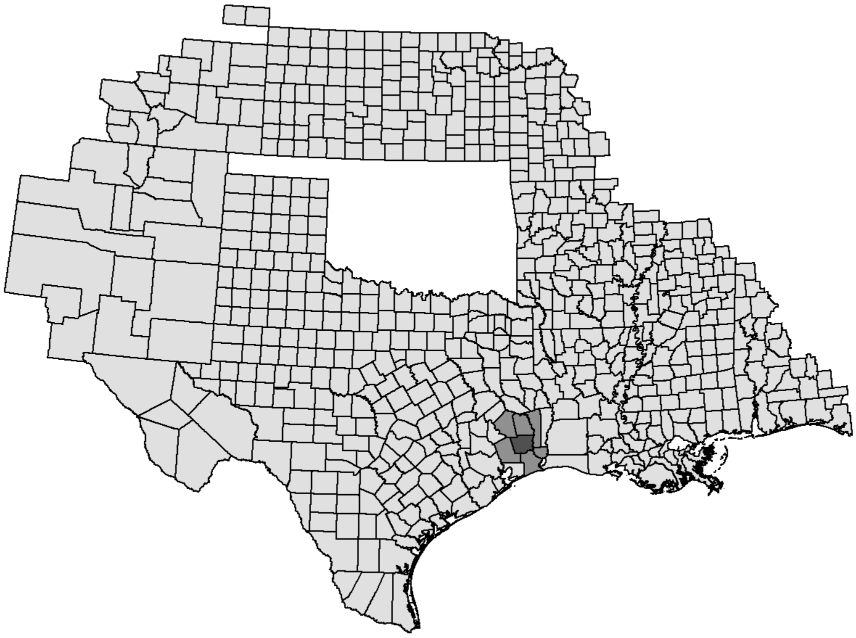

Figure 2.1: Map of Oil Discoveries 1900

Note: Oil abundant counties (dark grey), Neighbours of oil abundant counties (grey), Other counties (light grey)

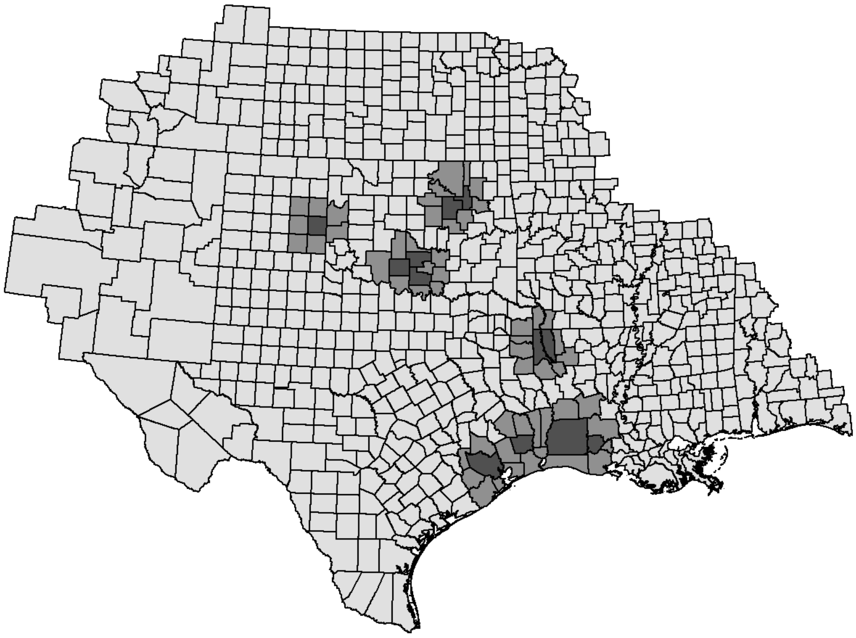

Figure 2.2: Map of Oil Discoveries 1910

Note: Oil abundant counties (dark grey), Neighbours of oil abundant counties (grey), Other counties (light grey)

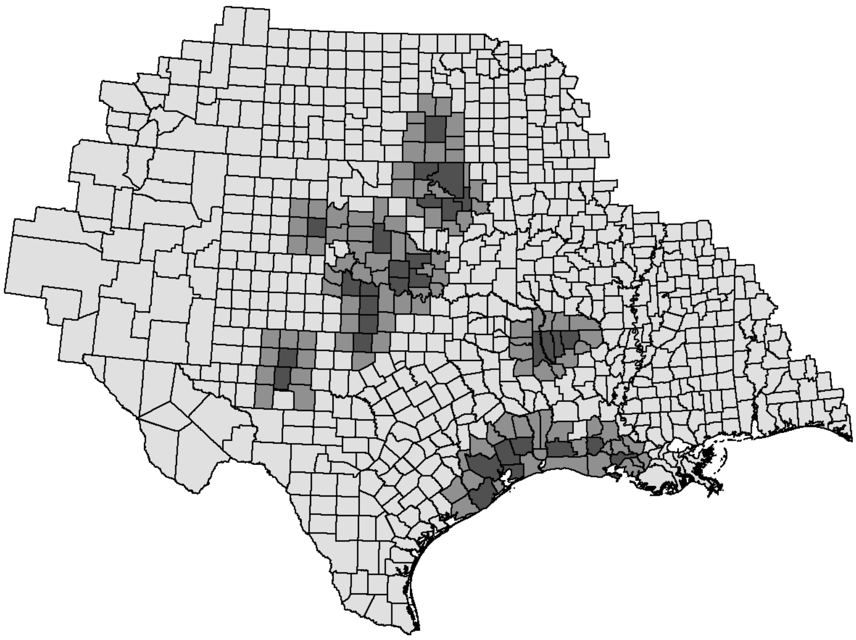

Figure 2.3: Map of Oil Discoveries 1920

Note: Oil abundant counties (dark grey), Neighbours of oil abundant counties (grey), Other counties (light grey)

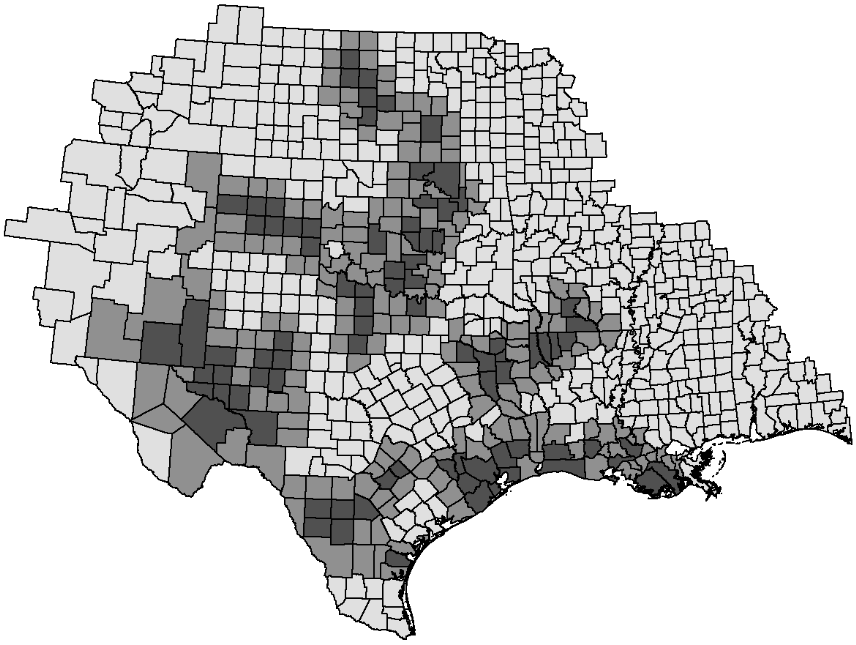

Figure 2.4: Map of Oil Discoveries 1930

Note: Oil abundant counties (dark grey), Neighbours of oil abundant counties (grey), Other counties (light grey)

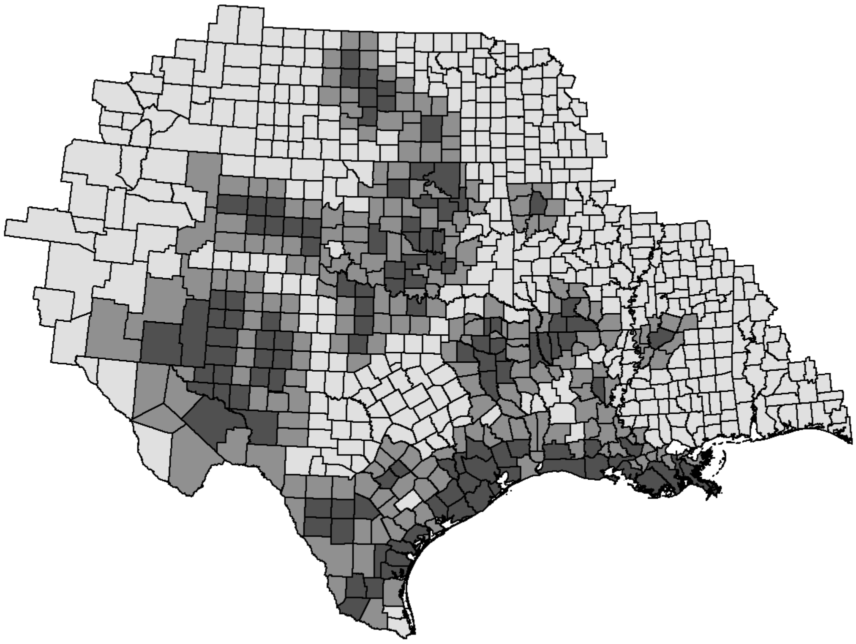

Figure 2.5: Map of Oil Discoveries 1940

Note: Oil abundant counties (dark grey), Neighbours of oil abundant counties (grey), Other counties (light grey)

	1900			1940		
	Mean	Standard deviation	Obs.	Mean	Standard deviation	Obs.
Agriculture LF Share	0.716	0.199	588	0.481	0.184	758
Manufacturing LF Share	0.038	0.054	588	0.062	0.057	758
Services LF Share	0.102	0.070	588	0.152	0.057	758
Oil Mining LF Share	0.000	0.001	588	0.016	0.045	758
Female LF Participation Rate	0.166	0.134	587	0.198	0.070	774
Male LF Participation Rate	0.934	0.045	588	0.866	0.141	774
Black population share	0.206	0.256	588	0.152	0.199	774
Urban population share	0.080	0.172	588	0.214	0.232	774
Marriage share, all women	0.692	0.092	587	0.727	0.039	774
Marriage share, women aged 16 to 25	0.495	0.150	581	0.516	0.071	774

Table 2.1: Summary statistics

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Ln(population)	Female pop share	Share of the LF employed in			
			Oil	Manuf.	Agriculture	Services
Discovered Oil Field	0.351***	-0.003	0.052***	0.008**	-0.083***	0.013**
	(0.068)	(0.003)	(0.007)	(0.004)	(0.017)	(0.006)
Observations	3,594	3594	3,577	3,577	3,577	3,577
Mean Dep Var		0.492	0.010	0.078	0.494	0.167

Number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 2.2: Broad economic effects of oil discoveries

VARIABLES	(1)	(2)	(3)	(4)
	Share of the Labour Force employed in			
	Oil Mining	Manufacturing	Agriculture	Services
Discovered Oil Field	0.065***	0.012***	-0.088***	0.004
	(0.008)	(0.004)	(0.016)	(0.004)
Mean Dep Var	0.012	0.087	0.536	0.087

Number of observations 3,577, number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.3: Sectoral shifts by gender: Men

VARIABLES	(1)	(2)	(3)	(4)
	Share of the Labour Force employed in			
	Oil Mining	Manufacturing	Agriculture	Services
Discovered Oil Field	0.001***	-0.009*	-0.063**	0.052*
	(0.000)	(0.005)	(0.030)	(0.028)
Mean Dep Var	0.001	0.045	0.329	0.479

Number of observations 3,475, number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.4: Sectoral shifts by gender: Women

VARIABLES	(1) LF Part. Rate Female	(2) Empl. Rate Female	(3) LF Part. Rate Male	(4) Empl. Rate Male
Discovered Oil Field	-0.003	-0.004	0.008	-0.005
	(0.010)	(0.008)	(0.007)	(0.004)
Observations	3,590	2,171	3,594	2,231
Mean Dep Var	0.246	0.938	0.907	0.941
Years	1900-1940	1910, 30, 40	1900-1940	1910, 30, 40

Number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.5: The effect of oil on labour force participation rates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES	Average ln (wage income)		Average wage inc. percentile rank		Average occupational earnings score			
Panel A	All sectors							
Discovered Oil Field	0.263***	0.047	7.427***	1.045	5.188**	-1.392*	4.246***	0.531
	(0.080)	(0.071)	(2.289)	(1.383)	(2.119)	(0.734)	(0.918)	(0.607)
Mean Dep Var	6.268	5.730	41.159	36.199	38.018	26.979	33.019	21.197
Panel B	Excluding individuals working in the oil mining sector							
Discovered Oil Field	0.224**	0.044	6.179**	0.978	4.174*	-1.447*	2.958***	0.497
	(0.087)	(0.071)	(2.452)	(1.384)	(2.227)	(0.738)	(0.771)	(0.607)
Mean Dep Var	6.251	5.728	40.729	36.162	37.632	26.956	32.683	21.182
Sample	Male LF	Female LF	Male LF	Female LF	Male LF	Female LF	Male LF	Female LF
Observations	758	758	758	758	758	758	3,577	3,478
Clusters	758	758	758	758	758	758	774	774
Year	1940	1940	1940	1940	1940	1940	1900-1940	1900-1940

Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 2.6: Oil discoveries and wages/earnings

VARIABLES	(1) Female LF Participation Rate	(2)
Discovered Oil Field	0.013	-0.008
	(0.013)	(0.009)
Sample	Single Women	Married women
Observations	3,543	3,590
Mean Dep Var	0.483	0.139

Number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)
*** p<0.01, ** p<0.05, * p<0.1

Table 2.7: The effects of oil on female LFP by marriage status

VARIABLES	(1)	(2) Share of women married	(3)
Discovered Oil Field	0.002	0.015**	0.015*
	(0.004)	(0.007)	(0.008)
Sample	All Women	Women $\leq$ 30	Women $\leq$ 25
Observations	3,590	3,582	3,573
Mean Dep Var	0.690	0.594	0.501

Number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)
*** p<0.01, ** p<0.05, * p<0.1

Table 2.8: Oil discoveries and marriage rates

VARIABLES	(1) LF Part. Rate	(2) Share of Female LF in services	(3) Share of Female LF in agriculture	(4) Average occ. earnings score	(5) LF Part. Rate	(6) Share of Female LF in services	(7) Share of Female LF in agriculture	(8) Average. occ. earnings score
Disc Oil Field	-0.000	0.019	-0.038	0.328	0.004	0.085**	-0.088**	1.217***
	(0.007)	(0.019)	(0.026)	(0.884)	(0.021)	(0.038)	(0.036)	(0.439)
Sample	White	White	White	White	Black	Black	Black	Black
Observations	3,590	3,452	3,452	3,455	2,954	2,686	2,686	2,709
Clusters	774	774	774	774	746	711	711	718
Mean Dep Var	0.184	0.483	0.188	31.899	0.459	0.475	0.495	7.318

Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.9: The effects of oil discoveries on women by race

VARIABLES	(1) Female pop share	(2) LF Part Rate	(3) Employment Rate	(4) Average occ. earnings score
Panel A	People born in state of residence			
Disc Oil Field	-0.000	0.010	-0.005	3.191***
	(0.002)	(0.008)	(0.004)	(0.642)
Observations	3,572	3,557	2,206	3,539
Panel B	People born outside of state of residence			
Disc Oil Field	-0.005	0.005	-0.007	5.913***
	(0.004)	(0.008)	(0.007)	(1.725)
Observations	3,586	3,583	2,221	3,566
Sample	All	Men	Male LF	Male LF

VARIABLES	(5) Share of Male LF employed in Oil Mining	(6) Manufacturing	(7) Services	(8) Agriculture
Panel A	People born in state of residence			
Disc Oil Field	0.038***	0.007	0.004	-0.061***
	(0.005)	(0.005)	(0.003)	(0.013)
Observations	3,538	3,538	3,538	3,538
Panel B	People born outside of state of residence			
Disc Oil Field	0.105***	0.015**	0.005	-0.126***
	(0.014)	(0.006)	(0.006)	(0.028)
Observations	3,566	3,566	3,566	3,566

Number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.10: Comparing state natives and state migrants, part 1

VARIABLES	(1) LF Part Rate	(2) Share Married	(3) Employment Rate	(4) Average occ. earnings score
Panel A		People born in state of residence		
Disc Oil Field	-0.006	0.015*	-0.007	0.654
	(0.012)	(0.008)	(0.007)	(0.622)
Observations	3,557	3,533	2,114	3,373
Panel B		People born outside of state of residence		
Disc Oil Field	0.002	0.009	0.001	0.268
	(0.090)	(0.016)	(0.016)	(1.027)
Observations	3,573	3,456	2,094	3,374
Sample	All Women	Women $\leq$ 25	Female LF	Female LF

VARIABLES	(5) Oil Mining	(6) Share of Female LF employed in Manufacturing	(7) Services	(8) Agriculture
Panel A		People born in state of residence		
Disc Oil Field	0.001***	-0.004	0.053	-0.066*
	(0.000)	(0.005)	(0.032)	(0.032)
Observations	3,370	3,370	3,370	3,370
Panel B		People born outside of state of residence		
Disc Oil Field	0.002***	-0.022***	0.060*	-0.068*
	(0.001)	(0.006)	(0.029)	(0.036)
Observations	3,365	3,365	3,365	3,365

Number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.11: Comparing state natives and state migrants, part 2

2.7 Appendix A: Theoretical results

An equilibrium of the economy described in section 2.2 is characterized by the system of equations (2.8) to (2.17):

$$L_M^o + L_M^c = L_M \tag{2.8}$$

$$L_W^c + L_W^s = L_W^{LF} \leq L_W \tag{2.9}$$

$$N_c = \bar{N}_c \tag{2.10}$$

$$\frac{\beta}{\alpha}\frac{L_M^c}{L_W^c} = \frac{w_W}{w_M} \tag{2.11}$$

$$\frac{\beta N_c}{(1-\alpha-\beta)L_W^c} = \frac{w_W}{r_c} \tag{2.12}$$

$$p_c = \frac{1}{A_c}\left(\frac{w_M}{\alpha}\right)^{\alpha}\left(\frac{w_W}{\beta}\right)^{\beta}\left(\frac{r_c}{1-\alpha-\beta}\right)^{1-\alpha-\beta} \tag{2.13}$$

$$p_o \leq \frac{w_M}{A_o} \tag{2.14}$$

$$p_s = \frac{w_W}{A_s} \tag{2.15}$$

$$(1-\gamma)\left[w_W L_W^{LF} + w_M L_M + r_c N_c\right] = w_W L_W^s \tag{2.16}$$

$$\gamma^{\gamma}(1-\gamma)^{1-\gamma}\frac{w_W}{p_c^{\gamma} p_s^{1-\gamma}} \geq \bar{u} \tag{2.17}$$

where beyond previously established notation w_i are the wages of workers of gender $i \in \{W, M\}$, r_c is the rental rate of specific capital while L_W^{LF} denotes the number of women in the labour force (which due to the presence of the outside option can be fewer than the total population of women).

Equations (2.8) to (2.10) are market clearing conditions in the various factor markets. Equations (2.11) and (2.12) are optimality conditions from the maximization problem of consumption good producers. Equations (2.13) to (2.15) are zero profit conditions for producers of the consumption good, oil and services respectively. Equation (2.16) is an optimality condition from the consumer's utility maximization problem. Equation (2.17) represents the participation condition of women into the labour force.

2.7.1 Low Oil Productivity

It can be shown that for A_o small enough such that (2.14) does not bind, the condition in (2.5) guarantees that (2.17) binds[8]. Therefore we begin our analysis with very small A_o where we know that (2.14) does not bind and (2.17) binds.

[8] This can be shown by contradiction. Assume there is an equilibrium where no men work in oil, $\bar{u}$ obeys (2.5) and (2.17) does not bind. This leads to a contradiction.

In such an equilibrium we have that:

$$L_M^{c*} = L_M \tag{2.18}$$

$$L_M^{o*} = 0 \tag{2.19}$$

$$w_W^* = \frac{p_c \bar{u}^{\frac{1}{\gamma}}}{\gamma[(1-\gamma)A_s]^{\frac{1-\gamma}{\gamma}}} \tag{2.20}$$

where (2.20) is obtained by plugging (2.15) in (2.17) and rearranging. With these results in place the group of equations (2.11) to (2.13) represent a system of three equations with three remaining unknowns. We first divide (2.11) by (2.12) and rearrange to obtain:

$$r_c^* = \frac{(1-\alpha-\beta)L_M}{\alpha \bar{N}_c} w_M \tag{2.21}$$

We plug (2.20) and (2.21) into (2.13) and then solve for w_M yielding:

$$w_M^* = \alpha (A_c)^{\frac{1}{1-\beta}} p_c \left[\frac{\beta\gamma[(1-\gamma)A_s]^{\frac{1-\gamma}{\gamma}}}{\bar{u}^{\frac{1}{\gamma}}} \right]^{\frac{\beta}{1-\beta}} \left(\frac{\bar{N}_c}{L_M} \right)^{\frac{1-\alpha-\beta}{1-\beta}} \tag{2.22}$$

This configuration is only an equilibrium if we can confirm that (2.14) does not bind. Plugging (2.22) into (2.14) yields us with the condition

$$A_o < \underbrace{\alpha (A_c)^{\frac{1}{1-\beta}} \frac{p_c}{p_o} \left[\frac{\beta\gamma[(1-\gamma)A_s]^{\frac{1-\gamma}{\gamma}}}{\bar{u}^{\frac{1}{\gamma}}} \right]^{\frac{\beta}{1-\beta}} \left(\frac{\bar{N}_c}{L_M} \right)^{\frac{1-\alpha-\beta}{1-\beta}}}_{A_1^*} \tag{2.23}$$

Thus if condition (2.23) is met the configuration above is an equilibrium featuring no oil production and some women not working. Moreover, for any $0 \leq A_o < A_1^*$ increases in A_o have no effect on any of the equilibrium quantities or prices (we have the exact same equilibrium in this range).

2.7.2 Oil Production Begins

For

$$A_o > \alpha (A_c)^{\frac{1}{1-\beta}} \frac{p_c}{p_o} \left[\frac{\beta\gamma[(1-\gamma)A_s]^{\frac{1-\gamma}{\gamma}}}{\bar{u}^{\frac{1}{\gamma}}} \right]^{\frac{\beta}{1-\beta}} \left(\frac{\bar{N}_c}{L_M} \right)^{\frac{1-\alpha-\beta}{1-\beta}}$$

Oil production begins. However, given the continuity of our set-up, we expect that in the neighbourhood of the threshold it should still be the case that some women don't work, whichever the effect of increased A_o is on female labour force participation. We thus search for equilibria where $A_o > A_1^*$ and (2.17) binds.

In such an equilibrium we have

$$w_M^* = A_o p_o \tag{2.24}$$
$$w_W^* = \frac{p_c \bar{u}^{\frac{1}{\gamma}}}{\gamma[(1-\gamma)A_s]^{\frac{1-\gamma}{\gamma}}}$$

Plugging these last two equations into (2.13) and solving for r_c yields:

$$r_c^* = (1-\alpha-\beta)(A_c p_c)^{\frac{1}{1-\alpha-\beta}} \left(\frac{\alpha}{A_o p_o}\right)^{\frac{\alpha}{1-\alpha-\beta}} \left[\frac{\beta\gamma[(1-\gamma)A_s]^{\frac{1-\gamma}{\gamma}}}{p_c \bar{u}^{\frac{1}{\gamma}}}\right]^{\frac{\beta}{1-\alpha-\beta}} \tag{2.25}$$

Plugging (2.20) and (2.25) into (2.12) and solving for L_W^c yields:

$$L_W^{c*} = (A_c p_c)^{\frac{1}{1-\alpha-\beta}} \left(\frac{\alpha}{A_o p_o}\right)^{\frac{\alpha}{1-\alpha-\beta}} \left[\frac{\beta\gamma[(1-\gamma)A_s]^{\frac{1-\gamma}{\gamma}}}{p_c \bar{u}^{\frac{1}{\gamma}}}\right]^{\frac{1-\alpha}{1-\alpha-\beta}} \tag{2.26}$$

Plugging (2.20), (2.24), (2.25) and (2.26) into (2.16), noting (2.9) and solving for L_W^{LF} yields:

$$L_W^{LF*} = \frac{w_W^* L_W^{c*} + (1-\gamma)A_o p_o L_M + (1-\gamma) r_c^* \bar{N}_c}{\gamma w_W^*} \tag{2.27}$$

Differentiating L_W^{LF*} with respect to A_o and signing the derivative yields:

$$\frac{\partial L_W^{LF*}}{\partial A_o} < 0 \text{ if}$$
$$A_o < \underbrace{\frac{\alpha(A_c p_c)^{\frac{1}{1-\beta}} \beta^{\frac{1-\alpha}{1-\beta}} [\beta + (1-\gamma)(1-\alpha-\beta)\bar{N}_c]^{\frac{1-\alpha-\beta}{1-\beta}}}{[\beta(1-\alpha-\beta)(1-\gamma)L_M]^{\frac{1-\alpha-\beta}{1-\beta}} p_o w_W^{*\frac{\beta}{1-\beta}}}}_{A_2^*} \tag{2.28}$$

Conversely, we have that

$$\frac{\partial L_W^{LF*}}{\partial A_o} > 0 \text{ if}$$
$$A_o > \underbrace{\frac{\alpha(A_c p_c)^{\frac{1}{1-\beta}} \beta^{\frac{1-\alpha}{1-\beta}} [\beta + (1-\gamma)(1-\alpha-\beta)\bar{N}_c]^{\frac{1-\alpha-\beta}{1-\beta}}}{[\beta(1-\alpha-\beta)(1-\gamma)L_M]^{\frac{1-\alpha-\beta}{1-\beta}} p_o w_W^{*\frac{\beta}{1-\beta}}}}_{A_2^*} \tag{2.29}$$

Two observations are in order at this stage. One is that (2.17) is guaranteed to bind between A_1^* and A_2^*. The imposition of restriction (2.5) guarantees that (2.17) binds at A_1^* and in the neighbourhood (via continuity). Moreover between A_1^* and A_2^* increases in A_o are associated with declines in female labour force participation, such that if (2.17) binds at A_1^* it will also bind at A_2^* and all the intermediate points between. The second observation is that as the size of the service sector gets arbitrarily small ($\gamma \to 1$) the magnitude of A_2^* goes to infinity, which shows that in the absence of the services sector, any increase in oil sector productivity beyond the threshold at which oil production begins brings about a decline in female labour force participation (the result of Corollary 1).

2.7.3 A Sufficiently Large Oil Boom Brings All Women into the Labour Force

All that is left to show is that there exists and $A_3^* > A_2^*$ such that beyond A_3^* all women are drawn into work. We will prove this by contradiction. Assume no such A_3^* exists. That means that for any $A_3^* > A_2^*$ we have that (2.17) binds. Because $A_3^* > A_2^* > A_1^*$ oil is produced in equilibrium, which means that equations (2.20) and (2.24) to (2.27) hold. For $A_3^* > A_2^*$ we know that female labour force participation is increasing. We then proceed to take the limit of the number of women in the labour force, given by (2.27) when A_3^* goes to infinity. We are allowed to do this because by assumption (2.27) holds for any A_3^*. But from (2.27) we have that:

$$\lim_{A_3^* \to \infty} L_W^{LF*} = \infty \tag{2.30}$$

But by assumption, the female population of the county (L_W) is finite. We have reached a contradiction which concludes the proof.

2.8 Appendix B: Further results and robustness checks

In this appendix, we present several additional findings and several checks of the robustness of our key findings. One first additional investigation concerns the exact nature of the service sector increase from which females are benefitting. We have shown that the share of females employed in the service sector increases. This could happen due to a variety of reasons: As manufacturing, refining and other capital-intensive sectors grow, demand for more advanced services such as banking, insurance, and business services like accounting and advertising might grow. On the other hand, income effects could lead to growth in the demand for personal services and entertainment. Of course, the types of jobs that would arise would be quite different in both scenarios. Our baseline classification includes all service sector industries and thus cannot inform us of which type of services experience particular growth. In table 2.12, we therefore disaggregate the service sector into three narrower categories: Personal services and Entertainment (this includes for example individuals working in private households, in hotels, in laundering, cleaning or dyeing services and in entertainment venues), Finance, Business, and Repair Services (e.g. banking, insurance, advertising, accounting, auto repair services), and professional services (e.g. medical, legal and educational services). Overall, we had found in section 2.4 that the share of women employed in services increases by 5.2 percentage points after oil discoveries. As our results in table 2.12 show, the lion's share of this increase - 4.5 percentage points or more than 85% - comes from the personal services and entertainment category, whereas the other two categories display only very small increases. This suggests that the main reason for the blossoming of service sector jobs are income effects and the resulting increased demand for personal and recreational services. Inter-industry demand linkages that would lead to growth in professional and business services seem to be, at least in terms of female employment, less relevant.

In addition, in this section we also implement a number of robustness checks that aim to provide further support for the notion that the effects we find are indeed caused by the discovery of major oil fields.

As a first check, we show that before oil discoveries, oil-rich counties were not systematically different from counties without major oil deposits. To do so, we run a series of cross-sectional regressions for 1900:

$$y_c = \beta OilRich_c + X_c'\gamma + s_s + u_c \tag{2.31}$$

where *OilRich* is 1 if a county sits on top of a large oil field an 0 otherwise.

As can be seen in table 2.13, across 10 different outcome variables, we do not find substantial differences between the two groups of counties. Oil- and non-oil counties have essentially the same male and female labour force participation rate, the same distribution of the labour force over our 4 sectors of interest, and very similarly-sized populations. The only statistically significant difference arises in occupational earnings scores, where it seems that oil-rich counties on average have worse-paying occupations. However, the difference is only about one and a half percentile ranks.

Overall, the two groups of counties thus look very balanced. These findings are reassuring as they imply that the oil wealth "treatment"can be seen as nearly randomly assigned and thus the setting we study constitutes a valid natural experiment and the common trend assumption needed for the validity of our difference-in-differences design is likely to hold.

To assess this key assumption more rigorously, we augment our baseline panel specification by replacing our variable of interest with a full set of leads and lags of the date of oil discovery relevant for each oil-rich county. The specification we estimate is thus of the form:

$$y_{ct} = \alpha_c + \tau_t + \sum_{j \in \{-30,-20,-10,10,20,30\}} \beta_j DiscoveredOilField_{c,t+j} + X_c' \gamma_t + u_{ct} \tag{2.32}$$

where the set of dummies $DiscoveredOilField_{c,t+j}$ code for whether an oil field is to be discovered $20-30$ years from period t, $10-20$ years from period t, $10-0$ years from period t or was discovered $0-10$ years prior to period t, $10-20$ years prior to period t or more than 20 years prior to period t, with the omitted reference category being represented by discoveries that occur more than 30 years after the reference period t. All the remaining variables and controls retain their meanings from specification 2.6 in section 2.3. We show graphical results from this specification for population and the labour force shares in our four major sectors in figures 2.6 to 2.10. Results for our remaining variables of interest are shown in tables 2.14 to 2.17.

Reassuringly, our results indicate that there is no evidence that oil-rich counties display systematically different characteristics before the discovery of oil, with virtually all the leading dummies (that indicate oil discoveries in the future) having no significant impact on any of our outcome variables of interest.

The findings related to the lagging dummies (that indicate time elapsed since discovery) largely confirm the results from our main specifications in the previous section: While there is no effect on either male or female labour force participation rates, sectoral reallocation is active. An interesting, but also intuitive further result is that the reallocation of male labour into the oil industry takes place immediately, while the increase in manufacturing needs a bit longer to materialize. In the case of the importance of services for female employment, it seems that oil counties originally display lover values and then catch up after oil discoveries take place. Yet another pattern is displayed by population, which increases sharply after oil discoveries and then continues to grow. All in all, our results from the lead and lag analysis are consistent with the results emerging from our main specifications. They support the view that the systematic differences we observe between oil-rich and baseline counties do indeed appear after oil discoveries and can be attributed to oil abundance in a causal way. A more parametric robustness check for our assumption of common trends is to include county-specific linear time trends, which is done in 2.18 and 2.19. Again, our basic conclusions remain unchanged. The one major exception to this is that the reallocation of men into the manufacturing sector becomes weaker and statistically insignificant, which can be explained by the fact that oil discoveries induce a post-treatment trend which is partly picked up by the linear trend.

Another related robustness check that also probes our identifying assumption involves performing our analysis on a limited sample from which all counties without oil deposits have been dropped. This is a further guard against the possibility that oil and non-oil counties might have been on different trends even in the absence of oil discoveries. It is a very demanding empirical exercise, as it involves deriving identification only from the time variation in oil field discoveries, as well as dropping more than three quarters of our observations. We do it for our key variables of interest- employment and labour force participation rates as well as the sectoral composition of the labour force and occupational earnings scores by gender, the female population share and the share of women aged 25 and below who are married. Results are shown in tables 2.20 and 2.21. Overall, the findings of this robustness test provide a convincing validation of our previous results: having a discovered oil field is found to be associated with a reallocation of men from agriculture into oil and (to a lesser extent) manufacturing, while women increasingly work in service sector jobs. Especially men experience an increase in occupational earnings scores, which can help to explain the increase in young women's marriage propensity. Overall, both male and female labour force participation rates stay unchanged. Generally, most results from our baseline specification survive unqualified, and the point estimates of the coefficients obtained by estimating over this limited sample are very similar to those obtained in the previous section. The two exceptions are the sex ratio, where we here find a small decrease in the female population share of less than one percentage point, and the female occupational earnings scores, which in this sample display a positive evolution. However, the point estimate for women is substantially below the one of males, indicating that our conclusion of an increased gender pay gap remains the same.

Another concern with our findings is the possibility of spatial spillovers. While workforces should be less mobile in our period of analysis than nowadays, there is still a possibility that workers commute short distances and that thus counties close to oil counties also get partially treated. In addition, there might be smaller oil fields that are not in our dataset close to larger ones, such that again we would observe some spatial spillovers, a problem already noted by Michaels (2011). To address this concern of spatial spillovers, we repeat our basic specification from section 2.4, but drop all county x year cells pertaining to counties which border a county with a discovered oil field, but have not (yet) discovered an oil field of their own in that given year. The results of this exercise are shown in tables 2.22 and 2.23. Again, our results from before are confirmed; if anything, we find slightly larger point estimates, consistent with minor spatial spillovers. Our basic conclusions remain the same as before: Oil does not change female labour force participation, but relatively more women move from agriculture into services, and younger women are more likely to get married. Men, on the other hand, experience substantial increases in earnings that are due to sectoral reallocation from agriculture to oil mining and manufacturing.

So far, all of our robustness checks have dealt with our panel difference-in-differences specification. In table 2.24, we perform two robustness checks for the 1940 cross-sectional specification which we had to employ for our wage regressions. The first (columns 1 and 3) involves adding an additional dummy for "having an undiscovered oil field"to our specification. With this we are testing whether oil-rich counties in general appear to be different in the 1940 cross-section. Naturally, oil that is not yet discovered should not affect economic outcomes, so if we observe any sizeable coefficients on this dummy, we would have to be worried. Fortunately for the credibility of our cross-sectional results, this is not the case: including a dummy for an undiscovered oil field barely changes the point estimate on our variable of interest, and the estimated coefficient for undiscovered oil fields is usually insignificant, albeit in the case of women (column 3) not very precisely estimated. In columns 2 and 4, we perform a placebo test: We drop all counties that had discovered oil fields in 1940 and compare the remaining oil-rich counties (whose oil wealth had not yet been discovered) to counties without oil. Again, if our identifying assumptions are valid, we would expect to see no effect of the undiscovered oil field dummy, and again this is the case, which provides further reassurance for our cross-sectional estimates.

Finally, one potential worry is that the timing of earlier oil discoveries might be spuriously correlated with the arrival of the boll weevil. This beetle decimated the South's cotton harvests in the late 19th and early 20th century, which has been shown to have had severe effects on local economies (Lange, Olmstead and Rhode, 2009), including a reduction of female labour force participation (Ager, Brueckner and Herz forthcoming). If the arrival or prevalence of the weevil were correlated with oil wealth or oil discoveries, this could therefore be cause for concern with our results. To examine this potential concern, we follow Pei, Pischke and Schwandt (2017), who suggest that instead of including an indicator for weevil presence as control in each of our regressions, an easier and more powerful test can be obtained by putting the weevil's presence on the left hand side of or our regression equation. We do so in table 2.25, using county-level data on the weevil's first arrival from Lange, Olmstead and Rhode (2009). We focus on 1900-1930, since the weevil plague subsided after that, and employ three different specifications. In column 1, we use our simple cross-sectional specification to analyse whether oil wealth is correlated with whether the weevil ever reached a given county at all. As can be seen, this is not the case. In columns 2 and 3, we instead use our standard difference-in-differences framework to analyse whether oil discoveries are associated with the arrival of the weevil. In column 2, we code the left hand side variable as 1 if the weevil has already reached the county and 0 if not, and similarly for oil discoveries on the right hand side. In column 3 instead, both our left hand and main right hand side variable are 1 if there has been a weevil arrival (or oil discovery) within the last 10 years. In both cases, we do not find a statistically significant effect of oil discoveries on weevil presence. The economic shocks due to the weevil plague are therefore no confounder for our analysis.

	(1)	(2)	(3)
VARIABLES	Share of the female LF employed in		
	Personal serv. and entertainment	Finance, Business, and repair serv.	Professional serv.
Discovered Oil Field	0.045	0.004	0.003
	(0.027)	(0.003)	(0.006)
Mean Dep Var	0.331	0.024	0.125

Number of observations 3,475, number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.12: Dissecting the female service sector increase

VARIABLES	(1) Labour Force Participation Rate Female	(2) Male	(3) Occ. earnings score Female	(4) Male	(5) Urban Pop Share
Oil Rich	0.001	0.004	-1.674*	-1.569	-0.071
	(0.016)	(0.005)	(0.924)	(2.444)	(0.067)
Observations	587	588	549	588	588
Mean Dep Var	0.213	0.930	15.121	27.509	0.233

VARIABLES	(6) Share of the Labour Force employed in Oil mining	(7) Manufacturing	(8) Services	(9) Agriculture	(10) ln(pop)
Oil Rich	-0.000	0.004	-0.008	0.026	-0.023
	(0.000)	(0.008)	(0.021)	(0.055)	(0.102)
Observations	588	588	588	588	588
Mean Dep Var	0.000	0.051	0.128	0.63	

Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.13: Balancing checks in 1900

VARIABLES	(1) Female pop share	(2) LF Part Rate	(3) Employment Rate	(4) Average occ. earnings score
-30 to -21	0.000	0.009	0.004	0.360
	(0.003)	(0.006)	(0.008)	(0.760)
-20 to -11	-0.000	0.012*	0.018	-0.759
	(0.004)	(0.007)	(0.011)	(0.809)
-10 to -1	0.004	0.017*	0.031***	-1.158
	(0.004)	(0.009)	(0.011)	(0.950)
0 to 10	-0.005	0.020*	0.018	3.055***
	(0.005)	(0.010)	(0.011)	(1.058)
11 to 20	0.005	0.026**	0.014	4.198***
	(0.006)	(0.013)	(0.012)	(1.507)
20+	0.008	0.005	0.019	3.717**
	(0.008)	(0.034)	(0.013)	(1.512)
Sample	All	Men	Male LF	Male LF
Observations	3,594	3,594	2,231	3,577

Number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.14: Leads and lags analysis, part 1

VARIABLES	(1)	(2)	(3)	(4)
	Share of Male LF employed in			
	Oil Mining	Manufacturing	Services	Agriculture
-30 to -21	0.003	0.007	0.006	0.006
	(0.003)	(0.006)	(0.005)	(0.013)
-20 to -11	-0.002	-0.004	-0.001	0.026*
	(0.004)	(0.010)	(0.004)	(0.015)
-10 to -1	-0.006	0.004	-0.002	0.039**
	(0.005)	(0.012)	(0.005)	(0.019)
0 to 10	0.061***	0.008	0.001	-0.049**
	(0.010)	(0.011)	(0.005)	(0.021)
11 to 20	0.064***	0.021*	0.008	-0.077***
	(0.090)	(0.012)	(0.006)	(0.023)
20+	0.043***	0.033**	0.012	-0.069**
	(0.090)	(0.014)	(0.090)	(0.028)

Number of observations 3,577, number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)
*** p<0.01, ** p<0.05, * p<0.1

Table 2.15: Leads and lags analysis, part 2

VARIABLES	(1) LF Part Rate	(2) Share Married	(3) Employment Rate	(4) Average occ.
-30 to -21	-0.016	0.006	0.006	0.361
	(0.018)	(0.011)	(0.010)	(0.664)
-20 to -11	-0.008	0.011	-0.010	0.448
	(0.025)	(0.013)	(0.015)	(0.730)
-10 to -1	-0.005	0.008	0.015	0.281
	(0.025)	(0.015)	(0.015)	(0.868)
0 to 10	-0.015	0.026*	0.000	1.073
	(0.026)	(0.015)	(0.016)	(0.843)
11 to 20	-0.005	0.019	-0.008	0.745
	(0.028)	(0.017)	(0.016)	(1.058)
20+	0.001	0.024	0.010	-0.295
	(0.033)	(0.024)	(0.018)	(1.147)
Sample	All Women	Women $\leq$ 25	Female LF	Female LF
Observations	3,590	3,573	2,171	3,478

Number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)
*** p<0.01, ** p<0.05, * p<0.1

Table 2.16: Leads and lags analysis, part 3

	(1)	(2)	(3)	(4)
VARIABLES		Share of Female LF employed in		
	Oil Mining	Manufacturing	Services	Agriculture
-30 to -21	-0.000	0.001	-0.029	0.023
	(0.000)	(0.006)	(0.055)	(0.049)
-20 to -11	-0.000	-0.006	-0.047	0.048
	(0.000)	(0.007)	(0.079)	(0.072)
-10 to -1	-0.000	0.001	-0.043	0.035
	(0.000)	(0.008)	(0.072)	(0.065)
0 to 10	0.001***	-0.008	-0.004	-0.015
	(0.000)	(0.008)	(0.078)	(0.071)
11 to 20	0.001	-0.013	0.032	-0.045
	(0.001)	(0.010)	(0.085)	(0.079)
20+	0.004***	-0.019	-0.006	0.001
	(0.001)	(0.013)	(0.101)	(0.093)

Number of observations 3,475, number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.17: Leads and lags analysis, part 4

	(1)	(2)	(3)	(4)
VARIABLES	Female pop share	LF Part Rate	Employment Rate	Average occ. earnings score
Disc Oil Field	-0.013***	0.014	-0.007	5.048***
	(0.004)	(0.011)	(0.010)	(0.953)
Sample	All	Men	Male LF	Male LF
Clusters	774	774	774	774
Observations	3,594	3,594	2,231	3,577
	(5)	(6)	(7)	(8)
VARIABLES	Share of Male LF employed in			
	Oil Mining	Manufacturing	Services	Agriculture
Disc Oil Field	0.074***	0.001	0.002	-0.096***
	(0.015)	(0.006)	(0.005)	(0.021)

Number of observations 3,577, number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.18: Results when including county-specific linear time trends, part 1

VARIABLES	(1) LF Part Rate	(2) Share Married	(3) Employment Rate	(4) Average occ. earnings score
Disc Oil Field	-0.018	0.029*	0.002	1.265
	(0.012)	(0.016)	(0.016)	(0.908)
Sample	All Women	Women ≤ 25	Female LF	Female LF
Observations	3,590	3,573	2,171	3,478
Clusters	774	774	774	774

VARIABLES	(5) Oil Mining	(6) Share of Female LF employed in Manufacturing	(7) Services	(8) Agriculture
Disc Oil Field	0.001	0.007	0.043*	-0.072**
	(0.000)	(0.005)	(0.023)	(0.027)
Observations	3,475	3,475	3,475	3,475
Clusters	774	774	774	774

Number of observations 3,475, number of clusters 774. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.19: Results when including county-specific linear time trends, part 2

VARIABLES	(1) Female pop share	(2) LF Part Rate	(3) Employment Rate	(4) Average occ. earnings score
Disc Oil Field	-0.007**	-0.001	-0.016**	3.523***
	(0.003)	(0.007)	(0.007)	(0.941)
Sample	All	Men	Male LF	Male LF
Observations	771	771	487	766
Clusters	171	171	171	171

VARIABLES	(5) Oil Mining	(6) Share of Male LF employed in Manufacturing	(7) Services	(8) Agriculture
Disc Oil Field	0.063***	0.005	0.003	-0.074***
	(0.010)	(0.005)	(0.004)	(0.017)

Number of observations 766, number of clusters 171. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.20: Results when dropping all counties without an oil deposit, part 1

VARIABLES	(1) LF Part Rate	(2) Share Married	(3) Employment Rate	(4) Average occ. earnings score
Disc Oil Field	-0.005	0.014	-0.009	1.148*
	(0.011)	(0.011)	(0.009)	(0.638)
Sample	All Women	Women ≤ 25	Female LF	Female LF
Observations	770	765	464	733
Clusters	171	171	171	171

VARIABLES	(5) Oil Mining	(6) Share of Female LF employed in Manufacturing	(7) Services	(8) Agriculture
Disc Oil Field	0.000	-0.005	0.036	-0.050
	(0.001)	(0.005)	(0.028)	(0.031)

Number of observations 733, number of clusters 171. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)
*** p<0.01, ** p<0.05, * p<0.1

Table 2.21: Results when dropping all counties without an oil deposit, part 2

VARIABLES	(1) Female pop share	(2) LF Part Rate	(3) Employment Rate	(4) Average occ. earnings score
Disc Oil Field	-0.006	0.013*	-0.001	4.277***
	(0.003)	(0.007)	(0.005)	(1.042)
Sample	All	Men	Male LF	Male LF
Observations	3,027	3,027	1,775	3,015
Clusters	741	741	722	741

VARIABLES	(5) Oil Mining	(6) Share of Male LF employed in Manufacturing	(7) Services	(8) Agriculture
Disc Oil Field	0.071***	0.010**	0.002	-0.086***
	(0.010)	(0.004)	(0.004)	(0.018)

Number of observations 3,015, number of clusters 741. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)
*** p<0.01, ** p<0.05, * p<0.1

Table 2.22: Results when excluding neighbouring counties without discovered oil wealth, part 1

VARIABLES	(1) LF Part Rate	(2) Share Married	(3) Employment Rate	(4) Average occ. earnings score
Disc Oil Field	-0.006	0.022**	0.004	0.424
	(0.012)	(0.010)	(0.011)	(0.757)
Sample	All Women	Women ≤ 25	Female LF	Female LF
Observations	3,023	3,007	1,717	2,918
Clusters	741	740	710	738

VARIABLES	(5) Oil Mining	(6) Share of Female LF employed in Manufacturing	(7) Services	(8) Agriculture
Disc Oil Field	0.001***	-0.010*	0.052*	-0.062**
	(0.000)	(0.006)	(0.027)	(0.028)

Number of observations 2,915, number of clusters 738. Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.23: Results when excluding neighbouring counties without discovered oil wealth, part 2

VARIABLES	(1)	(2) Average ln(annual wage income)	(3)	(4)
Discovered Oil Field	0.258***		0.038	
	(0.083)		(0.072)	
Not Yet Discovered	-0.042	-0.034	-0.072	-0.054
	(0.102)	(0.104)	(0.090)	(0.0903)
Counties	All	Omit disc oil fields	All	Omit disc oil fields
Sample	Male LF	Male LF	Female LF	Female LF
Observations	758	643	758	643

Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.24: Robustness checks 1940 cross section

VARIABLES	(1) Weevil ever present	(2) Weevil already present	(3) Weevil arrived within the last 10 yrs
Oil rich	-0.029		
	(0.022)		
Disc Oil Field		-0.039	
		(0.036)	
Discovery within the last 10 yrs			0.055
			(0.043)
Sample	Cross-section	Panel 1900-1930	Panel 1900-1930
Observations	758	2,820	2,820
Clusters	758	773	773

Standard errors, clustered at the county level, in parentheses. Standard errors and p-values are based on bias adjusted standard errors and effective degrees of freedom corrections as in Young (2016)

*** p<0.01, ** p<0.05, * p<0.1

Table 2.25: Robustness with respect to the boll weevil plague

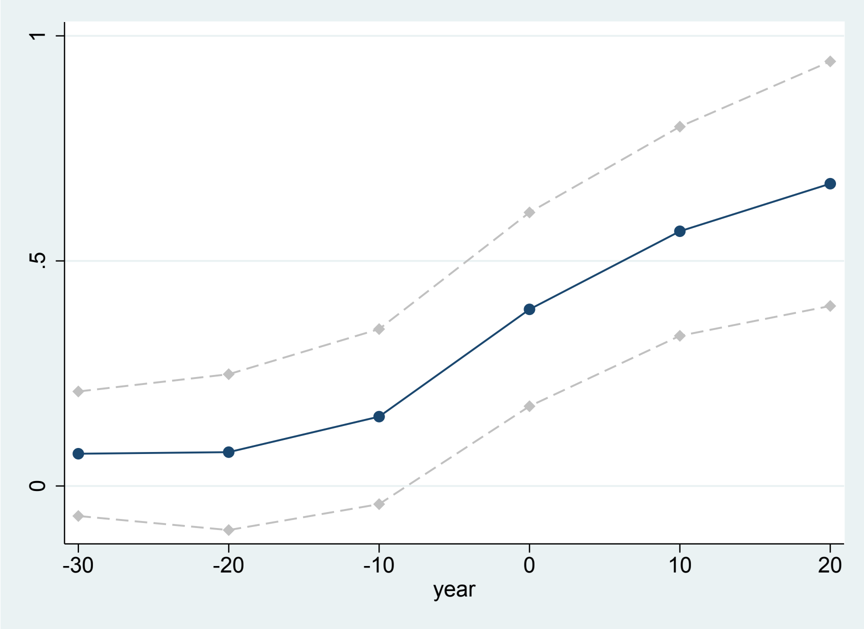

Figure 2.6: Leads and lags for log population

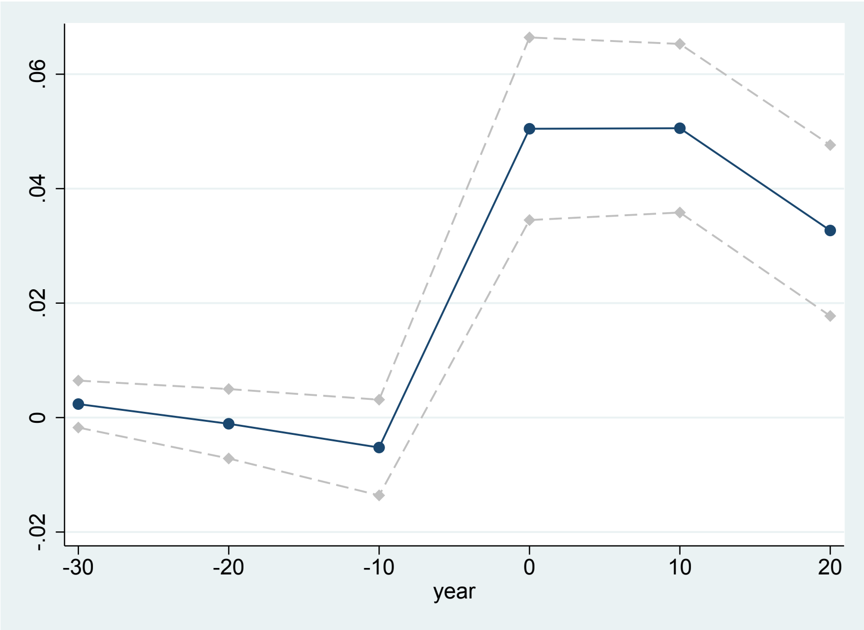

Figure 2.7: Leads and lags for the LF share employed in oil mining

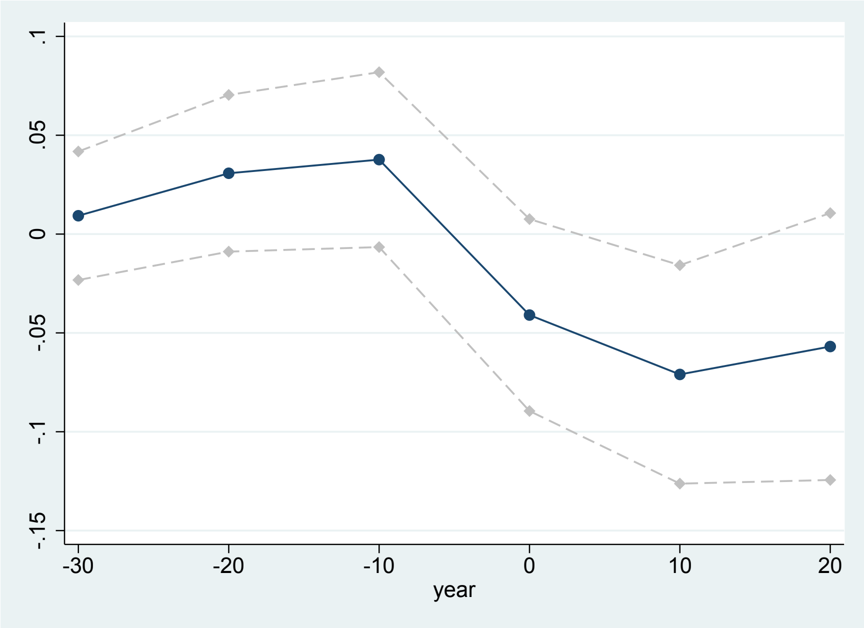

Figure 2.8: Leads and lags for the LF share employed in agriculture

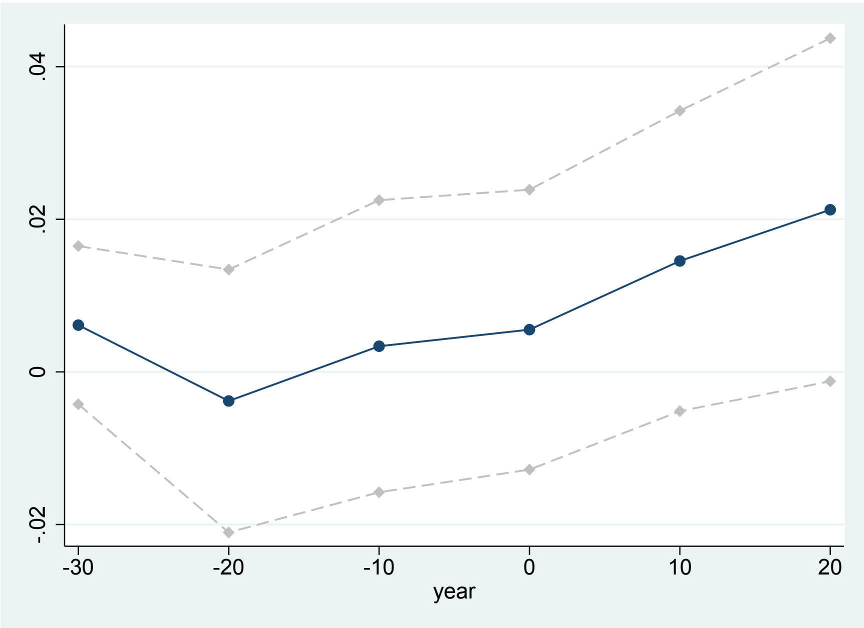

Figure 2.9: Leads and lags for the LF share employed in manufacturing

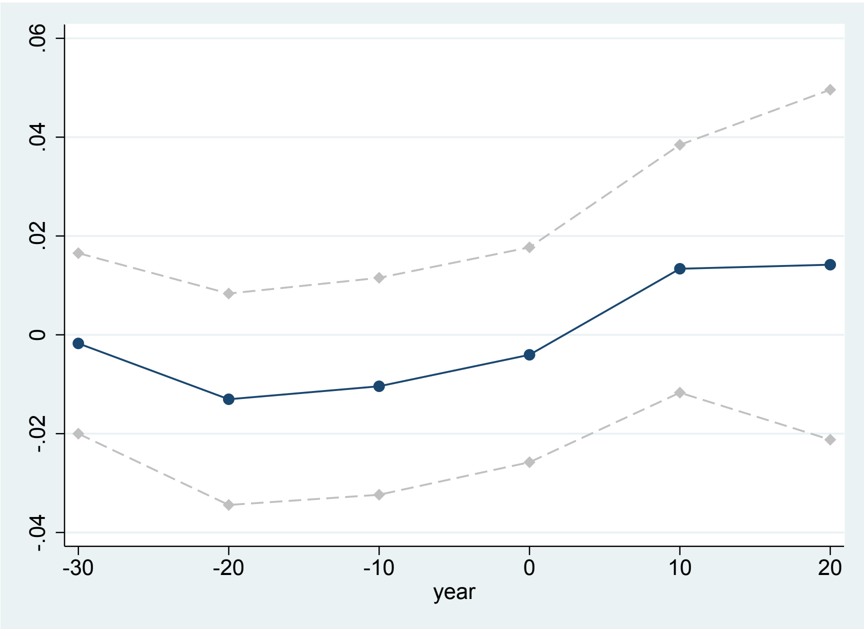

Figure 2.10: Leads and lags for the LF share employed in services

Chapter 3

Voting Behaviour and Public Employment in Nazi Germany[1]

3.1 Introduction

To what extent can governments use their economic means to favour their supporters or to punish their adversaries? While a large body of empirical literature has successfully established the economic value of political connections for firms, the evidence for voters or more aggregated units of observation is much more scarce. In this paper, I try to fill this gap by analysing whether cities benefit from having voted for the "right", i.e. the winning political party. In particular, the meteoric rise of the German National Socialist party in the 1930s, its seizure of power in 1933 and the subsequently enacted racial and political discrimination, programmes of large-scale public investment, and expansion of the armed forces create a quasi-experimental situation that allows to identify the causal effect of a city's vote share on subsequent public employment. Between 1928 and 1933, the Nazi party grew from being one of many small and unimportant radical parties to representing the largest fraction in the parliament, making Adolf Hitler chancellor in January 1933 and, together with a coalition partner, achieving a parliamentary majority in March of the same year. In the subsequent years, the Nazi government increased the armed forces and enacted several laws that prevented Jews and political opponents from holding public office. These policy measures together with the extremely rapid rise of the National Socialists create a unique possibility to estimate whether Hitler's government used public employment also in a way that favoured those cities that had helped him come

[1] I thank Steve Pischke, Guy Michaels, Michel Azulai, Sascha Becker, Florian Blum, Abel Brodeur, Robin Burgess, Marta De Philippis, Georg Graetz, Kilian Huber, Yu-Hsiang Lei, Andrei Potlogea, Robin Schädler, Pedro Souza, Daniel Sturm, Nico Voigtländer, several anonymous referees, and seminar participants at the LSE Labour Work in Progress Seminar for their helpful comments and suggestions. I also thank Michael Beaney from the LSE Language centre for proofreading an earlier version of this chapter.

to power.

On a firm-level, the value of political connections has been demonstrated convincingly by several papers. Fisman (2001) shows that rumours about the health of the Philippine dictator Suharto had a particularly strong influence on the share prices of firms that were politically connected to Suharto's regime. Similar positive effects of being politically connected have been found by Johnson and Mitton (2003) for Malaysia, Khwaja and Mian (2005) for Pakistan and Jayachandran (2006) for the United States of America. Other studies compare companies across countries: Faccio et al. (2006) show that around the globe, politically connected firms are more likely to be bailed out, while Faccio (2006) finds that political connections occur particularly often both in more corrupt and in more transparent countries. Of particular relevance for this paper is the study by Ferguson and Voth (2008), who show that firms that had (directly or through their executives) supported the German National Socialists prior to their seizure of power experienced particularly high stock market returns during the first two months of the Nazi regime: between January and March 1933, connected firms outperformed non-connected ones by between 5 and 8%.

The potential benefits of political connections for individual voters have been analysed less, particularly due to data restrictions: while political connections of firms' executives and firms' donations are often public, the average voter's political affiliations and convictions are mostly neither known to the government nor to the researcher, and hence cannot be analysed. One notable exception is the recent study by Hsieh et al. (2011), who document evidence that Venezuelan voters who had signed a petition calling for a referendum against Hugo Chavez were subsequently subject to drops in both earnings and employment. The peculiarities of this referendum, where signers had to sign not only with their name, but were also required to provide their address and birth date, allowed Hsieh et al. to identify the signers and to match them with data from the Venezuelan Household Survey. However, such detailed data on political affiliations are usually not available. One way out is to look at more aggregated units of observation such as cities, regions or electoral districts. Anderson and Tollison (1991), for example, present empirical evidence that US states with "influential" congressmen and senators (as measured by their tenure and their committee memberships) received more public funds during the New Deal era. Levitt and Snyder (1995) analyse the spending patterns of Federal programmes on a congressional district level and find that the Democratic majorities in the late 1970s have favoured districts with higher Democratic vote shares. Hodler and Raschky (2014) look at regions and show that in autocratic regimes, the birth regions of political leaders benefit more from foreign aid than others. Other evidence for regional favouritism in autocratic regimes is presented by Do et al. (forthcoming).

In this paper, I use cities as a "middle ground" between individual outcomes and larger units of aggregation. My paper adds to the existing literature by using a novel dataset to analyse whether cities with higher vote shares for the German National Socialists in 1933 experienced higher levels of public employment between then and 1939.

A closely related paper to mine is the contemporaneous study by Voigtländer and Voth (2016) who analyse the effect of public spending - in particular highway construction during the early years of the Nazi regime - on support for the NSDAP. They find that highway construction increased support for the NSDAP through signalling government "competence". Methodologically, by examining the opposing direction of causality, my paper complements the study by Voigtländer and Voth (2016). Their finding of a causal effect of public spending on NSDAP support also suggests a potential issue of reverse causality in a simple regression of public spending on NSDAP vote shares. This is one reason why I employ an instrumental variables strategy. In addition to this potential problem of reverse causality, several previous studies (most recently King et al. 2008) have highlighted the importance of the post-1929 economic crisis for the NSDAP's electoral results. While I control for city fixed effects and time-varying effects of several control variables, differential unobserved impacts of the economic crisis would likely lead to differences in public employment and could also be correlated with the 1933 Nazi vote share. In order to address both concerns, I employ a standard two-stage least squares estimation. As instrumental variable, I use the vote share of the "Economic Association" in 1912, a party alliance that tried to attract similar voters to the NSDAP later on.

I find that cities with higher NSDAP vote shares indeed had higher relative growth of public employment after 1933: for every additional percentage point in the vote share, the number of public employment positions increased by around 2.5 percent. When measured relative to the total population, a one standard-deviation increase in the 1933 vote share led to an increase in the share of public employment of 45% of a standard deviation in the year 1925. The findings are robust to including or excluding cities that underwent substantial changes in their population and territory during the period of observation, to using the 1930 or 1932 elections instead of the 1933 one as explanatory variable, and to employing different definitions of "public employment" as outcome variables. Taken all together, my findings indicate a significant positive effect of having voted for the National Socialists for cities, thus providing evidence that the Nazis did indeed use economic policy and public employment policies to reward more loyal cities or punish disloyal ones. In a broader context, this is further evidence that governments can have and use the ability to reward their voters or punish their adversaries, although some caveats apply to the representativeness of Nazi

Germany.
The structure of the remainder of this paper is as follows: section 3.2 presents the historical background of Hitler's rise to power and the National Socialist economic policy between 1933 and 1939, while section 3.3 discusses the data and identification strategy used in the analysis. The results and robustness checks are presented in section 3.4, and section 3.5 concludes.

3.2 Historical Background

In the early 1930's, Hitler's National Socialist German Workers' Party (*Nationalsozialistische Deutsche Arbeiterpartei*, NSDAP) experienced a meteoric rise from being one of many small parties in Weimar Germany to the strongest faction in the national parliament, the Reichstag. After an unsuccessful putsch in Bavaria in 1923, the party was banned and could only run for the national election in May 1924 by being the junior partner in an alliance with the German Völkisch Freedom Party (*Deutschvölkische Freiheitspartei*), another nationalist and anti-Semitic party in Weimar Germany. The two parties received 6.5% of the votes and also ran together in the December 1924 election, albeit under the new name of National Socialist Freedom Movement (*Nationalsozialistische Freiheitsbewegung*). This time, the alliance only achieved a vote share of 3%. Soon afterwards, the two parties separated and the NSDAP was re-formed in 1925. In 1928, it ran for the first time under this name at a national election, winning only 2.6% of the votes and 12 seats in the parliament. (Falter 1991, Chapter 2.1 and 2.2, Falter, Lindenberger, and Schumann 1986, Chapter 1.3) In the following years, the NSDAP changed its appearance and, benefiting from the deep recession that befell Germany in the wake of "Black Friday"(1929), grew stronger and stronger.[2] In September 1930, the National Socialists gained 18.3% of all votes, a share that they managed to even double two years later, when they came out of the July 1932 election with 37.4%, making them the strongest faction in the Reichstag. They and the Communists held more than half of all seats in the Reichstag, rendering it impossible to form a coalition of democratic parties with a parliamentary majority. As a result, the chancellors had to rely more and more on the authority and legislative powers of the president via so-called "emergency decrees". After the demise of 3 chancellors (Heinrich Brüning, Franz von Papen and Kurt von Schleicher) within half a year, the associates of President Hindenburg managed to convince him to appoint Hitler to head the government, which happened on January 30, 1933. (Kolb 2005, Part C) The new chancellor was still far from being a dictator. At the time of his appointment, Hitler, like his predecessors von Papen and von Schleicher, had

[2] For a recent review of key aspects of Germany's economy at the time of the crisis, see Ritschl (2013).

no parliamentary majority. However, Hindenburg soon dissolved the Reichstag, and in the elections that followed in March, the NSDAP won 43.9% of the votes. Together with its coalition partner, the national conservative German National People's Party (*Deutschnationale Volkspartei*, DNVP), the National Socialists now also had a majority in the parliament. Subsequently, the Enabling Act (*Ermächtigungsgesetz*) was passed, giving legislative powers to the executive branch of the government. In the following months, Hitler used these powers to put the German states under the rule of centrally appointed "Commissars"(a process commonly known as "coordination"or *Gleichschaltung*), to ban trade unions and to pressure all other parties until they dissolved. By July 1933, the NSDAP was the only remaining party in Germany and with the death of President Hindenburg in 1934, the last remaining non-Nazi source of power died, and Hitler and his party had now control over every aspect of government. (Kershaw 1999, ch. 10-12)

Not surprisingly, economic policy was an important item on the agenda of the newly-appointed chancellor. Already in May 1932, the NSDAP had demanded an "immediate economic programme"(*Wirtschaftliches Sofortprogramm*) in order to address the issue of unemployment. In particular, the party advocated increasing employment through large public investments that were, at least in parts, supposed to be financed through debt. (Barkai 1988, p.42) In the following years massive military and non-military investment ensued. Full employment was achieved by 1936, a success that the general public attributed largely to Hitler. (Abelshauser 1998) However, modern econometric analyses suggest Germany's economic upswing was not caused by Hitler's policies, but largely due to market forces and an international economic recovery that would have benefited the country also in the absence of a Nazi government. Public expenditure grew in the process of rearmament, but it often crowded out private demand. (Ritschl 2002, 2003, Weder 2006)

During the same time, the Nazis also enacted several important policies in order to redesign the public sector according to their ideas. In 1933, the Law for the restoration of the professional civil service (*Gesetz zur Wiederhestellung des Berufsbeamtentums*) was passed. It allowed the dismissal of "non-Aryan"or politically "unreliable"civil servants.[3] A similar law was enacted for lawyers, and as far as "non-Aryan"professionals were concerned, both laws were made stricter with a reform of the Citizenship Act in 1935 that precluded Jews from holding public office. While these laws, ceteris paribus, led to a decrease in public employment, employment in the armed forces increased. Within two and a half years, the strength of the German army increased fourfold to around 400,000 men in autumn 1935. While part of this had been achieved by integrating police

[3]See Waldinger (2010, 2012) for some economic consequences of such dismissals.

units into the armed forces, the increase in size between 1933 and 1934 was to a large extent due to volunteers. The officer corps alone increased between October 1933 and October 1935 by nearly 3,000 men. With the re-introduction of national conscription in October 1935, the expansion of the army was advanced further. The navy and the air force, a completely new formation, experienced similar increases. (Deist 2003, ch. II) Obviously, this increased military force also required new installations. In this respect, it is noteworthy that the city of Coburg, a small town in Northern Bavaria with very high vote shares for the Hitler movement[4] that was labelled "the first Nazi town"in a book by Hayward and Morris (1988), experienced a substantial amount of public construction after the Nazi's seizure of power. In 1934, several new military barracks were built, followed by a regional centre for the Hitler Youth in 1937. Other projects, such as a new monument to remember soldiers killed in action or a "thingstead", were planned, but never realised. (Nöth 2006)

In addition, the new Nazi government also greatly changed the composition of public employees. The dismissal of Jewish or politically opposing civil servants is one prominent example. Women were also affected by this. While it is a myth that the Nazis drove women out of the labour force en masse, they did take action against women in the upper ranks of civil and professional service jobs. Women were excluded, for example, from the judiciary, the bar, and (with few exceptions) the highest levels of public service jobs. At the same time, however, the number of women working in low-level office jobs increased. (Stephenson 2001, ch. 3)

Finally, and particularly importantly for this study, it was Hitler's explicit desire that deserving party veterans would be rewarded for their loyalty by giving them attractive positions in the local administration and related branches. In Heilbronn, for example, several men were hired by the municipal administration solely based on account of their long-lasting party membership, and one position seems to have been created exclusively for one such party stalwart. Another applicant was told by the NSDAP district head that "...as an old party member, you have a salary that is not commensurate with your contribution in establishing the 3rd Reich. I will immediately look for a suitable position for you."(Schlösser 2003) The importance of party affiliation and patronage in the distribution of public sector jobs can also be seen in the fact that civil servants and teachers are particularly prevalent among the "March Converts", i.e. among those that joined the Nazi party after it had already come to power. Falter (2013) notes that the number of civil servants in the party more than quadrupled in 1933 and suspects that many of these entries were motivated by

[4]In 1933, for example, the NSDAP received 55.8%, compared to the overall national result of 43.9%.

the desire to secure or advance public sector careers.[5]
It is easily conceivable that giving local administration jobs to party veterans and removing political opponents from office would have increased the number of public employment jobs in NSDAP strongholds (where there presumably were more long-standing party members) relative to cities whose electorate had withstood the new regime (where there were more opponents and thus more people that potentially could be dismissed). Local folklore gives an example for another form of (presumed) punishment of opposing cities. According to popular belief, the independent city of Lübeck[6] lost its independence due to its opposition to Hitler. Allegedly, the town council in 1932 had prevented Hitler from speaking within the city's borders, and Hitler took revenge in 1937 by revoking the city's independent status and making it part of Schleswig-Holstein. While the overall credibility of this story is rather dubious (see Pressemitteilung der Stadt Lübeck 2012), its existence alone suggests that people believed that Hitler's policy was driven by such thoughts.
Taken together, the clearly discriminating purpose of the Nazis' public employment laws and the anecdotal evidence above suggests that the Nazis, once in power, might have used public employment and the appointments of public servants to reward cities and regions that had been loyal to them and to punish those that had been reluctant until the very last. If this were the case, one would expect to find a relative increase in public employment for cities with high NSDAP vote shares. Such a rewarding behaviour could be due to political patronage, which has been shown to be operational in other settings as well (e.g. Xu 2017 for the case of the British Empire, Gagliarducci and Manacorda 2016 for nepotism in modern Italy). Hitler's desire to give public sector jobs to long-standing party members and similar decrees in road construction (Silverman 1998, p.188) indicate that this was also a motive in Nazi Germany. Ideological considerations also might have played a role. Erdmann (1969, p.115ff.) points out that the Nazis tried to fabricate the illusion of a long-standing National socialist tradition, creating a propagandistic cult around the early days of the party and the events that the movement's "old guard" had lived through. Loyal supporters were given honorary medals like the "blood medal" (*Blutorden*; see Schmitz-Berning 1998, p.117f.) for participants of Hitler's putsch in 1923, and cities would be given (or adopt with official consent) honorary titles such as "the capital of the movement" (Munich, see Schmitz-Berning 1998, p.296f.) or

[5]According to Schoenbaum (1966/1997, ch. 7), the traditional central bureaucracy proved somewhat resilient to party patronage, with obvious exceptions such as the ministries headed by Goebbels and Göring. However, he also notes that the local level saw a closer union of state administration and party: In 1935, for example, around 20% of all State and local offices were occupied by party members that had joined the party before 1933. Among city offices, this share even amounted to 47%.

[6]Three *Länder* in the Weimar Republic were merely city-states: Bremen, Hamburg and Lübeck, all of them former Hansa cities.

"First National Socialist town in Germany"(Coburg, see Sandner 2000, p.157). Celebrating long-standing loyalty and National Socialist tradition thus appears frequently in Nazi propaganda. Finally, rewarding core voters as opposed to marginal ones can be utility-maximizing for candidates in some models of voting games, for example when candidates need loyal "activists"(Lindbeck and Weibull 1987) or if core voters are more responsive and less risky investments than swing ones (Cox and McCubbins 1986). The evidence presented by Hsieh et al. (2011), Hodler and Raschky (2014) and Do et al. (forthcoming) also highlights the importance of political and regional favouritism in distributional politics, particularly in autocratic regimes and countries with worse institutions.[7] However, it should be noted that a priori, it is also conceivable that public spending could be particularly increased in more disloyal regions in order to "buy support"from former adversaries or marginal voters and thus stabilise the regime in its early days. An emerging body of economic literature has shown that local government spending has a positive causal effect on support for the government (see for example Manacorda, Miguel, and Vigorito 2011 and Litschig and Morrison 2012). Particularly important for the context of my study, Voigtländer and Voth (2016) find that areas traversed by newly-built motorways reduced their opposition to the Nazi regime between 1933 and 1934. If the NSDAP were distributing public funds and jobs in a way to broaden its support base, one would expect to find a relative decrease in public employment for cities with high NSDAP vote shares, or a relative increase for cities with low NSDAP vote shares.

3.3 Empirical Strategy

3.3.1 Data and Summary Statistics

In order to evaluate whether the Nazis allocated more public sector jobs to cities with high Nazi vote shares, I collected data on the number and fraction of people working in such jobs from the German Censuses of Occupation. Administered in 1925, 1933, and 1939, they contain fairly detailed data about the number of people working in different occupations and types of jobs. Unfortunately, the definitions of jobs and occupations and the method of counting them vary somewhat over time. The 1939 census, for example, reports separate numbers for people working in the public administration and armed forces, in teaching occupations, church-related occupations, in legal or economic counselling, and in the entertainment industry. The 1925 census, on the other hand, groups all of those occupations together, while the 1933 census has a slightly finer categorisation

[7] Related to this literature, but of less importance for the setting I study, Burgess et al. (2015) document the role of ethnic favouritism in autocratic regimes.

that at least separates the entertainment industry from administration, armed forces, church, and teaching. The ideal measure for the purpose of my analysis would be the 1939 census category of people working in public administration and the armed forces, but in order to obtain a consistent measure of "public employment"for all three censuses, I am forced to group several occupations, following the widest definition of the 1925 census. Because of this, my measure of public employment does not only include occupations related to the administration and the armed forces, but also teaching professions, artists and other entertainment professions, and church-related professions. For more details, I refer to Appendix A. While this wider definition introduces additional noise into my outcome variable, table 3.1 provides evidence that public administration proper is still the key driver of my measure. In Panel A, I use the narrower 1939 employment categories to decompose my aggregated 1939 measure. As can be seen, in 1939 nearly two thirds of the people working in public employment according to my wide definition worked in occupations belonging to public administration and the armed forces. Typical occupations in this narrow category in 1939 include officers, non-commissioned officers and long-serving privates in the *Wehrmacht*, civil servants in various levels of public administration, and the police (see Appendix A). In Panel B, I decompose my aggregated measure in 1933 into the two components that the 1933 census reports separately, and again it can be seen that artists and entertainers are only a very small part.[8] Finally, and most importantly, Panel C shows that the variation that my aggregate measure uses is extremely similar to the variation in the narrower categories of public employment according to the 1933 and 1939 definitions: In 1939, my aggregated measure of public employment is extremely highly correlated with the number of workers in the narrower 1939 census category of public administration, armed forces, and judicature (both numbers normalised by city population) and even more so with the somewhat broader 1933 category. Thus, while aggregating several census categories introduces noise, the resulting variable still seems to be driven to a very large extent by public administration proper. To address further concerns, I will show that using only 1933 and 1939 and the narrower definition of public employment according to the 1933 census leads to similar results, as does using the number of civil servants (*Beamte*) across all sectors, which is available for 1933 and 1939. I will also show cross-sectional results for 1939 only, using the narrower categories of the 1939 census.

In my baseline results, I will use the data from 1925, 1933 and 1939 and thus the wider, but internally consistent public employment definition described above. My main outcome variable is the log of the number of public administration jobs (denoted *logadmin*). Since the German cities experienced considerable

[8]The fraction of the entertainment category in 1933 is much larger than in 1939 since the 1939 census groups some artists (e.g. actors) with teachers.

population growth between 1925 and 1939, I also examine the ratio of public administration jobs to total population (*adminpop*) and the ratio of public administration jobs to the labour force (*adminshare*). While the latter measure is robust to population changes, it is subject to another type of criticism: public investment is likely to also have increased employment in other sectors (for example due to increased spending on construction or military equipment). In some cities, growth in these industries might have outpaced growth in public employment, leading to a decrease in *adminpop* or *adminshare*, even though the city might still have been benefiting from increased public employment. Put differently, a relative increase in *adminpop* or *adminshare* might just mean that the rest of the city's economy was doing relatively poorly, thereby increasing the share of public sector jobs. However, this is not a concern for *logadmin*, which measures absolute, rather than relative, increases in public employment. Both the absolute and the relative measures therefore have their respective advantages and disadvantages. In practice, both lead to similar results, as will be shown below.

The advantage of using public sector employment, on the other hand, is that most of the respective job categories are under the direct control of the central or regional government (e.g. the number of officers and non-commissioned officers in the armed forces) and less constrained by the pre-existing regional industry, making them a more natural outcome measure that also has closer resemblance to the individual-level outcome measure employed by, for example, Hsieh et al. (2011). Data about this measure of public employment is available for nearly 300 cities; In particular, it is available for all cities with more than 20,000 inhabitants and for some smaller ones that happened to be independent cities, not belonging to any other administrative district (*Kreisfreie Städte*). Figure 3.1 gives an impression of the geographical distribution of public employment growth over the sample period. Specifically, it maps quintiles of the difference in *logadmin* between 1939 and 1925. As can be seen, the increases in public employment are spread relatively evenly across the country, with perhaps a small concentration of strong increases in Central Germany.

The main explanatory variable of interest is the NSDAP vote share in the election of March 1933. For this, I use the extensive database on social and electoral variables of Weimar Germany compiled by Dirk Hänisch and Jürgen Falter. This database also contains other socioeconomic variables that might be of interest when analysing NSDAP vote shares. In particular, I include the Jewish share of a city's population in 1925 and the unemployment rate at the time of the census in 1933.[9] In addition to these socioeconomic variables, I also control for

[9]For Berlin, the data in the database are on the level of the city's administrative districts. I created an aggregated measure for Berlin by adding all districts and boroughs belonging to it. In order to assess the validity of this aggregation, I compared the aggregated population to

longitude, latitude, an indicator for being in the Rhineland and an indicator for being a Nazi gau capital. These variables are included to control for potential geographic determinants of public, especially military employment. According to articles 42-44 of the Versailles treaty, Germany was not allowed to maintain or construct fortifications or assembly troops on the left (Western) bank of the Rhine or within 50km of its right (Eastern) bank. In 1936, Hitler violated this stipulation by "reoccupying"the Rhineland with armed forces. If the Rhineland is also correlated with voting behaviour, this could create a spurious correlation between public employment and vote shares. To avoid this concern, I control for a city lying in the "Rhineland"as defined by the Versailles Treaty. Similarly, the territorial reorganisation of the former German states (*Länder*) into new units called *Gaue* might also have led to differential changes in public employment in the gau capital cities. If these capital cities had also been more likely to vote for the Nazis, again a spurious correlation might arise. I therefore also include an indicator for whether a city was a gau capital in 1938. (Das Buch der deutschen Gaue 1938)

One potential problem is the question whether a city in 1925 was the same city in 1933 and 1939, as many German cities underwent changes in their territory and population, acquiring smaller surrounding towns and villages, merging with other cities and the like. The prime example for this is Wilhelmshaven, which more than quadrupled its population between 1933 and 1939 due to the acquisition of the neighbouring city of Rüstringen. Similar mergers occurred in the Ruhr area in 1928-1930. In order to avoid problems due to these territorial restructurings, I excluded all those cities which experienced a substantial enlargement in their population between 1910 and 1925, 1925 and 1933 or 1933 and 1939.[10]

In addition, I use voting data from the 1912 Reichstag election, for which I have city-level data for all cities that had more than 10,000 inhabitants in 1910. These were obtained from the official election results, published by the Statistisches Reichsamt in 1913. Ultimately, I end up with a sample of 246 cities for the three census years 1925, 1933 and 1939. The 1933 census was administered on June 16, four and a half months after Hitler had become chancellor, but still before his large-scale rearmament programmes had begun. I therefore usually

the one from the censuses in 1925 and 1933. Some differences exist, but they are well below 5%.

[10]In particular, for all cities whose population growth between 1910 and 1925 or 1925 and 1933 or 1933 and 1939 exceeded the mean growth rate by more than one standard deviation, I analysed whether this large population growth was due to territorial gains or changes that made the city grow by more than 25% alone. If this was the case, I excluded the city. For details, see Appendix B. As an alternative measure, I simply excluded all cities whose population growth between either 1910 and 1925, 1925 and 1933 or 1933 and 1939 exceeded the mean growth by more than one standard deviation. The results are not sensitive to this, as shown in section 3.4.2.

treat it as a pre-NSDAP year but will show in section 3.4.2 that my results are not driven by this.

Table 3.2 shows summary statistics of the explanatory and explained variables. As can be seen, both the number and shares of public employees increased from 1925 to 1933, and then decreased again. The number of public employees is higher in 1939 than in 1925, whereas their shares out of the total population or labour force are lower. Given the large amount of public investment and the substantial increase of the German armed forces between 1933 and 1939, this might seem surprising. The most likely explanation is that public employment was driven up between 1925 and 1933 by general employment measures, since it was already used as a means of fighting unemployment before the Nazis came to power. Hitler's predecessors von Papen and von Schleicher had already made credit and funds for public employment measures available (Golecki 1986, p.XXXIV-XL), and Blaich (1970) uses the example of Ludwigshafen to show how cities themselves tried to fight the economic crisis by employing otherwise unemployed workers in the construction of roads and sewerage systems. This pattern of a strong increase as a reaction to the economic crisis makes it even more difficult to uncover the causal effect of the NSDAP vote share using a standard OLS approach. Cities where more of these emergency projects were carried out could experience an increase in public employment between 1929 and 1933, followed by a decrease until 1939, when full employment had rendered these emergency measures obsolete. As an alternative measure of public employment, I also use the number of civil servants (*Beamte*) across all sectors. Data for this variable exists for 1933 and 1939, and this measure (which should measure more permanent jobs, but across a wider spectrum of sectors) shows a raw increase even between 1933 and 1939. The mean NSDAP vote share in my sample of cities is 41.8%, very close to the national average of 43.9%. The 1912 Economic Association was a much smaller party. Across my whole sample, it averaged 1.5% of all votes, but with a sizeable dispersion- in 191 cities it did not receive any votes, while four cities recorded EA vote shares greater than 20%.

3.3.2 The "Economic Association"and its voters

It is unlikely that a simple OLS regression, even after controlling for city fixed effects and control variables, uncovers the causal effect of Nazi support on subsequent public employment. The main endogeneity concern in such a regression is due to the economic crisis in the years following 1929. For example, cities that were more adversely affected by the crisis might have been differentially prone to vote for the NSDAP in 1933 and they might also have been subsequently those with different public employment shares. The NSDAP vote share would then be correlated with the error term, and as a consequence, the estimate of β in such a regression will be inconsistent. Another potential problem could arise if public employees themselves are more or less likely to vote for the NSDAP, creating a reverse causality problem. Finally, Voigtländer and Voth (2016) have shown that local highway construction caused increased support for the NSDAP government, suggesting an additional reverse causality problem: Voting behaviour might affect public spending, but public spending (or expected public spending) might also affect voting behaviour.

In order to address these issues of potential endogeneity, I instrument the 1933 NSDAP vote share by the vote share of another party, the "Economic Association"(*Wirtschaftliche Vereinigung*, henceforth EA) in the 1912 election. The EA was an alliance of several smaller parties, most notably the "Christian-Social Party"(*Christlich-Soziale Partei*) and the "German-Social Party"(*Deutschsoziale Partei*). Most of these parties had conservative, nationalist platforms that denounced both socialism and capitalism and tried to attract the votes of middle-class voters particularly in Protestant and rural areas. In addition, both the "Christian-Social Party" and the "German-Social Party" were openly anti-Semitic. (Gräfe 2012, Bergmann 2012) The constituent parties of the EA were not strong, and the alliance only obtained a few seats in the 1912 election. However, there are strong parallels between the voters that the EA tried to attract, and the voters that in 1933 voted for the NSDAP.

While the NSDAP had started out using anti-capitalist and socialist rhetoric and catering to the preferences of blue-collar voters, it markedly changed its approach as a result of its disappointing results in 1928. After 1928, the party focussed more on rural areas and presented itself less as a radical force against capitalism but rather as an ultra-nationalist, conservative party that advocated law and order and the fight against the treaty of Versailles. The aim was to attract more middle-class voters that heretofore had been repulsed by the party's more proletarian agenda. (Stachura 1978) This transformation was a successful one; By 1933, the NSDAP had become, in the words of Jürgen Falter (1991, p.372), "a people's party with a middle-class belly"(*eine Volkspartei mit Mittelstandsbauch*), in which the middle classes were the largest fraction. Thus, after 1928, the NSDAP presented itself more as an ultra-nationalist party for the middle-class, with a particular focus on rural and Protestant voters, trying to attract the very voters that the EA before World War I had tried to attract, as well as sharing its antisemitism.[11] Because of this, the vote share of the EA in the 1912 election and the NSDAP vote shares after 1928 are significantly positively related (this is also shown more formally in the first-stage results below), and the former can be used as an instrument for the latter.

[11] After 1930, the NSDAP toned down its antisemitism considerably (see for example Voigtländer and Voth 2012). Still, it remained, in the words of Herbert (2000, p.18f.) "a receptacle"for Anti-Jewish elements.

The basic idea of this instrument is to use variation in NSDAP vote shares that is not due to the economic crisis post 1929, but due to persistent political attitudes of the local population such as antisemitism or extreme nationalism.[12] In order to be a valid instrument, the 1912 EA share has to satisfy the exclusion restriction. In particular, the identifying assumption of this strategy is that the 1912 EA share does not have an effect on public employment outcomes later on, other than through affecting support for the NSDAP. Several aspects make the 1912 EA vote share attractive in this respect. Firstly, dating more than 20 years prior to the 1933 election, use of the 1912 EA share should not be susceptible to reverse causation problems due to either public spending in the Weimar era, expected public spending by a potential NSDAP government, or the voting behaviour of public employees. Secondly, using a vote result prior to the economic crisis that started in 1929 allows to purge the 1933 vote shares of any factors due to this crisis. One remaining concern, however, is that there might still be unobserved factors that are correlated both with the 1912 EA share and with the 1933 NSDAP share and that might also be relevant for the evolution of public employment over time. Cities with a high EA share might be fundamentally different from those with low or no EA votes, and these differences might also affect public employment patterns. The absence of such differences cannot be tested or proved in a formal sense. However, I can at least examine whether the instrument is correlated with the level and evolution of relevant variables before 1933. In table 3.3, I run a cross-sectional regression for 1925, relating the three public employment outcomes (the ratio of public employment to either the population or the labour force, and the natural logarithm of public employment) to the 1912 EA vote share. Odd numbered columns show simple bivariate correlations, even numbered ones include controls. Since I cannot include fixed effects in the cross-sectional regressions, I include the natural logarithm of a city's population as an additional control to make sure I compare cities of similar size. On the other hand, I omit the 1933 unemployment rate from these 1925 regressions. There is clearly no relationship between public employment in 1925 and the 1912 EA vote share. Regardless of the in- or exclusion of controls, the estimates are statistically insignificant and of small magnitudes. For example, according to the specification with controls, increasing the 1912 EA vote share by one full standard deviation would decrease the 1925 public administration jobs to population ratio by only 2% of a standard deviation. While the absence of a correlation between my instrument and the level of public employment in 1925 is reassuring, the identifying assumption of my strategy is not that the

[12]It is for this reason that I do not control for measures of long-term antisemitism that Voigtländer and Voth(2012) have constructed and used. If antisemitism led to increased votes for the EA in 1912, this is "good"variation that I want to use in my estimates, and not remove it by controlling for it.

1912 EA vote share is uncorrelated with the counterfactual *level* of public employment, but with its counterfactual *changes*. To assess this, In columns 1-3 of table 3.4 I show the results of a "placebo test", examining whether cities with different 1912 EA vote shares experienced different evolutions of public employment between 1925 and 1933. Specifically, I regress my 3 outcomes on the interaction between the 1912 EA share and an indicator for the year 1933, an indicator for the year 1933, city fixed effects and my control variables, each of which is interacted with an indicator for the year 1933. The results from this exercise are encouraging. The 1912 EA vote share is never significantly associated with the development of public employment between 1925 and 1933, and the point estimates are small. For example, increasing the 1912 EA vote share by one standard deviation is associated with an increase in administration jobs as a ratio to total population of less than 1% of a standard deviation of the 1925 administration job ratio. Even more encouraging, in columns 4-6, I examine whether the 1912 EA share is correlated with the evolution of a city's economy as measured by the employment shares of three broadly defined sectors and again do not find any relationship. Between 1925 and 1933, cities with high or low 1912 EA vote shares did not experience different evolutions of employment in either agriculture, manufacturing or commerce. Finally, in columns 7-9, I look at the occupational class composition of a city's economy, measured by the respective shares of blue-collar employees, white-collar employees, and self-employed. Again, there is no relationship between voting for the Economic Association and changes in a city's occupational class structure.[13]

Another important question is whether cities with different 1912 EA vote shares were affected differentially by the economic crisis after 1929. I provide a coarse assessment of this in table 3.5, where I correlate unemployment during the crisis with voting for the EA in 1912. A slight negative bivariate correlation is present between the unemployment rate in 1933 and the 1912 EA vote share, but this is not robust to the inclusion of controls. In addition, when looking at changes in unemployment by regressing the difference in the logs of unemployment in 1932 and 1930 on the 1912 vote share, I again do not detect any sizable or statistically significant relationship.

The previous results strengthen the instrument's case for exogeneity: The 1912 EA share is not correlated with the level or evolution of public employment between 1925 and 1933. It is also not associated with changes in the city's sectoral or occupational class structure. Finally, there is also no evidence that cities with high EA vote shares experienced different evolutions of unemployment during

[13]For columns 4-9, it should be noted that the measures employed differ for the 1925 and 1933 census, as the 1925 census does not contain the number of people employed in a given sector, but the number of people employed in a given sector and their family members. However, such uniform differences in measurement should be absorbed by the year fixed effect.

the economic crisis in the early 1930's. It is still possible that cities with high and low EA votes differed along other dimensions, but the previous results seem to indicate that such differences, if present, were not very consequential in economic terms.

Based on this, I will employ a standard two-stage least square estimation procedure. The basic model of interest is

$$y_{it} = \beta \cdot NSDAPshare33_i \cdot D1939_t + \tau \cdot D1939_t + \gamma' \cdot X_i \cdot D1939_t + c_i + u_{it} \quad (3.1)$$

where y_{it} represents the outcome, i.e. either *logadmin* , *adminshare* or *adminpop*, in city i and year t, where t can be 1925, 1933 or 1939. $NSDAPshare33_i$ denotes the NSDAP vote share in the 1933 election in city i and $D1939_t$ is a dummy which is 1 for the year 1939. The city fixed effects c_i will account for time-invariant city characteristics that have a constant effect on public employment over time. X is a vector of time-invariant city-level controls, discussed in the previous section. To address the endogeneity problem arising from the likely correlation between the 1933 NSDAP vote share and the economic crisis, the variable of interest in equation 3.1 will be replaced by its prediction based on the same controls and the 1912 EA share:

$$NSDAPshare33_i \cdot D1939_t = \eta \cdot EAshare12_i \cdot D1939_t + \theta \cdot D1939_t + \chi' \cdot X_i \cdot D1939_t + \xi_i + \epsilon_{it} \quad (3.2)$$

These predicted values will then be used as regressors of interest in the second stage:

$$y_{it} = \beta \cdot (\widehat{NSDAPshare33}_i \cdot D1939_t) + \tau \cdot D1939_t + \gamma' \cdot X_i \cdot D1939_t + c_i + u_{it} \quad (3.3)$$

3.4 Results

3.4.1 Baseline estimates

Before turning to the instrumental variable estimation described in equations 3.2 and 3.3, table 3.6 presents the results of simple OLS estimations of equation 3.1. As discussed in the previous section, these regressions are based on the very strong assumption that the 1933 NSDAP vote share is exogenous, i.e. uncorrelated with the error term. This is not very likely, with the economic crisis after 1929 being a prime confounder. An extensive body of literature has analysed the reasons for the NSDAP's rapid electoral successes, which socioeconomic groups were more likely to vote for the National Socialists and why they did so. (see, among others, Frey and Weck 1981, Falter et al. 1983, Falter et al. 1985, Falter 1991, van Riel and Schram 1993, Stögbauer 2001, King et al. 2008, Adena et al.

2015, Koenig 2015, Spenkuch and Tillmann 2016, Satyanath, Voigtländer and Voth forthcoming). While some disagreement about certain issues still exists, there is a clear consensus that the economic crisis that affected Germany in the early 1930's was a prime driver of National Socialist vote shares. If cities that were affected more by the crisis were differentially likely to vote for the NSDAP in 1933, the OLS estimates here would be biased. The sign of the bias is unclear a priori: if cities that were affected more by the crisis were more industrialised and therefore more strongly connected to the communist parties, and public employment increased in these cities as a response to the crisis, the resulting OLS estimate could be downwardly biased. A similar negative bias could arise if public servants were more likely to vote for the NSDAP and cities with more public employment experienced slower growth in public employment (for example since they were less affected by the crisis and thus did not need large-scale investment programmes, or at least only smaller ones). On the other hand, if cities that were affected more by the crisis turned towards the Nazis, it is easily conceivable that an upward bias might arise.[14] With these caveats in mind, the estimates of β in table 3.6 are negative, but typically not very sizable and not significantly different from zero. Taken at face value, this would mean that a city's NSDAP vote share in 1933 had no or a slightly negative effect on the city's public employment share in 1939. This could mean that if anything, instead of "favouring" loyal cities, the new government tried to "buy support" from more resistant cities, for example in an attempt to stabilise its power in the early days of the regime. However, due to the reasons above, OLS estimates should be viewed with caution and I turn next to the instrumental variable (IV) estimates discussed in section 3.3.2. By using an instrumental variable that significantly predates the 1929-1932 economic crisis, this approach avoids any correlation between the NSDAP vote shares and the intensity of the crisis.

The 2SLS estimates based on using the 1912 vote share for the "Economic Association" as instrumental variable for the 1933 NSDAP vote share are presented in table 3.7. As can be seen from the first stage result, the 1912 EA vote share and the 1933 NSDAP vote share are indeed strongly and, as one would expect, positively related. The first-stage results indicate that a one percentage point increase in the 1912 EA vote share increases the 1933 NSDAP vote share by around 0.35 percentage points. Thus, the translation between 1912 EA voters and 1933 NSDAP voters is not one-to-one, which is not surprising, given that some members of the Economic Alliance joined other parties after 1918, in

[14] The exact relationships between the economic crisis in the late 1920s and the rise of the NSDAP are still debated in the literature. The most recent study by King et al. (2008) finds that the most adversely affected groups reacted differently in their voting behaviour: While the "working poor" such as self-employed shopkeepers and professionals increasingly voted for the Nazis, the unemployed turned towards the communists. A priori, it is therefore not clear how adverse effects of the economic situation would correlate with the NSDAP vote share.

particular the German National People's Party, Hitler's coalition partner from January 1933 (Bergmann 2012).

Turning to the main results in Panel B, addressing the potential endogeneity of the 1933 vote shares uncovers a positive and significant effect of voting for the NSDAP on subsequent public employment. An increase in the 1933 NSDAP vote share of one percentage point would increase the number of public sector jobs by around 2.3%, a quite substantial increase. When looking at public employment as a share of population, the results indicate that a one percentage point increase in the 1933 NSDAP vote share is associated with a 0.1 percentage point increase in the ratio of public employment to total population. Put differently, an increase of one standard deviation in the 1933 vote share leads to an increase of around 45% of a standard deviation in terms of the 1925 public employment share. The results for public employment as a share of the labour force are similar in magnitude.[15] All these results are relative results- they show that cities with greater NSDAP vote shares experienced a relative increase in public employment compared to cities with lower Nazi support. As the summary statistics in section 3.3.1 show, both the number of public employment jobs and its shares are lower in 1939 compared to 1933. In this sense, my results show that public employment fell less for cities with greater NSDAP support. Below I will also show results based on the number of civil servants, which at least in raw numbers increased between 1933 and 1939. Taken all together, the results from table 3.7 show that high 1933 NSDAP vote shares led to a subsequent relative increase in public sector jobs, both in absolute numbers and in ratios of the population. This pattern thus would not be consistent with the Nazi government buying support from opposing cities, but rather rewarding its strongholds via public employment.

Table 3.8 provides further evidence that the estimates in table 3.7 are based on government discrimination, rather than other economic forces at work. Here, I repeat the analysis above, but this time using the metal industry, a sector that contracted during the 1929 crisis and expanded during the pre-war build-up, but was not under direct government control, so I would not expect to find an effect here.[16] The estimates are all not significantly different from zero. In the

[15] As explained in section 3.3.1, the 1912 EA vote share has a somewhat unusual distribution with many zeros and a right tail of four cities with very high values. This distribution could mask some heterogeneity in the effects, which I further explore in Appendix C. Cities with no 1912 EA vote share are important for the precision of the estimate, but tend to dampen the effect size. Because of this, the 55 cities with nonzero EA vote share display larger effect sizes, but lower precision. The four cities with EA vote shares above 20% are influential for the overall results but less so when focussing on only cities with positive instrument values. Removing both extremes results in a smaller sample and less precise estimation, but does not change the sign of the results.

[16] In the case of *adminpop* and *adminshare*, a slight negative coefficient could arise mechanically, as an increase in the share of public employment has to be compensated by all other sectors.

case of *adminpop* and *adminshare*, the point estimates are relatively sizable, but very imprecisely estimated, whereas for *logadmin* also the point estimate would indicate a very limited negative effect of 1%.

Overall, my findings indicate that cities that had voted more for the Nazis experienced a relative increase in the number and importance of public sector jobs. As the historical record in section 3.2 shows, there are several potential channels that could explain this. On the one hand, the NSDAP could have simply allocated public funds or public employment positions to its strongholds or taken them away from opposing cities. The case of Lübeck losing its independence would be an example, but it is most likely only local folklore. Another, potentially more important channel is political patronage. Patronage has been shown to be operational in other settings (e.g. Xu 2017 for the case of the British Empire, Gagliarducci and Manacorda 2016 for nepotism in modern Italy), and it is likely that it was also important in Nazi Germany. As described above, the Nazi government passed several laws and decrees to redesign the German civil service, removing Jews, political opponents and women from office and giving jobs to party veterans. This changed the composition of the public sector, but it could also have easily changed its distribution across cities: Cities that had opposed the NSDAP before probably had more political opponents in public jobs that were then removed, whereas Nazi strongholds probably had more party stalwarts that were rewarded, for example by creating new positions for them. In this sense, my results are likely an understatement of the change in the civil service, since they only show changes in the size of the public sector, and not in its composition. This is one caveat that has to be borne in mind. Another caveat is that part of the increased importance of public sector jobs could be due to supply factors rather than increased demand. The 1939 census groups the armed branches of SA and SS together with the regular *Wehrmacht* forces, and in cities leaning more towards the Nazis, more people might have decided to join these outfits (or the armed forces in general). While such differential supply-side behaviour is possible, the actual allocation of men to units and thus barracks would still be a decision to be made by the respective organisation and thus the government or one of its branches.

3.4.2 Robustness

This section addresses several potential concerns with the findings from the IV regressions. As explained in Section 3.1, there were several mergers and restructurings of cities between 1925 and 1933. I tried to exclude all cities whose population growth was mostly driven by territorial enlargement. Still, a certain arbitrary element remains- when are territorial changes so important that a city is no longer comparable over time? In table 3.9, I repeat the analysis of table

3.7 for two different and somewhat “extreme”samples: In columns 1-3, I do not drop any cities, while columns 4-6 exclude all cities whose growth between either 1910 and 1925, 1925 and 1933 or 1933 and 1939 exceeded the respective mean by more than one respective standard deviation. For convenience, the first stage results are omitted and only the respective F statistic is displayed. Compared to the main results in table 3.7, standard errors increase a bit both when including more cities (columns 1-3) and when including fewer (columns 4-6). In the latter case, the point estimates are very similar to the main specification, whereas not dropping any cities if anything seems to lead to slightly larger point estimates. Overall, however, the qualitative conclusions from before do not change.

The validity of using the 1933 NSDAP vote share as explanatory variable might also be questioned on several grounds. First of all, throughout the analysis, I have treated 1939 as the only year after the NSDAP rise to power, considering 1925 and 1933 to be pre-treatment years. The reason for this is that the first large-scale public investment programme by the National Socialists, the Rheinhardt programme, started in June 1933, few days before the 1933 census data were collected, and the expansion of the armed forces occurred even later. Against this definition, however, one can argue that the National Socialists came to power already in January 1933 and already had a parliamentary majority after March 1933. As shown by Ferguson and Voth (2008), shareholders immediately reacted to this change in power, and so might have other economic variables. Moreover, the Nazis started a major prosecution of their political enemies after the Reichstag Fire in February 1933, which might also have had implications for public employment. Similarly, the “Law for the restoration of the professional civil service”was also already passed in April 1933. Hence, treating 1933 as a pre-treatment year might be problematic, though if anything, it would most likely bias my estimates towards 0.

In table 3.10, I examine the effect of dropping the potentially confounded year 1933, only comparing 1925, a clear pre-treatment year, to 1939, a clear post-treatment year. The results are again very similar to the baseline results and if anything slightly larger, which would be in line with the 1933 numbers already being slightly contaminated by the Nazi rise to power. However, any such contamination appears to be very limited.

Another potential concern with the 1933 election is that it was not the election that brought Hitler into power, it was only the one that gave him a parliamentary majority. Secondly, since the election happened after the Reichstag Fire and the subsequent prosecution of Communists, it is questionable whether this election was really a free one. In table 3.11, I therefore redo the analysis of table 3.7, but use different elections as main explanatory variables, in particular those in September 1930, July 1932 and November 1932. While data for the 1930 election are available for all cities in my sample, the results for the 1932 elections were unfortunately only reported at the district level. I can therefore run these regressions only on a limited sample that includes cities that were also a district at the same time (*Stadtkreise*, as opposed to cities that were part of a *Landkreis*), which decreases the sample size by around one third. Still, the results from table 3.11 confirm the previous results, both in terms of sign and magnitude and indicate that there is a positive relationship between voting shares for the NSDAP and subsequent public employment. In addition, the first stage F statistics show that the relationship between the 1912 EA vote shares and the later NSDAP vote shares is stronger in 1932 than in 1930. This is consistent with the NSDAP becoming more and more attractive for the nationalist lower middle class voters to whose preferences the constituent parties of the Economic Association had catered in Imperial Germany 20 years earlier.

Another potential problem is the wide definition of public employment that I am using throughout the analysis. As mentioned in section 3.3.1, creating a consistent measure for 1925, 1933 and 1939 requires aggregating several occupations, some of which are clearly not part of "public employment". This will increase the noise in my outcome variable and thus impair the precision of my estimates, but if these other jobs also react to the Nazi rise to power, it could introduce biases. In table 3.12, I therefore examine the robustness of my results with regards to different outcome variables. Such variables unfortunately do not exist for 1925, but for both 1933 and 1939. In columns 1 and 2, I therefore show results for my baseline outcome for these two years only. For brevity, I focus on *adminpop* and *logadmin* as outcome variables, omitting *adminshare*, which is very highly correlated with *adminpop*. As can be seen, the results when using 1933 and 1939 are very similar to the main results in table 3.7, which is further evidence that in spite of the late census, 1933 is still a valid pre-treatment year. In columns 3 and 4 then, I use a more narrow public employment definition based on the 1933 census definition that at least omits the cinema, theatre and other entertainment occupations. The results from using this definition are very similar to those in columns 1 and 2. In columns 5 and 6, I employ a very different outcome variable. Instead of trying to get at the public administration *sector*, I here use the occupational *class* of civil servants (*Beamte*) across all sectors.[17] Again, results are very similar: Increasing the 1933 NSDAP vote share by one percentage point would lead to an increase in the number of civil servants by 2.7%, very much in line with the 2.3% in column 2 of this table and column 3 in table 3.7.

An even narrower definition of public employment is possible when using data

[17] I cannot use 1925 for this regression, since the 1925 census aggregates civil servants with all white-collar employees. The 1933 and 1939 civil servant numbers differ slightly in that the former omits the most top-level civil servants, but this should be a minor difference.

from 1939 only. In table 3.13, I show results from cross-sectional regressions for the five narrow employment categories that together form my employment definition that is consistent across the three censuses. These regressions cannot control for fixed effects and therefore rely on stronger identifying assumptions. As before, to at least partly remedy this, I include the natural logarithm of a city's population as additional control. In Panel A, I show the respective ratio of each of the five categories to total population, i.e. the equivalent of *adminpop*. As can be seen, the overall increase in *adminpop* is driven mostly by public administration and the armed forces. Legal and economic counselling and teaching occupations show at most a very slight increase. Entertainment does not respond at all, whereas church-related occupations even decrease. Panel B, which uses the natural logarithm, by and at large confirms the previous findings. However, the estimate for public administration becomes relatively imprecise and loses statistical significance. Still, the point estimate of a 3.1% is very similar and if anything slightly larger than the previous ones. Overall, as already argued in section 3.3.1, public administration and the armed forces appear to be the prime drivers of my measure of public employment. Taken together, the results in tables 3.12 and 3.13 show that while the aggregation of several occupational classes is undesirable from a conceptual point of view, it seems to have little implications for the sign or magnitude of my findings.

3.5 Conclusion

Between 1928 and 1933, the NSDAP developed from a small and unimportant party in Weimar Germany into the strongest party in the German parliament, bringing its leader Adolf Hitler to the head of the German government by January 1933 and gaining a parliamentary majority by March of the same year. Subsequently, Hitler used this power not only to concentrate all political competences among his followers, but also to enact large public investment and rearmament programmes and to overhaul the civil service. In this paper, I document evidence that the public employment policies during the early Nazi era were not ideologically colour blind. Using a city's 1912 vote share of the Economic Alliance, a small party in Imperial Germany that catered to similar voters as the late NSDAP, as an instrumental variable for the 1933 NSDAP vote share, I find that the latter had a positive and significant effect on subsequent public employment at the city level. A one percentage point increase in the 1933 vote share caused the number of public employment jobs to grow by around 2.5 percent, a finding which is not driven by cities undergoing territorial changes, by the inclusion or exclusion of the potentially already contaminated census year of 1933, or by the potentially worrisome and unfree 1933 election shares and is also robust to using different definitions of public employment as outcome variables. The results of this study also shed additional light on the ability of governments to use economic policy as a means to reward and protect their voters and supporters, and/or to punish their political adversaries. Thereby, it adds to a vast body of literature that has documented such behaviour on a firm-level and, to a certain extent, also for individuals. Of course, some cautionary remarks apply. In particular, Germany's Nazi government had incomparable powers to any modern democratic government. Being freed of the constraints usually posed by a parliamentary opposition, judicial review by courts and a free press, it seems reasonable to assume that the National Socialists' ability to reward a city's loyalty was substantially larger than in most countries at most times. On the other hand, my measure of public employment only captures changes in its size, not its composition and therefore potentially understates the extent of rewarding for loyal followers.

Anecdotal evidence suggests that one important mechanism behind my finding is the preferred allocation of public sector jobs to loyal party stalwarts. Several questions remain in this respect for future research: Can this behaviour also be documented at the individual level? Were there other channels through which public employment was targeted towards cities favouring the Nazis? Are there any long-term effects of the increased public employment in the 1930's, i.e. did the economic reward for the cities survive the Second World War and persist longer than the Nazi government, or did they disappear when the favoured party members lost office? And, in a broader context, what are the welfare implications of such favouring behaviour? In these respects, there is substantial scope for further research.

3.6 Tables and Figures

Share of the 1939 subcategories in the overall definition of public employment, 1939		
	Mean	Std. deviation
Public Administration, Armed Forces, Judicature	64.935	13.351
Teaching and artistic occupations	23.325	8.654
Church-related occupations	7.079	8.242
Legal and Economic Counselling	3.094	1.388
Entertainment (without artists)	1.567	1.050

Share of the 1933 subcategories in the overall definition of public employment, 1933		
	Mean	Std. deviation
Public Administration, Armed Forces, Church, Education 1933	92.099	3.637
Entertainment	7.901	3.637

Correlation of baseline public employment definition with narrower definitions in 1939 (all measures normalised by city population)	
Public Administration, Armed Forces, Church, Education (Census def. 1933)	1.000
Public Administration, Armed Forces, Judicature (Census def. 1939)	0.956
Observations	246

Table 3.1: Assessing outcome variable

	Mean	Standard deviation	Minimum	Maximum
Public Employment Level 1925	2,643.358	4,556.106	199	38,703
Public Employment Level 1933	3,260.264	5,788.368	234	48,687
Public Employment Level 1939	2,843.760	4,484.459	170	42,090
Civil Servants Level 1933	2,764.179	4,619.024	173	35,786
Civil Servants Level 1939	3,334.577	5,093.963	179	43,037
Public Empl. as a percentage share of population 1925	3.929	1.917	1.121	10.844
Public Empl. as a percentage share of population 1933	4.285	1.941	0.882	10.408
Public Empl. as a percentage share of population 1939	3.639	1.548	0.840	9.134
Public Empl. as a percentage share of the labour force 1925	8.548	4.481	1.857	24.113
Public Empl. as a percentage share of the labour force 1933	9.572	4.544	2.520	25.125
Public Empl. as a percentage share of the labour force 1939	8.277	3.921	1.927	20.494
Population 1925	67,770.809	111,899.740	11,782	700,222
Population 1933	74,590.573	124,042.812	12,089	756,605
Population 1939	81,164.516	128,917.766	12,641	829,318
Labour Force 1925	32,748.565	56,838.156	5,440	358,477
Labour Force 1933	34,801.268	60,752.118	5,832	379,032
Labour Force 1939	37,865.756	63,288.081	6,227	432,082
NSDAP Vote Share 1933	41.883	8.765	13.654	60.395
Economic Association vote share 1912	1.463	4.493	0.000	37.683
Unemployment rate 1933	22.281	6.038	5.622	39.227
Jewish population share 1925	0.849	0.807	0.012	5.491
Observations	246			

Table 3.2: Summary statistics

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES			Public employment			
	Ratio to Population		Ratio to Labour Force		Natural Logarithm	
EA share 1912	0.003	0.009	-0.002	0.004	-0.013	0.000
	(0.030)	(0.022)	(0.067)	(0.052)	(0.010)	(0.005)
Controls		X		X		X

Cross-sectional results for 1925 with 246 cities. Controls are the natural logarithm of city population, longitude, latitude, an indicator for the Rhineland, an indicator for being a gau capital, and the Jewish population share in 1925. Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.3: Check for different levels in public employments in 1925

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
VARIABLES	Public employment			Share of the LF employed in			Share of workers classified as		
	Ratio to Pop.	Ratio to LF	Natural Log	Commerce	Manuf.	Agric.	Blue collar	White collar	Self empl.
EA share 1912 · D1933	0.005	0.036	0.000	0.043	-0.049	-0.001	0.006	0.034	-0.041
	(0.007)	(0.031)	(0.002)	(0.058)	(0.043)	(0.011)	(0.034)	(0.027)	(0.025)

Panel data results for 1925 and 1933 with 246 cities and 492 observations. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1933 as well as interactions of an indicator for 1933 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, latitude and an indicator for being in the Rhineland.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3.4: Check for different trends in public employment and other sectors before 1933

VARIABLES	(1) Unemployment rate 1933	(2)	(3) $\Delta ln(Unemployed)_{1932-1930}$	(4)
EA share 1912	-0.124*	-0.066	0.002	0.000
	(0.068)	(0.056)	(0.002)	(0.003)
Controls		X		X
Observations	246	246	201	201

Cross-sectional results for 1925. Controls are longitude, latitude, an indicator for the Rhineland, an indicator for being a gau capital, and the Jewish population share in 1925. Columns 1 and 2 also control for the natural logarithm of city population. Robust standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.5: Relationship between the instrument and the economic crisis

	(1)	(2)	(3)
VARIABLES	Public Employment		
	Ratio to Pop.	Ratio to LF	Natural Log
NSDAP share 1933 · D1939	-0.007	-0.009	-0.002
	(0.011)	(0.025)	(0.003)

Panel data results for 1925, 1933 and 1939, 246 cities and 738 observations. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, latitude and an indicator for being in the Rhineland.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.6: OLS estimates

	(1)	(2)	(3)
Panel A: First stage	Dep. variable: NSDAP vote share 1933		
EA share 1912 · D1939		0.339***	
		(0.093)	
F-stat first stage		13.21	
Panel B: 2SLS estimation	Public Employment		
	Ratio to Pop.	Ratio to LF	Natural Log
NSDAP share 1933 · D1939	0.098**	0.266**	0.023*
	(0.049)	(0.122)	(0.013)

Panel data results for 1925, 1933 and 1939, 246 cities and 738 observations. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, latitude and an indicator for being in the Rhineland.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.7: IV estimates

VARIABLES	(1)	(2)	(3)
	Metal industry Employment		
	Ratio to Pop.	Ratio to LF	Natural Log
NSDAP share 1933 · D1939	-0.195	-0.401	-0.010
	(0.149)	(0.325)	(0.014)
F-stat first stage		13.21	

Panel data results for 1925, 1933 and 1939, 246 cities and 738 observations. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, latitude and an indicator for being in the Rhineland.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.8: Metal industry

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Public Employment					
	Ratio to Pop.	Ratio to Lab. Force	Natural Log	Ratio to Pop.	Ratio to Lab. Force	Natural Log
NSDAP share 1933 · D1939	0.130**	0.317**	0.020*	0.103*	0.273*	0.025
	(0.062)	(0.143)	(0.012)	(0.060)	(0.150)	(0.016)
Observations	843	843	843	663	663	663
Number of cities	281	281	281	221	221	221
F-stat first stage	14.46	14.46	14.46	10.37	10.37	10.37

Panel data results for 1925, 1933 and 1939. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, latitude and an indicator for being in the Rhineland. Columns 1-3 do not drop any cities that underwent size changes during the period of observation, columns 4-6 exclude all cities whose growth between either 1910 and 1925, 1925 and 1933 or 1933 and 1939 exceeded the respective mean by more than one respective standard deviation.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.9: Robustness of the IV estimate: City size

	(1)	(2)	(3)
VARIABLES		Public Employment	
	Ratio to Pop.	Ratio to LF	Natural Log
NSDAP share 1933 · D1939	0.105**	0.319**	0.023*
	(0.048)	(0.128)	(0.013)
F-stat first stage		13.15	

Panel data results for 1925 and 1939, 246 cities and 492 observations. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, latitude and an indicator for being in the Rhineland.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.10: Robustness: Dropping 1933

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Public Employment								
	Ratio to Population	Ratio to Labour Force	Natural Log	Ratio to Population	Ratio to Labour Force	Natural Log	Ratio to Population	Ratio to Labour Force	Natural Log
NSDAP share 1930 · D1939	0.158*	0.429*	0.037						
	(0.090)	(0.225)	(0.024)						
NSDAP share July 1932 · D1939				0.090**	0.235**	0.025**			
				(0.038)	(0.105)	(0.010)			
NSDAP share Nov 1932 · D1939							0.103**	0.268**	0.028**
							(0.046)	(0.129)	(0.012)
Observations	738	738	738	480	480	480	483	483	483
Number of cities	246	246	246	160	160	160	161	161	161
F-stat first stage		7.200			10.11			9.591	

Panel data results for 1925, 1933 and 1939. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, latitude and an indicator for being in the Rhineland.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.11: Robustness: Different elections

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Public empl. definition 1925		Public empl. definition 1933		Civil servants	
	Ratio to pop	natural log	Ratio to pop	natural log	Ratio to pop	natural log
NSDAP share 1933 · D1939	0.090*	0.023*	0.094*	0.023	0.118**	0.027**
	(0.051)	(0.014)	(0.054)	(0.014)	(0.053)	(0.012)
F-stat first stage			13.15			

Panel data results for 1933 and 1939, 246 cities and 492 observations. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, latitude and an indicator for being in the Rhineland.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.12: Robustness: Outcome variables

VARIABLES	Public Admin.	Teaching	Legal Counselling	Church	Entertainment
Panel A	Ratio of respective occupation group to total population				
NSDAP share 1933 · D1939	0.120*	0.012	0.004**	-0.037***	0.001
	(0.070)	(0.011)	(0.002)	(0.012)	(0.001)
Panel B	Natural Logarithm				
NSDAP share 1933 · D1939	0.031	0.010	0.004**	-0.071**	0.004
	(0.027)	(0.012)	(0.002)	(0.030)	(0.030)
F-stat first stage	12.35				

Cross-sectional results for 1939 with 246 cities. Controls are the natural logarithm of city population, longitude, latitude, an indicator for the Rhineland, an indicator for being a gau capital, the unemployment rate in 1933, and the Jewish population share in 1925. Robust standard errors in parentheses.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3.13: Outcome categories in 1939

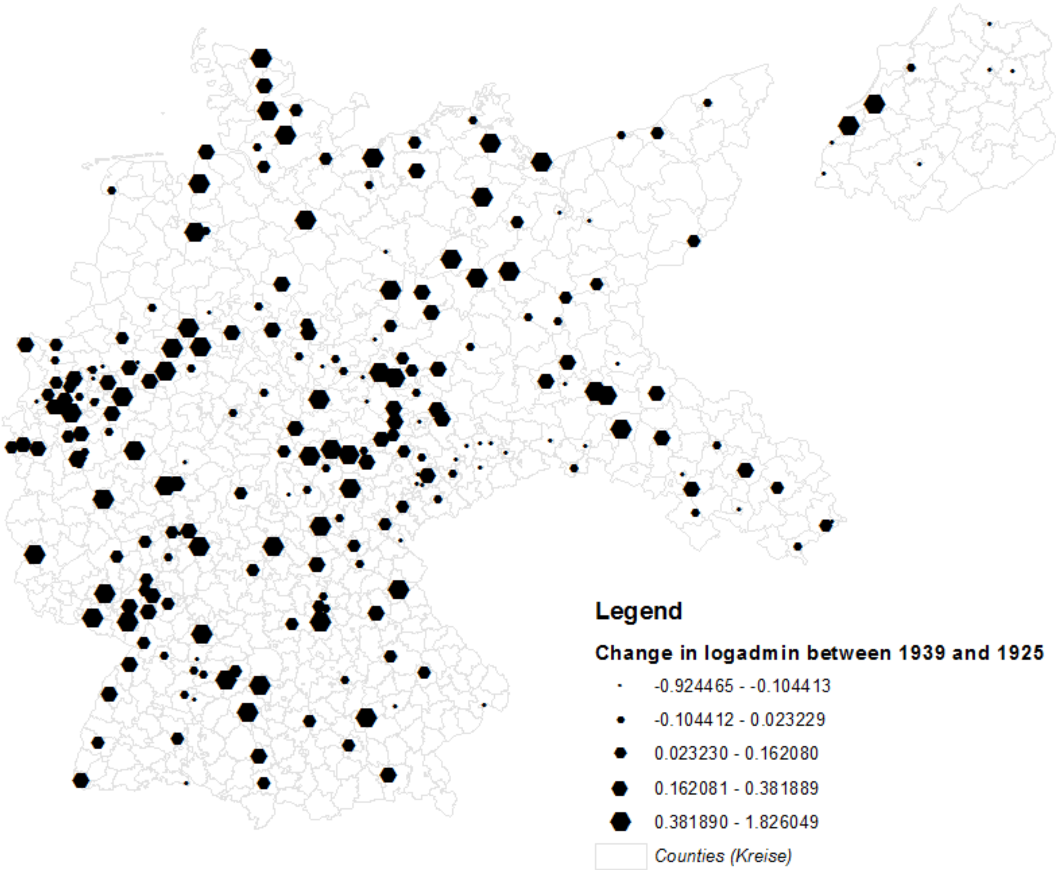

Figure 3.1: Quintiles of the change in *logadmin* between 1939 and 1925. County boundaries are based on shapefiles for 1938 Germany from MPIDR (2011).

3.7 Appendix A

The 1925 census of occupations groups the following occupations together (Occupation group "D"): Administration, Armed Forces, Church, Free professions (*Verwaltung, Heerwesen, Kirche, freie Berufe*). In 1933, some of these groups are reported separately: Occupation group 51 of the 1933 census reports the number of people working in jobs related to Administration, Armed Forces, Church, Education and others (*Verwaltung, Wehrmacht, Kirche, Bildung, Erziehung usw.*), while occupation group 54 deals with occupations connected to Theatre, Cinemas and Movie Recording, Broadcasting, Music, Sports and Showmen (*Theater, Lichtspiele und Filmaufnahme, Rundfunkwesen, Musikgewerbe, sportliche und Schaustellungsgewerbe*). Clearly, the latter group is not in the focus of my analysis. However, since these professions are contained in Occupation Group D of the 1925 census, I also included them for 1933 and added the Occupation Groups 51 and 54 of the respective census.
The 1939 census makes even finer distinctions: Occupation Group 61 deals with Administration and Armed Forces (*Berufe der öffentlichen Verwaltung und Rechtspflege, der Wehrmacht usw*), Occupation Group 62 with teaching professions and artists (*Lehr- und Bildungsberufe, künstlerische Berufe*). Church-related professions are reported in Group 63 (*Berufe der Kirche, Mönche und Nonnen*), while Group 64 contains professions related to legal counselling (*Berufe der Rechts- und Wirtschaftsberatung*). Group 68 finally contains the entertainment industry (*Berufe des Unterhaltungsgewerbes (außer Künstler)*). Again, a better measure would be to only count groups 61 and maybe 62 and 64, but due to the reporting schemes in 1925 and 1933, I added up the number of people working in groups 61-64 and 68.
Which different sectors and occupations does this aggregate measure contain? The following gives an overview of the different subcategories in the 3 censuses. Unfortunately, apart from 1939, these subcategories are not reported separately at the city level.

- 1925
  - Occupation Group D: Administration, Armed Forces, Church, Free professions (*Verwaltung, Heerwesen, Kirche, freie Berufe*)
    - W137. National, state, district and municipal administration, judicature (if in civil service position), penal system (*Reichs-, Landes-, Bezirks- und Gemeindeverwaltung, Reichtspflege (soweit in beamteter Stellung) und Strafvollzug)*
    - W138. Army and navy, army and navy administration (incl. military hospitals) (*Heer und Marine, Heeres- und Marineverwaltung (einschl. Militärlazarette)*)

        * W139. Church, divine service, mission, institutions and associations for religious purposes (*Kirche, Gottesdienst, Mission, Anstalten und Vereine für religiöse Zwecke*)
        * W140. Education, instruction, libraries, scientific institutes and art collections (*Bildung, Erziehung, Unterricht, Büchereien, wissenschaftliche Institute und Kunstsammlungen*)
        * W141. Legal counselling and representation of interests (*Rechtsberatung und Interessenvertretung*)
        * W142. Artists, private scholars, authors (*Künstler, Privatgelehrte, Schriftsteller*)
        * W143. Theatres and operas, music business (*Theater und Opernhäuser, Musikgewerbe*)
        * W144. Cinemas (*Lichtspielwesen*)
        * W145. Broadcasting (*Rundfunkwesen*)
        * W146. Sports business, horse rental, showman business (*Sportliche Gewerbe, Pferdeverleihung, Schaustellungsgewerbe*)

- 1933
    - Occupation Group 51: Administration, Armed Forces, Church, Education and others (*Verwaltung, Wehrmacht, Kirche, Bildung, Erziehung usw.*)
        * 511. National, state and municipal administration, public judicature (*Reichs-, Landes- und Gemeindeverwaltung, öffentliche Rechtspflege*)
        * 512. Wehrmacht (incl. army and navy administration, naval yard, army and navy hospitals etc.) (*Wehrmacht (einschl. Heeres- und Marineverwaltung, Marinewerft, Heeres- und Marinelazarette usw.)*)
        * 513. Church, institutions and associations for religious purposes (*Kirche, Anstalten und Vereine für religiöse Zwecke*)
        * 514. Education, instruction (*Bildung, Erziehung, Unterricht*)
        * 515. Visual arts, free literary and scientific activity (*Bildende Kunst, freie schriftstellerische und wissenschaftliche Betägigung*)
        * 516. Legal and economic counselling, representation of interests (*Rechts- und Wirtschaftsberatung, Interessenvertretung*)
        * 517. Residential business (administration incl. allocation) (*Wohngewerbe (Verwaltung einschl. Vermittlung*)
    - Occupation Group 54: Theatre, Cinemas and Movie Recording, Broadcasting, Music, Sports and Showmen (*Theater, Lichtspiele und Fil-*

maufnahme, Rundfunkwesen, Musikgewerbe, sportliche und Schaustellungsgewerbe)

- 1939
  - Occupation Group 61: Occupations in public administration and judicature, Wehrmacht etc. (*Berufe der öffentlichen Verwaltung und Rechtspflege, der Wehrmacht usw*)
  - Occupation Group 62: Teaching and instruction occupations, artistic occupations (*Lehr- und Bildungsberufe, künstlerische Berufe*)
  - Occupation Group 63: Occupations in the church, monks and nuns (*Berufe der Kirche, Mönche und Nonnen*
  - Occupation Group 64: Legal and economic counselling (*Rechts- und Wirtschaftsberatung*)
  - Occupation Group 68: Entertainment occupations (excl. artists) (*Berufe des Unterhaltungsgewerbes (außer Künstler)*)

To give a sense of what occupation group 61 in 1939 measures, the following professions were reported separately at a coarser level (state or Prussian province): Judges and attorneys (*Richter und Staatsanwälte*), trainee judges and civil servants (*Regierungs- und Gerichtsrefendare*), Bailiffs and executory officers (*Gerichtsvollzieher, Vollstreckungsbeamte*), Officers (Wehrmacht) (*Offiziere (Wehrmacht)*), Non-commissioned officers and long-serving privates (Wehrmacht) (*Unteroffiziere und langdienende Mannschaften (Wehrmacht)*), Leaders and trainee leaders of the Reich Labour Service (*Reichtsarbeitsdienstführer und -anwärter*), Police (*Polizei- und Gendarmerieoffiziere*, *Vollzugsbeamte der Ordnungspolizei (ohne Offiziere)* and *Vollzugsbeamte der Sicherheitspolizei (ohne Offiziere)*), fire brigade (without engineers and technicians) (*Feuerwehrleute (ohne Ingenieure und Techniker)*), Leaders and men of the armed units of SA and SS (*Führer der bewaffneten Einheiten der SS und der SA* and *Männer der bewaffneten Einheiten der SS und der SA*), civil servants in administration (if not included in other occupations) (*Verwaltungsbeamte (soweit nicht in anderen Berufen nachgewiesen)*), clerical assistants (also department managers) and similar employees in public administration and judicature, in the NSDAP, in the public administration of labour and the economy etc (if not included in other occupations) (*Sachbearbeiter (auch Dienststellen - und Abteilungsleiter u. ähnl. Angestl. in der öffentlichen Verwalt. u. Rechtspflege, in der NSDAP, in der öffentl. Arbeits- und Wirtschaftslenkung usw. (soweit nicht in anderen Berufen)*).

The largest of these categories are the Wehrmacht and the civil servants in public administration. In the whole of Germany, Occupation Group 61 comprised 1,074,571 members in 1939. Of those, 328,819 were in the *Wehrmacht*, and 403,019 were *Verwaltungsbeamte*. Another 86,848 people belonged to the broad category of *Sachbearbeiter (auch Dienststellen - und Abteilungsleiter u. ähnl. Angestl. in der öffentlichen Verwalt. u. Rechtspflege, in der NSDAP, in der öffentl. Arbeits- und Wirtschaftslenkung usw. (soweit nicht in anderen Berufen)*. Many of those were public administration employees that were not in the civil service (i.e. *Angestellte* as opposed to *Beamte*). Various police categories accounted for 125,222 members, and the leaders of the Reich Labour Service for another 50,767. The armed branches of SA and SS were relatively small outfits, amounting to only 29,882 or less than 3% of the whole measure.

For the regressions where I use the number of civil servants (*Beamte*) as outcome variables, I use the 1933 and 1939 census. In 1933, the respective occupational class is called "civil servants and soldiers, excluding those in senior positions"(*Beamte und Soldaten (ohne die in leitender Stellung)*.[18] In 1939, the respective category is called "civil servants"(*Beamte*).

In table 3.4, I analyse a city's social class composition, using the categories of blue collar, white collar and self employed. Blue collar is defined to contain labourers (*Arbeiter*), helping family members (*Mithelfende Familienangehörige*), and domestic workers (*Hausangestellte*). The white collar category contains the census categories of employees (*Angestellte*) and civil servants (*Beamte*), which are reported separately in 1933 and grouped together in 1925. Self-employed is a census category in itself (*Selbstständige*). In 1933, it also contains high-ranking civil servants and officers (*Beamte und Offiziere in leitender Stellung*). As described in the text, the 1925 census does not report the number of people working in each broad category, but the number of workers and their dependants, whereas the 1933 census reports the number of people working in each category. However, such uniform differences in measurement should be absorbed by the year fixed effect.

3.8 Appendix B

In order to address the problem caused by city mergers and restructurings, I analysed all cities whose growth rate between either 1910 and 1925, 1925 and 1933 or between 1933 and 1939 exceeded the respective mean growth rates by more than one standard deviation. For those cities, I examined whether they grew by 25% or more alone because of enlargements. Details about which cities or villages were added to the respective cities were obtained from Wikipedia unless stated otherwise. The names, population data and sources for the cities

[18]High-ranking civil servants and soldiers are reported with the self-employed, which is why they are missing from my measure.

are given below.
From 1925 to 1933, 25 cities exceeded the mean growth rate by more than one standard deviation. 19 of them were dropped for the following reasons.
Beuthen's population in 1925 stood at 62543. Newly added districts had a total population of 26080 in 1925 according to the Statistisches Jahrbuch Deutscher Städte 1928. Hence, Beuthen grew by 40% alone due to these acquisitions. Similarly, Bielefeld (population in 1925: 86062) received incorporations totalling a 1925 population of 27893 (Statistisches Jahrbuch Deutscher Städte 1932), representing a growth of more than 32%. Bochum (population in 1925: 211249) was enlarged through several rounds of incorporations that, according to the Statistisches Jahrbuch Deutscher Städte 1929 and 1931, totalled 156462 and meant that it was dropped from the dataset as well. Duisburg incorporated several cities and towns in 1929, including the large city of Hamborn. According to the Statistisches Jahrbuch Deutscher Städte 1930, Hamborn alone led to a population growth of more than 46%. Essen grew by 161977 people or nearly 35% relative to its baseline level of 470525 in 1925. (Statistisches Jahrbuch Deutscher Städte 1931). Gelsenkirchen incorporated Buer and other cities in 1928, leading to growth of nearly 60% (Statistisches Jahrbuch Deutscher Städte 1930) Hagen's incorporations were as large as 43900 or 44% of its 1925 population. Herne incorporated the towns of Börnig, Sodingen, Cray, Oestrich, Bladenhort and Holthausen. The 1925 census gives the following numbers for Börnig, Sodingen and Holthausen, respectively: 7979, 8198, 5942. The other villages are not listed and hence must have been smaller than 2000 inhabitants. Still, even without them, the three larger ones totalled 22119 people, which represents a 32% increase in population for Herne. Hindenburg in 1927 acquired several surrounding towns and municipalities, growing by nearly 69% (Statistisches Jahrbuch Deutscher Städte 1928). Luenen acquired Brambauer and parts of another town, Derne. Brambauer alone led to growth of around 56%. Neustrelitz, a town of 12260 inhabitants in 1925, was merged with Strelitz, thereby gaining 4687 inhabitants as of 1925, or 38%. Oberhausen incorporated several surrounding entities, totalling 84466 according to the Statistisches Jahrbuch Deutscher Städte 1931, or nearly 80% of the city's 1925 population. The towns of Lennep and Lüttringhausen (together 27826 according to the 1925 census) were added to Remscheid, making it grow by 36%. Rheine had a 1933 population of 17732. According to the homepage of the administrative district of Münster (of which the city is a part), the city acquired additional territory in 1929 that made its population grow by about 10000 inhabitants. Bad Salzemen (9998) and Frohse (2064, both numbers according to the 1925 census) were added to Schönebeck, which as a consequence grew by 56%. Solingen's 1925 population was more than doubled by the acquisition of Gräfrath, Höhscheid, Ohligs and Wald, totalling

83799 inhabitants (Census of population 1925). Several towns were incorporated into Wiesbaden, making its 1925 population of 102737 grow by 30684 or nearly 30% according to the Statistisches Jahrbuch Deutscher Städte (1928). Witten's population in 1925 stood at 45295. Annen, Stockum, Düren and parts of Langendreer and Bommern were added to this. While Düren is missing from the 1925 census list and hence must have had less than 2000 inhabitants, Annen and Stockum had 1925 populations of 17822 and 3196, respectively. Zweibrücken received the villages of Bubenhausen and Ernstweiler. Bubenhausen's population as of 1925 was 3817, or 24% of Zweibrücken's in the same year. For Ernstweiler, the census contains no population data. However, even under a very conservative assumption of only 200 inhabitants, the two acquisitions would exceed the 25% threshold, so Zweibrücken was also dropped.

Six cities were not dropped, although they experienced substantial territorial gains.

Dortmund, with a 1925 population of 321743, received additional incorporations totalling 70491 according to the Statistisches Jahrbuch Deutscher Städte 1931, or 22%. Similarly, Eschweiler received the surrounding villages of Nothberg, Hastenrath and Scherpenseel. Nothberg and Hastenrath are listed in the 1925 census as having populations of 2176 and 2187, while Scherpenseel had less than 2000 inhabitants. Even under the conservative assumption that it was exactly at this cut-off, the sum of the three gains would total only 6363, or 24% of Eschweiler's 1925 population. Ellguth-Zabrze (2205), Sosnitza (6453), Richtersdorf (3661) and Zernik (2083, all figures from the 1925 census) were made part of Gleiwitz, making its 1925 population grow by 17.5%. Heilbronn experienced substantial population growth between 1925 and 1933, but I could not find any evidence for territorial gains. Mainz acquired Bretzenheim (5692), Weisenau (6637), Ginsheim (4611), Bischofsheim (5438) and Gustavsburg (below 2000, all figures from 1925 census). Even if Gustavsburg's population had been at 2000, this would have resulted in growth of 22.5% relative to the 1925 level. Euren, Biewer, Kürenz, Olewig and a part of Pallien were made part of Trier (1925 population: 58140). The 1925 census gives the population of Euren and Kürenz as 3248 and 4268, respectively; Biewer, Pallien and Olewig are not listed and hence must have been smaller than 2000 inhabitants. However, even under the most conservative assumption that they each had exactly 2000 inhabitants, the sum of the added populations would only reach 23% of Trier's 1925 population.

Between 1933 and 1939, 15 cities exceeded the mean growth rate by more than one standard deviation. Seven of them were dropped.

In a large-scale reorganisation, the cities of Altona, Wandsbek and Harburg-Wilhelmsburg were added to Hamburg (1129307). Their population as of 1933 stood at 400818. Potsdam (1933 population: 73676) acquired several surrounding towns, including Nowawes (1933 population: 29229). Radebeul (1933 population: 12949) was merged with Kötschenbroda (1933 population: 18909). Weingarten (8385 according to the census of occupation 1933) was incorporated into Ravensburg (18930) in 1939, making the latter grow by 44%. Stolberg (17394) acquired parts of Büsbach, Eilendorf and Eschweiler, whose total is given as 12199 by the census 1933. In a curious reorganization, Rüstringen (48562 in 1933 according to the census) was added to Wilhelmshaven (1933: 28016). Zweibrücken was dropped already because of its large growth between 1925 and 1933.

Eight cities were not dropped.

For Neubrandenburg, Oranienburg and Swinemünde, I could not find any evidence of territorial gains. Cuxhaven acquired Groden (1678), Westerwisch and Süderwisch (864), Stickenbüttel (644), Duhnen (725) and Neuwerk with Scharhörn (63), totalling nearly 18% of its 1933 population of 22234. Dessau incorporated Rosslau, Jonitz and Naundorf. The latter had been excorporated just before the 1933 census and were then reincorporated in 1935. Their respective populations according to the 1933 census stood at 12845, 1721 and 527, which represents a growth of around 19% relative to Dessau's 1933 population of 78634. Landau acquired Queichsheim and Mörlheim, totalling 3013 inhabitants or 18% of Landau's 1933 population (all data from the 1933 census). Suhl (15477) acquired Heinrichs. Heinrichs' population as of 1925 was 2895, which would mean a growth of 18.7%. Even if Heinrichs experienced further growth between 1925 and 1933, it is very unlikely that it would exceed the 25% threshold, so I did not drop Suhl. Wittenberg incorporated Teuchen and Labetz in 1938. Both towns are not listed in the 1925 census and hence together cannot have exceeded 4000 inhabitants in 1925. Given Wittenberg's 1925 population of 23457, the two towns fell considerably short of the 25% threshold in 1925, and it is highly unlikely that they grew so fast as to exceed it in 1933, when Wittenberg's population stood at 24480.

Between 1910 and 1925, 25 cities exceeded the mean growth rate by more than one standard deviation. 13 of them were dropped.

Berlin experienced a massive increase in area and population due to the Greater Berlin Act of 1920. Gera (1910 population according to the census: 49276) acquired a vast number of surrounding towns and villages. Four of them alone (Debschwitz, Untermhaus, Pforten and Zwötzen) had a combined 1910 population of 23967, leading Gera to be dropped. Greiz was enlarged by the acquisition of Pohlitz, Dölau and several smaller villages. The two former alone had a combined population of 6025, enlarging Greiz's 1910 population of 23245 by more than 25%. Hirschberg with its 1910 population of 20564 acquired several smaller towns and Kunnersdorf/Cunnersdorf according to Salomon and Stein (1928), which in 1910 had a population of 5411, making the city grow by more than 25% alone. Osternburg and Eversten were added to Oldenburg, boosting that city's population by more than 66% at 1910 levels. Pirna's population in 1910 stood at 19525. Between then and 1925, several towns and villages were incorporated into it, and the incorporation of Copitz and Neundorf alone added nearly 45% of the city's 1910 population to it. Similarly, Riesa incorporated Gröba, Oberreussen and Weida. While Oberreussen had less than 2000 inhabitants in 1910, Gröba and Weida had 4471 and 2119, respectively, or 43% of Riesa's 1910 population of 15287. Waldenburg incorporated several minor districts and villages and Altwasser, which by itself increased Waldenburg's population by 88% in 1910 terms. Wattenscheid was considerably enlarged after 1926. While the Hänisch-Falter database contains data for the enlarged city in 1925, the 1910 census and 1912 election results refer to the original, small city only, which was therefore dropped. The same holds for Castrop-Rauxel, which was created in 1926 through a merger of Castrop, Rauxel and other municipalities. Bochum, Essen and Luenen were already dropped due to their enlargements between 1925 and 1933 or 1933/39.

12 Cities were not dropped.

In the case of Ahlen, Bottrop, Datteln, Gladbeck, Herten, Marienburg, Recklinghausen and Schneidemühl, I did not find any evidence for territorial acquisitions, their growth seems to have been purely organic. Dortmund acquired Deusen, Dorstfeld, Eving, Huckarde, Kemminghausen, Lindenhorst, Rahm, Wischlingen, Brackel and Wambel, of which Deusen, Kemminghausen, Rahm and Wischlingen had fewer than 2000 inhabitants in 1910. Even under the extreme assumption that they had exactly 2000 inhabitants, the total growth due to the acquisition of all 10 towns would have amounted to only 23%, so Dortmund was not dropped. Similarly, Hannover acquired Linden, but thereby growing only by 24%. Similarly, Schweinfurt incorporated Oberndorf, but this only represented a growth of around 15% at 1910 levels. The most difficult case is Regensburg. Its population in 1910 stood at 52624. Between then and 1925, it acquired Stadtamhof (4369) and Steinweg (3575) as well as 5 villages that had fewer than 2000 inhabitants in 1910. If these 5 villages had a total population of more than 5212 inhabitants, Regensburg's inorganic growth would have exceeded 25% and I would have dropped the city. However, in the respective district of Oberpfalz, the 1910 census gives the average population of all villages below 2000 inhabitants as 395, so the 5 villages combined would have had to exceed this average by more than a factor of 2.5 to reach 5212 inhabitants, which seems unlikely. I therefore decided to not drop Regensburg.

3.9 Appendix C

In this appendix, I present several additional results. In my empirical strategy, I use several control variables detailed in section 3.3.1. Control variables are not of interest per se for my study. Moreover, in the absence of a credible identification strategy for them, their regression coefficients identify mere conditional correlations without a causal interpretation. Because of this, I have omitted them from the main regression tables. In table 3.14, however, I show the basic fixed-effects regression results from table 3.6, including the coefficients for all control variables. Table 3.15 does the same for the baseline IV estimates of table 3.7.

Conditional on all the other included variables, latitude and the Jewish population share in 1925 do not appear to be correlated with public employment. Cities in the East experienced a relative decrease of public employment, while the dummy for the Rhineland is actually negative. Given the military reoccupation of the Rhineland after 1935, this is surprising, but on the other hand, the Rhineland and longitude are strongly negatively correlated, so part of this effect could potentially be captured by the longitude coefficient already. Cities with a greater unemployment rate in 1933 experience relative increases in public employment jobs in the ratio specifications (*adminpop* and *adminshare*), but not in the log specification, which indicates that such cities experienced a relative population decline. Interestingly, NSDAP gau capitals experienced declines in public employment. One explanation is that most gau capitals were already regional or state capitals before the Nazis had come to power and thus already had sizeable public employment shares before, such that they on average experienced less growth.

Panel A of table 3.15 shows the correlations of the control variables with the 1933 NSDAP vote share. The Jewish population share is positively correlated with voting for the Nazis, as is being in the East and North. Being in the Rhineland has a strong negative effect on the 1933 NSDAP vote share, which is not surprising given the region's Catholicism and the findings of Spenkuch and Tillmann (2016). Gau capitals and the 1933 unemployment rate, on the other hand, are not found to have statistically significant effects on voting for the Nazis. The negative point estimate for the unemployment rate is in line with Falter et al. 1985, who find a negative correlation between the unemployment rate and the 1933 NSDAP vote share in a sample of precincts. Turning to the IV estimates in panel B, most of the estimates are similar to those in table 3.14: Latitude turns slightly negative, while the effect of the unemployment rate is again apparently due to differential population changes and thus absent in the *logadmin* specification. Cities in the West still are found to have experienced growth in the public sector, whereas Gau capital cities are negatively associated

with public employment growth. The one major change of the IV controls relative to the OLS controls is that the Rhineland coefficients are now considerably larger in value, sometimes even positive, but generally insignificant. Again, it seems that the public employment effect of being in the Rhineland is to a large extent already captured by the longitude effect.
A further robustness check is motivated by figure 3.1, which shows a slight concentration of large public employment increases in Central Germany. Because of this, controlling for longitude or latitude in a linear form might not be sufficient. Table 3.16 shows results when additionally controlling for the square of longitude and latitude (interacted with a dummy for Post-1939). If anything, allowing for a nonlinear effect of geography strengthens my results: Point estimates increase, the first stage becomes more powerful, and the standard errors decrease slightly.
In table 3.17, I repeat the robustness check of table 3.9 for the "Placebo check" in table 3.8. Generally, the point estimates increase in absolute value, but their precision either stays relatively unchanged (columns 1-3) or decreases (columns 4-6). As a result, there is still no statistically significant effect to be seen, but the relative imprecision of the estimates makes it hard to draw any firm conclusions.
Finally, as mentioned in sections 3.3.1 and 3.4.1, the instrumental variable I use has a somewhat unusual distribution: In 191 cities, the EA did not run or did not receive any votes in 1912. In the remaining 55 cities, the mean vote share was 3.12% (standard deviation 6.54). Four cities recorded particularly high EA vote shares of more than 20%, the extreme case being Siegen with more than 37%. The spike at 0 and the presence of potentially influential variables on the right tail of the instrument's distribution may mask some heterogeneity in the effect. This is further explored in table 3.18. Focussing on *adminpop* and *logadmin* as outcome variables, I show IV results for several subsamples, and show how the IV coefficient depends on its two constituents, the reduced form (effect of the 1912 EA share on the outcome) and the first stage. Columns 1 and 5 show the results for the baseline of 246 cities, reproducing the main results from table 3.7. In Columns 2 and 6, I focus on the 55 cities that had nonzero EA vote shares in 1912. Naturally, the precision decreases when dropping 75% of the sample, and the first stage F statistic is weakened. The size of the first stage coefficient decreases, indicating that among the cities where the EA ran in 1912, the translation from EA vote shares to 1933 NSDAP vote shares is smaller. Because of this and a slight increase in the reduced form, the effect size for these cities increases. In columns 3 and 7, I keep the cities with no EA vote share, but drop the four cities with EA vote shares larger than 20. They represent around 1.5% of all cities, but more than 7% of all those with values larger than 0. Now, the reduced form decreases, while the first stage becomes less precise, but larger in absolute value, both of which leads to smaller effect

sizes. Finally, in columns 4 and 8, I drop both the four cities with large EA vote shares and all the cities with no EA voting. The first stage coefficient in this sample is very close to the baseline, but the reduced form is larger, leading to overall larger effect sizes. Overall, the cities with zero EA votes in 1912 are important for the precision of the estimate. However, they tend to decrease the reduced form and increase the coefficient of the first stage, thereby dampening the effect. The four cities with very large EA vales are quite influential in the overall regression, increasing the reduced form and decreasing the first stage and thus increasing overall estimates, as can be seen when comparing columns 1 and 5 to 3 and 7. Among cities with positive EA vote shares, those four cities are more influential for the first stage, tilting its slope closer towards 0 (columns 2/6 vs. 4/8). Finally, as columns 4 and 8 show, when focussing on a very limited sample of only cities that have positive EA shares below 20%, i.e. excluding both extremes, precision is low, but the qualitative conclusions remain the same as in the full sample.

VARIABLES	(1)	(2)	(3)
	Public Employment		
	Ratio to Pop.	Ratio to LF	Natural Log
NSDAP share 1933 · D1939	-0.007	-0.009	-0.002
	(0.011)	(0.025)	(0.003)
Jewish pop share 1925 · D1939	-0.006	0.100	0.038
	(0.090)	(0.223)	(0.026)
Unemployment rate 1933 · D1939	0.038***	0.067**	-0.007*
	(0.014)	(0.031)	(0.004)
Latitude · D1939	0.004	-0.055	0.025**
	(0.055)	(0.120)	(0.013)
Longitude · D1939	-0.107***	-0.209**	-0.024***
	(0.038)	(0.087)	(0.008)
DRhineland · D1939	-0.420*	-0.924*	-0.126**
	(0.220)	(0.483)	(0.060)
D(Gau capital)· D1939	-0.507**	-1.224**	-0.107**
	(0.208)	(0.481)	(0.052)
D1939	0.142	3.500	-0.710
	(2.808)	(6.091)	(0.632)

Panel data results for 1925, 1933 and 1939, 246 cities and 738 observations. Robust standard errors, clustered at the city level, in parentheses. All regressions also control for city fixed effects.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.14: OLS estimates

	(1)	(2)	(3)
Panel A: First stage	Dep. variable: NSDAP vote share 1933		
EA share 1912 · D1939		0.339***	
		(0.093)	
Jewish pop share 1925 · D1939		1.595**	
		(0.764)	
Unemployment rate 1933 · D1939		-0.127	
		(0.082)	
Latitude · D1939		0.771**	
		(0.330)	
Longitude · D1939		0.808***	
		(0.203)	
DRhineland · D1939		-4.269**	
		(1.649)	
D(Gau capital)· D1939		0.076	
		(1.316)	
D1939		-4.487	
		(16.370)	
F-stat first stage		13.21	
Panel B: 2SLS estimation	Public Employment		
	Ratio to Pop.	Ratio to LF	Natural Log
NSDAP share 1933 · D1939	0.098**	0.266**	0.023*
	(0.049)	(0.122)	(0.013)
Jewish pop share 1925 · D1939	-0.193	-0.389	-0.007
	(0.127)	(0.312)	(0.034)
Unemployment rate 1933 · D1939	0.054***	0.109**	-0.003
	(0.018)	(0.043)	(0.005)
Latitude · D1939	-0.090	-0.300*	0.002
	(0.076)	(0.178)	(0.020)
Longitude · D1939	-0.179***	-0.396***	-0.042***
	(0.054)	(0.135)	(0.013)
DRhineland · D1939	0.038	0.272	-0.014
	(0.357)	(0.842)	(0.097)
D(Gau capital)· D1939	-0.541**	-1.313**	-0.115**
	(0.220)	(0.536)	(0.049)
D1939	1.033	5.829	-0.493
	(3.446)	(7.902)	(0.814)

Panel data results for 1925, 1933 and 1939, 246 cities and 738 observations. Robust standard errors, clustered at the city level, in parentheses. All regressions also control for city fixed effects.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.15: IV estimates

VARIABLES	(1)	(2)	(3)
		Public Employment	
	Ratio to Pop.	Ratio to LF	Natural Log
NSDAP share 1933 · D1939	0.097**	0.260**	0.026**
	(0.043)	(0.107)	(0.012)
F-stat first stage		16.03	

Panel data results for 1925, 1933 and 1939, 246 cities and 738 observations. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, the square of longitude, latitude, the square of latitude, and an indicator for being in the Rhineland.

*** p<0.01, ** p<0.05, * p<0.1

Table 3.16: IV estimates when controlling for second-order polynomials in longitude and latitude

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Metal industry Employment					
	Ratio to Pop.	Ratio to LF	Natural Log	Ratio to Pop.	Ratio to LF	Natural Log
NSDAP share 1933 · D1939	-0.204	-0.439	-0.017	-0.236	-0.494	-0.018
	(0.137)	(0.302)	(0.015)	(0.186)	(0.406)	(0.016)
Observations	843	843	843	663	663	663
Number of cities	281	281	281	221	221	221
F-stat first stage	14.46	14.46	14.46	10.37	10.37	10.37

Panel data results for 1925, 1933 and 1939. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population in 1925, the unemployment rate in 1933, longitude, latitude and an indicator for being in the Rhineland. Columns 1-3 do not drop any cities that underwent size changes during the period of observation, columns 4-6 exclude all cities whose growth between either 1910 and 1925, 1925 and 1933 or 1933 and 1939 exceeded the respective mean by more than one respective standard deviation.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3.17: Robustness of the estimates for the metal industry: Robustness to city size

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: 2SLS estimates	Ratio of public employment to population				Natural logarithm of public employment			
NSDAP share 1933 · D1939	0.098**	0.257*	0.048	0.175	0.023*	0.086**	0.010	0.064**
	(0.049)	(0.147)	(0.049)	(0.118)	(0.013)	(0.040)	(0.012)	(0.033)
Panel B: Reduced Form	Ratio of public employment to population				Natural logarithm of public employment			
EA share 1912 · D1939	0.033**	0.048***	0.029	0.065*	0.008*	0.016***	0.006	0.024***
	(0.014)	(0.016)	(0.028)	(0.036)	(0.004)	(0.005)	(0.007)	(0.009)
Panel C: First stage	NSDAP vote share 1933							
EA share 1912 · D1939	0.339***	0.187**	0.606***	0.373**	0.339***	0.187**	0.606***	0.373**
	(0.093)	(0.088)	(0.143)	(0.165)	(0.093)	(0.088)	(0.143)	(0.165)
F-stat first stage	13.21	4.564	17.87	5.123	13.21	4.564	17.87	5.123
Observations	738	165	726	153	738	165	726	153
Number of cities	246	55	242	51	246	55	242	51
Sample	All cities	%EA > 0	%EA < 20	0<%EA < 20	All cities	%EA > 0	%EA < 20	0<%EA < 20

Panel data results for 1925, 1933 and 1939. Robust standard errors, clustered at the city level, in parentheses. All regressions control for city fixed effects, an indicator for 1939 as well as interactions of an indicator for 1939 with an indicator for being a gau capital, the Jewish population share in 1925, the unemployment rate in 1933, longitude, latitude, and an indicator for being in the Rhineland.
*** p<0.01, ** p<0.05, * p<0.1

Table 3.18: Subsamples depending on the IV value

Conclusion

This thesis used three different historical settings to shed light on economic questions. In the first chapter, we studied the causal relationship between trade and economic growth. In order to avoid problems like reverse causality, path dependence and confounding factors such as institutions, we made use of one of the first major trade expansions in human history: the systematic crossing of open seas in the Mediterranean at the time of the Phoenicians. We constructed a measure of maritime connectedness of points along the Mediterranean coast and argued that this measure of connectedness is a valid instrument for trade. We then estimated a reduced form relationship between our instrument and economic development, which we measured by way of the presence of archaeological sites. We found that a clear positive relationship between geographic connectedness and local archaeologic site density after 750BC, when traversing the open sea had become regular. We also found evidence for some importance of sea trade already during the Bronze Age, but our results in this respect were weaker and less consistent across specifications. Our results also highlighted that the positive effects of connectedness for settlement density persisted during the Hellenic and Roman periods. Further evidence for a causal effect of trade on growth was provided by extending our strategy to the world scale and finding a positive relationship between maritime connectedness and population density in the year 1AD.

In the second chapter, we studied how a male-biased labour demand shock affects female labour market involvement. Again, we made use of variation provided by geography and history. During the first half of the 20th century, many large oil fields were discovered in the American South and Southwest. Since the oil extraction industry is heavily male-dominated, we treated these discoveries as male-biased shifters of labour demand. Combining data on the location and discovery dates of such oil fields with labour market data from the US census allowed us to implement a difference-in-differences research design at the county level. Doing so, we found a zero net effect of oil discoveries on female labour force participation due to two opposing channels. Oil discoveries indeed bring about increased demand for male labour and therefore higher male wages. This increases the supply of "marriageable "men in the local economy, and we accordingly found an increase in young women's propensity to marry, which depresses their labour supply. However, oil discoveries also increase demand for services and thus female labour. Overall, female labour force participation therefore remains unchanged. While previous research (e.g. Ross 2008) had shown a negative correlation between resource wealth and female labour force participation, our results demonstrated that such a negative association can be offset by service sector growth.

The third chapter further contributed to the literature on the value of political connections. I analysed whether the German National Socialists, once in power, preferentially allocated public employment jobs to cities where they had obtained higher vote shares in 1933. I argued that in a regression of public employment on the 1933 Nazi vote share, the latter is likely to be endogenous. I then proposed to address this endogeneity by way of an instrumental variable strategy, using the 1912 vote share of another antisemitic party as instrument for the 1933 Nazi vote share. Doing so, I found a significant positive effect of Nazi votes on public employment.

Bibliography

Abelshauser, Werner. 1998. Germany: Guns, Butter and Economic Miracles. In: Mark Harrison (ed.). *The Economics of World War II*. Cambridge: Cambridge University Press.

Abulafia, David. 2011. *The Great Sea: A Human History of the Mediterranean.* London: Penguin; New York: Oxford University Press.

Acemoglu, Daron, David H. Autor and David Lyle. 2004. Women, War, and Wages: The Effect of Female Labour Supply on the Wage Structure at Midcentury. *Journal of Political Economy* 112: 497-551.

Acemoglu, Daron, Simon Johnson, and James Robinson. 2005. The Rise of Europe: Atlantic Trade, Institutional Change, and Economic Growth. *American Economic Review* 95: 546-579.

Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa and Ekaterina Zhuravskaya. 2015. Radio and the Rise of the Nazis in Prewar Germany. *Quarterly Journal of Economics* 130: 1885-1939

Ager, Philipp, Markus Brueckner, and Benedikt Herz. Forthcoming. The Boll Weevil Plague and Its Effect on the Southern Agricultural Sector, 18891929. *Explorations in Economic History.*

Agouridis, Christos. 1997. Sea Routes and Navigation in the Third Millennium Aegean. *Oxford Journal of Archaeology* 16: 1-24.

Akbulut, Rahsan. 2011. Sectoral Changes and the Increase in Women's Labour Force Participation. *Macroeconomic Dynamics* 15:240-64.

Albanesi, Stefania, and Claudia Olivetti. 2016. Gender Roles and Medical Progress. *Journal of Political Economy* 124: 650-695.

Algaze, Guillermo. 2008. *Ancient Mesopotamia at the Dawn of Civilization. The Evolution of an Urban Landscape.* Chicago: Chicago University Press.

Alesina, Alberto F., Paola Giuliano, and Nathan Nunn. 2013. On the Origins of Gender Roles: Women and the Plough. *Quarterly Journal of Economics* 128: 469-530.

Allcott, Hunt and Daniel Keniston. 2014. Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America. NBER Working Paper 20508.

Anderson, Gary M. and Robert D. Tollison. 1991. Congressional Influence and Patterns of New Deal Spending, 1933-1939. *Journal of Law and Economics* 34: 161-175.

Angrist, Joshua and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics*. Princeton and Oxford: Princeton University Press.

Aragon, Fernando M., and Juan P. Rud. 2013. Natural Resources and Local Communities: Evidence from a Peruvian Gold Mine. *American Economic Journal: Economic Policy* 5: 1-25.

Aragon, Fernado M., Juan P. Rud and Gerhard Toews. 2016. Resource Shocks, Employment, and Gender: Evidence from the Collapse of the UK Coal Industry. Working Paper.

Arnaud, Pascal. 2007. Diocletian's Prices Edict: The Prices of Seaborne Transport and the Average Duration of Maritime Travel. *Journal of Roman Archaeology* 20: 321-336.

Ashenfelter, Orley and James Heckman. 1974. The Estimation of Income and Substitution Effects in a Model of Family Labour Supply. *Econometrica* 42:73-85.

Ashraf, Quamrul, and Oded Galor. 2011a. Cultural Diversity, Geographical Isolation, and the Origin of the Wealth of Nations. NBER Working Paper 17640.

Ashraf, Quamrul, and Oded Galor. 2011b. Dynamics and Stagnation in the Malthusian Epoch. *American Economic Review* 101: 2003-2041.

Attanasio, Orazio, Hamish Low, and Virginia Sanchez-Marcos. 2008. Explaining Changes in Female Labour Supply in a Life-Cycle Model. *American Economic Review* 98: 1517-42.

Ballard, Robert D. Lawrence E. Stager, Daniel Master, Dana Yoerger, David Mindell, Louis L. Whitcomb, Hanumant Singh, and Dennis Piechota. 2002. Iron Age Shipwrecks in Deep Water off Ashkelon, Israel. *American Journal of Archaeology* 106: 151-168.

Bagnall, Roger et al. (eds.) 2014. *Pleiades: A Gazetteer of Past Places*, `http://pleiades.stoa.org`

Barkai, Avraham. 1988. *Das Wirtschaftssystem des Nationalsozialismus: der Historische und Ideologische Hintergrund*. Extended edition. Frankfurt am Main: Fischer.

Becker, Gary S. 1981. *A Treatise on the Family* Cambridge and London: Harvard University Press.

Bergmann, Werner. 2012. Christlich-soziale Partei (Deutschland). In: Wolfgang Benz (ed.). *Handbuch des Antisemitismus, Band 5: Organisationen, Institutionen, Bewegungen* Berlin: De Gruyter Saur.

Bester, C. Alan, Timothy G. Conley, and Christian B. Hansen. 2011. Inference with Dependent Data Using Cluster Covariance Estimators. *Journal of Econometrics* 165: 137-151.

Black, Dan, Terra McKinnish, and Seth Sanders. 2005. The Economic Impact of the Coal Boom and Bust. *The Economic Journal* 115: 449-76.

Blaich, Fritz. 1970. Möglichkeiten und Grenzen Kommunaler Wirtschaftspolitik während der Weltwirtschaftskrise 1929-1932. *Archiv für Kommunalwissenschaften* 9: 92-108.

Bleakley, Hoyt, and Jeffrey Lin. 2012. Portage and Path Dependence. *The Quarterly Journal of Economics* 127: 587-644.

Bosker, Maarten, and Eltjo Buringh. 2017. City Seeds: Geography and the Origins of the European City System. *Journal of Urban Economics* 98:139-157.

Boustan, Leah Platt and William J. Collins. 2013. The Origin and Persistence of Black-White Differences in Women's Labour Force Participation In: Leah Platt Boustan, Carola Frydman, Robert A. Margo (eds). *Human Capital in History: The American Record*. Chicago: University of Chicago Press.

Braudel, Fernand. 2001. *The Mediterranean in the Ancient World*. London: Penguin Books.

Broodbank, Cyprian. 2006. The Origins and Early Development of Mediterranean Maritime Activity. *Journal of Mediterranean Archaeology* 19.2: 199-230.

Broodbank, Cyprian. 2013. *Making of the Middle Sea*. London: Thames and Hudson Limited.

Buckley, Patrick D. 2003. Essays in Applied Microeconomics: Signaling, Marriage Hazard Rates, and Ability Sorting. PhD diss., Department of Economics, Massachusetts Institute of Technology.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria and Gerad Padró i Miquel. 2015. The Value of Democracy: Evidence from Road Building in Kenya. *American Economic Review* 105: 1817-1851.

Carrington, William J. 1996. The Alaskan Labour Market during the Pipeline Era. *Journal of Political Economy* 104: 186-218.

Caselli, Francesco, and John W. Coleman. 2001. The U.S. Structural Transformation and Regional Convergence: A Reinterpretation. *Journal of Political Economy* 109: 584-616.

Caselli, Francesco, and Guy Michaels. 2013. Do Oil Windfalls Improve Living Standards? Evidence from Brazil. *American Economic Journal: Applied Economics* 5: 208-38.

Comin, Diego, William Easterly, and Erick Gong. 2010. Was the Wealth of Nations Determined in 1000 BC? *American Economic Journal: Macroeconomics* 2: 65-97.

Corden, Max W., and Peter J. Neary. 1982. Booming Sector and De-industrialisation in a Small Open Economy. *The Economic Journal* 92: 825-48.

Cox, Garz W. and Matthew D. McCubbins. 1986. Electoral Politics as a Redistributive Game. *Journal of Politics* 48: 370-389.

Cucchi, Thomas, Jean Denis Vigne, and Jean Christophe Auffray. 2005. First Occurrence of the House Mouse (Mus musculus domesticus Schwarz & Schwarz, 1943) in the Western Mediterranean: a Zooarchaeological Revision of Subfossil Occurrences. *Biological Journal of the Linnean Society* 84.3: 429-445.

Cunliffe, Barry. 2008. *Europe Between the Oceans. 9000 BC - AD 1000.* New Haven: Yale University Press.

Das Buch der Deutschen Gaue: Fünf Jahre Nationalsozialistische Aufbauleistung. 1938. Bayreuth: Gauverlag Bayerische Ostmark.

Davis, Donald, and David Weinstein. 2002. Bones, Bombs, and Break Points: The Geography of Economic Activity. *American Economic Review* 92: 1269-1289.

Deist, Wilhelm. 2003. The Rearmament of the Wehrmacht. In: *Germany and the Second World War, Vol. 1: The Build-Up of German Aggression.* Oxford: Oxford University Press.

Deutscher Städtetag (ed.) 1928, 1929, 1930, 1931, 1932. *Statistisches Jahrbuch Deutscher Städte: Amtliche Veröffentlichung des Deutschen Städtetages.* Jena: Fischer.

Diamond, Jared M. 1997. *Guns, Germs, and Steel: The Fates of Human Societies* . New York: W. W. Norton.

Dixon, John, Johnston R. Cann, and Colin Renfrew. 1965. Obsidian in the Aegean. *The Annual of the British School at Athens* 60: 225-247.

Dixon, J. E., J. R. Cann, and Colin Renfrew. 1968. Obsidian and the Origins of Trade. *Scientific American* 218.3: 38-46.

Do, Quy-Toan, Andrei A. Levchenko, and Claudio Raddatz. 2016. Comparative Advantage, International Trade, and Fertility.*Journal of Development Economics* 119: 48-66

Do, Quoc-Anh, Kieu-Trang Nguyen and Anh N. Tran. Forthcoming. One Mandarin Benefits the Whole Clan: Hometown Favoritism in an Authoriatarian Regime. *American Economic Journal: Applied Economics.*

Donaldson, Dave. Forthcoming. Railroads of the Raj: Estimating the Impact of Transportation Infrastructure. *American Economic Review.*

Donaldson, Dave, and Richard Hornbeck. 2016. Railroads and American Economic Growth: a "Market Access" Approach. *Quarterly Journal of Economics* 131: 799-858.

Erdmann, Jürgen. 1969. *Coburg, Bayern und das Reich 1918-1923.* Coburg: Ro ßteutscher.

Faccio, Mara. 2006. Politically Connected Firms. *American Economic Review* 96: 369-386.

Faccio, Mara, Ronald W. Masulis, and John J. McConnell. 2006. Political Connections and Corporate Bailouts. *Journal of Finance* 61: 2597-2635.

Falter, Jürgen W. 1991. *Hitlers Wähler.* München: Beck.

Falter, Jürgen W. 2013. Die "Märzgefallenen" von 1933. Neue Forschungsergebnisse zum Sozialen Wandel Innerhalb der NSDAP-Mitgliedschaft während der Machtergreifungsphase [1998]. *Historical Social Research*: 25 (Supplement): 280-302.

Falter, Jürgen W. and Dirk Hänisch. 1990. *Wahl- und Sozialdaten der Kreise und Gemeinden des Deutschen Reiches von 1920 bis 1933.* GESIS Datenarchiv, Köln, ZA8013 Datenfile Version 1.0.0 doi: 10.4232/1.8013.

Falter, Jürgen W., Thomas Lindenberger, and Siegfried Schumann. 1986. *Wahlen und Abstimmungen in der Weimarer Republik.* München: Beck.

Falter, Jürgen W., Andreas Link, Jan-Bernd Lohmöller, Johann de Rijke, and Siegfried Schumann. 1983. Arbeitslosigkeit und Nationalsozialismus- Eine Empirische Analyse des Beitrages der Massenerwerbslosigkeit zu den Wahlerfolgen der NSDAP 1932 und 1933. *Kölner Zeitschrift für Soziologie und Sozialpsychologie* 35: 525-554.

Falter, Jürgen W., Andreas Link, Jan-Bernd Lohmöller, and Johann de Rijke. 1985. Hat Arbeitslosigkeit Tatsächlich den Aufstieg des Nationalsozialismnus bewirkt? (Did Unemployment Really Cause the Rise of National Socialism?). *Jahrbücher für Nationalökonomie und Statistik* 200: 121-136.

FAO/IIASA, 2010. Global Agro-ecological Zones (GAEZ v3.0). FAO, Rome, Italy and IIASA, Laxenburg, Austria

Ferguson, Thomas and Hans-Joachim Voth. 2008. Betting on Hitler- The Value of Political Connections in Nazi Germany. *Quarterly Journal of Economics* 123: 101-137.

Fernandez, Raquel. 2011. Does Culture Matter?. In *Handbook of Social Economics*, vol. 1A, ed. Jess Benhabib, Matthew O. Jackson, and Alberto Bisin. Amsterdam: North Holland.

Fernandez, Raquel. 2013. Cultural Change as Learning: The Evolution of Female Labour Force Participation over a Century. *American Economic Review* 103: 472-500.

Fetzer, Thiemo R. 2015. Fracking Growth. CEP Discussion Paper No. 1278, London School of Economics.

Feyrer, James. 2009. Trade and Income–Exploiting Time Series in Geography. NBER Working Paper 14910.

Fisman, Raymond. 2001. Estimating the Value of Political Connections. *American Economic Review* 91: 1095-1102.

Frankel, Jeffrey A., and David Romer. 1999. Does Trade Cause Growth? *American Economic Review* 89: 379-399.

Frey, Bruno S. and Hannelore Weck. 1981. Hat Arbeitslosigkeit den Aufstieg des Nationalsozialismus Bewirkt? (Did Unemployment Lead to the Rise of National Socialism?) *Jahrbücher für Nationalökonomie und Statistik* 196: 1-31.

Furchtgott-Roth, Diana, and Andrew Gray. 2013. The Economic Effects of Hydrofracturing on Local Economies: A Comparison of New York and Pennsylvania. Growth and Prosperity Report no.1, Manhattan Institute for Policy Research, New York.

Gaddis, Isis, and Janneke Pieters. 2017. The Gendered Labour Market Impacts of Trade Liberalization: Evidence from Brazil. *Journal of Human Resources* 52: 457-489.

Gagliarducci, Stefano and Marco Manacorda. 2016. Politics in the Family: Nepotism and the Hiring Decisions of Italian Firms. Working Paper.

Galor, Oded, and Ömer Özak. 2016. The Agricultural Origins of Time Preference. *American Economic Review* 106: 3064-3103.

Goldin, Claudia. 1977. Female Labour Force Participation: The Origin of Black and White Differences, 1870 and 1880. *Journal of Economic History* 37: 87-108.

Goldin, Claudia. 1990. *Understanding the gender wage gap: An Economic History of American Women.* New York: Oxford University Press.

Goldin, Claudia. 2006. The Quiet Revolution That Transformed Women's Employment, Education, and Family. *American Economic Review, Papers and Proceedings* 96: 1-21.

Goldin, Claudian and Lawrence F. Katz. 2002. The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions. *Journal of Political Economy* 100: 730-70.

Golecki, Anton. 1986. Das Kabinett von Schleicher. In *Akten der Reichskanzlei Weimarer Republik.* Boppard am Rhein: Harald Bodt Verlag.

Gräfe, Thomas: Deutschsoziale Partei. In: Wolfgang Benz (ed.). *Handbuch des Antisemitismus, Band 5: Organisationen, Institutionen, Bewegungen* Berlin: De Gruyter Saur.

Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. 2005. Engines of Liberation. *Review of Economic Studies* 72: 109-33.

Hayward, Nicholas F. and David S. Morris. 1988. *The First Nazi Town.* Aldershot: Avebury.

Herbert, Ulrich. 2000. Extermination Policy: New Answers and Questions about the History of the 'Holocaust' in German Historiography. In: Ulrich Herbert (ed.). *National Socialist Extermination Policies: Contemporary German Perspectives and Controversies.* Oxford: Berghahn.

Hibbs, Douglas A. and Ola Olsson. 2004. Geography, Biogeography, and Why Some Countries Are Rich and Others Are Poor. *Prodeedings of the National Academy of Sciences* 101: 3715-3720.

Hodler, Roland and Paul A. Raschky. 2014. Regional Favoritism. *Quarterly Journal of Economics* 129: 995-1033.

Horden, Peregrine, and Nicholas Purcell. 2000. *The Corrupting Sea: a Study of Mediterranean History.* Oxford: Wiley-Blackwell.

Hsieh, Chang-Tai, Edward Miguel, Daniel Ortete, and Francisco Rodriguez. 2011. The Price of Political Opposition: Evidence from Venezuela's Maisanta. *American Economic Journal: Applied Economics* 3: 196-214.

Hubatsch, W. and T. Klein (eds.) 1975 ff. *Grundriß der Deutschen Verwaltungsgeschichte.* Marburg.

Humphreys, Macartan, Jeffrey D. Sachs, and Joseph E. Stiglitz. eds. 2007. *Escaping the Resource Curse.* New York, NY: Columbia University Press.

Jayachandran, Seema. 2006. The Jeffords Effect. *Journal of Law and Economics* 49: 397-425.

Jelnov, Pavel. 2016. The Marriage Age U-shape: What Can we Learn from Energy Crises. Working Paper.

Johnson, Simon and Todd Mitton. 2003. Cronyism and Capital Controls: Evidence from Malaysia. *Journal of Financial Economics* 67: 351-382

Jones, Larry E., Rodolfo E. Manuelli, and Ellen R. McGrattan. 2015. Why Are Married Women Working so Much? *Journal of Demographic Economics* 81: 75-114.

Kang, Alice. 2009. Studying Oil, Islam, and Women as if Political Institutions Mattered. *Politics & Gender* 5: 560-68.

Kershaw, Ian. 1999. *Hitler 1889-1936: Hubris.* London: Penguin.

Khwaja, Asim Ijaz and Atif Mian. 2005. Do Lenders Favor Politically Connected Firms? Rent Provision in an Emerging Financial Market. *Quarterly Journal of Economics* 120: 1371-1411.

King, Gary, Ori Rosen, Martin Tanner, and Alexander F. Wagner. 2008. Ordinary Economic Voting Behaviour in the Extraordinary Election of Adolf Hitler. *Journal of Economic History* 68: 951-996.

Klasen, Stephan, and Janneke Pieters. 2015. What Explains the Stagnation of Female Labour Force Participation in Urban India? *World Bank Economic Review* 29: 449-478.

Knappett, Carl, Tim Evans, and Ray Rivers. 2008. Modelling Maritime Interaction in the Aegean Bronze Age. *Antiquity* 82.318: 1009-1024.

Koenig, Christoph. 2015. Loose Cannons- War Veterans and the Erosion of Democracy iin Weimar Germany. Working Paper.

Kolb, Eberhard. 2005. *The Weimar Republic*, translated 2nd. ed. New York: Routledge.

Kotsadam, Andreas, and Anja Tolonen. 2016. African Mining, Gender, and Local Employment. *World Development*: 325-339.

Lange, Fabian, Alan L. Olmstead and Paul W. Rhode. 2009. The Impact of the Boll Weevil, 1892-1932. *Journal of Economic History*, 69: 685-717.

Levitt, Steven D. and James M. Snyder Jr. 1995. Political Parties and the Distribution of Federal Outlays. *American Journal of Political Science* 39: 958-980.

Lindbeck, Assar and Jörgen W. Weibull. 1987. Balanced-Budget Redistribution as the Outcome of Political Competition. *Public Choice* 52: 273-297

Litschig, Stephan and Kevin Morrison. The Impact of Intergovernmental Transfers on Education Outcomes and Poverty Reduction. *American Economic Journal: Applied Economics* 5: 206-240.

Manacorda, Marco, Edward Andrew Miguel, and Andrea Vigorito. 2011. Government Transfers and Political Support. *American Economic Journal: Applied Economics* 3: 1-28.

Marchand, Joseph. 2012. Local Labour Market Impacts of Energy Boom-Bust-Boom in Western Canada. *Journal of Urban Economics* 71: 165-74.

Martinez, Luis R. 2016. Sources of Revenue and Government Performance: Evidence from Colombia. Working Paper.

McEvedy, Colin 1967. *The Penguin Atlas of Ancient History.* Hamondsworth: Penguin Books Ltd.

McEvedy, Colin, and Richard Jones. 1978. *Atlas of World Population History.* Hamondsworth: Penguin Books Ltd.

Michaels, Guy. 2011. The Long Term Consequences of Resource-Based Specialisation. *The Economic Journal* 121: 31-57.

Michaels, Guy, and Ferdinand Rauch. Forthcoming. Resetting the Urban Network: 117-2012. *Economic Journal.*

Minnesota Population Center. *National Historical Geographic Information System*: Version 11.0 [Database]. Minneapolis: University of Minnesota. 2016. http://doi.org/10.18128/D050.V11.0.

MPIDR [Max Planck Institute for Demographic Research] and CGG [Chair for Geodesy and Geoinformatics, University of Rostock]. 2011. *MPIDR Population History GIS Collection* (partly based on Hubatsch and Klein 1975 ff.). Rostock.

Bezirksregierung Münster. 2013. *Stadt Rheine- Kurzinformation.* http://www.bezreg-muenster.de, retrieved April 2, 2013.

Ngai, Rachel L., and Barbara Petrongolo. Forthcoming. Gender Gaps and the Rise of the Service Economy. *American Economic Journal: Macroeconomics.*

Norris, Pippa. 2010. Perhaps Petroleum Perpetuates Patriarchy? A Response and Critique to Ross. Unpublished manuscript. Harvard Kennedy School of Government.

Nöth, Stefan. 2006. Die Stadtentwicklung Coburgs seit 1920. In: Peter Morsbach and Otto Titz (eds.) *Stadt Coburg.* München: Karl M. Lipp.

Nunn, Nathan, and Diego Puga. 2012. Ruggedness: The Blessing of Bad Geography in Africa. *Review of Economics and Statistics* 94: 20-36.

Olivetti, Claudia. 2013. "The Female Labour Force and Long-run Development: The American Experience in Comparative Perspective", Working Paper no. 19131, National Bureau of Economic Research, Cambridge, MA.

Pascali, Luigi. Forthcoming. The Wind of Change: Maritime Technology, Trade and Economic Development. *American Economic Review.*

Pei, Zhuan, Jörn-Steffen Pischke and Hannes Schwandt. 2017. Poorly Measured Confounders are More Useful on the Left Than on the Right. NBER Working Paper 23232.

Penn Energy Research 2000. Oil and Gas Journal Data Book. Tulsa, OK: Pennwell Publishing Co.

Penn Energy Research 2006. Oil and Gas Journal Data Book. Tulsa, OK: Pennwell Publishing Co.

Papyrakis, Elissaios, and Reyer Gerlagh. 2007. Resource Abundance and Economic Growth in the United States. *European Economic Review* 51: 1011-39.

Pratt, Joseph A. 1980. *The Growth of a Refining Region.* Greenwich, Connecticut: Jai Press.

Redding, Stephen J., and Daniel M. Sturm. 2008. The Costs of Remoteness: Evidence from German Division and Reunification. *American Economic Review* 98: 1766-97.

Redding, Stephen, and Anthony J. Venables. 2004. Economic Geography and International Inequality. *Journal of International Economics* 62: 53-82.

van Riel, Arthur and Arthur Schram. Weimar Economic Decline, Nazi Germany Recovery, and the Stabilization of Political Dictatorship. *Journal of Economic History* 53: 71-105.

Ritschl, Albrecht. 2002. Deficit Spending in the Nazi Recovery, 1933-1938. *Journal of the Japanese and International Ecoomies* 16: 559-582.

Ritschl, Albrecht. 2003. Hat das Dritte Reich Wirklich eine Ordentliche Beschäftigungspolitik Betrieben? *Jahrbuch für Wirtschaftsgeschichte* 2003/I: 125-140.

Ritschl, Albrecht. 2013. "Reparations, Deficits, and Debt Default: the Great Depression in Germany". In: Nicholas Crafts and Peter Fearon (eds.). *The Great Depression of the 1930s.* Oxford: Oxford University Press.

Ross, Michael L. 2008. Oil, Islam and Women. *American Political Science Review* 102: 107-23.

Ross, Michael L. 2012. *The Oil Curse: How Petroleum Wealth Shapes the Development of Nations.* Princeton, NJ: Princeton University Press.

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. 2015. Integrated Public Use Microdata Series: Version 6.0 [dataset]. Minneapolis: University of Minnesota. http://doi.org/10.18128/D010.V6.0.

Sachs, Jeffrey D. and Andrew M. Warner. 1999. The Big Rush, Natural Resource Booms And Growth. *Journal of Development Economics* 59: 43-76.

Sachs, Jeffrey D. and Andrew M. Warner. 2001. The Curse of Natural Resources. *European Economic Review* 45: 827-38.

Salomon, Bruno and Erwin Stein. 1928. *Die Riesengebirgskreise.* Berlin: Deutscher Kommunalverlag.

Sandner, Harald. 2000. *Coburg im 20. Jahrhundert.* Coburg: Neue Presse.

Satyanath, Shanker, Nico Voigtländer and Hans-Joachim Voth. Forthcoming. Bowling for Fascism: Social Capital and the Rise of Nazi Party. *Journal of Political Economy.*

Schiller, Karl. 1936. *Arbeitsbeschaffung und Finanzordnung in Deutschland.* Berlin: Junker und Dünnhaupt.

Schlösser, Susanne. 2003. Die Heilbronner NSDAP und Ihre "Führer". Eine Bestandsaufnahme zur nationalsozialistischen Personalpolitik auf Lokaler Ebene und Ihren Auswirkungen "vor Ort". In Christhard Schrank and Peter Wanner (eds.). *heilbronnica 2- Beiträge zur Stadtgeschichte*. Quellen und Forschungen zur Geschichte der Stadt Heilbronn 15.

Schmitz-Berning, Cornelia. 1998. *Vokabular des Nationalsozialismus*. Berlin: deGruyter.

Schoenbaum, David. 1997. *Hitler's Social Revolution*. New York, London: W.W. Norton (Reprint of original work published in 1966).

Sherratt, Susan, and Andrew Sherratt. 1993. The Growth of the Mediterranean Economy in the Early First Millennium BC. *World Achaeology* 24: 361-378.

Silverman, Dan P. 1998. *Hitler's Economy: Nazi Work Creation Programs, 1933-1936*. Cambridge: Harvard University Press.

Spenkuch, Jörg and Philipp Tillmann. 2016. Special Interests at the Ballot Box? Religion and the Electoral Success of the Nazis. Working Paper.

Stachura, Peter D. 1978. Der Kritische Wendepunkt? Die NSDAP und die Reichtsagswahlen vom 29. Mai 1928. *Vierteljahreshefte für Zeitgeschichte* 26: 66-99.

Stadt Lübeck (Pressemitteilung). 2012. *1. April 1937: Der Staat Lübeck Endet*, http://www.luebeck.de/aktuelles/presse/pressedienstarchiv, retrieved January 15, 2012.

Statistisches Reichsamt (ed.) 1913. *Statistik des Deutschen Reiches: Die Reichtstagswahlen von 1912*, vol. 250 (Neue Folge). Berlin.

Statistisches Reichsamt (ed.) 1915. *Statistik des Deutschen Reiches: Die Volkszählung im Deutschen Reiche am 1. Dezember 1919*, vol. 240 (Neue Folge). Berlin.

Statistisches Reichsamt (ed.) 1927-28. *Statistik des Deutschen Reiches: Die Berufliche und Soziale Gliederung der Bevölkerung in den Ländern und Landesteilen 1925*, vol. 403-405 (Neue Folge). Berlin.

Statistisches Reichsamt (ed.) 1928. *Statistik des Deutschen Reiches: Die Bevölkerung des Deutschen Reiches nach den Ergebnissen der Volkszählung 1925*, vol. 401 (Neue Folge). Berlin.

Statistisches Reichsamt (ed.) 1934. *Statistik des Deutschen Reiches: Amtliches Gemeindeverzeichnis für das Deutsche Reich*, vol. 450 (Neue Folge). Berlin.

Statistisches Reichsamt (ed.) 1935-1936. *Statistik des Deutschen Reiches: Die Berufliche und Soziale Gliederung der Bevölkerung des Deutschen Reichs 1933*, vol. 453-56 (Neue Folge). Berlin.

Statistisches Reichsamt (ed.) 1942. *Statistik des Deutschen Reiches: Volks-, Berufs- und Betriebszählung vom 17. Mai 1939: Die Berufstätigkeit der Bevölkerung in den Reichsteilen*, vol. 557 (Neue Folge). Berlin.

Stephenson, Jill. 2001. *Women in Nazi Germany*. Harlow: Longman.

Stögbauer, Christian: The Radicalisation of the German Electorate: Swinging to the Right and the Left in the Twilight of the Weimar Republic. *European Review of Economic History* 5: 251-280.

Talbert, Richard JA, ed. 2000. *Barrington Atlas of the Greek and Roman World: Map-by-map Directory*. Princeton, Oxford: Princeton University Press.

Temin, Peter. 2006. Mediterranean Trade in Biblical Times. In Ronald Findlay et al., eds. *Eli Heckscher, International Trade, and Economic History*. Cambridge: MIT Press, 141-156.

Temin, Peter. 2013. *The Roman Market Economy*. Princeton: Princeton University Press.

Voigtländer, Nico and Hans-Joachim Voth. 2012. Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany. *Quarterly Journal of Economics* 127: 1339-1392.

Voigtländer, Nico, and Joachim Voth. 2013. How the West 'Invented'Fertility Restriction. *American Economic Review* 103: 2227-64.

Voigtländer, Nico and Hans-Joachim Voth. 2016. Highway to Hitler. Working Paper.

Waldinger, Fabian. 2010. Quality Matters: The Expulsion of Professors and the Consequences for PhD Student Outcomes in Nazi Germany. *Journal of Political Economy* 118: 787-831.

Waldinger, Fabian. 2012. Peer Effects in Science: Evidence from the Dismissal of Scientists in Nazi Germany. *Review of Economic Studies* 79: 838-861.

Weder, Mark. 2006. Some Observations on the Great Depression in Germany. *German Economic Review* 7: 113-133.

Whitehouse, David, and Ruth Whitehouse. 1975. *Archaeological Atlas of the World*. London: Thames and Hudson.

Wright, Gavin. 1986. *Old South, New South*. New York:Basic Books.

Xu, Guo. 2017. The Costs of Patronage: Evidence from the British Empire. Working Paper.

Young, Alwyn. 2016. Improved, Nearly Exact, Statistical Inference with Robust and Clustered Covariance Matrices using Effective Degrees of Freedom Corrections. Working Paper.