Company Towns: Local Governments under Industry Dominance and Decline

Elizabeth Mitchell Elder

August 13, 2021

Abstract

Recent evidence suggests that local economies grow more slowly in areas dependent on resource extraction. Building on literature on the "resource curse," I argue the consequences are also political: local governments in areas dependent on a resource industry shape public policy to the demands of that industry, leading to underinvestment in government capacity and public goods. Focusing on the case of Eastern Coal Country in the 20th century United States, I present evidence that the coal industry hindered the growth of local government capacity. Many former coal-producing areas are now in a state of profound social and economic decline; an implication of these results is that this decline may stem not only from recent job losses, but also from the coal industry inhibiting investment in human capital and infrastructure that could have better prepared these areas for the resources' inevitable exhaustion.

* Contact at e.mitchell.elder@princeton.edu

1 Introduction

Literature on the "political resource curse" suggests that places with economies centered on resource extraction develop less effective, less democratic political institutions. Judged against these expectations, the United States is considered a success story: despite heavy resource extraction in the early 1900s, husbanding of the profits through effective institutions and a diversified national economy allowed the country to prosper and democratization to continue (David and Wright, 1997; Wright, 1990).

The areas of the United States providing the coal that fueled its 20th century prosperity—in general, the coal fields of Appalachia and the Illinois Basin—show little evidence of this success. Coal-producing areas are poor and underdeveloped compared to their neighbors. Their economies were buffeted by the fluctuations of coal prices for 150 years and have now been debilitated by the near-disappearance of coal employment. Few coal-producing counties have successfully developed other industries to replace mining. Why did these areas, which contributed so much to the country's economy, remain so poor?

I argue that coal-producing areas in Appalachia and the Midwest have been subject to a political resource curse. Local governments in mining areas are lower in capacity than their neighbors because their economies historically depended on resource extraction. The profits earned from mining accrued mainly to relatively few elites; the gains to be had from mining incentivized these elites to invest in controlling local governments. Mining interests received many policy concessions from local governments, including low taxes on mining activity. As a result, local governments in coal country were and remain chronically underfunded and underdeveloped compared to their neighbors.

This paper proceeds in three parts. First, I join qualitative and historical evidence with granular descriptive data to characterize the political economy of Eastern coal country; I asses how this case fits within the resource curse paradigm, generate predictions for how coal's dominance affected local political institutions, and verify the plausibility of these predictions using qualitative work from the region.

Second, I test the hypothesis that areas of Appalachia and the Illinois Basin with economies dominated by coal mining developed smaller, lower-capacity local governments. Using a

differences-in-differences design, I demonstrate that coal-dominated counties' local governments fell behind their neighbors in size starting when the coal industry arrived. A complementary descriptive analysis suggests that while coal-dominated areas' governments caught up to other areas' in public employment by the postwar period, they continued to collect less in tax revenue and spend less money into the 21st century. This suggests that, in coal country as in resource-dependent places around the world, dependence on extractive resources leads to weak government institutions.

Finally, I present a case study of a county in Eastern Kentucky which illustrates how the economic and political dynamics of coal counties described in the first section of the paper produced the lower government employment and tax rates demonstrated in the second. This section also traces the consequences of coal industry dominance and hampered local governments for policy outcomes, from the height of the industry's dominance to the present, in the area of healthcare.

The political resource curse is a particularly stark example of a broader problem: when localities' economies are dominated by an elite group with coherent interests, local governments must cater to those interests—even when the elites' demands keep local governments weak and poorly-funded, as may be the case when the dominant elite is an industry group with more interest in preventing taxes and interference than in the provision of services. Local governments, which represent smaller constituencies with a less diverse ecosystem of economic and political interest groups than higher levels of government, are vulnerable to capture by a single powerful interest in a way that higher levels of government are not. The results presented here suggest that the potential harms of this capture go beyond control over policy outcomes—capture by economic interests can undermine the capacity of institutions of local government.

2 Theory

Countries with economies that depend on resource extraction are, on average, poorer than countries that depend on other sectors or have more diversified economies. This pattern was first empirically documented by Sachs and Warner (1995), and it has since spurred a large

body work by an ever-growing set of researchers (see Frankel, 2010; Van der Ploeg, 2011; Badeeb et al., 2017, for reviews). Some disagreement remains, but a consensus has developed that the economies of resource-dependent countries grow more slowly than others. This is the "resource curse."

Why might this be? Initially, explanations centered on purely economic consequences of resource extraction, like its effects on currency value and investment in other sectors of the economy (Sachs and Warner, 1995). However, most analysts now agree that there is a political dimension to the resource curse.

Resource extraction pumps money into the pockets of whomever—be they corporate investors, government officials, or rebel groups—can claim a share of the profits. This money can affect political processes in a number of ways; for example, it can buy support for authoritarian (or democratic) regimes, bribe government officials for favorable policy, or spur conflict over territorial control (see for example Dunning, 2008; Brollo et al., 2013; Ross, 2015, for a review). Different effects occur in different contexts, but on average, countries with resource-dependent economies have more corrupt and less democratic governments than other countries (Deacon and Rode, 2015). This is the "political resource curse."

In the late nineteenth and early 20th centuries, the United States was heavily engaged in resource extraction. It was responsible for 39% of the world's coal production in 1913, and was the world's leading producer of many other natural resources (David and Wright, 1997). The United States seems to have avoided the resource curse; it was and remains prosperous and democratic, perhaps because the national economy remained diversified even during its period of greatest resource extraction. However, the regions that produced the country's coal at its resource-extracting peak were largely dependent on mining, and they remain poorer and less developed than their neighbors. It was here, I argue, that the resource curse had its bite: while the country overall did not suffer an economic or political resource curse, coal-producing areas did.

I argue that coal production and coal producers hampered the growth of local political institutions in coal-producing areas of the United States throughout the 20th century. Incentivized by large profits accruing to a small set of mining investors, elites invested in building influence over governments in coal-producing areas. This elite influence produced

corrupt and weak local governments, as well as policy gains for mining interests like low tax rates, limited regulation, and cooperation from police in labor disputes.

What is more, the corruption and weakening of political institutions while coal interests were most powerful has had consequences for modern politics. Local governments lack the funding and capacity to solve important problems, like adapting their economies to a world with negligible coal employment and dealing with a massive influx of opioid drugs. As a consequence, I argue, residents of formerly coal-producing regions fare worse than their neighbors on indicators of well-being today.

2.1 The Case: Eastern Coal Country

Six states accounted for 83% of the United States' coal production in 1919: Pennsylvania, West Virginia, Kentucky, Illinois, Indiana, and Ohio (U.S. Census Bureau, 1922). Though coal production today has moved westward, for the first half of the 20th century, the heart of coal country was in Appalachia and the Illinois Basin.

Before coal production became large-scale and mechanized, coal-producing areas were generally sparsely populated by subsistence farmers. For example, Logan County, West Virginia became one of the nation's largest coal producers in the 1900s, with 75% of its non-farm labor force employed in the coal industry in 1920. In 1880, however, the coal industry had not yet arrived. There was one incorporated town; the county's total population was 7,300, and its local governments employed two people. Between 1880 and 1920, as the region's coal industry developed, Logan County's population grew by more than 5 times to more than 41,000—more than live there today.

Figure 1 summarizes characteristics of the region's coal-intensive counties, defined as counties with more than 20% of their non-farm labor force working in the coal industry in 1920, as compared to all other counties in the region between 1850 and 2000. Coal counties lagged behind others in terms of population and property value in the late 1800s. By the early 1900s, the median coal county had surpassed the median non-coal county in population, and though property values still lagged behind, the gap had closed somewhat. Coal and non-coal counties had converged further on both measures by the late twentieth century.

The period of coal counties' greatest population growth coincides with the period of

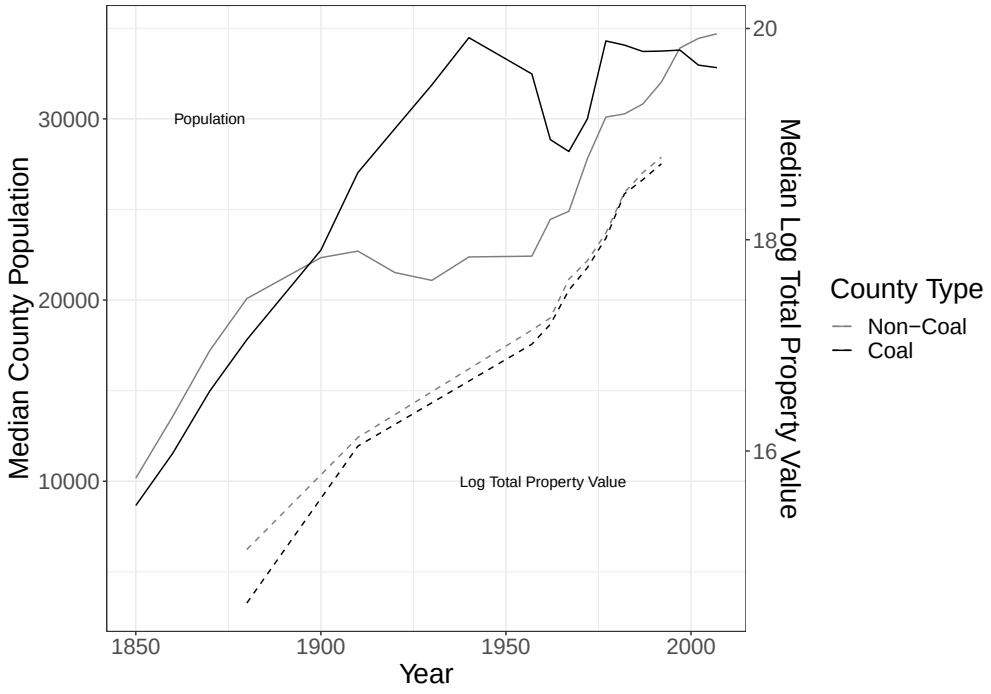

Figure 1: **Trends in Coal and Non-Coal Counties.** The solid lines represent population in the median coal (coal accounts for 20% or more of non-farm employment) and non-coal (all others) counties over time. Dotted lines represent log total property value in the median coal and non-coal counties.

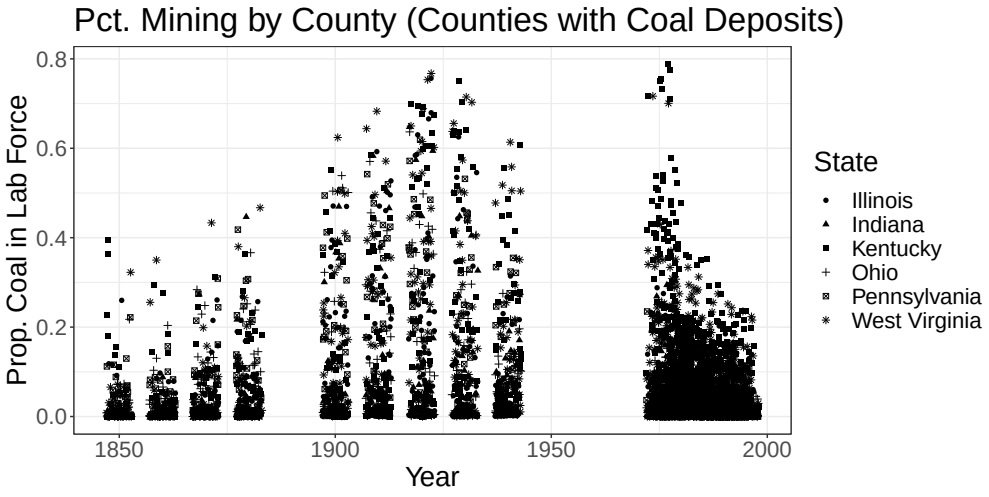

Figure 2: **Trends in County Coal Employment.** Each point represents the proportion of a county's non-farm labor force employed in the coal industry in a year. Dots are jittered along the horizontal axis (time). Only counties with some coal deposits are included.

highest coal employment, the decades around 1920. Figure 2 shows the proportion of the labor force employed in the coal industry between 1850 and 1995[1] in counties with some coal deposits. Though a handful of counties had coal-focused economies already by the 1880s, the period between 1900 and 1930 saw the largest number of counties focus their economies on mining. Some of these places remained invested in coal into the 1970s, but by the end of the twentieth century, even the most coal-focused places saw 20% or less of their workforce in the industry.

The growing populations of coal mining areas were settled in distinctive ways. Even as these areas became more prosperous and populous, they remained rural in character; coal miners and their families often lived near the mines where they worked, in remote areas with few other options for employment (Green, 2011). Many miners—60% in West Virginia, up to 80% in Eastern Kentucky—lived in "company towns," settlements built and owned by the coal companies to house mine workers. These towns were often unincorporated and remote, far from centers of population and with limited access to transportation networks.

As Figure 2 suggests, though coal mining was already occurring in Appalachia and the

[1]In 1940 and earlier, data is from the United States Census; in 1975 and later, data is from the Quarterly Census of Employment and Wages.

Illinois basin by the mid-1800s, the industry changed and grew dramatically between 1870 and 1920. Coal lands in this region had been primarily owned by private individuals, and mining was done primarily by small private companies. After the Civil War, corporations (especially iron, steel, and railroad corporations and their subsidiaries) quickly accumulated the rights to mine land throughout the region (Goldberg and Power, 1972; Leistritz and Voelker, 1975; Klesta and Others, 2016). Mining became more widespread and intensive, and ownership of land and mineral rights became heavily concentrated.

Thorough surveys of land ownership in Appalachia (Appalachian Land Ownership Task Force, 1983) and the Illinois Basin (Smith et al., 1978) in the 1970s found that land ownership was far more concentrated in mining areas than other parts of the region and that, for many reasons, large mining corporations paid less in property taxes than other landowners. These patterns were already quite old by the time of these surveys; the accumulation of property and corporate consolidation by mining interests in the late 19th and early 20th centuries meant that "absentee landlords"—the out-of-town, -state, or -country corporations that own much of the land and mineral rights in Eastern Coal Country—were a longstanding figure in the area's politics by the 1950s (Lutz, 1978). Legacies of this process are still present today: in five of West Virginia's heaviest coal-mining counties, ten corporations owned more than half the private land as of 2013 (Spence et al., 2013).

To summarize, areas of Appalachia and the Illinois basin that would become major coal producers were generally smaller and less developed than other areas before the industry arrived. With the arrival of the coal industry came massive population growth, often settled in coal towns outside of population centers, and increasingly concentrated ownership of land and mineral rights.

2.2 The Resource Curse in Eastern Coal Country

As is often the case in resource-rich areas, many coal-producing counties in Appalachia and the Illinois Basin developed economies specializing in coal mining, to the exclusion of other industries (Sachs and Warner, 1995). This led to a slowing of economic growth in coal-specialized counties (James and Aadland, 2011; Matheis, 2016), a classic expression of the resource curse.

How might an economy dependent on mining have affected governing institutions in eastern coal country? Literature on the resource curse offers insight into the relationship between industry and government in places with resource-dependent economies. This case differs from most contexts covered by the resource curse literature in two key ways: all the coal wealth in this region is privately owned, not (partly or fully) state-owned, as in the case of many oil-producing parts of the world; and governments in coal-producing areas do not suffer from an overabundance of mining-related revenue that weakens their fiscal discipline or loosens their ties to the electorate. On the contrary, as I will show later, mining wealth and extraction was taxed unevenly (if at all) until the second half of the twentieth century, so tax revenue was quite low. However, other features of this case fit well with mechanisms described by past work on the resource curse.

First, mining specialization is associated with the concentration of economic power (Isham et al., 2005). Resource extraction requires capital investment, but not an especially large or skilled labor force[2], and mined resources are generally exported to other areas for processing and use. Therefore, the profits of resource industries flow fairly narrowly to investors and ownership of mining companies, not broadly across a large labor force or multiple sectors of the economy. The more consolidated the mining sector in an area, the more concentrated the economic power in the area in the hands of relatively few mining elites.

Second, because resource extraction tends to dominate the economies of areas that contain resource deposits (Sachs and Warner, 1995), mining industries are more likely to be the dominant economic interest in their jurisdiction than are other industries. Governments in resource-rich areas often depend on the jobs and money mining companies bring. In return, mining companies depend on governments in resource-rich areas for favorable policy, since they generally cannot move their operations in search of a better policy environment.

To summarize, literature on the resource curse predicts two key patterns should characterize politics in coal-dominated areas of eastern coal country. First, economic—and by extension, political—power should be concentrated in the hands of relatively few mining

[2]Coal mining in the United States through the first half of the twentieth century relied on a reasonably large but low-skilled labor force; other resource industries and modern coal mining rely on smaller but higher-skilled labor forces.

elites. Second, because of the mutual dependence between local governments and mining companies in their jurisdiction, both benefit when local governments create policy favorable to the coal industry.

Historical evidence from eastern coal country corroborates the idea that coal elites were especially economically powerful in their areas. As discussed above, ownership of land and mining rights was concentrated in relatively few corporations (Gaventa, 1982). Wages for mine workers were relatively low, and there were few options for alternative employment in mining areas. In coal towns, especially when unincorporated, mining companies had substantial control over residents' housing, the prices of goods, and service provision (Green, 2011). Even in more conventionally-structured towns, though, the upper and middle classes were either directly involved in mining ventures or had commercial interests that aligned with mining (Gaventa, 1982).

The ability of mining companies to extract policy gains from governments in resource-rich areas has precedent in many nations worldwide; within the United States, the story of coal interests' pursuit of influence on state governments is instructive (Rode, 2013). Through financial means ranging from campaign contributions to bribery, mining interests have produced some of the country's most corrupt state governments and have secured substantial policy gains. Consistent regulation of worker safety and environmental damage came only after the industry's peak of power and employment had passed. Mineral extraction is still not taxed in Pennsylvania, and elsewhere a near-complete lack of tax on coal extraction in heavy producer states remained until the late 20th century.

2.3 Local Governments and Coal Industry Influence

While individual reports of the coal industry's influence on governments abound, this paper aims to causally identify and quantitatively measure the industry's effects. How might government accession to coal companies' interests be visible in aggregate data? I make two moves to make this problem more tractable: I focus on local governments as the unit of analysis, and my analyses focus on the readily-measurable outcomes of government employment, revenue, and spending.

Though many important policy domains in this area are regulated at the state level,

studying localities has key advantages. First, it is easier to find comparable coal-producing and non-coal producing places: notoriously coal-dependent West Virginia is substantially unlike other states, but there are local areas throughout Appalachia and the Illinois Basin that are similar to each other except for their coal endowment. These comparable local areas serve as leverage in understanding the causal effect of the coal industry on local governments.

Second, local governments in many coal-producing areas were relatively weak before the coal industry arrived. Literature on political dimensions of the resource curse suggests that the curse operates most strongly—or even exclusively—in places with weak institutions before the arrival of extractive industries. Well-developed institutions can resist capture. Coal-endowed areas are often rugged and were often sparsely populated before the industry arrived, with relatively few towns and incorporated areas, making them good candidates to observe to the effects of resource-industry dominance on government capacity.

Third, the coal industry depended on local governments' cooperation in important policy domains. For example, labor history documents that coal companies used local police forces as private armies in labor disputes, or at least depended on them to turn a blind eye towards their violent actions against labor organizers (Fishback, 1995). The suppression of union organizing was a top priority for the industry, and the willingness with which police forces cooperated with mining ownership in these disputes suggests a close relationship between the two.

What relationship ought we expect between coal industry dominance and local government employment, revenue, and spending? Coal corporations, like most corporations, preferred to pay less in local taxes. In areas where the coal industry accounted for much of the tax base, I expect that this led to lower local government revenues and smaller local government institutions. Though this may seem in some ways to be counter to coal elites' preferences—corporations need infrastructure for transport, management needs schools for their children, etc.—I argue that in places where a single industry is the primary source of economic activity, industry elites would rather construct public goods themselves than pay the tax revenue to allow local governments to do so.

This preference is clearest in the case of unincorporated coal towns. Coal corporations' reasons for constructing coal towns were myriad, but their prevalence reflects that companies

would rather construct towns from scratch, paying the upfront costs to build public services and infrastructure and governing them themselves, than cede land, tax money, and control to an elected local government to construct them. Even outside these enclaves, I argue, in places where coal was the dominant industry, corporations preferred to provide infrastructure themselves over allowing increased taxation.

I therefore expect that local governments in more coal-dependent places were smaller and worse-funded than local governments in non-coal places. Because of coal elites' relative economic power, aversion to paying taxes, and preference for weaker local governments, I argue that coal elites were willing and able to limit government revenue and size. This kept local governments in coal areas small when the industry was most powerful.

Understanding coal interests' effects on local taxation is especially important for the long-term development of local governments in coal-producing areas. When the coal industry was most dominant, localities largely depended on the property taxes they collected for funding (Peterson et al., 1981); cities and counties had few (or no) other options for revenue. If coal interests were able to keep property tax revenues low, local governments would have been less capable of building strong and effective institutions during a crucial period for their growth and professionalization. This early disadvantage in capacity would limit local governments' ability to grow once the industry has left, taking its tax revenue with it.

2.4 Historical Evidence for Local Government Capture

In the empirical sections that follow, I use quantitative analyses to test whether developing a coal-dependent economy leads to smaller local governments. These analyses build on—and are no substitute for—extensive work by regional specialists on three crucial questions: what did coal companies want from local governments? Through what means did they seek to get it? And finally, did they succeed?

Coal companies' policy demands centered on two areas: minimal taxation on land and mineral rights, and assistance from the criminal justice system in suppressing labor organizing and other forms of threatening activism[3]. The desire to pay less in taxes is all but

[3]Limiting the passage and enforcement of regulations surrounding worker safety and the environment was another priority, but these matters are generally handled by state and national governments.

universal among corporations. The focus on enlisting the police and courts to arrest, jail, and prosecute union organizers, environmental activists, and journalists is less universal, but it is well documented in coal country (Gaventa, 1982; Caudill, 1962; Fishback, 1995; Kiffmeyer, 1998).

John Gaventa's classic *Power and Powerlessness: Quiescence and Rebellion in an Appalachian Valley* provides a thorough account of the means through which the coal industry has pursued these interests. Gaventa offers many examples of overt attempts to win conflicts, indicative of the "first face of power:" running coal-affiliated candidates in local elections, buying votes or punishing the "wrong" votes through threats to company-provided housing or benefits, repeatedly appealing tax appraisals to lower their property's assessed value (see also Appalachian Land Ownership Commission, 1983), and of course, the aforementioned criminal justice-assisted reprisals against the outspoken.

More subtly, Gaventa describes the cultivation of a local middle and upper class with interests aligned with the industry. In areas dominated by the coal industry, the prosperity of all residents depends on the industry's success. Local lawyers, middle managers, and small business owners have the financial resources, social capital, and professional skills that would allow them to engage in electoral or legal challenges to the industry's dominance, but have little to no economic incentive to do so. Policy concessions to the coal industry, which might raise objections over special treatment from other local industries in areas with a more diverse economy, go unchallenged. Conflicts over issues like land ownership and taxation, then, rarely even rise to the level of political contention.

So, were coal companies able to count on local governments in coal areas for their desired policies? There seems to be little question in literature on Eastern Coal Country. As one expert said in an interview, "in the three-legged stool of business, government, and society, business and government are really just one leg" in the region[4]. There is a widespread sense, substantiated by local evidence (Smith et al., 1978, Appalachian Land Ownership Task Force, 1983), that coal companies pay little in taxes, and that this has led to underinvestment in roads, schools, and healthcare (Gaventa, 1982; Caudill, 1962). In fact, in the 1970s, the West Virginia State Tax Department took the official position that "the coal industry's support

[4]Author interview with Thomas J. Kiffmeyer, 5/20/2021.

of local government and schools, through property taxes, has not been realistic given the extent of the industry's mineral and fee property holdings" (as quoted by the Appalachian Land Ownership Task Force, 1983). The many instances of cooperation from the courts and law enforcement with coal interests suggest the companies' interest prevail in that arena as well.

This evidence provides valuable support for the notion that coal companies have the will and the means to succeed in keeping local governments underfunded. In the analyses that follow, I seek to complement this evidence in three ways: by comparing local government size in counties across the entire region of Appalachia and the Illinois basin; by studying the timing of the emergence of the gap in government size between coal-and non-coal counties; and by employing quantitative methods to rule out alternative explanations for the size of coal areas' governments.

3 Empirics

3.1 Data and Measures

To test the hypothesis that a coal-dependent economy led to weaker government institutions in the coalfields of the Eastern United States, I draw on data on economic and political conditions at the county level. Counties are a mutually exclusive and exhaustive unit of geography—that is, all land in this region is part of one and only one county—and are small enough in the states of interest that county-level data captures fairly local conditions. I limit my analysis here to the Eastern coal country region discussed above, which includes the six states of Illinois, Indiana, Kentucky, Ohio, Pennsylvania, and West Virginia.

I measure a county's dependence on coal using employment data drawn from the full-count Censuses from 1850-1940, the period covering the peak of coal employment in this region. For each county, at each Census, I divide the number of people employed in the coal mining industry by the total number of non-farm, non-domestic laborers in the county. This measure captures the extent to which coal provides the livelihoods of county residents and therefore reflects the power of coal, relative to the area's other industries, in the local

economy.

I measure the size of local government institutions using two kinds of data. First, using the same 1850-1940 full-count Census data, I calculate the number of people in each county whose industry of employment is classified as "local government administration." This includes people whose full time job (or the job that provided most of their income) was provided directly by the local government, including police officers, firefighters, lamplighters, clerks, assessors, judges, and elected officials, but excluding postal workers and teachers.

Second, the Census of Local Governments provides information on the revenues from property taxes, expenditures (total and by category), and employees of many local governments in the United States over the period from 1957 to 2002. Earlier versions of the survey were conducted in 1880 and 1912. These data also measure the total value of all property in each county for each year; since the revenue measure captures revenues from property taxes, this directly measures the county's tax base. I use the property tax revenue, expenditures, and employment of local governments, summed across all local governments in a county, as measures of local government strength.

The three sections that follow all assess the relationship between coal industry dominance and local government size. The first presents simple descriptive statistics on the size of local governments in coal and non-coal counties over time. The second presents a difference-in-differences analysis, showing that counties where the coal industry developed abruptly diverged from non-coal counties in government size at the time the industry arrived. The third extends these results by showing the relationship between coal industry size and local government employment, revenue, and spending throughout the 20th century.

3.2 Government Growth in Coal and Non-Coal Counties

Local governments grew substantially between 1850 and 1940, and even more dramatically after World War II. A first step in understanding the relationship between a coal-dependent economy and government size is to compare the trajectories of local government growth among counties with and without sizeable coal industries.

To simplify this comparison, I limit the sample to two types of counties: those with large coal industries in 1920 (defined as having 20% or more of their non-farm labor force working

in the coal industry; n=107) and those without appreciable coal industries (less than 1% in the industry; n=270).

I capture government size using the number of local government employees in each county. However, coal and non-coal counties are located in different states with different roles for local governments, and the groups also differ in population over time. To account for this flexibly, I measure government size using the residuals of an OLS regression of the logged number of government employees on a three-way interaction between county's logged population, a set of state fixed effects, and a set of period fixed effects. This gives a measure of the county's local government size, after accounting for differences in population and state between coal and non-coal counties over time.

Figure 3 presents the residualized government employment measure in the median coal and median non-coal county between 1850 and 2002. Early in this period, the median coal county employed fewer people in government than the median non-coal county. This gap widens steadily after 1860, with coal counties' governments employing fewer people; the gap is largest around 1920, and it all but closes as the twentieth century ends. Throughout the 20th century, coal counties' local governments employed fewer people than non-coal counties, after accounting for differences in state and population.

This figure suggests that counties which became heavy coal producers diverged from other places in government employment around the time the coal industry was becoming most powerful, from the early 1900s through the 1920s. However, this is insufficient to conclude that the growing coal industry was the cause of the divergence. Areas could have selected into specializing in coal production for many reasons, and this selection process, rather than the specialization itself, could have produced differences in governments between coal and non-coal counties. The next section takes up this concern.

3.3 Differences in Differences Analysis

As the coal industry expanded over the late 19th and early 20th centuries, some coal-producing areas experienced steady and gradual growth in mining employment. In other areas, the industry arrived abruptly. For example, in 1910, nineteen of the 10,500 residents of Harlan County, Kentucky worked in the coal industry. By 1920, 5,267 worked in coal

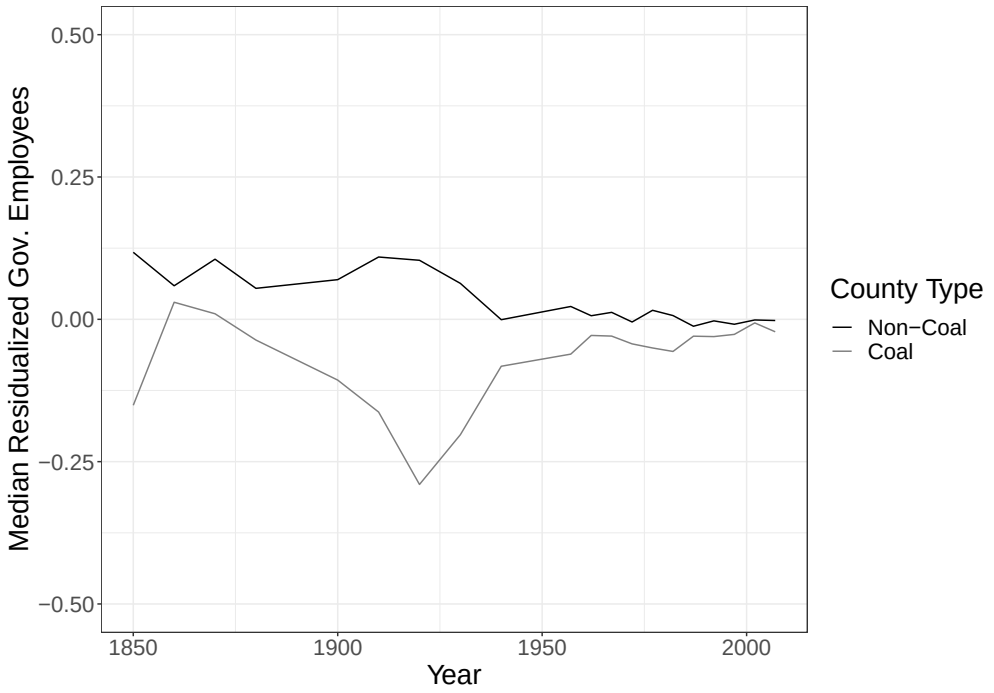

Figure 3: **Government Employment in Coal and Non-Coal Counties Over Time.** On the y-axis is a measure of government employment after accounting for state and population differences. The x-axis shows time. Lines represent residualized government employment in the median coal and non-coal county in each year.

mining, representing 69% of the county's non-farm labor force.

In all, there are 23 counties in Appalachia and the Illinois Basin that developed dominant coal industries over a short time period. Each of these counties, between two consecutive Censuses in the 1850-1940 period, went from less than 5% of their labor force in coal mining to at least 20%, and eventually reached at least 40%. In all cases, coal became the county's largest industry at some point in (and often throughout) the period between 1880 and 1940. For the purposes of this analysis, this group of 23 counties serves as my "treatment group": they began untreated (i.e. no coal industry) and, at an identifiable point in time, received the "treatment" of developing a large, dominant coal industry.

To assess the relationship between coal specialization and local government size, I compare changes in pre- and post-treatment government size in these treated counties to changes in government size at the same time for counties that did not receive the treatment; that is, a differences in differences design. I define the untreated counties as those in which no more than 5% of the labor force worked in the coal industry at any point between 1850 and 1940 (n=280)[5].

The outcome of interest in this analysis is the same as in the previous section: a measure of the total number of local government employees in each year, residualized for state, population, and time differences. However, to capture the relationship between these variables and government size absent industry intervention, the regression model used to produce the outcome here is trained only on control counties. I expect that after the coal industry arrives in treated counties, governments in treated counties (as measured with this residualized figure) will grow more slowly than governments in control counties.

The key identifying assumption of this analysis is that of "parallel trends:" the potential outcomes under the control condition of both groups of counties move in parallel over the time period being studied. Here, this means that if the treated counties had never developed a coal industry, their governments would have grown at the same rate as the governments of control counties. Within states and years, and accounting for population size, I argue this assumption is plausible; as long as they are being compared to similarly-populated

[5]Section 0.5 in the Supplemental Information includes a map of the treatment and control counties in this analysis.

control units with the same state-delegated responsibilities, there is no reason to expect coal counties' governments to have grown differently from their neighbors absent intervention. I test below whether treatment and control counties' governments followed parallel trends before treatment occurred, and I construct several alternative control groups that make this assumption even more plausible.

The coal industry arrived in each of the treated counties between 1870 and 1920. The earliest time period in which outcome data is available is 1850, allowing for the observation of pre-treatment trends in the outcome even for the earliest-treated units. Outcome data is available at 5 or 10 year increments through 1997 (except the 1940-1957 period), and to observe how treated and control units' governments grew over the course of the 20th century, I test for differences at every point following treatment. To account for the different timing of treatment receipt and the measurement of effects at multiple time points after treatment, I implement a difference in differences design with staggered treatment assignment and multiple time periods.

In some of the treated counties, the coal industry did not survive long after the downturn of the 1920s, and in most, coal employment accounted for less than 10% of the labor force by the 1980s. However, I consider treatment status here to be irreversible—that is, once a county is "treated" by developing a dominant coal industry, it remains treated in all following years, regardless of the industry's size in later periods. If I am correct that the presence of a dominant coal industry in the early 20th century shaped the development of local governments durably, even after the industry leaves, it cannot be considered comparable to a unit that never received the treatment at all.

I implement this analysis in `R` using the `did` package and the method described in Callaway and Sant'Anna (2020). Callaway and Sant'Anna's method breaks the treated observations into groups based on when each unit is treated and calculates treatment effects for each time at which outcomes are observed. It then combines these group-time effects into weighted average treatment effects at each time point for which outcomes are observed for all the groups in my data, where the weights depend on the number of observations in each group[6].

[6]This differs from a conventional two-way fixed effects method for estimating effects from a difference in

Figure 4 presents the results. Each point represents the average treatment effect for groups at each time pre- and post-treatment; the bars represent bootstrapped 95% confidence intervals clustered at the county level. The x-axis shows the time relative to treatment. The leftmost two points, then, test for differences in differences between treatment and control groups 20 and 10 years before treatment occurs. These estimates are insignificant, suggesting trends in government employment in treatment and control units are not diverging before treatment is received.

The black coefficients in Figure 4 represent the average treatment effect across groups at each time point after treatment. The borderline significant negative coefficient at time 10 suggests that treated units, between the Census ten years before being treated and the Census ten years after being treated, gained less in residualized government employment than did control units over those same years. The estimates from further after treatment compare the change in treated and control units between their final pretreatment period and the period elapsed since treatment occurred. Though the statistical significance of the estimates vary, coefficients remain negative until around 40 years after treatment, at which point estimates generally become close to 0.

A pooled test of the treatment effect over the periods from 0 to 40 years post-treatment suggests a significant effect of $-.30$ $(-0.53, -0.06)$ points. This represents a one standard deviation decrease in government size in treated units relative to control units. To put this in perspective, consider two Pennsylvania counties that differed by .3 points on the same scale in 1900: Berks and Lancaster, which both had about 159,000 residents at the time. Berks County's governments employed 117 people, while Lancaster's employed 156. This difference of about 25 employees per 100,000 residents is similar to the estimated treatment effect of the emergence of a large coal industry on a coal county.

Results available in the Supplemental Information repeat these tests excluding counties without any coal deposits and excluding counties with slaveowners to ensure the comparisons

differences design in three ways: first, it disaggregates estimates at different lengths of treatment exposure; second, it removes comparisons between already-treated units and those becoming treated; and third, it weights group-time effects by the number of observations, but not the variance in treatment assignment. See Callaway and Sant'Anna (2020) for further details.

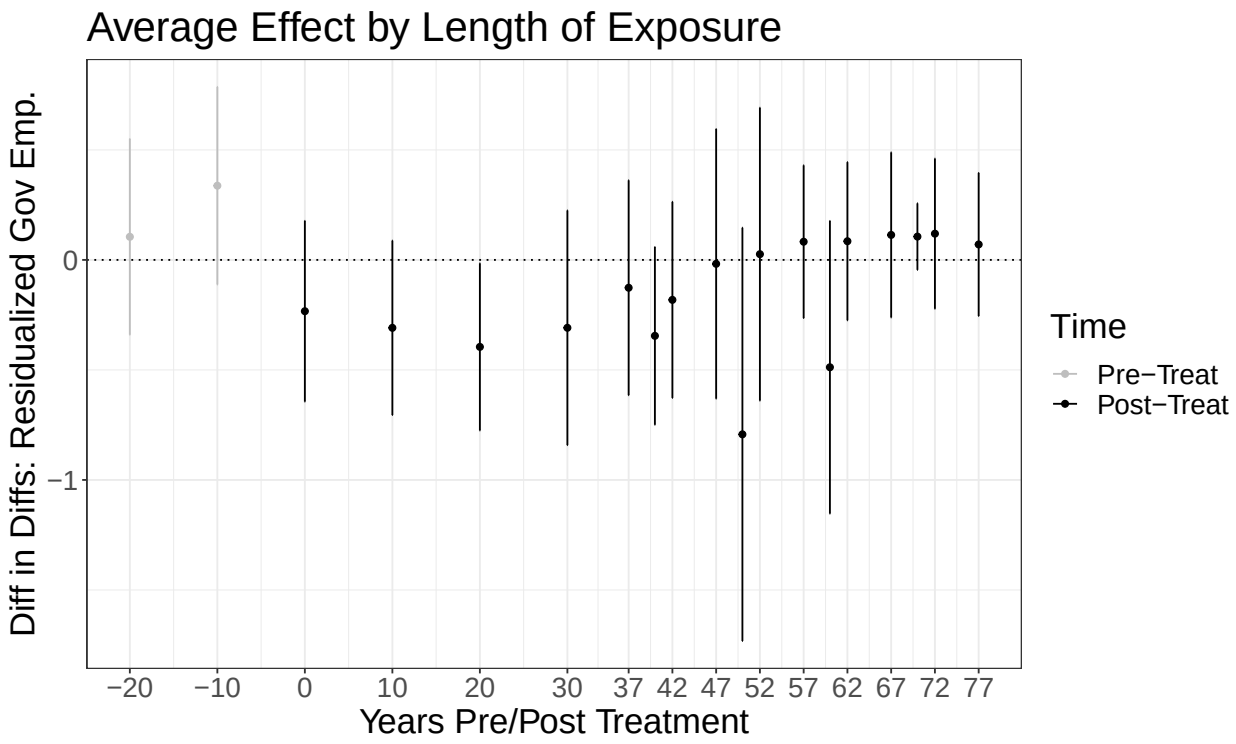

Figure 4: **Differences in Differences: Government Employment in Coal and Non-Coal Counties Over Time.** On the y-axis is the difference between treated and control groups in the difference in government employment from the previous time period. The measure of government employment is logged and residualized for state and population differences. The x-axis shows time relative to treatment; -20 represents 20 years before receiving treatment, while 20 represents 20 years afterwards. Bars represent bootstrapped 95% confidence intervals.

do not reflect other factors that could have set treated and control counties on different trajectories of government spending around this time. The results are robust to these different samples.

These results suggest that the arrival of a dominant coal industry depressed the growth of local governments in coal-producing areas. Several decades after the arrival of the industry, coal-dominated counties' government employment had returned to near parity with employment other places. Given the substantial decline in coal employment over the second half of the twentieth century, it is possible the dwindling of the coal effect is because the coal industry had left these areas. However, for decades after the industry arrived, coal mining areas showed distinctively slow government growth for their population size and states.

3.4 Extension: Correlational Results

In this section, I move away from the goal of causal inference and test the descriptive claim that counties with coal-dependent economies have weaker governments by a broader set of measures. The differences in differences analysis in the previous section was enabled by the long-term panel data available on government employment; data is more limited for other important features of local governments, including property tax revenue and expenditures. The difference-in-differences analysis also limited the sample of coal counties to 20 units with a distinctive set of useful features. A broader descriptive analysis can explore the relationship between a continuous measure of coal industry size in all counties and local governments in all the region's counties.

I test the expectation that coal-producing counties have weaker local governments by regressing measures of local government capacity on my labor force measure of coal dependence. I use three measures of local government capacity: number of government employees, property tax revenue collected, and total expenditures, collected at intervals between 1880 and 1992.

I regress each government outcome on the proportion of the county's nonfarm, nondomestic labor force employed in the coal mining industry. I include state fixed effects to account for between-state differences in government structure, and to scale the outcomes to the population and funding base of the area, I control for logged population and logged total

property value in the year of each government outcome. All outcome measures are logged. Observations are weighted by their populations in 1880. The regression models are weighted least squares with standard errors calculated using the HAC adjustment proposed by Kelly (2020) to account for spatial autocorrelation between observations.

The control variables I include in these models to make the local government outcomes more comparable across counties—population, total property value, and state—complicate the interpretation of these relationships. For example, there is suggestive evidence that coal companies influenced local officials to ensure their property was assessed at lower-than-appropriate values for tax purposes (Caudill, 1962); controlling for property value could therefore be suppressing some of the relationship between coal dependence and tax revenue. However, given the often-substantial differences in population and property value between coal- and non-coal counties, the well-documented relationship between resource dependence and prosperity (James and Aadland, 2011), and the differences in local government institutions between coal-heavy and less-coal-heavy states, I argue these state, property value, and population measures must be included as covariates.

For dates between 1880 and 1940, when the coal industry employed substantial proportions of the labor force in many counties in the region, I predict government outcomes in a year using coal employment from the most recent census. This captures the contemporary relationship between coal dependence and local government at a particular time. For outcomes measured starting in the 1950s, after coal employment has become negligible in most places, I predict government outcomes using coal employment in 1920, the industry's peak. This captures the longer-run relationship between coal dependence and local government after the industry's strength has waned.

Figure 5 presents the result for government employment. Each point represents the coefficient on the proportion of county labor force in the coal industry in a regression of the logged number of government employees on coal employment, logged total population, logged property value, and state fixed effects. From around 1900-1940, counties more dependent on the coal industry had local governments that employed fewer people. By the 1950s, counties that had been more dependent on coal had similarly-sized local governments to others, though in some years the relationship remains negative and significant.

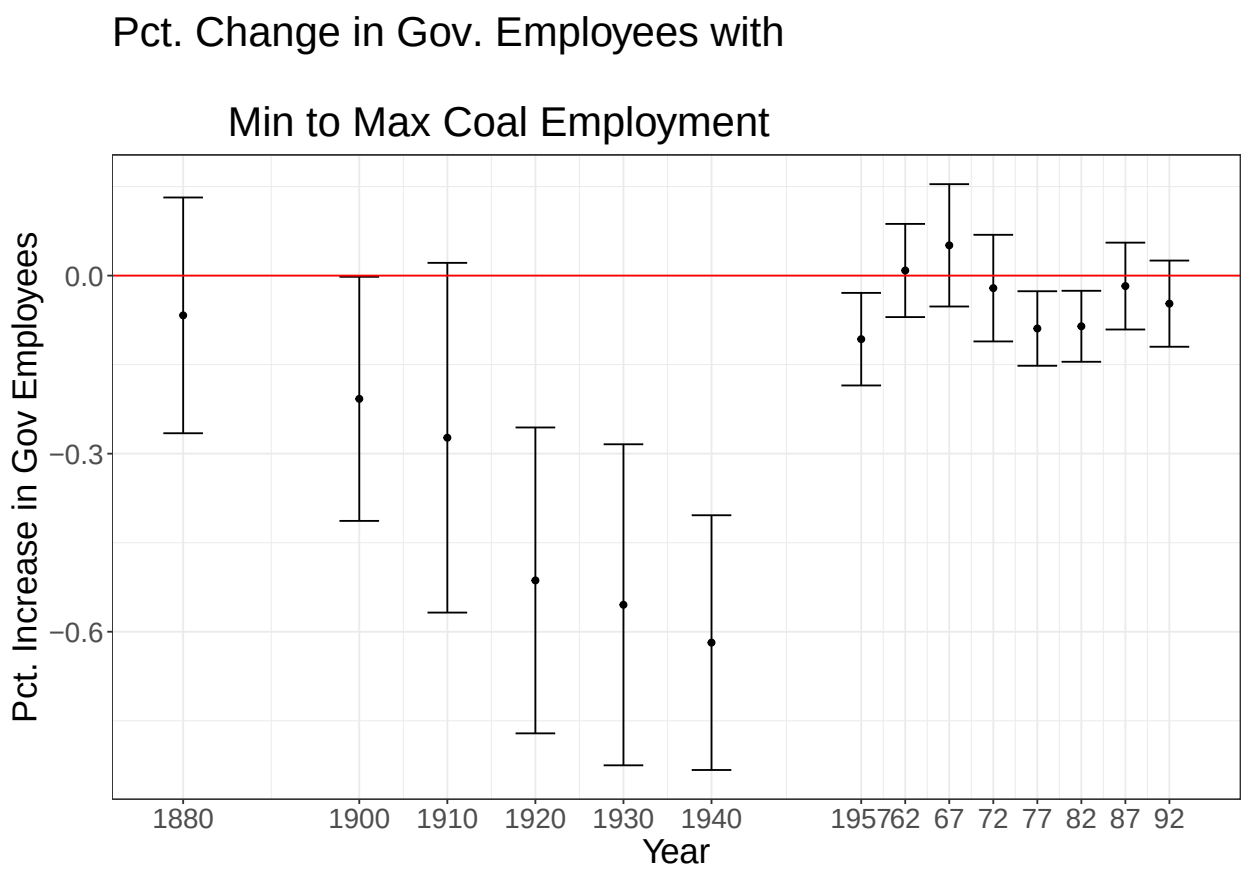

Figure 5: **Coal Employment and Government Employment.** Points represent the coefficient on the proportion of county labor force in the coal industry in a regression of the logged number of government employees on coal employment, logged total population, logged property value, and state fixed effects. Bars represent 95% confidence intervals with standard errors calculated with HAC adjustment for spatial clustering.

The results in Figure 5 echo those in the differences in differences analysis in the previous section: in the initial years after coal became powerful, coal counties had smaller local governments, but coal and non-coal counties converged in government employment as time passed and the industry employed fewer people.

Figure 6 presents the relationship between coal employment and property tax revenue. Again, the results suggest a negative association: counties with employment more concentrated in the coal industry collect less revenue from property taxes, even after accounting for their population size, property tax base, and state. Unlike in the case of government

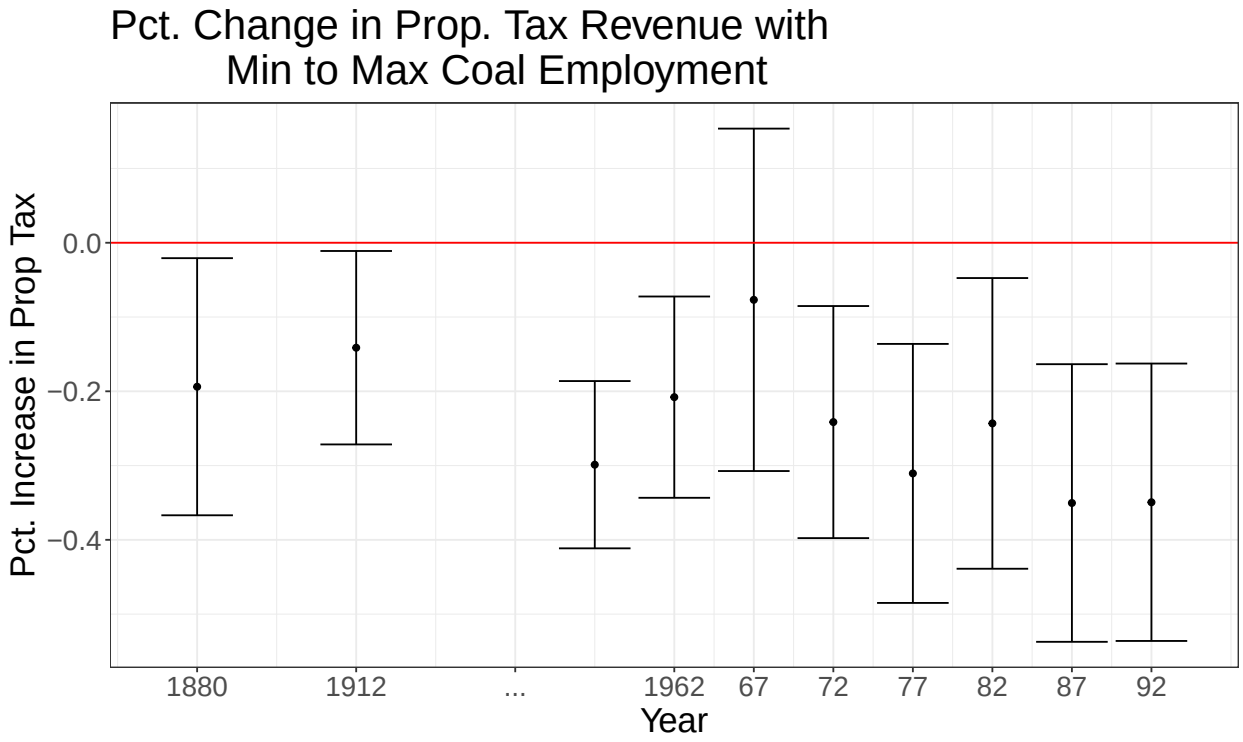

Figure 6: **Coal Employment and Property Tax Revenue.** Points represent the coefficient on the proportion of county labor force in the coal industry in a regression of the logged amount of government revenue from property taxes on coal employment, logged total population, logged property value, and state fixed effects. Bars represent 95% confidence intervals with standard errors calculated with HAC adjustment for spatial clustering.

employment, however, this relationship persists into the late 20th century, long after coal employment in the region declined.

Finally, Figure 6 presents the relationship between coal employment and total local government spending. This outcome measure is available at a more limited set of time points, but when it is available, the results are consistent with those from the models of property tax revenue: even after coal employment has declined, former coal-dependent counties' local governments spend less money.

Taken together, these results suggest that local governments in areas with more of their labor force employed by the coal industry tend to employ fewer people, collect less revenue, and spend less money. Patterns diverge in the later time periods between government

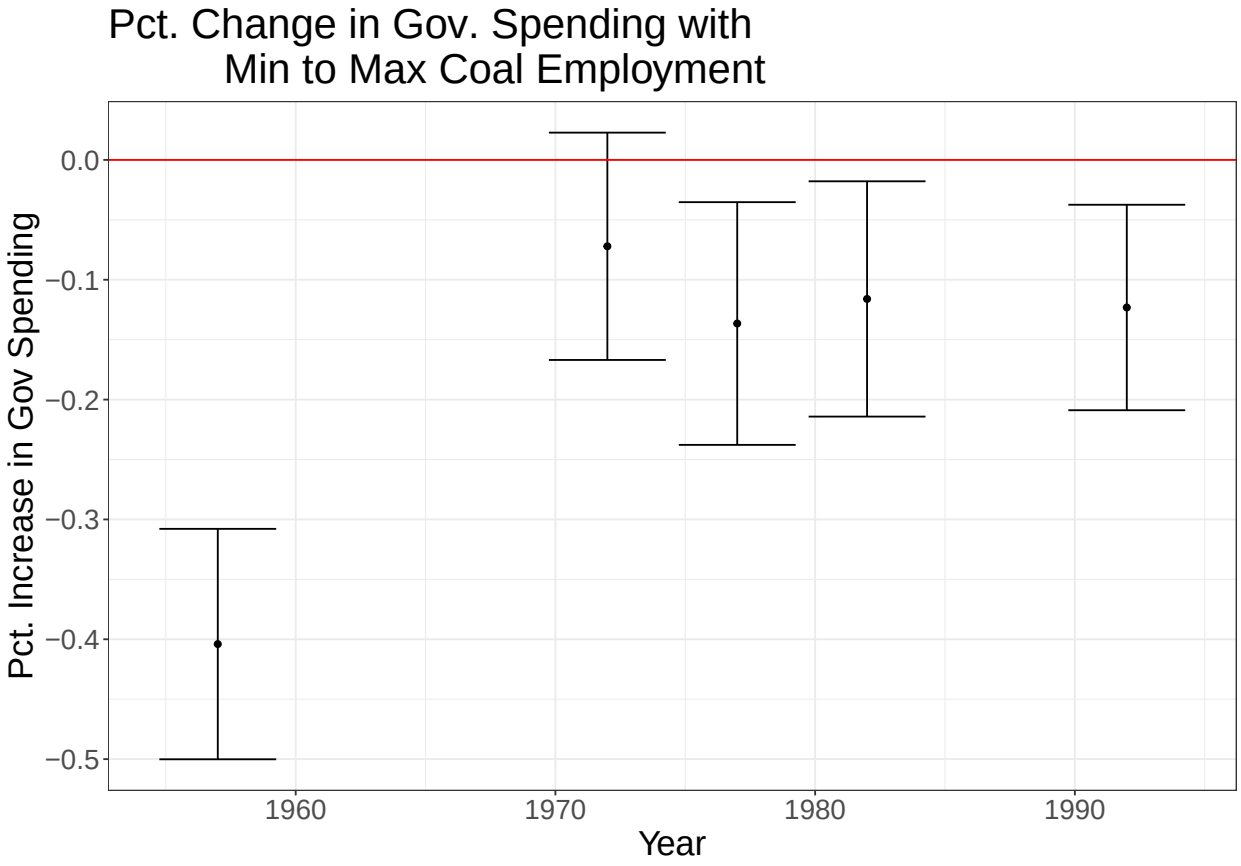

Figure 7: **Coal Employment and Government Spending.** Points represent the coefficient on the proportion of county labor force in the coal industry in a regression of the logged amount of government spending on coal employment, logged total population, logged property value, and state fixed effects. Bars represent 95% confidence intervals with standard errors calculated with HAC adjustment for spatial clustering.

employment and government spending and revenue: though employment in coal counties recovers after the industry declines, more coal-dependent counties continue to collect less in revenue decades afterwards.

4 Case Study: Linking Historical Coal Dominance and Modern Well-Being

The results presented thus far suggest that the coal industry suppressed the growth of local governments in Eastern coal country. To connect these findings to the dynamics described in the theory section, I now briefly describe the history of a place that illustrates the interplay between coal companies, local governments, and public policy. Perry County, Kentucky rapidly transformed from an isolated region of subsistence farmers to one of the country's largest coal producers in the early twentieth century. The coal industry's influence prevented the county's governments, funded at similar rates to those of their neighbors before the industry arrived, from fully benefiting from this economic growth.

Perry County, Kentucky was formed in 1821. The 1850 Census counted 2975 residents, only 19 of which performed market labor; the rest were farmers or did not work outside the home. Perry County remained primarily agricultural until 1912, when it was first connected to the regional freight railroad system. In the 1910 Census, just before the railroads arrived, the county had grown to 11,255 residents, but 221—less than 2% —worked outside a farm or home. Just two of those 221 market labor residents listed coal as the industry of their primary occupation.

The county's local governments were small during this agricultural period, with only a handful of full-time employees, but they were broadly similar in size and funding to those of in neighboring areas. In 1912, on the eve of the coal industry's emergence, Perry County's local governments collected property taxes only slightly less than the state median; in 1880, their taxes were somewhat higher than the typical Kentucky county.

Meanwhile, sixty miles north in Paintsville, a teacher named John C. C. Mayo had received financial backing from a Pittsburgh family to fund the purchase of land and mineral

rights in Eastern Kentucky. Mayo traveled throughout the region in the 1890s and 1900s acquiring property for his backers, including substantial tracts of mineral rights in Perry County, with the eventual goal of attracting railroad companies to Eastern Kentucky. The land and rights Mayo acquired were combined with land purchased or claimed by other out-of-state-funded speculators to form the Kentucky River Coal Corporation in 1915. With the arrival of the railroads a few years earlier, the KRCC was ready to lease its land to coal operators, launching the mining industry in Perry County (Appalachian Land Ownership Task Force, 1983; Kentucky River Properties, 2021).

4.1 Establishing Industry Dominance

By the 1920 Census, Perry County's population had more than doubled in ten years to 26,040. Sixty-three percent of its non-farm labor force now worked in the coal industry, up from only two residents ten years before. Coal dominated the county's economy for decades to follow. The industry still accounted for 63% of Perry County's labor force in 1930, before falling to 49% in 1940 and 43% in 1975, less than 20% in the early 1980s, and less than 10% in the 1990s. Though few miners are left today, the coal industry continues with new technology: as recently as 2008, coal production rose to more than 17 million tons, just a few thousand off the county's all-time peak.

As the mining industry in Perry County took off in the 1920s, the Kentucky River Coal Corporation recorded 143,000 acres of mineral holdings in Eastern Kentucky. Because of the area's tax policies, however, it is unlikely the KRCC's substantial mineral assets were reflected in a substantial tax bill. Historically, taxes on unmined resources were set through a negotiation between owners and local assessors; coal owners would file a return stating the value of their coal, the assessor could reevaluate the property, the owner could appeal the reevaluation, and "haggling" would continue until an agreement was reached (Vasek Jr, 1985). Because coal companies prioritized building influence over local assessors in coal-heavy areas, the agreements were generally quite favorable to owners (Caudill, 1962, Appalachian Land Ownership Task Force, 1983). When mining rights were owned by a different entity than the surface property, local authorities often had difficulty locating the mining rights owners; in these cases, unmined minerals went completely untaxed (Vasek Jr, 1985; Parrish,

2011).

While coal interests were working with local governments to keep their tax bills low, they were also avoiding local governments altogether when possible. As of 1936, there were only two incorporated towns in Perry County, containing less than a quarter of the county's population (WPA report). Outside of the county seat of Hazard, many of the area's settlements were coal towns or coal camps (Rennick, 2000), unincorporated areas without town governments in which mining companies were responsible for the provision of housing, infrastructure, education, and healthcare. Unlike in other regions, where growing populations moved to urban centers with larger local governments, population gains in early 20th century Perry County were directed to these unincorporated, or at least rural, areas.

4.2 Government Size, Revenue, and Spending

As local governments in the region were beginning to grow and professionalize in the early 20th century, Perry County began to fall behind its neighbors. In the early 1900s, the county's governments collectively employed 3-6 people for every 10,000 residents, quite close to the state median. Perry County had fallen behind the median government employment by 1930. In 1940, its local governments employed 20 people per 10,000 county residents compared to a state median of 25, and by 1957, they employed 135 to a state median of 178[7].

The KRCC continued to accumulate land in the region (though at an incremental pace, compared to the rush in the late 1800s and early 1900s), and by 1978, when the Appalachian Land Ownership Task Force surveyed property records in Central Appalachia, the company owned 29,685 acres of surface land and 26,272 acres of mineral rights in Perry County alone—accounting for 13.5% of the total land in the county, and an even larger share of the mineral rights. The company's holdings in the county were assessed at a value of $2,979,500, only

[7]This trend reversed in the 1970s and 80s, with Perry County local governments employing more people than the median Kentucky county. Federal aid, beginning with Great Society programs under Lyndon Johnson, and state transfers began to bolster local government coffers at this point, and literature on Eastern Kentucky suggests that patronage networks developed to use this money to employ local officeholders' relatives and supporters in a growing array of local government positions. See Kiffmeyer (1998); Eller (2008).

1.7% of the property value in the county. That year, they paid a total of $19,044 in property taxes on the land, a rate of .006%, and no local taxes on the mineral rights or value of unmined coal[8]. This is less than the overall 1% property tax rate in the county that year and amounts to about 1% of the property taxes collected in the county.

The 1978 tax year marked the last gasp of a state effort to make the taxation of unmined minerals more lucrative, uniform, and insulated from local political pressure (Vasek Jr, 1985). In 1976 the state legislature, motivated by the prospect of garnering tax revenue from out-of-state (i.e. nonvoting) landowners, passed a state property tax of 31.5c per $100 of unmined coal and used state resources to assess its value. This state assessment process grew the total assessed value of unmined coal in the state from $109 million in 1976 to $276 million in 1978 (Parrish, 2011)—a clear indication that the local assessment process had produced undervaluations highly favorable to coal owners.

However, the 1976 legislation also forbade local governments from taxing unmined coal as property, eliminating even the patchy revenue counties had previously been able to collect from these resources. A portion of the revenue from this state tax was redistributed to local governments. Any relief this redistributed revenue provided was short-lived: in 1978, the state legislature reduced the tax rate on unmined coal to .1c per $100 of value and returned assessment responsibility to the local level, a move so clearly influenced by the mining industry that one observer called it "reminiscent of the 'good old days' of King Coal" (Collins, 2015). Today, the revenue from this astonishingly low tax rate is not seen as enough to fund the assessment process (Parrish, 2011).

Perry County's local governments, then, have largely not collected tax revenue from their unmined mineral resources. The surface property owned by mining companies is also often taxed at lower rates than other land in Eastern Kentucky because it has relatively few improvements, because companies take advantage of tax breaks intended for agricultural land, and because the leverage and resources of large corporations allow for repeated appeals and negotiation over property valuations (Appalachian Land Ownership Task Force, 1983).

This low taxation of this valuable property limited governments' revenue. Recall that in

[8]Because the KRCC largely leased its rights to coal operators rather than mining itself, it would also not have been responsible for the state severance tax on coal mining enacted in 1972 (Roenker, 2013)

1880 and 1912, Perry County's property tax revenues were similar to or higher than other counties in the state. From the 1950s through the 1990s, on the other hand, Perry County's local governments collected less in local revenue than their peers. In 1978, for example, the median Kentucky county's local governments collected $50.40 in property taxes per resident per $100,000 of property. Perry County collected $30.46, 87th of the state's 120 counties. The county's local governments were heavily subsidized by state and federal transfers: seventy-one percent of their funding came from intergovernmental transfers in 1967, as compared to a median of 55% in Kentucky and 37.5% in the region overall. In property tax revenue and in employment, Perry County's local governments began to fall behind their neighbors after the coal industry arrived.

4.3 Coping with Industry Decline: The Case of Healthcare

How did Perry County's smaller and less-funded local governments affect the provision of public services? I will focus here on the issue of health. In the days when coal employed most of the county's residents, company-provided health care kept people fairly healthy. However, as these company-provided services dried up along with the large coal companies after World War II, the public health system was unable to cope, and health outcomes quickly declined.

In 1936, when around half of Perry County's labor force was employed by the coal industry, a Works Progress Administration (WPA) survey assessed the area's health services. It described the then-ten-year-old county health department relatively favorably, noting its growing vaccination program. The county contained two hospitals with 100 beds and 37 doctors (many of whom worked for coal companies in coal camps) for a population of around 42,000, both of which fell below national targets, but not dramatically so. The death rate in the county was lower than the state average, despite a higher than average rate of waterborne illnesses due to pollution in the water supply—a common problem in areas where coal companies were responsible for water supply (Boone and Lee, 1947). Around the same time, the infant mortality rate in Perry County was 44.9 per 1,000 live births, lower than the state average at the time and substantially lower than the rate in other Eastern Kentucky coal counties (Boone and Lee, 1947).

In the 1930s, Perry County's coal town residents were healthier than their neighbors;

many other mining areas had far worse outcomes. For example, in the average Kentucky town where mines were surveyed by Boone and Lee (1947), infant mortality rates were 59 per 1,000 live briths, 30% higher than in Perry County. These poor conditions prompted the the United Mine Workers of America's Welfare and Retirement Fund, which paid for medical care to coal workers using a fee collected from mine operators, to fund the construction of ten hospitals in southeastern Kentucky and western Virginia–one in Hazard, the county seat of Perry County.

The new mining-funded hospitals, christened the Miners Memorial Hospital Association, quickly ran into financial trouble. The mechanization of the coal industry after 1950 diminished the number of patients eligible for the hospitals' care, and drops in coal prices diminished the industry profits that funded the hospitals (Mulcahy, 1993). The UMWA also cut back dramatically on its coverage of miners' healthcare in this period. By the early 1960s, the Association was attempting to sell four branches of the hospital system, including the one in Hazard. Federal intervention and funding kept the hospitals from closing, but this multifaceted disruption to the security of area miners' livelihoods had already triggered unrest.

Violent protests broke out throughout the region in 1962 in response to the cuts in health coverage and the prospective hospital closure, though the focus of the protests gradually broadened to include the economic concerns of miners and poor residents more generally. The unofficial center of these protests was Perry County; protesters would gather in Hazard and travel to different non-union mines in the region. Working through both armed gangs and the local criminal justice system, coal owners and their middle class allies mobilized against the protests (Black, 1990). One 500-person caravan was dispersed by the Perry County sheriff, himself a coal owner (Eller, 2008). When a county judge was seen by some as insufficiently tough on the protesters, a local political action committee called the Citizen's Committee for Law and Order was formed to defeat this and future pro-labor candidates for local government offices in the county, and on election day, seven organizers were arrested to prevent them from organizing for the judge's reelection (Black, 1990; Cantrell, 1987). Local newspapers portrayed the protesters and organizers as outsiders and communists, minimizing sympathy with their demands.

Eventually, the movement shifted its activities away from protesting and towards lobbying and organization over the use of federal antipoverty funds. These funds were largely controlled by local government officials who used them to supplement county budgets and cultivate patronage networks through jobs and welfare. Attempts by activists to use federal programs to fund activities that promoted collective action or heightened the economic independence of residents were met with serious resistance by local power brokers (Eller, 2008; Kiffmeyer, 1998; Black, 1990; Gaventa, 1982). Due to then-President Kennedy's concern about Appalachian poverty, local efforts succeeded in securing federal assistance to keep open all ten branches of the hospitals opened by the UMWA fund, but the growing numbers of uninsured miners and their families remained unaddressed.

How well was the county's public health system equipped to cope with the growing number of residents losing access to medical care? The county's strikingly low property tax revenue, alongside its history of private health systems provided through coal companies or funded by unions, manifested in low spending on public health. In 1962, Perry County spent $1.29 per resident on public health and hospitals, far lower than the $2.34 spent by the median Kentucky county. Even as funding from the sympathetic Johnson administration began to bolster Appalachian local governments' budgets over the next several years, spending in Perry County generally remained below the median Kentucky county.

By the 1970s, mortality statistics show that Perry County had one of the highest age-adjusted death rates in the state. A brief decline in mortality in the late 1970s was followed by two straight decades of death rates in the top quintile of the Kentucky counties. Between 1980 and 2014, life expectancy in the median county in the region rose by 4 years; in Perry County, it fell by 5 months.

One culprit for the county's high mortality rate is the growing use of opioid drugs. In 2017, Perry County led the nation in hospitalizations for opioid abuse. Drug use has led to other problems, including an outbreak of Hepatitis C. A lack of resources for testing and treatment have impeded the county's efforts to stem the disease. The county department of public health—now part of the combined Kentucky River District Health Department, which serves seven counties in the region—has relied on outside grants to fund needle exchange programs, testing, and staff to help with treatment.

Today, the hospital originally opened by the United Mine Workers of America's Welfare and Retirement Fund remains the only hospital in Perry County. A 2017 report suggests that residents see the local health system as understaffed and overwhelmed, with $\frac{1}{3}$ unsatisfied with their county's access to health care services (Community and Economic Development Initiative of Kentucky, 2019).

To summarize, Perry County began as a remote, sparsely populated agricultural area. Ownership of its land and mineral rights was quickly consolidated around of the turn of the century, and it developed a coal industry large enough to dominate its economy almost overnight. The county's local governments were typically sized and funded before the coal industry arrived but fell behind after it became dominant. Coal companies funded adequate healthcare for many residents while they thrived, but as mine-backed healthcare disappeared, the health of Perry County residents declined and public health infrastructure was insufficient to prevent climbing mortality.

The history of Perry County illustrates how coal companies kept their property tax bills low by building unincorporated coal towns, exploiting and lobbying state-level tax policy, and influencing local processes of assessment and taxation. It illustrates how the coal industry influenced the use of the funding local governments did have, using actors in the justice system to retaliate against threats. And it illustrates the subtler problems with relying on private funding for public goods like healthcare systems, as public-sphere actors struggle to cope with the inevitable departure of those private resources.

5 Conclusion

In this paper, I have presented a combination of descriptive and causal analyses that document the relationship between a coal-dominated economy and local government capacity in the case of Eastern Coal Country in the United States. Drawing on literature on the resource curse, I predicted that localities with economies focused on coal extraction would have smaller, weaker local governments than comparable localities in different economic circumstances.

Comparable coal and non-coal counties employed similar numbers of people in local

government until around 1880, the point at which the coal industry began to expand dramatically in this region. By the late nineteenth and early 20th centuries, counties with larger proportions of their labor force employed in the coal industry tended to have local governments that employed fewer people and collected less revenue. By the postwar period, when coal employment had declined dramatically, local governments in coal areas no longer employed fewer people, but they continued to collect less in revenue and spent less money.

A differences in differences analysis of government employment suggests a causal interpretation of these patterns. Using a sample of counties in which the coal industry developed suddenly at an identifiable time, I demonstrate that coal and non-coal counties, on a similar trajectory of government employment before the industry arrived, diverged after the industry arrived. This suggests that the arrival of a large coal industry caused local government growth to slow.

These results align with past work showing that dependence on resource extraction can lead to weakened government institutions. Applying literature on the resource curse to the case of coal in Appalachia and the Illinois Basin, I have argued that profits from resource extraction accrue to relatively few industry elites, who can invest these gains in extracting preferred policy from governments in mining-heavy areas. These preferred policies include low taxes and high autonomy for the industry, which constitute an overall preference for weaker local governments in mining areas. The results presented here suggest that in this context, mining elites were successful.

The analyses presented here focus on measures of government capacity related to government size: the number of employees, government revenue, and government spending. Data availability constraints prevent me from testing a key mechanism in the story relating industry power to government size: corruption, in the form of mining interests paying off elections and elected officials in return for policy gains. Anecdotal accounts suggest extensive corruption in coal mining areas (Caudill, 1962), but data on corruption cases against local government officials is not available at a granular enough level to test this prediction quantitatively. This is an important step for future work.

Coal is not the only industry that dominated local economies in this region. In other parts of Appalachia and the Midwest, substantial portions of counties' labor forces were

employed in steel mills, automobile factories, and other manufacturing establishments. These industries also constructed company towns and avoided taxation by local governments. In other work, I test whether the model of industry dominance and local government explored here extends to cases beyond resource extraction.

Today, in many former coal-producing areas east of the Mississippi, mining is drawing to a close. Few of these places have developed diverse economies in the wake of coal employment's decline; economic, educational, and health outcomes in many are poor. A possible implication of the results presented here is that this state of affairs was not inevitable. The coal industry hindered the growth of local governments, perhaps preventing investment in human capital and infrastructure that could have better prepared these areas for the inevitable point when the coal was exhausted. Future work should examine more closely the link between local government capacity and the ability to weather industrial decline.

References

Acharya, A., M. Blackwell, and M. Sen (2018). *Deep roots*. Princeton University Press.

Appalachian Land Ownership Task Force (1983). *Who Owns Appalachia? Landowership and its Impact*. University of Kentucky.

Badeeb, R. A., H. H. Lean, and J. Clark (2017). The evolution of the natural resource curse thesis: A critical literature survey. *Resources Policy 51*, 123–134.

Black, K. (1990). The roving picket movement and the appalachian committee for full employment, 1959-1965: A narrative. *Journal of the Appalachian Studies Association 2*, 110–127.

Boone, J. T. and R. Lee (1947). *A Medical Survey of the Bituminous-coal Industry: Report of the Coal Mines Administration*. US Government Printing Office.

Brollo, F., T. Nannicini, R. Perotti, and G. Tabellini (2013). The political resource curse. *American Economic Review 103*(5), 1759–1796.

Callaway, B. and P. H. Sant'Anna (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.

Cantrell, D. (1987). Interview with Bruce Stephens, Jr.

Caudill, H. (1962). *Night Comes to the Cumberland: A Biography of a Depressed Area*. Little Brown, Boston.

Collins, T. (2015). *Selling the State: Economic Development in Kentucky*. Macomb, Illinois: Illinois Institute for Rural Affairs.

Community and Economic Development Initiative of Kentucky (2019). Hazard ARH Regional Medical Center 2019 Community Health Needs Assessment. Technical report, Appalachian Regional Healthcare, Hazard, Kentucky.

David, A. P. and G. Wright (1997). Increasing Returns and the Genesis of American Resource Abundance. *Industrial and Corporate Change 6*(2), 203–245.

Deacon, R. T. and A. Rode (2015). Rent seeking and the resource curse. In *Companion to the political economy of rent seeking*. Edward Elgar Publishing.

Dunning, T. (2008). *Crude democracy: Natural resource wealth and political regimes*, Volume 7. Cambridge University Press Cambridge.

Eller, R. (2008). *Uneven ground: Appalachia since 1945*. University Press of Kentucky.

Fishback, P. (1995). An alternative view of violence in labor disputes in the early 1900s: The bituminous coal industry, 1890–1930. *Labor History 36*(3), 426–456.

Frankel, J. A. (2010). The natural resource curse: A survey. *NBER Working Paper* (w15836).

Gaventa, J. (1982). *Power and powerlessness: Quiescence and rebellion in an Appalachian valley*. University of Illinois Press.

Goldberg, E. F. and G. Power (1972). Legal Problems of Coal Mine Reclamation: A Study in Maryland, Ohio, Pennsylvania, and West Virginia.

Green, H. (2011). *The company town: the industrial Edens and satanic mills that shaped the American economy*. Basic Books.

Isham, J., M. Woolcock, L. Pritchett, and G. Busby (2005). The varieties of resource experience: natural resource export structures and the political economy of economic growth. *The World Bank Economic Review 19*(2), 141–174.

James, A. and D. Aadland (2011). The curse of natural resources: An empirical investigation of us counties. *Resource and Energy Economics 33*(2), 440–453.

Kiffmeyer, T. J. (1998). From self-help to sedition: The appalachian volunteers in eastern kentucky, 1964-1970. *The Journal of Southern History 64*(1), 65–94.

Klesta, M. and Others (2016). Eastern Kentucky: A Region in Flux. *Forefront 7*(1).

Leistritz, F. L. and S. W. Voelker (1975). Coal Resource Ownership: Patterns, Problems, and Suggested Solutions. *Natural Resources Journal 15*(4), 643–662.

Lutz, P. F. (1978). Governor Marland's Political Suicide: The Severance Tax. *West Virginia History 40*(1), 13–27.

Matheis, M. (2016). Local economic impacts of coal mining in the united states 1870 to 1970. *The Journal of Economic History 76*(4), 1152–1181.

Mulcahy, R. (1993). Health care in the coal fields: The miners memorial hospital association. *The Historian 55*(4), 641–656.

Parrish, J. M. (2011). Taxation of Unmined Coal in Kentucky. *Ky. J. Equine Agric. & Nat. Resources L. 4*, 231.

Peterson, P. E. et al. (1981). *City limits*. University of Chicago Press.

Rennick, R. M. (2000). Perry County- Post Offices. Technical report, County Histories of Kentucky.

Rode, A. (2013). Institutions and the resource curse within the united states. Technical report, Working paper.

Roenker, J. (2013). Kentucky's Coal Severance Tax. Technical report, Legislative Research Commission, Frankfort, Kentucky.

Ross, M. L. (2015). What have we learned about the resource curse? *Annual Review of Political Science 18*, 239–259.

Sachs, J. D. and A. M. Warner (1995). Natural resource abundance and economic growth. Technical report, National Bureau of Economic Research.

Smith, J., D. Ostendorf, and M. Schechtman (1978). Who's mining the farm.[land ownership in illinois].

Spence, B., C. Kunkel, E. Schewel, T. Boettner, and L. Martin (2013). Who owns west virginia. *West Virginia Center on Budget & Policy 13*.

Suryanarayan, P. and S. White (2021). Slavery, reconstruction, and bureaucratic capacity in the american south. *American Political Science Review 115*(2), 568–584.

U.S. Census Bureau (1922). *Mines and Quarries: General Report and Analytical Tables and Reports for States and Selected Industries*, Volume 11. Government Printing Office.

Van der Ploeg, F. (2011). Natural resources: curse or blessing? *Journal of Economic literature 49*(2), 366–420.

Vasek Jr, S. J. (1985). The impact and desirability of taxing unmined coal interests in the same manner as other real property. *J. Min. L. & Pol'y 1*, 221.

Wright, G. (1990). The origins of american industrial success, 1879-1940. *The American Economic Review*, 651–668.

Supplemental Information

0.1 Map of Coal Employment

Figure 8 shows the spatial distribution of coal employment in the six-state region of interest in this paper in 1920. The counties in which coal employed the largest proportion of people

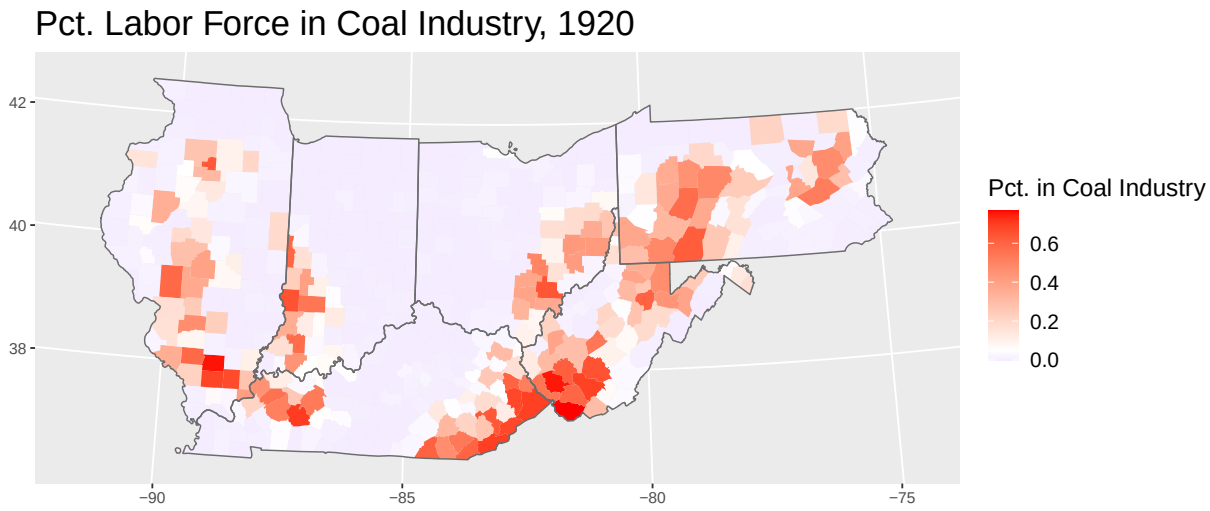

Figure 8: **Map of Coal Employment.** The color of each county represents the proportion of its non-farm, non-domestic labor force employed in the coal industry in 1920.

are concentrated in two bands: the coal deposits of the Illinois basin, stretching across Illinois into Western Kentucky and north into Western Indiana; and Appalachia, stretching from the anthracite coalfields of northeastern Pennsylvania to the southwest through West Virginia and Ohio into Eastern Kentucky.

0.2 Instrumental Variables Analysis

The descriptive results presented in Section 3.4 suggest that counties more dependent on coal employment had smaller local governments. The differences in differences analysis in Section 3.3 supports a causal interpretation of this relationship in the case of government employment: for 30-40 years after the coal industry arrived, coal dependence depressed the growth of the local government workforce. However, the descriptive results for local

government revenue and spending suggest a longer-term effect, with coal areas' governments remaining poorer well into the late 20th century.

Is the long-term association between coal employment and government revenue/spending indicative of a causal relationship? The data available for these measures do not allow for a differences in differences design in this case. However, an instrumental variables analysis can provide a test of an important alternative explanation: that underdeveloped counties select into specializing in coal production, and this selection explains why these counties' governments are poorer later on.

Counties need substantial coal deposits to develop an economy that specializes in coal mining; the tonnage of coal deposits in a county is therefore an excellent predictor of the proportion of a county's workforce employed in coal ($F = 299$ in 1920). For a measure of coal deposits, I rely on the United States Geographical Survey's USCOAL database. USCOAL was an effort to measure the amount of coal resources present in all deposits in the United States. Compiled from a variety of surveys in sources between the 1950s and 1980s, USCOAL contains, when possible, an estimate of the total tonnage of coal originally present in each deposit; otherwise, it contains an estimate of the coal remaining in the ground in the 1950s. Aggregated to the county level, this data provides an estimate of the total coal resources available in a county.

In terms of relevance, coal deposits are a strong instrument for coal employment. The other key criterion for a good instrument is the exclusion restriction: are coal deposits associated with local government size only through their relationship with coal employment? It seems plausible that coal in the ground cannot affect local governments other than through the mining industry that develops to extract it.

One potential exclusion restriction concern is that in the United States, coal-endowed areas tend to be relatively rugged and therefore remote. Because of spatial clustering in these variables, variation in coal deposits could be shared with variation in ruggedness, which could violate the exclusion restriction if ruggedness affects government size other than through the coal industry—a plausible concern. Though coal deposits are associated with rugged terrain in this region, this association disappears when accounting for state differences in ruggedness, so the inclusion of state fixed effects in these regressions should address this

concern.

By instrumenting for coal-industry dependence with coal deposits, I am only using the variation in coal dependence that can be explained by tonnage of coal deposits when estimating the relationship between coal dependence and local government size. This removes the possibility that some coal-endowed counties choose to specialize in coal, while others don't, in ways that could explain the relationship between coal dependence and local government size.

To estimate the relationship between local government size and the instrumented measure of coal dependence, I regress the three measures of local government size discussed in the main text—number of government employees, amount of property tax revenue, and total amount of spending—on state fixed effects, population size, property tax base, and the proportion of a county's nonfarm labor force employed in the coal industry, instrumented by the tonnage of coal deposits in the county. As in the models in the main text, all observations are weighted by 1880 population size, and standard errors are calcluated using the HAC adjustment proposed by Kelly (2020). From 1880-1940, the independent variable is measured in the same year as the outcome and controls; starting in 1957, the independent variable is coal employment in 1920.

Results are presented in Figure 9, Figure 10, and Figure 11. Each figure presents the coefficient on coal employment in the instrumental variables model described above, as well as the OLS results from the main text for comparison. The coefficients in the instrumental variables models are generally less precisely estimated than the OLS coefficients (especially in earlier years, when the coal deposit instrument is weaker), but in most cases, the two are similar in magnitude; in a substantial majority of cases, the IV coefficients are within the confidence interval of the OLS estimates.

In general, the results of the instrumental variables models point to the same conclusion as those of the OLS results discussed in the main text. Counties with more coal employment had local governments with fewer employees through 1940, but not consistently afterwards; they collected less revenue consistently over time; and they spent less money in the 1950s, though perhaps not later on.

These results suggest that the descriptive relationship between coal employment and local

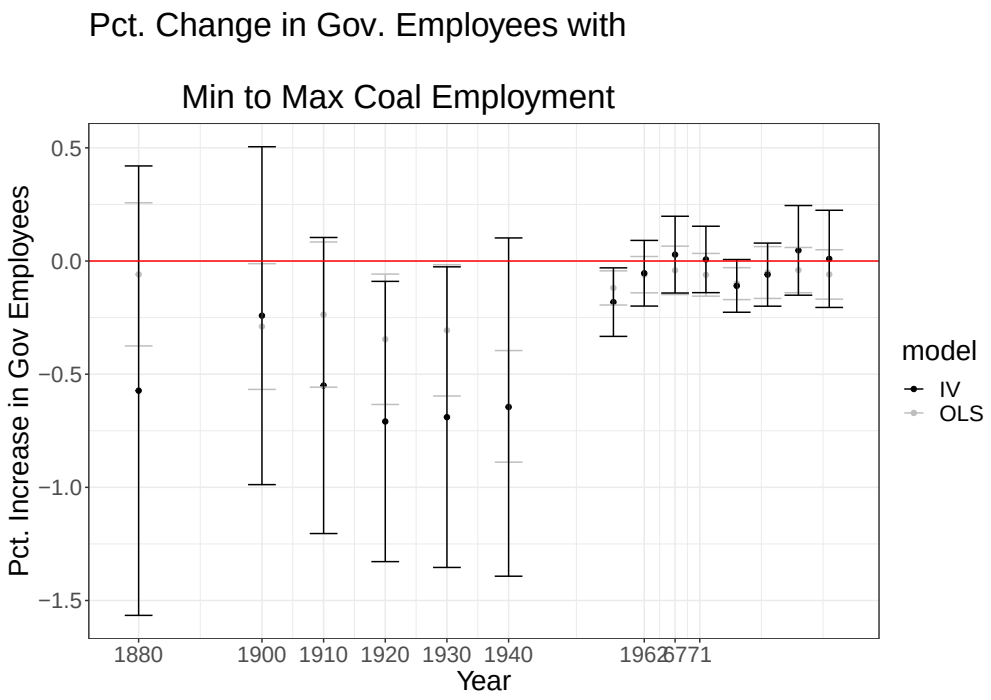

Figure 9: **IV Results: Employment.** Points and lines in black represent coefficients and 95% confidence intervals on coal employment, instrumented using coal deposits, in predicting the number of government employees in a county, controlling for state, population, and property value. Gray points and lines represent the OLS regression estimates presented in the main text for comparison.

government size cannot be explained by coal-endowed counties selecting into specializing in resource extraction for reasons that lead to smaller local governments.

0.3 Difference in Differences: Effects by Cohort

The difference in differences analysis presented in the main text shows the average effect of coal industry onset on government employment over time, combining information across five different "treatment cohorts:" that is, groups of counties that developed a coal industry in a particular year. Figure 12 shows estimates for each cohort separately.

The trend reflected in the averaged results is visible in each group, with the exception of the 1900 cohort, though the standard errors and exact estimates differ. Before treatment onset, there is no clear pattern of differences in trends between treatment and control counties.

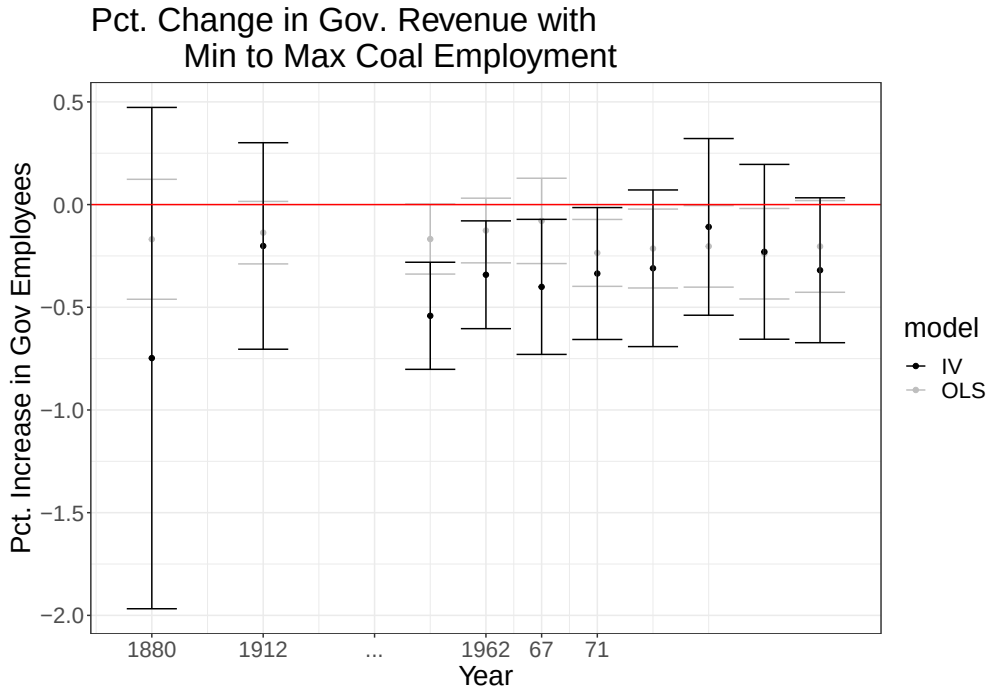

Figure 10: **IV Results: Revenue.** Points and lines in black represent coefficients and 95% confidence intervals on coal employment, instrumented using coal deposits, in predicting the total property tax revenue for local governments in a county, controlling for state, population, and property value. Gray points and lines represent the OLS regression estimates presented in the main text for comparison.

Starting in the first year of treatment and continuing for 30-40 years, governments in treated counties grow more slowly than governments in control counties. This trend disappears, and in some cohorts reverses, as distance from treatment increases.

0.4 Difference in Differences: Robustness Checks

The differences-in-differences analysis presented in Section 3.3 of the main text uses all counties with a coal employment proportion that never exceeds 5% as the comparison group for the 23 treated counties of interest. Though the government employment trends pre-treatment appear similar in the treated and control groups, it remains possible that the trends in the potential outcomes in the treated and control groups under control diverged at the time of treatment. To check the plausibility of this concern, I now repeat the differences-

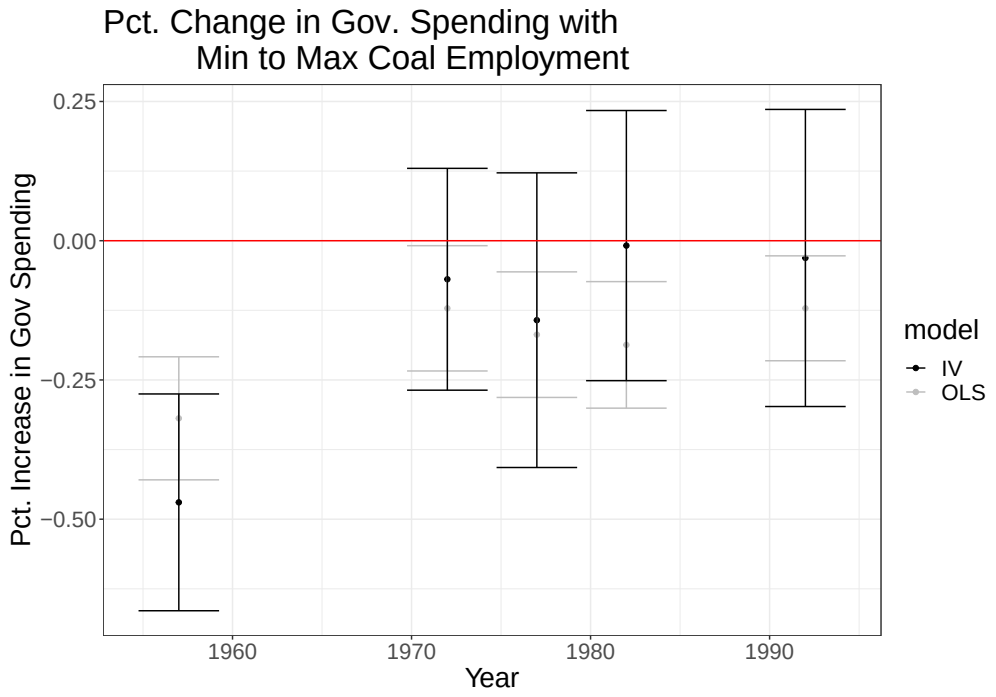

Figure 11: **IV Results: Spending.** Points and lines in black represent coefficients and 95% confidence intervals on coal employment, instrumented using coal deposits, in predicting the total spending by all local governments in a county, controlling for state, population, and property value. Gray points and lines represent the OLS regression estimates presented in the main text for comparison.

in-differences analysis using more limited samples that address potential differences between treated and control units[9].

One potential concern is slavery. Kentucky and West Virginia, the two heaviest coal-mining states in the sample, were also the only two states in which slavery was common. Local governments in slaveholding areas were changing in the late 1800s in the processes of Reconstruction and Redemption (Suryanarayan and White, 2021), and it is possible that the paths of slave- and non-slave areas diverged at similar times to the coal industry's arrival. I therefore repeat the analysis using only counties with no slaves counted in the 1860 Census[10]

[9]See Section 0.5 for a map of the treatment counties and those used as comparison counties here and in the main text.

[10]This eliminates a number of control counties in Kentucky and West Virginia, as well as 6 of the 23 treatment counties.

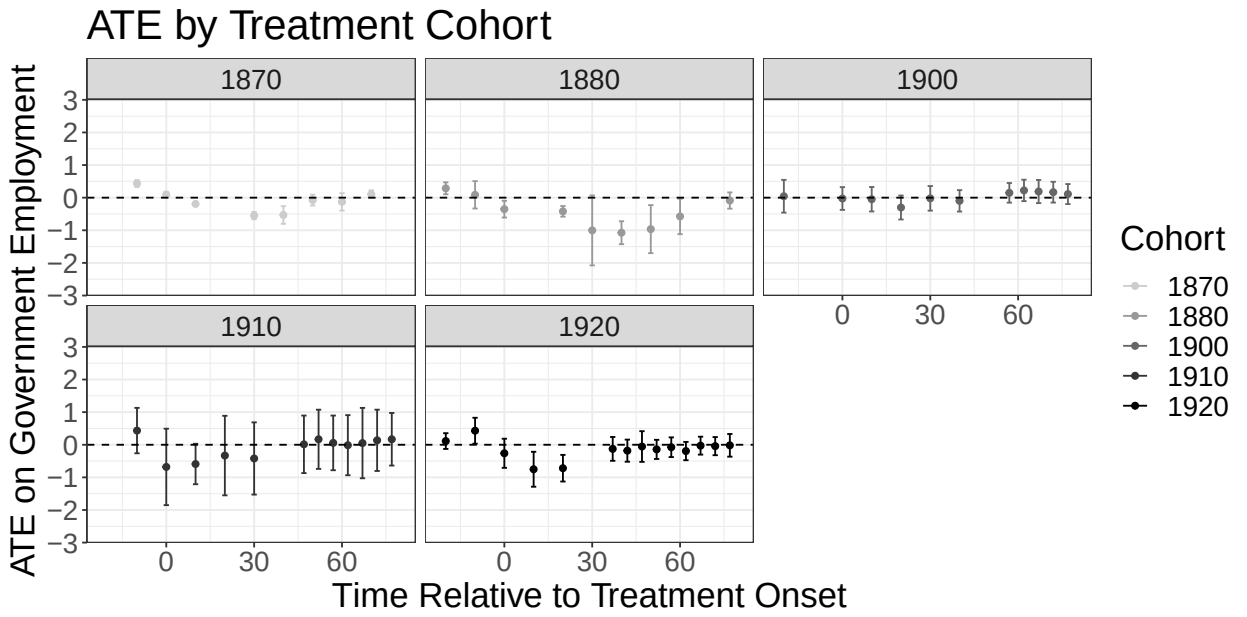

Figure 12: **Differences in Differences: Government Employment in Coal and Non-Coal Counties Over Time, By Treatment Cohort.** On the y-axis is the difference between treated and control groups in the difference in government employment from the previous time period, with different estimates for each group of counties based on the year in which the industry arrived. The measure of government employment is logged and residualized for state and population differences. The x-axis shows time relative to treatment; -20 represents 20 years before receiving treatment, while 20 represents 20 years afterwards. Bars represent bootstrapped 95% confidence intervals.

using data from Acharya et al. (2018). The results are presented in Figure 13; they are quite similar to the results using the full sample.

Another potential concern is that coal-endowed areas in this region are often rugged and remote. These areas may have been set on different trajectories than the rest of the region as industrialization progressed around the turn of the century. To address this concern, I repeat the differences-in-differences analysis using only counties with some amount of coal deposits, using data from the USGS (control n = 41; further details can be found in the Instrumental Variables section of the SI). The results can be found in Figure 14. Again, results are similar to those in the full sample.

Finally, the differences in differences analysis presented in the main text relies on a group

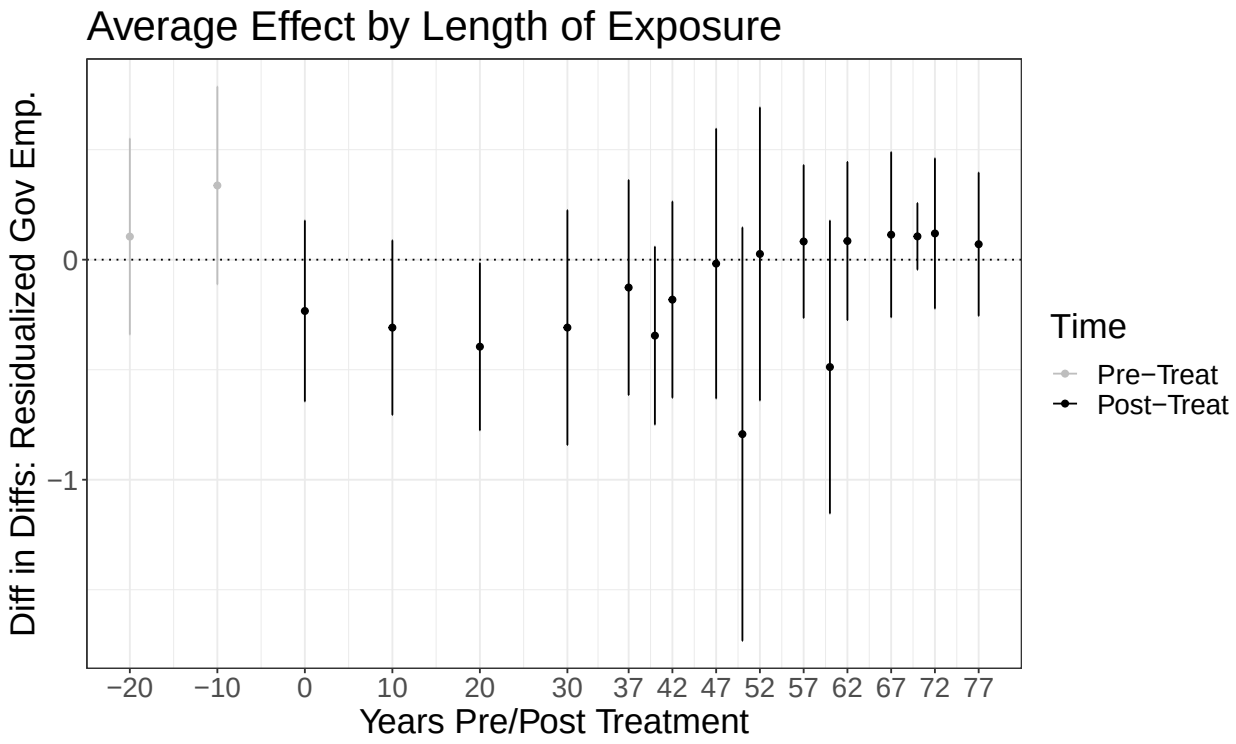

Figure 13: **Differences in Differences: Government Employment in Coal and Non-Coal Counties Over Time, Non-Slave Counties Only.** On the y-axis is the difference between treated and control groups in the difference in government employment from the previous time period. The measure of government employment is logged and residualized for state. time, and population differences. The x-axis shows time relative to treatment; -20 represents 20 years before receiving treatment, while 20 represents 20 years afterwards. Bars represent bootstrapped 95% confidence intervals.

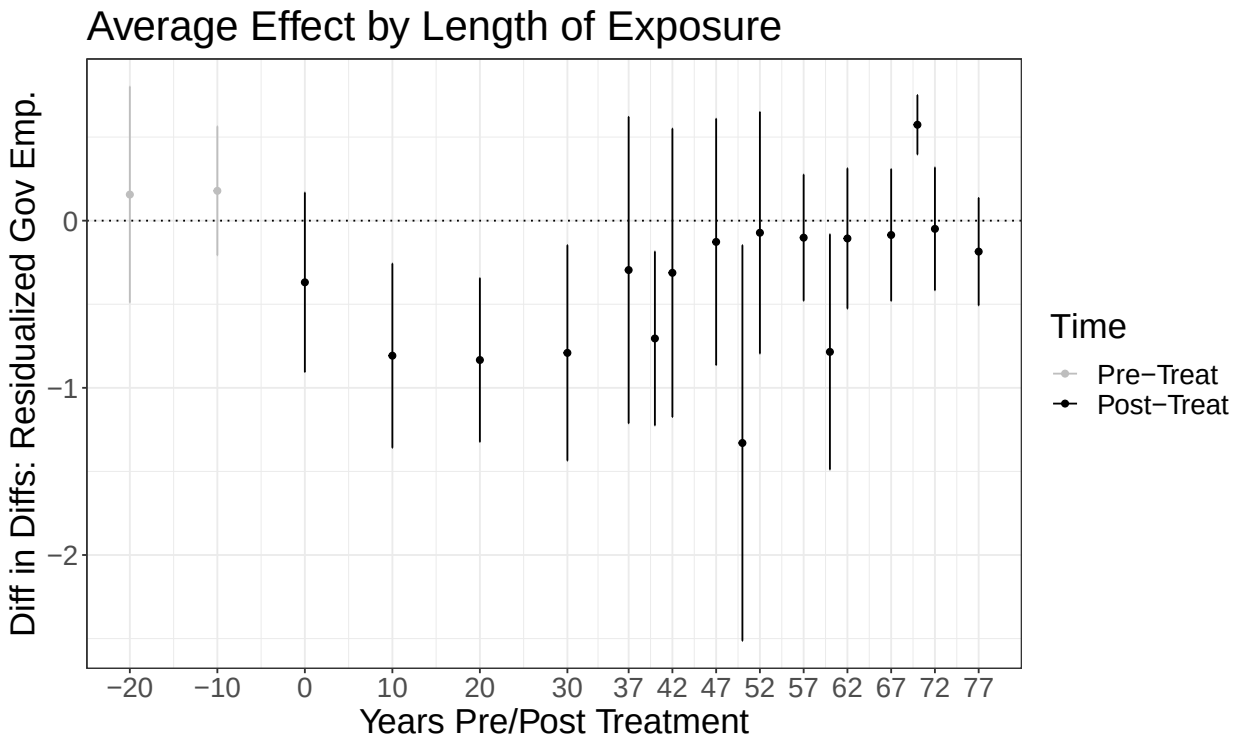

Figure 14: **Differences in Differences: Government Employment in Coal and Non-Coal Counties Over Time, Coal-Endowed Counties Only.** On the y-axis is the difference between treated and control groups in the difference in government employment from the previous time period. The measure of government employment is logged and residualized for state and population differences. The x-axis shows time relative to treatment; -20 represents 20 years before receiving treatment, while 20 represents 20 years afterwards. Bars represent bootstrapped 95% confidence intervals.

of 23 counties that quickly developed dominant coal counties. This quick development allows me to precisely identify the time at which these coal counties were first "treated" with the industry, which is useful for analysis. However, limiting the analysis to this small a group of counties makes the analysis vulnerable to the concern that only one or a few counties could be driving the result.

To address this concern, I repeat the pooled analysis discussed the main text (combining periods 0-40 years after treatment) 23 times, each time leaving out one of the treated counties. The results of these analyses are presented in Figure 15. Where the full sample pooled analysis ATT was $-.30$, the "leave-one-out" analyses range from $-.37$ to $-.24$. This suggests that the estimated treatment effect is robust to slight variations in the treatment sample, and a single treated observation is not producing the observed results.

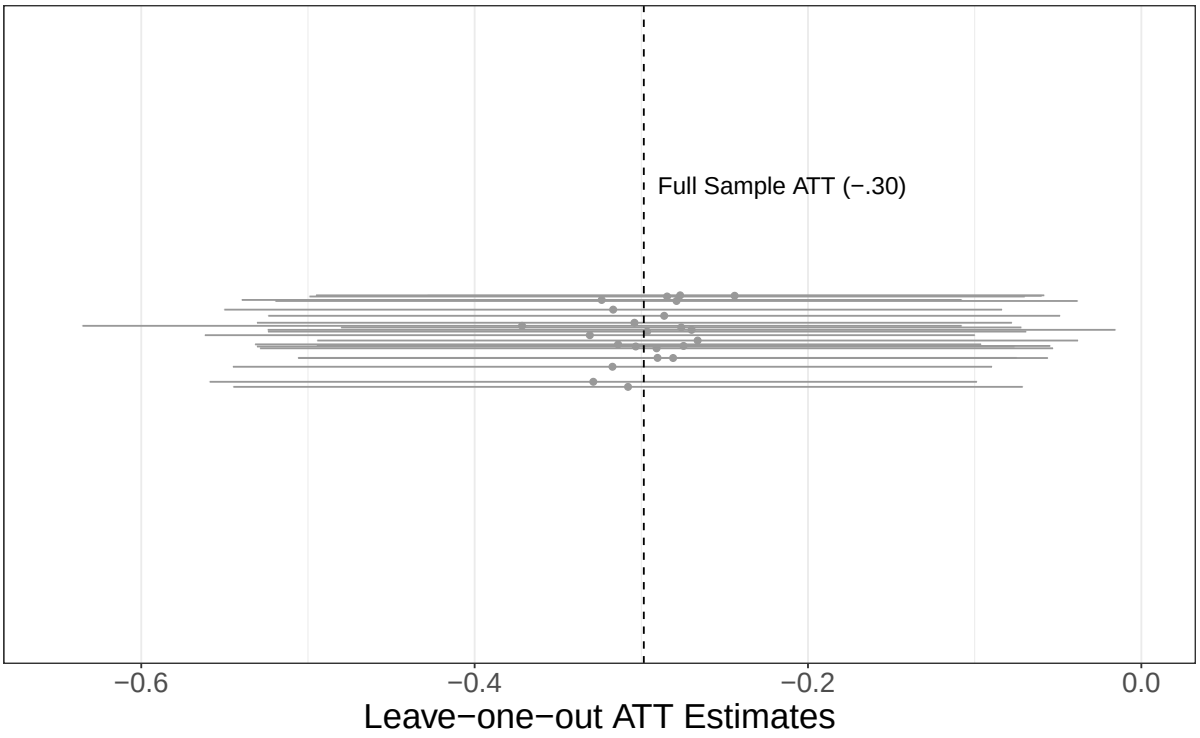

Figure 15: **Differences in Differences: Government Employment in Coal and Non-Coal Counties Over Time, Coal-Endowed Counties Only.** On the y-axis is the difference between treated and control groups in the difference in government employment from the previous time period. The measure of government employment is logged and residualized for state and population differences. The x-axis shows time relative to treatment; -20 represents 20 years before receiving treatment, while 20 represents 20 years afterwards. Bars represent bootstrapped 95% confidence intervals.

0.5 Characteristics of Treated and Comparison Counties: Differences in Differences

First, to clarify the choice of treatment counties for the difference in difference analysis, Figure 16 shows the trajectories of coal employment over time in all counties in the region in which coal employed at least 40% of the labor force at any point between 1850 and 1995.

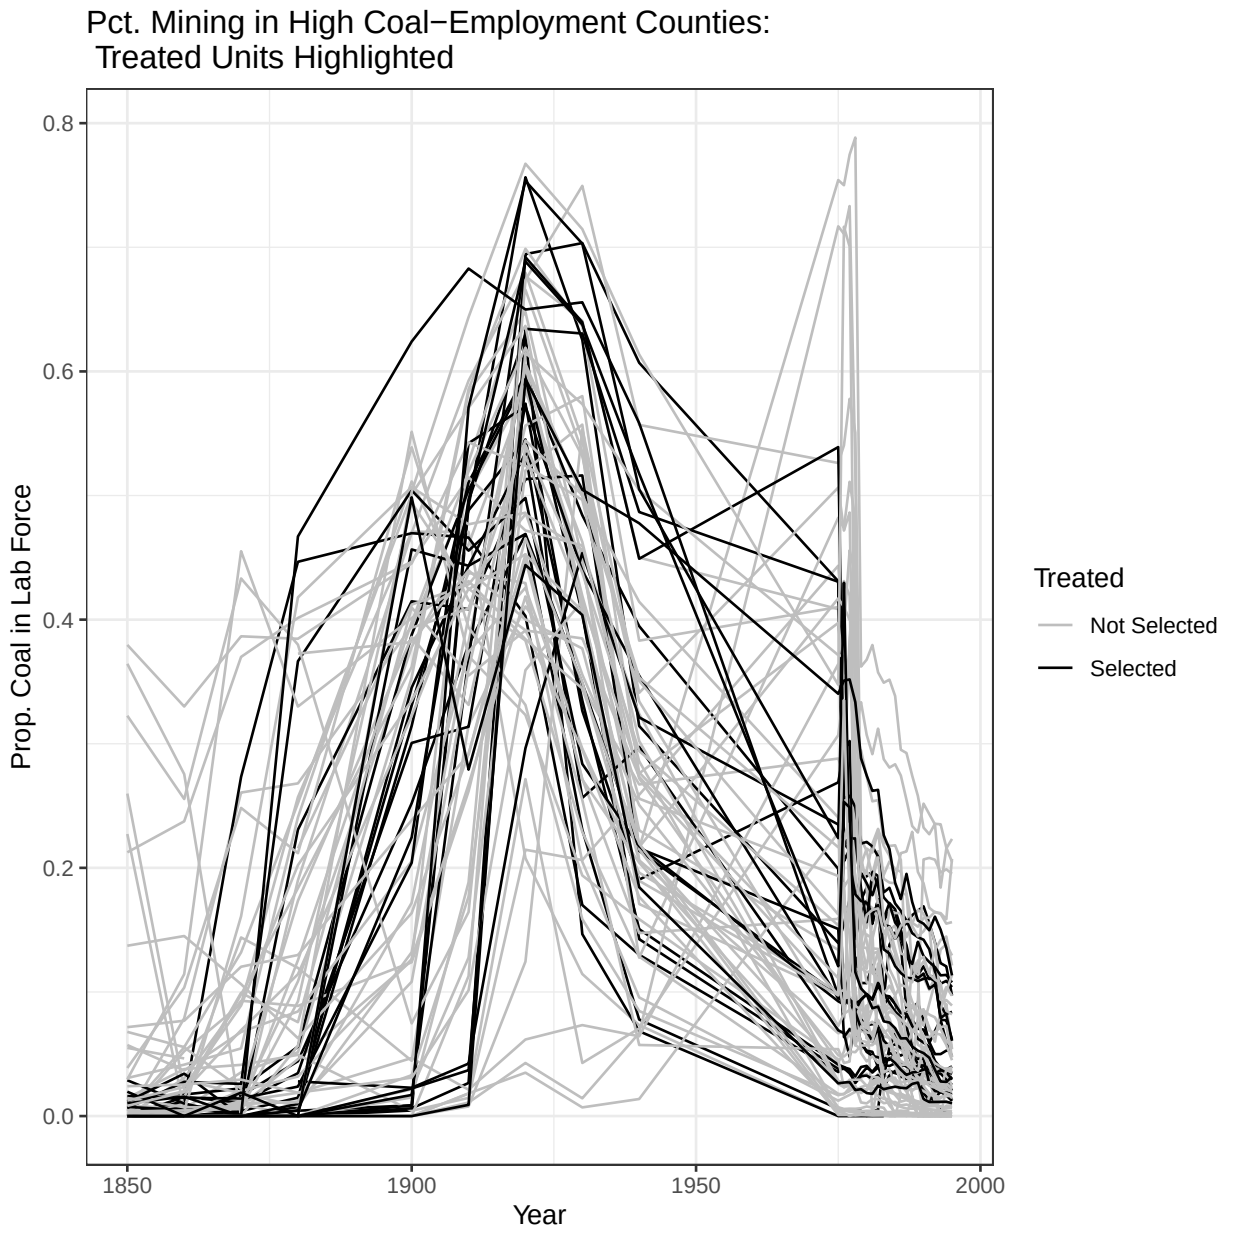

Figure 16: **Coal employment in high-coal counties.** Lines show the proportion of each county's labor force in the coal industry at each year, for all counties with at least 40% of their labor force in the coal industry at some point over this period. Treatment counties in DiD analysis are highlighted in black.

The treatment counties used in the difference in differences analysis are shown in black, while all other high-coal-employment counties are shown in gray. The treatment counties stand out for the nearly-vertical lines indicating the sudden (i.e. between two decennial

Censuses) emergence of a large coal industry in the late 1800s or early 1900s. The coal industries in these treated counties peaked and declined at different times, though most were highest between 1900 and 1920.

High-coal-employment counties that were not used as treatment comparisons were rejected for one of three reasons. First, they did not allow for the comparison of pre-treatment trends, because data were missing or because their industries arrived too early. Second, their coal industries developed too gradually, not allowing the identification of a single time period at which the industry arrived. Third, counties whose industries developed after 1940 were rejected because of a lack of coal employment data between 1940 and 1975, which makes it impossible to identify when the industry arrived. I dropped a final county because its coal industry employed a substantial proportion of its labor force in only one Census, reasoning that this time would have been too short to allow a durable effect of the industry on governance.

Selected characteristics of the counties in the treated group and control group, as well as all counties not used in either, can be found in Table 17. The median treated county was substantially smaller than control and other counties in 1870, but had surpassed them in population by 1920; the groups had similar populations by 1970. No group's median county had a noticeable proportion of its non-farm labor force working in coal in 1870, but the groups diverged in 1920 and afterwards. The number of government employees per 10,000 residents in the median treated and other counties remain lower than in the median control county throughout the period. Finally, the table reports the proportion of the median county's labor force employed by its largest industry, regardless of what the largest industry was; the median control county's employment was less concentrated than treated and other counties, especially in 1920.

Next, to illustrate the comparisons being made in the differences in differences analysis in Section 3.3 in the main text, as well as the robustness checks in the previous section, Figure 18 maps the counties in the "treatment" and comparison groups.

In Figure 18, the red counties are the 23 "treatment" counties in the difference in difference analysis; all are places that rapidly developed coal industries at an identifiable point in time between 1870 and 1920. All the counties in various shades of blue were included as con-

Group:	Control	Other	Treated
N	280.00	221.00	23.00
Population			
1870	14040.00	13119.00	7575.00
1920	22573.50	23979.00	30819.00
1970	28765.00	25721.00	28315.00
Pct. Working in Coal			
1870	0.00	0.01	0.00
1920	0.00	0.13	0.57
1975	0.00	0.02	0.10
Gov. Emp. Per 10k Residents			
1870	5.50	4.52	3.16
1920	8.55	6.87	5.52
1972	339.30	321.02	320.94
Prop. in Largest Industry			
1880	0.07	0.09	0.09
1920	0.13	0.25	0.58

Figure 17: Characteristics of the median countries in the control group, treated group, and "other", counties in neither group.

trol counties in the analysis presented in the main text; these are the counties in the region that never employed more than 5% of their workforce in coal between 1850 and 1940. The non-shaded counties had middling coal employment or coal employment that grew gradually, and they are not included as comparisons in any of the difference in difference analyses.

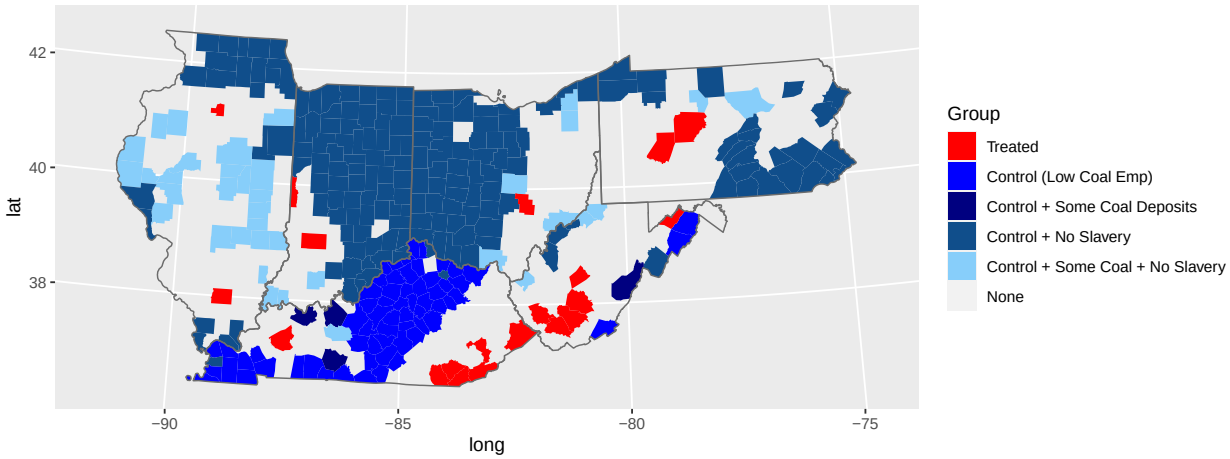

Figure 18: **Counties Used in Difference in Differences comparisons.**

The different shades of blue represent different categories of counties based on the variables used in the robustness checks in the previous section: slaves per household (from Acharya et al., 2018) and presence of coal deposits. Counties of the color labeled "Control (Low Coal Emp)", mostly located in non-coal-producing parts of Kentucky, were only included as control counties in the main text and were eliminated from both robustness checks for slaveholding and a lack of coal deposits.

Many counties in Indiana and Ohio (and several elsewhere) had no coal deposits, but were not slaveholding areas, so they were included as controls in the analysis eliminating slave areas. A handful of places in Kentucky and West Virginia had coal deposits but held slaves, so they were included in the test limited to coal-endowed counties. Finally, a group of counties largely in Illinois had both coal deposits and no slaves and were eligible for both robustness-test comparisons.

0.6 Maps of Relevant Features

This section contains maps of the 6-state region of interest in this paper. Each map colors counties based off of some characteristic of theoretical interest. I first show a map of the key independent variable of interest, the percent of a county's labor force employed by the coal industry, at coal employment's peak in 1920. The maps that follow illustrate how this key variable is related to other characteristics of coal- and non-coal counties.

Figure 19 shows the proportion of each county's non-farm labor force employed in the coal industry in 1920. The two major stretches of coal employment are located in the Illinois basin mining region, stretching from western Kentucky north into Illinois and Indiana, and Appalachia, reaching from eastern Kentucky north through Ohio and West Virginia into Pennsylvania.

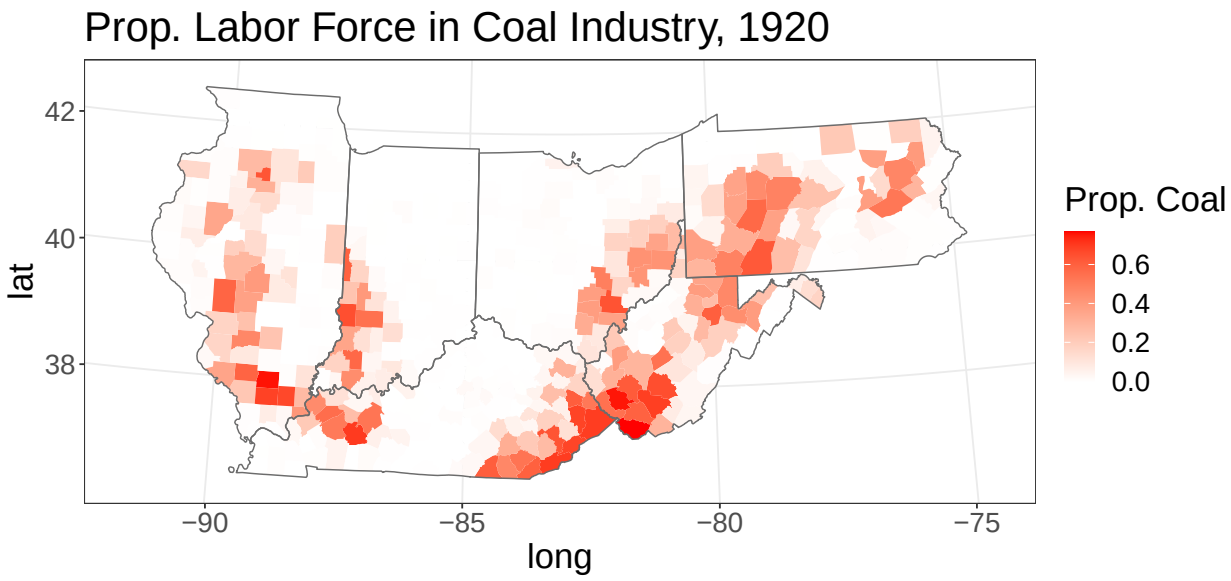

Figure 19: **Percent of non-farm labor force in coal industry in 1920.**

Figure 20 shows the original (i.e. before mining activity) tonnage of coal deposits in each county, as estimated by the USGS in the mid-to-late 20th century (see the Instrumental Variables section of the appendix for more information). The distribution of resources largely tracks with the pattern of coal employment.

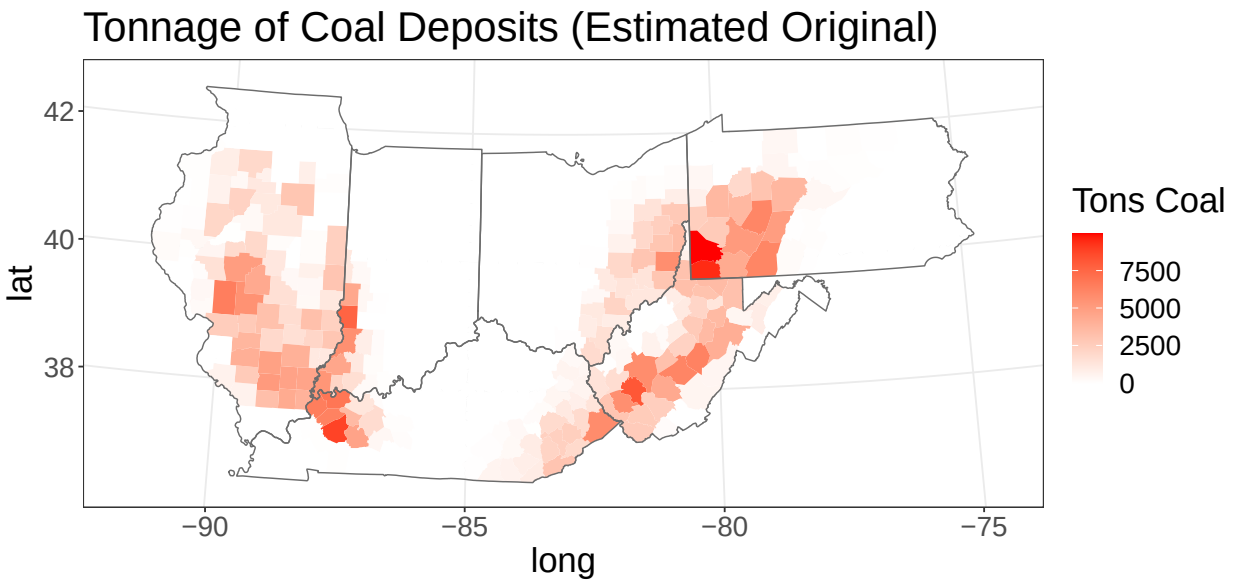

Figure 20: **Tonnage of coal deposits.**

Figure 21 shows counties' access to two important forms of transportation—railroads and navigable waterways—in 1860, before the large-scale coal industry had arrived in much of the region (from Acharya et al., 2018). Counties are coded as having rail access if at least one railroad passed through them, and counties are coded as having water access if they contained a steamboat-navigable river or canal, or access to a Great Lake. Though much of West Virginia and eastern Kentucky were inaccessible at this time, much of the rest of the region, including other areas that would develop strong coal economies, was equipped with some access to transportation.

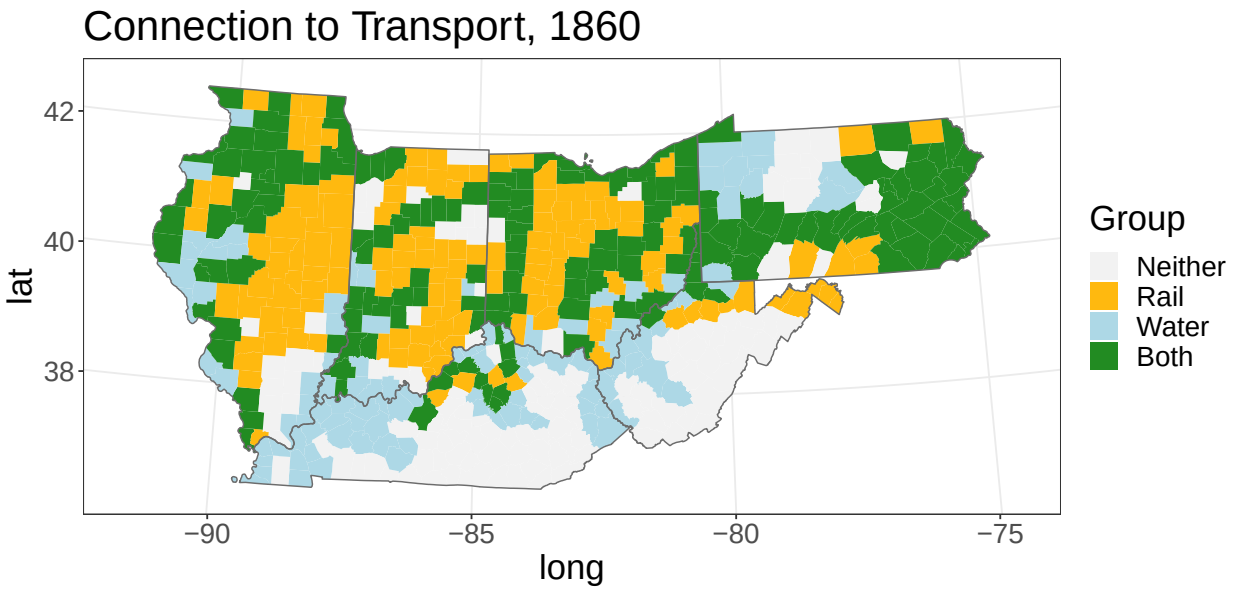

Figure 21: **Access to rail and water transportation.**

Figure 22 shows the ruggedness of each county's terrain, measured by the standard deviation of altitudes across the county (from Acharya et al., 2018). Some have attributed the low development of coal mining areas to their rough terrain. As this map shows, the coalfields of Appalachia do have more variation in altitude than other areas of the region, but coal-producing areas in the Illinois Basin are quite flat.

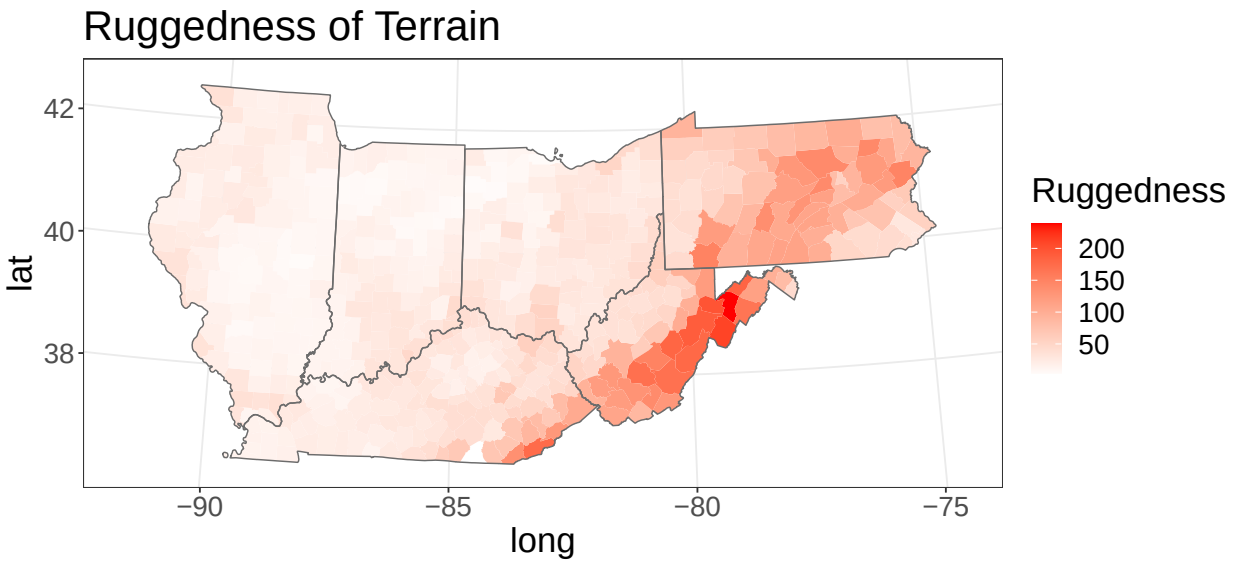

Figure 22: **Ruggedness of terrain.**