Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America

Hunt Allcott and Daniel Keniston*

March 13, 2013

Abstract

PRELIMINARY FIRST DRAFT

The rise in oil and gas prices and drilling activity in the past decade has caused economists and policymakers to reconsider whether natural resource production benefits producer economies or instead creates a "Natural Resource Curse." We use confidential establishment-level data from the US Census of Manufactures and Longitudinal Business Database to estimate the effects of expansions and contractions of the oil and gas sector on growth since the early 1970s. Our approach combines cross-county variation in oil and gas supply with large time series variation in production activity. Oil and gas booms increase growth rates in producer counties by 60 to 80 percent relative to non-producer counties, and a necessary condition for the resource curse is satisfied: local wages increase by 0.3 to 0.5 percentage points per year during a boom. Nevertheless, manufacturing growth is positively associated with natural resource booms. Manufacturing employment and output both rise, while productivity does not, suggesting that at least in the rural counties we study, manufacturing firms benefit from increases in local demand.

We thank Yunmi Kong and Pingting Wei for superb research assistance and Julia Garlick for help in data preparation. We are grateful for feedback from Dan Black, Pat Kline, and seminar participants at Cornell and NYU. Thanks also to Randy Becker, Allan Collard-Wexler, Jonathan Fischer, Todd Gardner, Cheryl Grim, Javier Miranda, and Justin Pierce for their advice and help in using U.S. Census Data, and Jean Roth for advice on the Current Population Survey.

*Allcott: NYU Department of Economics and NBER. Email: hunt.allcott@nyu.edu. Keniston: Yale Department of Economics and the Cowles Foundation. Email: daniel.keniston@yale.edu.

1 Introduction

Williams County is a rural area in western North Dakota that until recently was best known for being the host of the Miss North Dakota pageant. Between 2006 and 2011, however, the county went from producing $382 million of oil and gas per year, or about $19,000 annually for each of its 20 thousand inhabitants, to producing $2.1 billion. This fast growth, brought on by high oil prices and improvements in drilling technologies, has dramatically changed the area.

"It's hard to think of what oil hasn't done to life in the small communities of western North Dakota," writes the New York Times Magazine (Brown 2013). "It has minted millionaires, paid off mortgages, created businesses; it has raised rents, stressed roads, vexed planners and overwhelmed schools ... It has forced McDonald's to offer bonuses and brought job seekers from all over the country - truck drivers, frack hands, pipe fitters, teachers, manicurists, strippers." Locals hope that unlike the previous boom of the 1970s and early 1980s, this boom "won't afflict the state with the so-called Dutch Disease in which natural-resource development and the sugar rush of fast cash paradoxically make other parts of the economy less competitive and more difficult to sustain."

Of course, oil and gas production has affected producer economies worldwide, not just in North Dakota: the economic histories of Canada, Iraq, the Netherlands, Nigeria, Qatar, Venezuela, and many other countries have been changed - perhaps in very different directions - by the booms and busts of the past 40 years. And the interest extends beyond the New York Times. Policy makers are asking how to respond to booms: for example, should they encourage development through low royalties and accommodating local development policies? Or should they discourage development and perhaps even ban new drilling technologies, as New York State has done? Planners and businesses want to know how producer economies will fare if and when the current boom ends. Economists have started to revisit these and related questions, both in the United States (Black, McKinnish, and Sanders 2005a, Carrington 1996, Michaels 2010, etc.) and internationally (Aragon and Rud 2011, Caselli and Michaels 2013, Collier and Goderis 2009, Dube and Vargas 2012, Harding and Venables 2013, Mehlum, Moene, and Torvik 2006, Sachs and Warner 1995, etc.).

In this paper, we ask *how do oil and gas booms and busts differentially affect the manufacturing sector in US counties that produce oil and gas vs. counties that do not?* We focus on the manufacturing sector because of concern about a local version of "Dutch Disease": an increase in labor demand from the non-tradable and natural resource sectors drives up local wages. This causes the manufacturing sector to contract, as it often sells output into national or international output markets with exogenous prices. In the long run, this slows growth if there are local learning-by-doing spillovers in the manufacturing sector, as suggested by Krugman (1987), Matsuyama (1992), and van Wijnbergen (1984). Dutch Disease is perhaps the most likely mechanism through which a "Natural Resource Curse" could afflict a developed economy with advanced environmental regulation and stable political institutions. This is an empirical question: it is not obvious that the basic conditions for Dutch Disease would be satisfied in a study of US counties. If manufacturing wages don't rise because migration softens the labor demand shock or because manufacturing workers are not closely substitutable with oil and gas workers, then the manufacturing sector is less likely to contract during a boom.

We also focus on a second channel through which natural resource booms could affect the manufacturing sector: positive productivity spillovers. This is inspired by recent empirical work documenting agglomeration economies: increases in productivity driven by an increased density of economic activity (Glaeser and Gottlieb 2009). Greenstone, Hornbeck, and Moretti (2010) show that the entry of large manufacturing establishments can generate positive spillovers to other nearby producers, while Bleakley and Lin (2012) and Kline and Moretti (2012) show that agglomerative effects can persist long after the initial source of comparative advantage has disappeared. This also is an empirical question: it is not obvious that oil and gas drilling would cause the same kind of productivity spillovers as new manufacturing plants, trade and financial hubs, or planned place-based policies.

Our study exploits extraordinary historical data on local economies in the United States. We gather a new county-level panel dataset of oil and gas production using market research data and historical records from state regulatory agencies. We exploit rich publicly-available data on employment, earnings, population, and hourly wages from the Regional Economic Information System and the Current Population Survey. We also use restricted-access establishment-level microdata from the U.S. Census of Manufactures and Longitudinal Business Database, which allow us to examine employment, output, and productivity and design nuanced tests of whether and why natural resource booms affect different types of firms.

We combine these data in what is effectively a difference-in-differences design, which estimates the difference in outcomes between oil and gas producer and non-producer counties in boom years vs. busts. We define counties as "treatment" counties if they produce oil or gas between 1969 and 2011, and "control" if they do not. We then test whether changes in outcomes differ between treatment and control as the oil and gas sector expands and contracts between the early 1970s and today. These expansions and contractions are very large: between 1974 and 1982, national oil and gas employment expands by 600,000 jobs. Between 1983 and 1992, about 400,000 of those jobs disappear. Then, between 2004 and 2011, national oil and gas employment again increases by 400,000. We both quantify and discuss the conceptual implications of treatment effect spillovers to control counties, for example when people migrate from a control county to work in a treatment county.

We can illustrate a simple version of our empirical strategy by returning to North Dakota. The black line on Figure 1 shows total state-level oil and gas employment from 1969 to 2011. Clearly visible are the boom and bust of the 1970s and 1980s and the second boom of the most recent decade. The dashed blue line shows the unconditional difference in means between log manufacturing employment in the 19 counties in the state that produce oil and gas and the 31 counties that do not. Manufacturing is declining in the treatment counties before the 1970s boom, but this decline is arrested and eventually reversed just as the resource boom nears its peak. As resource employment contracts beginning in the early 1980s, manufacturing employment also contracts in treatment relative to control. The more recent boom tells the same qualitative story: state-level oil and gas employment and relative manufacturing employment have both increased in every year since 2005. The data from this small case study in one state do not appear to be consistent with the Dutch Disease model. If anything, manufacturing growth appears to be *positively* associated with oil and gas booms.

Our nationwide regressions exploit a much larger sample size and condition on a battery of control variables and fixed effects. We find three key results. First, a necessary condition for Dutch Disease is satisfied: natural resource booms (busts) increase (decrease) growth rates in treatment counties by 60 to 80 percent, and attendant to a boom are increases in wages of between 0.3 and

0.5 percentage points per year. Wages per worker are two to five percent higher in treatment counties for an extended period, from 1978-1986. Second, however, Dutch Disease appears not to afflict the manufacturing establishments we study: just as in North Dakota, manufacturing growth is *positively* associated with resource booms. During booms (busts), manufacturing employment increases (decreases) by 0.7 percentage points per year, and sales revenues increase (decrease) by 1.8 percent per year. Furthermore, there is little evidence of negative association even for the most potentially-vulnerable industries: those that are labor intensive, produce goods for more distant markets, and are not immediately upstream or downstream of the oil and gas sector. Third, there is little evidence that resource booms cause manufacturing productivity to increase: most of our point estimates are not statistically different than zero. For estimates including all manufacturing establishments, our standard errors are tight enough to rule out that revenue productivity and output prices increase (decrease) by more than about 0.15 percent per year during natural resource booms (busts).

These basic results are remarkably robust to a variety of alternative empirical approaches, including different controls for pre-boom levels and trends, state-by-year fixed effects, trimming to a common support of propensity scores, and using a "doubly-robust" estimator with both regression controls and inverse probability weights (Robins, Rotnitzky, and Zhao 1995). Our interpretation is that at least some of the manufacturing plants in our sample of rural counties sell some output into local markets and thus benefit from the increase in local economic activity driven by the natural resource boom. This suggests a new model of how resource booms might affect the manufacturing sector: if manufacturing firms sell into both local markets and more competitive national and international markets, an increase in local demand from a resource boom coupled with an increase in labor input can increase total output, even if the firm is less competitive on national and international markets.

In the remainder of this first section, we discuss related literature and how we contribute. Section 2 provides a background on the modern evolution of the oil and gas sector. To fix ideas about the channels through which the natural resource sector can affect producer economies, Section 3 presents a model of an open economy with tradable, non-tradable, and resource sectors. Section 4 details our data, Section 5 outlines the empirical strategy, and Section 6 presents results. Section

7 concludes.

1.1 Literature

This paper combines the literature on the Natural Resource Curse with the urban and labor economics and industrial organization literatures on agglomeration and productivity spillovers. The Natural Resource Curse is the potentially counterintuitive idea that an endowment of valuable natural resources might reduce welfare; van der Ploeg (2011) provides an overview of the theoretical and empirical literature. There are many anecdotal counterexamples to the Resource Curse, such as Botswana and Norway (van der Ploeg 2011), and David and Wright (1997) and Wright and Czelusta (2007) argue that resource abundance has been a crucial ingredient in US economic growth. This has motivated econometric studies to understand the average effect of resource abundance, as well as sources of heterogeneity in that effect. Sachs and Warner (1995) show that natural resource abundance is indeed negatively correlated with economic growth, but some subsequent analyses arrive at different results when instrumenting for resource abundance (Brunnschweiler and Bulte 2009), including country fixed effects (Manzano and Rigobon 2001), or conditioning on the quality of institutions (Collier and Goderis (2009), Mehlum, Moene, and Torvik (2006)). Van der Ploeg (2011) argues that "the cross-country and panel data results are sensitive to changing the sample period, the sample of countries, or the definition of various explanatory variables.... The road forward might be to exploit variation within a country where variables that might confound the relationship between resources and macroeconomic outcomes do not vary and the danger of spurious correlation is minimized." This is the road that we follow.

In recent years, other analyses have followed this same road forward using cross-state regressions in the United States (James and James (2012), Papyrakis and Gerlagh (2007)) and in Spain (Domenech 2008), as well as more local sources of variation in developing countries (Aragon and Rud (2011), Caselli and Michaels (2013), Dube and Vargas (2012), and Monteiro and Ferraz (2012)). Three papers examine particular resource booms within the United States. Carrington (1996) studies the effects of the Trans-Alaska pipeline construction on the Alaskan economy, finding that labor supply was very elastic on the extensive and intensive margins, but there was little effect in sectors unrelated to the pipeline, and no long-run effect after construction finished. Black, McKinnish, and

Sanders (2005a) use difference-in-differences to analyze the boom and bust in coal production in Kentucky, Ohio, Pennsylvania, and West Virginia, using variation in initial coal production across counties. They find that the boom increased wages and decreased poverty, and that there were employment spillovers into local non-tradables sectors such as retail but no positive or negative spillovers to manufacturing. Michaels (2010) exploits cross-sectional variation in oil abundance across counties to study the long-term effects of resource abundance in the southern United States. He shows that after oil was discovered, oil abundant counties specialized in oil production, but this did not generate a resource curse: higher incomes increased population, which increased the provision of local public goods such as roads and airports, which in turn increased output in agriculture and manufacturing.

On paper differs from this existing work in several ways. First, our geographic scope is much broader - the entire United States instead of a region. This affords us a larger sample of counties, workers, and firms and thus more precise estimates, as well as parameter estimates that apply to a larger and thus more policy-relevant population. Second, our access to microdata from the U.S. Census of Manufactures and Longitudinal Business Database allows us to ask questions that no previous paper has been able to ask. For example, we can examine additional outcomes, such as total factor productivity, and particular sub-sets of firms, such as those that have labor-intensive production technologies or are upstream or downstream of the oil and gas sector.

A third way in which we differ is in our conceptual connection to the empirical literature on agglomeration and productivity spillovers. As Ellison and Glaeser (1997) point out, spillovers are particularly difficult to identify because both heterogeneous local input costs and spillovers can cause agglomeration. To identify spillovers, several recent papers have exploited natural experiments, including the siting of large manufacturing plants (Greenstone, Hornbeck, and Moretti 2010), portage sites (Bleakley and Lin 2012), and the boundaries of the Tennessee Valley Authority development region (Kline and Moretti 2012). Our analysis builds on this literature by exploiting the fact that natural resource booms and busts are similarly interesting and useful as natural experiments to identify spillovers. More broadly, this project builds on previous analyses of local economic shocks, which may be due to national-level sectoral trends (Bartik 1991, Blanchard and Katz 1992), military base closures (Hooker and Knetter 2001), or place-based economic development policies (Busso, Gregory, and Kline 2012).

2 Background: Evolution of the Oil and Gas Sector

In 1970, oil and gas markets were relatively stable. Prices had been steady and slightly declining for the entire post-war period. This suddenly ended in October 1973 when 10 Arab oil exporters imposed an embargo on the US in retaliation for its support of Israel during the Yom Kippur war, cutting oil production by 5 percent. The embargo lasted for five months, during which time oil prices in the U.S. more than doubled. Figure 2 presents "first purchase" oil prices - the price per barrel of the first inter-firm sale of oil produced in the US, as reported by the U.S. Energy Information Administration (EIA) - from 1960 to the present. Like all prices in this paper, these are in real 2010 dollars. Real prices remained on this higher plane until a second unexpected event, the 1979 Iranian Revolution. Protests that year, and further disruption caused by the Iran-Iraq war, caused Iranian production to drop by 70 percent. As a result, US oil prices rose from $30 in 1978 to $77 in 1981.

Figure 3 illustrates a second way in which the year 1970 was a turning point for the US oil industry: that was the year that oil production "peaked" and began to decline. This decline was monotonic for the first few years of the 1970s, until it was arrested by the supply-side response to the 1973 price shock. The dashed black line on Figure 3 shows the "rig count": the total number of oil and gas rigs in operation to drill new wells or work over existing wells, as reported by the U.S. EIA. The rig count rose from 976 in 1971 to 3970 in 1981.

This supply response, coupled with declining global demand caused by the recession of the early 1980s, caused prices to drop in two phases. First, after peaking in March 1981, they dropped steadily to $49 in 1985. Then, in the first six months of 1986, prices suddenly dropped another 60 percent. This in turn induced a bust in drilling activity, and the national rig count was below 1000 for all but two years between 1986 and 2002. In the past decade, high global demand has spurred a second boom of high prices and increased drilling activity.

While oil production entails large capital costs, it also requires significant labor input. Figure 4 shows total national employment in the oil and gas sector, as reported by the U.S. Bureau of

Economic Analysis through the Regional Economic Information System (REIS). Employment rose from under 400 thousand people in the early 1970s to over one million in the early 1980s. Closely following the peak in prices and rig counts, employment peaked in 1982, dropped in 1983, and held steady in 1984 and 1985 before dropping sharply in 1986 and declining steadily until 2002. Figure 4 also shows total US manufacturing employment, which is mildly procyclical through 2000 and then drops off significantly between 2001 and 2010. The vertical lines on the graph mark the change from SIC to NAICS industrial classification codes, which reclassifies some jobs away from both the oil and gas and manufacturing sectors. Using consistent definitions, employment in these two sectors did not change much between these two years.

This large time series variation in oil and gas drilling activity and employment demand has very different incidence across the cross section of counties. Figure 5 maps each county's oil and gas intensity, in units of average annual value of oil and gas production over 1969-2011 per 1969 inhabitant. Our empirical strategy exploits the interaction of the time series variation in Figure 4 with the cross-sectional geographical variation in Figure 5.

3 Model

3.1 Setup

To elucidate the set of mechanisms through which exogenous shocks to the resource sector might affect the other sectors of the economy, we outline a simple model of a small, open economy in the spirit of Matsuyama (1992). Extending Matsuyama, we include a local non-tradable good sector, for a total of three industries:

1. A resource sector, with productivity R, exogenously set price p_r, and employment n_r:
   - Output: $X_r = RF(n_r) \quad F(0) = 0,\ F'(\cdot) > 0\ F''(\cdot) < 0$
   - Productivity R changes exogenously with resource boom.
2. A tradable good sector, with productivity M, exogenously set price p_m and employment n_m:
   - Output: $X_m = MH(n_m) \quad H(0) = 0,\ H'(\cdot) > 0\ H''(\cdot) < 0$

- Productivity M, evolves over time according to learning-by-doing, with the functional form:

$$M_t = M_t + \delta X_{tm} + \gamma X_{tr}$$

In this equation, we have included time subscripts, which we suppress in other equations for legibility. A value of $\gamma > 0$ allows positive spillovers from the resource sector to have a direct effect on productivity of tradable goods.

3. A non-tradable sector, with productivity L, endogenously determined price p_l and employment n_l:

- Output: $X_l = LG(n_{tl}) \quad G(0) = 0,\ G'(\cdot) > 0\ G''(\cdot) < 0$
- Productivity L is constant

We assume that the regional economy is open to resource and tradable goods, but closed to non-tradables and labor. The model thus captures the short-term effects of the resource boom, before adjustments in population which might increase labor supply, or potential entry by new non-tradable producers.

In equilibrium the price of local goods adjusts to equilibrate non-tradable supply and demand

$$C_l = X_l = LG(n_l)$$

and the sum of employment in all three sectors must equal total regional labor supply, normalized to one

$$n_r + n_l + n_m = 1$$

Households supply labor inelastically, and have Cobb-Douglas preferences over tradable and non-tradable goods with contemporaneous utility function

$$U = \alpha \ln C_l + (1 - \alpha) \ln C_m$$

and budget constraint

$$w + \pi = p_l C_l + p_m C_m$$

with w being labor income and π being the household's share of profits from local firms. For simplicity, we assume that household income from profits comes only from the resource sector

$$\pi_r = p_r X_r - w n_r$$

and that firms from the tradable and non-tradable sectors are owned by individuals outside the region of interest and thus return no profits. Including these additional profit terms in the representative consumer's budget constraint would make little difference in the results since, as shown below, booms in the resource sector create offsetting effects in tradable and non-tradable profits. Including resource sector profits captures the role of royalty payments made by resource firms to landowners. The $\pi > 0$ term also ensues that a resource boom increases consumption of non-tradables instead of increasing wages and non-tradable goods prices proportionally.

The key endogenous outcomes in this model are the labor shares across industries, n_r, n_l, n_m. These are pinned down by equalizing the marginal product of labor equalized across sectors:

$$w = p_r R F'(n_r) = p_l L G'(n_l) = p_m M H'(n_m) \tag{1}$$

as well as the demands for tradable and non-tradable goods :

$$p_l C_l = \alpha (w + \pi) \tag{2}$$

$$p_m C_m = (1 - \alpha)(w + \pi) \tag{3}$$

3.2 Predictions

We use this model to derive some basic relationships between resource booms, captured by increases in the productivity of the resource sector R_t, and the non-resource sectors of the economy.

Prediction 1: Resource booms decrease contemporaneous local production of tradable goods:

Wages equilibrate returns to labor in resource and tradables manufacturing:

$$\frac{p_r}{p_m} \cdot \frac{R}{M} = \frac{H'(n_m)}{F'(n_r)}$$

An increase in R, or equivalently in p_r, increases the right-hand side, requiring an increase in labor in the resource sector (n_r) to compensate. This takes labor from the tradable goods industry, decreasing the output and employment in this sector. Thus an increase in R increases resource production both directly (through higher productivity) and indirectly, through attracting more labor to the resource sector. It also increases wages overall since marginal returns to labor increase in all sectors. This captures the fundamental mechanism behind the "Dutch Disease" phenomenon often described in the literature.

Prediction 2: Resource booms increase prices and production of non-tradable goods:

In contrast to zero-trade cost manufacturing sector, the model predicts positive growth in the non-tradable industry. To see this, we can re-write profits in the resource sector as:

$$\pi = \mu \cdot \left(p_r R n_r \left(\frac{F(n_r)}{n_r} - F'(n_r) \right) \right)$$

Differentiating

$$\frac{\partial \pi}{\partial R} = \mu n_r p_r \left(\left(\frac{F(n_r)}{n_r} - F'(n_r) \right) - R F''(n_r) \frac{\partial n_r}{\partial R} \right) > 0$$

shows that a local resource boom increases the rents that inhabitants of the region acquire from the natural resource sector, or alternatively in the contribution of the resource industry to aggregate local consumption. Combined with the increase in wages due to higher productivity in the resource sector, overall consumption $w + \pi$ increases, raising demand for local non-tradable goods. This increased demand not only raises prices but also increases production, with non-tradable output determined by the representative consumer's demand

$$LG\left(n_l\right) = \frac{\alpha\left(w+\pi\right)}{p_l}$$

and prices set to equalize the marginal product of labor and the wage:

$$p_l = w/LG'\left(n_l\right).$$

Combining these two equations and simplifying generates the condition determining production of the non-tradable goods sector,

$$\frac{G\left(n_l\right)}{G'\left(n_l\right)} = \alpha + \alpha\frac{\pi}{w} = \alpha + \alpha\left(\frac{F\left(n_r\right)}{F'\left(n_r\right)} - n_r\right) \tag{4}$$

It is straightforward to show that since $\partial n_r/\partial R > 0$ and $dn_l/dn_r > 0$, non-tradable labor and production must be increasing in the productivity of the resource sector,and prices of non-tradables are increasing as well. These effects are purely driven by increases in demand, and not by any change in the productivity of the non-tradable sector.

Prediction 3: With productivity spillovers, resource booms create long-term divergence between regions:

The potential decline of the tradable sector due to the resource boom has no particular welfare consequences if tradable productivity is taken to be exogenous. However, if this sector is particularly prone to productivity growth through learning by doing, then these sectoral shifts may have long-term consequences. To illustrate this, we consider two locations, A and B, where location A has a resource boom which increases the productivity of resource extraction in the first period, R_1^A. In the first period, the regions' other sectoral productivities, $M_1^{A,B}$and $L_1^{A,B}$ are the same. Their labor allocations, however, will differ. As we show above, an increase in R_1^A increases n_{1r}^A and decreases n_{1m}^A (manufacturing labor) relative to n_{1m}^B. Since tradable output depends only on labor, this too will be lower in region A.

$$X_{1m}^A = M_1^A H\left(n_{1m}^A\right) < X_{1m}^B$$

Due to learning-by-doing in manufacturing, the second period's tradable productivity will be lower in region B,

$$M_2^A = M_1^A + \delta X_{1m}^A + \gamma X_{1r}^A < M_2^B,$$

if $\delta\left(X_{1m}^A - X_{1m}^B\right) > \gamma\left(X_{1r}^B - X_{1r}^A\right)$.

In the second period, the resource boom ends, $R_2^A = R_2^B$, but tradable productivity remains different between regions. Taking the ratio of wages

$$\frac{F'\left(n_{2r}^A\right)}{H'\left(n_{2m}^A\right)} \cdot \frac{M_2^B}{R_2^B} \frac{R_2^A}{M_2^A} = \frac{F'\left(n_{2r}^B\right)}{H'\left(n_{2m}^B\right)}$$

shows that even after the resource boom is over, region A still has relatively more labor in the resource sector. As Matsuyama (1992) argues, even short-term differences in resource sector productivity may lead to long term divergence in tradable or manufacturing output.

Note, however, that our model suggests an alternative outcome is also possible. If spillovers from the resource sector are strong and/or the size of the boom is large enough $\left(X_{1r}^A \ggg X_{1r}^B\right)$, then long-term tradable output may be higher in areas experiencing resource booms. This basic intuition–that resource booms may have long term positive impacts on other sectors of the economy–underlies much of the literature on agglomeration and spillovers.

3.3 Tests

The model suggests a series of three empirical questions. First, how much do resource booms affect county-level aggregate outcomes: employment, earnings, population, and wages? These outcomes are of interest *per se* as measures of growth. The answer to this question also suggests the magnitude of effects on the manufacturing sector. For example, if resource booms have only small impacts on wages, then there is limited scope for Dutch Disease. Quick population migration to producer counties could limit wage impacts.

Second, how do resource booms contemporaneously affect employment and output in the manufacturing sector? If there is no short-run contraction in response to higher wages, then there is no loss of learning-by-doing potential. The effects of a boom might vary substantially for different types of manufacturing firms. In reality, tradability is a continuum determined by transport costs and economies of scale, and some firms might produce goods that are more highly tradable than others. Furthermore, different manufacturing sectors use more or less labor-intensive technologies, and increases in wages due to resource booms should cause the more labor-intensive sectors to contract more. Finally, productivity spillovers and local demand effects could be stronger for "linked" manufacturing firms - that is, firms that produce inputs to or purchase outputs from the natural resource sector.

Third, do resource booms affect the total factor productivity of manufacturing firms, either in the short-run or long run? Changes in the short-run productivity of the non-resource sector identify the direct productivity spillover effects (γ) of resource booms. Changes in the longer term productivity of the non-resource sectors identify the learning by doing effect (δ), conditional on the boom affecting manufacturing or non-tradable production in the short-run. The strength of these spillovers may depend on the degree to which firms are linked with the resource sector.

4 Data

In this section we describe the data at our disposal. We begin by presenting the data on oil and gas resources, labor and industry outcomes, and other data. We then define the sample population of treatment and control counties. Finally, we define "booms" and "busts."

4.1 Resource Data

We have constructed a new county-by-year panel dataset of oil production (in barrels) and natural gas production (in million BTUs) from the 1960s to the present. The original source of much of these data is an extraordinary database from a market research company called DrillingInfo that includes monthly production of oil and gas from all individual oil or gas wells in nearly all U.S. states. For our purposes, however, these data are incomplete: they do not include all states, and in some states where data are available, the data begin in the 1980s – after the oil and gas boom of the 1970s, which provides substantial identifying variation. Therefore, we have acquired county-level oil and gas production data from oil and gas regulatory agencies or severance tax authorities in nine additional states: Illinois, Indiana, Kansas, Kentucky, Louisiana, Michigan, Montana, Nevada, and New York. Appendix Table I gives more info on the sources of these data. In several cases, it is not possible to acquire county-level production data for all years back to 1969, but we do have state-level production data from the Energy Information Administration. In these cases, we impute county-level historical production by multiplying state-level production by the county's share of state production in the earliest year when it is observed.

As suggested by Figure 5, there is substantial spatial variation in the natural resource intensity of local economies in the US. There are 1051 counties with positive production, while more than 2000 counties have zero production. Within the counties that produce any oil, the interquartile range of oil production is substantial: 11 to 3404 barrels per person. There is also significant variation within states: with the exception of Louisiana, all states that produce oil or gas also have some counties that produce no oil or gas.

The extraction of natural resources from a location may itself be affected by many of the economic outcomes that we wish to study. To overcome this potential source of endogeneity, we are collecting a unique dataset on the initial stock of oil and gas resources in the United States. Our measure of a county's initial resource stock is the sum of unproven reserves measured by the U.S. Geological Survey (2013), plus the sum of proven reserves reported by the U.S. EIA (2013), plus the sum of observed extraction until the year when the reserves data are reported. In this draft, we do not yet have results using these data.

We construct coal production data analogously to the oil and gas data. We observe coal production for every mine in the United States from 1960 to the present using data gathered by the Bureau of Mines (for 1960-1976), EIA (for 1977), and the Mine Safety and Health Administration (for 1978-2012). These data are primarily gathered to evaluate worker safety, but mines are required to report several variables that are of great use for our project: production, number of employees, and number of hours worked. As with oil production, there is substantial variation. There are 421 counties with positive production, while 2658 counties have zero production. Within the counties that produce any coal, the interquartile range of coal production is 22.6 to 2714 tons per person. There is also significant variation within states: for example, in Wyoming, which has more production than any other state, nearly all of the production comes from one county. Just as we are constructing a dataset of initial stock of oil and gas, we are also constructing a dataset of the initial stock of coal for each county in the country.

4.2 Outcome Data

4.2.1 Public County-Level Data

Our primary source of annual data on employment, earnings, and population is the Regional Economic Information System (REIS).[1] The REIS is prepared by the U.S. Bureau of Economic Analysis using IRS tax records, unemployment insurance and social security payments, Census data, and other information. The data include total county-level employment and earnings for all counties in all years, as well as employment and earnings for one-digit sectors such as retail, manufacturing, and mining. However, the sectoral data are sometimes withheld, especially for smaller counties, to avoid disclosing identifiable firm-level information. We use REIS data from 1969 to 2011. Table 1 presents descriptive statistics for all data, beginning with the REIS.

In some specifications, we control with pre-1969 data from the County Data Books.[2] The County Data Books include 1960 and 1970 total population and total employment from the decennial Census of Population, as well as 1963 and 1967 manufacturing, retail, and mining employment from the Economic Census. As with the REIS, these sectoral data are sometimes withheld for confidentiality. We use county land area in square miles from the US Census.[3]

4.2.2 Current Population Survey

For some of our wage specifications, we use individual-level microdata from the Current Population Survey (CPS). The CPS is a monthly survey that includes 50 to 60 thousand households each month. Households are surveyed for four consecutive months, then ignored for eight months, and then surveyed again for the next four months. For each individual in the sampled household, we

[1] The REIS data are available from http://www.bea.gov/regional/.
[2] We downloaded the County Data Book datasets from ICPSR, series 7736.
[3] The land area data are available from http://quickfacts.census.gov/qfd/download_data.html.

observe gender, age, race, education level, state of residence, employer classification (government, private-sector, self-employed, or unemployed), and the employer's industry.

To increase precision we focus on questions that allow the direct computation of an hourly wage, rather than using annual data. For people who are paid by the hour, we use the answer to the question, "How much does [person] earn per hour?" For non-hourly employees, we divide weekly earnings ("How much does [person] usually earn per week at this job before deductions?") by weekly hours ("How many hours per week does [person] usually work at this job?"). From 1969-1987, these questions were asked on the May CPS. Beginning in 1979, these questions were also asked on each household's "outgoing rotation": the fourth and eighth interviews of the panel, which occur exactly 12 months apart. These data are available from 1979-2012 in the Merged Outgoing Rotation Group (MORG) database.[4]

We construct two datasets from the CPS. One we treat as a repeated cross section, simply combining all observations from the May CPS for 1977 and 1978 with all observations from the MORG beginning in 1979. (Before 1977, the public May CPS data do not include a complete set of state identifiers, and beginning in 1979, the MORG offers a larger sample of individuals answering the hourly wage questions.) The second is a panel based on the MORG, which includes each individual's change in hourly earnings in the 12 months between his or her two outgoing rotations. Since the CPS sampling frame is the household, not the individual, this panel includes only individuals who do not change residence in these two years. The CPS does not include unique individual identifying codes, so we use the approach of Madrian and Lefgren (1999) to match individuals between the datasets in each of the separate years.

4.2.3 Restricted-Access Census Data

Many of the outcomes we examine are drawn from confidential establishment-level data from the Census of Manufactures (CM) and the Longitudinal Business Database (LBD). The LBD data are derived from payroll taxes, and they comprise total number of employees and total wages paid for every business in the United States from 1976 to 2010. We observe the county where the

[4] The May CPS data we use can be downloaded from http://www.nber.org/data/cps_may.html, and the MORG data are available from http://data.nber.org/morg/annual/.

establishment is located, as well as its four-digit SIC code and six-digit NAICS code. The LBD also allows us to observe each establishment's year of entry and exit, provided that the event occurs between 1977 and 2009.

The Census of Manufactures includes establishment-level microdata for all manufacturing establishments in the United States. The data include book value of physical capital stock, number of employees, total wage bill, value of materials inputs, and total revenues. For employment and earnings, data for establishments that do not respond and small establishments with fewer than five employees is imputed from the LBD and marked as an "administrative record." For these establishments, we use the employment and earnings variables, but we do not use any other (imputed) variables. The CM microdata are available for 1963 and quinquennially (every five years) beginning in 1967, i.e. 1972, 1977, 1982, ..., 2007.

For about 6000 relatively-homogeneous products defined at the 7-digit SIC level, the CM also asks establishments to report both their physical production quantities and sales revenues. These are the data used by Foster, Haltiwanger, and Syverson (2008) to highlight the distinction between physical productivity and revenue productivity. We divide revenues by physical production to arrive at an establishment-by-product-by-year dataset of manufacturing output prices. We drop any reported prices that differ from the median by a factor of more than four.

4.3 Other Data

We use two additional datasets to examine subsamples of firms for which the resource boom may have particularly revealing effects. First, we classify industries as "highly tradable" using average shipment distances from the 2007 Commodity Flow Survey (CFS). The CFS is a sample survey of manufacturing, mining, wholesale, and selected retail and services establishments that gathers data on what commodity is shipped, the value, weight, and mode of transportation, and the destination. We use publicly-available data on average shipment distance by 3-digit NAICS industry, as shown in Table 2. A "highly tradable" goods industry is one with average shipment distance longer than 300 miles.

Second, we classify industries as upstream or downstream of the oil and gas sector using the Bureau of Economic Analysis Input-Output tables for 1987. An upstream industry is one that sells more than one percent of its output to the oil and gas sector, while a downstream industry is one in which the input cost share of oil and gas is larger than one percent. Tables 3A and 3B show industries classified as upstream or downstream, along with their oil and gas sector output shares and input cost shares.

4.4 Treatment and Control

Our analysis considers the continental United States, and assigns counties (and all establishments within each county) to "treatment" or "control." On rare occasions, counties will merge or split. In these cases, we construct pseudo-counties at the most disaggregated level at which data are observed for a consistent geographic area over the entire 1969-2011 period[5]. This gives a population of 3079 counties in the continental U.S.

Resource extraction is a fundamentally rural industry, and most urban counties have low or zero employment and output shares in oil and gas extraction while growing with very different trends compared to rural counties. We therefore limit the analysis to the 2,427 counties with 1969 population density of less than 100 people per square mile and less than 250,000 total inhabitants in 1969. This excludes urban and suburban counties: for example, San Mateo County, a suburban county which extends from Menlo Park, California to just south of San Francisco, had 1231 people per square mile in 1969 and 1604 in 2010.

This population of rural counties is divided into two groups using a simple rule: a county that produces any oil or gas in any year after 1969 is called a "treatment" county, and a county that does not produce any oil, gas, or coal after 1969 is a "control" county. We exclude 83 coal-only counties from both treatment and control because while such counties are not oil and gas producers, coal demand is directly affected by oil and natural gas price changes during the sample. The final sample includes 1065 treatment counties and 1362 control counties. For robustness checks, we also construct a "high treatment group" of counties that average more than $1000 in oil and gas production per 1969 inhabitant.

[5] For example, many independent cities in Virginia merge or split with adjacent counties, so data are reported separately for the city and county for some years and then combined for other years. In these cases, we combine the city and county for all years.

Table 1 shows both the means and standard deviations for the entire sample as well as the means in treatment and control. The groups differ statistically in 1969 only on manufacturing employment, which is about 20 percent higher in control counties but may be affected by the substantial fraction (22%) of nondisclosed data. In all regressions, we control for 1969 levels and pre-1969 trends in manufacturing employment and other variables. Because we exclude urban counties, the sample includes only 18 percent of 1969 US manufacturing employment. While this is a small share, it sums to 3.8 million workers.

The CPS data include state identifiers, but not county identifiers. Thus, in analyzing the CPS, we divide US states into treatment and control using a comparable approach. Treatment states are those that average more than $1000 annual oil and gas output per 1969 inhabitant, while control states are those that average less than $400. These figures were chosen because they are at large discontinuities in the distribution of per capita oil and gas output, meaning that assignment to treatment and control is not immediately sensitive to the exact choice of cutoff. States that would be in control but average more than 1.2 tons of annual coal production per 1969 inhabitant are excluded from the sample. These states are Arizona, Illinois, Indiana, Kentucky, Ohio, Pennsylvania, Tennessee, and Virginia. The 1.2 ton cutoff was set to include Tennessee, which has a natural resource boom in the 1970s driven by its increase in coal output. The ten treatment states are Colorado, Kansas, Louisiana, Montana, North Dakota, New Mexico, Oklahoma, Texas, Utah, and Wyoming. There are 26 control states.

4.5 Defining Booms and Busts

In this version of the paper, we employ a very simple difference-in-difference empirical strategy where we are interested in the interaction effect of being in a treatment county during a "boom year" or a "bust year." To do this, we must define booms and busts. We define a boom (bust) as a year in which there is a large increase (decrease) in national oil and gas sector employment. We define 1974-1982 and 2004-2011 as boom years, and we define 1983 and 1986-1992 as bust years. When analyzing quinquennial CM data, we define 1977, 1982, and 2007 as booms, and 1987 and 1992 as busts.

In alternative specifications, we could also define a boom as any year in which national oil and gas employment increases. We could also use the continuous measure of total national oil and gas employment, in an analogy to the Bartik (1991) approach of testing for local impacts associated with national level changes in industrial composition.

5 Empirical Strategy

In this section, we specify our estimating equations, all of which are fundamentally difference-in-differences estimators: they compare the difference in growth between treatment and control counties in boom vs. bust vs. stable years. We also discuss two conceptual issues: unconfoundedness and geographic spillovers.

5.1 Estimating Equations

5.1.1 Graphical

Before presenting formal results, we will plot the annual differences between treatment and control counties. This illustrates how the time series of treatment effects is associated with the time series of oil and gas booms and busts. In the equation below, Y_{ct} denotes outcomes such as employment, population, and earnings, c indexes counties, and t indexes time in years. T_c is the treatment indicator. Y_{0c} represents a vector of pre-1970 values of the outcome, which may be observed in different years depending on the data series. For example, when using natural log of population as our dependent variable, we control for 1969 population using REIS data and 1960 population using the 1960 Census, which we observe through the County Data Book. Including these two pre-treatment values means that we control for both the levels and trends in county population. The variable ϕ_{dt} represents a vector of Census division-by-year indicator variables.

The estimating equation is:

$$\ln Y_{ct} = \sum_{t=1970}^{2011} [\tau_t T_c + \gamma_{0t} \ln Y_{0c}] + \phi_{dt} + \varepsilon_{ct} \tag{5}$$

In this and all other specifications, we use robust standard errors and cluster by county, because the treatment indicator varies at the county level and errors may be serially correlated within county.

5.1.2 County-Level Aggregates

Our first sets of formal specifications use county level observations. Denoting Δ as the difference operator, we use annual differences in county-level outcomes $\Delta \ln Y_{ct}$ as independent variables. The variable B_t takes value 1 if year t is a boom year, -1 if a bust, and 0 if neither. The coefficient of interest is τ; this measures the average difference in growth in treatment counties vs. control counties during a boom year. The value $-\tau$ measures the expected reduction in growth during a bust. We combine across both booms and busts to maximize power, which is useful in some specifications.

$$\Delta \ln Y_{ct} = \tau T_c B_t + \phi_{dt} + \psi_c + \varepsilon_{ct} \tag{6}$$

5.1.3 Current Population Survey Wages

Increased wages is a necessary condition for Dutch Disease, so it is especially important to cleanly estimate the effects of natural resource booms on wages. While the REIS, CM, and LBD earnings per worker data provide useful measures of wages, they are not ideal. Ideally, we would estimate the change in cost per unit of labor input, that is, a quality-adjusted unit of labor over a unit of time. The earnings per worker measures could potentially suffer from composition effects, for example if a resource boom induces lower-education workers to enter the local workforce either by transitioning from unemployment or by migrating from elsewhere. They also do not measure total hours worked: earnings per worker could increase if employees work more hours, without any increase in unit labor costs.

We address these concerns by using the hourly earnings data from the Current Population Survey. One specification treats the CPS data as a repeated cross section and estimates difference-in-differences specifications using "Mincerian" controls for age, education, gender, and race. To estimate a comparable τ using repeated cross sections instead of annual differences, we construct a variable $\widetilde{B}_t = \sum_{y=1970}^{t} B_y$, which captures the cumulative net number of boom years between 1970 and year t. The coefficient of interest is the interaction of $\widetilde{B}_t$ with T_s, the state-level treatment indicator variable. The variable X_i denotes individual i's vector of demographic characteristics, and μ_m is a vector of eleven month indicator variables. The specification is:

$$\ln Y_{ismt} = \tau T_s \widetilde{B}_t + \kappa T_s + \phi_{dt} + \beta X_i + \mu_m + \varepsilon_{ismt} \tag{7}$$

A second CPS specification exploits the MORG panel, and uses individual i's change in wages between year $t-1$ and year t:

$$\Delta \ln Y_{ismt} = \tau T_s \cdot B_t + \kappa T_s + \phi_{dt} + \varepsilon_{ismt} \tag{8}$$

In both of these regressions, standard errors are robust and clustered by state.

5.1.4 Productivity and Prices

Our productivity and price regressions use establishment-level data from the Longitudinal Business Database and Census of Manufactures. The specifications are closely analogous to the CPS panel regression in Equation (8), as they exploit unit-level changes. In the estimating equation, f indexes establishments, and we also add λ_{nt}, the full interactions of two-digit SIC codes and years.

$$\Delta \ln Y_{fct} = \tau T_c B_t + \kappa T_c + \phi_{dt} + \lambda_{nt} + \varepsilon_{fct} \tag{9}$$

5.2 Conceptual Issues

5.2.1 Unconfoundedness

Throughout this paper, we make causal arguments: a resource boom *caused* outcomes to differ between treatment and control. Our argument requires unconfoundedness: no unobserved factors differentially affect treatment and control counties during resource booms. We help to substantiate this econometrically through controls for pre-levels and pre-trends in Equation (5) and county-specific trend effects in Equation (9).

In some alternative specifications presented below, we restrict attention to the subset of counties with overlap on propensity scores between treatment and control. To do this, we estimate a probit regression of treatment status with census division indicator variables and the log of 1960 and 1969 total employment and population, 1963 and 1969 manufacturing employment, and 1969 population density on the right hand side. The propensity score is the predicted value from this regression. There are 59 counties with propensity scores that do not overlap: 20 in control with propensity scores smaller than the smallest value in treatment, and 39 counties in treatment with propensity scores larger than the largest value in control.

Furthermore, in other alternative specifications, we both include the standard set of fixed effects and controls and re-weight treatment and control counties by inverse probability weights, making treatment and control groups balanced on propensity scores. This estimator is "doubly robust" in the sense of Robins, Rotnitzky, and Zhao (1995): it is consistent if either the propensity score model or the linear regression model is correctly specified. As we shall see, the results are remarkably consistent across these different specifications, which builds confidence around the unconfoundedness assumption.

We also substantiate the unconfoundedness assumption by examining three events: two booms and one bust between the 1970s and today. While a confounding linear trend between treatment and control could bias estimated effects from one boom, it would take a very particular confound to bias results from multiple events. The bulk of our results appear to hold for each of the three events we study, which further builds confidence in unconfoundedness.

5.2.2 Geographic Spillovers

Resource booms and busts could affect both treatment and control counties. Mechanically, the population that moves into treatment counties must be migrating from some other counties, either in our control group or the urban counties excluded from our sample. As Busso, Gregory, and Kline (2012) point out in the related context of place-based local economic development policies, our estimates at least partially reflect re-allocation of economic activity from one area to another. Producer states may redistribute some tax revenues to control counties within their states, and firms may expand in control counties to serve higher demand in nearby treatment counties.

As a result, we cannot estimate the aggregate national impacts of a resource boom. Instead, we estimate how resource booms differentially affect the potential outcomes of producer vs. non-producer counties. Our "treatment effects" reflect the difference in potential outcomes for one local area that decides to change from treatment to control. This is a relevant policy question: for example, New York state has banned fracking, a production technology that has augmented the recent oil and gas boom. States and counties can also have continuous policy choices, such as where to set tax rates and how quickly and extensively to provide complementary public goods such as schools, sewers, and roads. Our estimates are informative for local policy makers evaluating the costs and benefits of such decisions.

6 Results

In this section, we present empirical tests of the three basic questions suggested by the model. First, how do resource booms affect county aggregate outcomes: employment, earnings, population, and wages. Second, how do resource booms affect employment and output in the manufacturing and non-tradable sectors? Third, do resource booms affect manufacturing productivity?

6.1 County-Level Aggregate Effects

Figures 7A-7E present estimates of Equation (5) for employment, earnings, population, wages, and manufacturing wages using the REIS data. Each graph reflects the same basic pattern: these aggregate outcomes rise from the late 1970s through the early 1980s, then drop beginning in 1982 or 1983 and decline steadily until they increase again beginning in 2003 or 2004.

Figure 8 combines the employment, wage, and population coefficients on the same graph. This illustrates the dynamic effects of adjustment to a local economic shock analyzed by Blanchard and Katz (1992). As the resource sector expands, total employment increases and wages rise. Population adjusts more slowly, meaning that the short-run effects of a resource boom are to increase wages and decrease unemployment. However, within one to three years, people migrate in search of higher wages, and this migration puts downward pressure on wages. Had the employment demand increase from the boom flattened out, population eventually would adjust until wages equilibrate. In the early 1980s, this never happened: before population reached any newer equilibrium, the boom ended, and employment demand and wages dropped, eventually causing out-migration from the treatment counties. The pattern later begins to repeat itself with the boom of the early 2000s.

Table 4 presents the formal estimates of the effects of resource booms on aggregate outcomes. Column 2 contains the estimates of τ from Equation (6). Column 1 is a simpler version, excluding ϕ_{dt} and ψ_c and including only year indicator variables. Across all outcomes, the estimates are not very sensitive to these controls. Column 3 uses the "high treatment" group instead of the entire treatment group. Excluding the treatment counties that are not in "high treatment" is what reduces the sample size. In all cases, the effects are substantially larger, as would be expected.

Columns 4 and 5 are additional specifications that help to assess robustness to observed confounders. Column 4 includes only the sample of counties within the common support of propensity scores between treatment and control. Column 5 presents the "doubly-robust" estimator that uses the same sample as column 4 and the same right-hand-side variables as column 2, and additionally re-weights treatment and control by inverse probability weights.

There are three important takeaways from this table. First, natural resource booms and busts substantially affect growth in producer counties. For example, the descriptive statistics in the bottom panel of Table 1 show that county-level real earnings grow by an average of 1.69 percent per year between 1969 and 2011. In comparison, the coefficient estimates in the second row of Table 4 suggest that resource booms (busts) increase (decrease) annual growth by about 1.1 percentage points. In a resource-intensive "high treatment" county, a resource boom increases earnings growth by 1.8 percent, more than doubling the sample average growth rate. Second, the results are remarkably robust to alternative controls, trimming, and weighting.

The third takeaway is that the necessary condition for Dutch Disease is satisfied: resource booms are associated with large wage increases. Coefficient estimates in the fourth and fifth rows show that overall wages, and manufacturing sector wages, rise by 0.3 to 0.5 percent in each year that the oil and gas sector expands. If the manufacturing sector is selling into a national output market where prices are fixed, it should contract when faced with this increase in labor costs.

6.2 Manufacturing Sector Effects

Figure 10 presents estimates of Equation (5) with the log of county-level manufacturing sector output as the dependent variable. This graph is not as stark is the graphs of total employment. Two factors may account for this. First, the manufacturing sector might respond differently than the county economy as a whole. Second, the estimates are less precise because the REIS withholds data for some county manufacturing sectors to avoid non-disclosure.

The general pattern, however, appears to be comparable: manufacturing employment grows slightly in treatment relative to control at the end of the 1970s and into the early 1980s, then drops off by ten log points between 1982 and 1987. Relative employment recovers during the 1990s and through the resource boom of the early 2000s. There seems to be no evidence of Dutch Disease.

Figure 11 corroborates this, plotting the estimated conditional mean difference in changes of log manufacturing employment between treatment and control against the change in national oil and gas sector employment, for each year between 1970 and 2011. The best fit line slopes upward, meaning that manufacturing growth in treatment counties is positively associated with increases in employment demand in the oil and gas sector.

Table 7 presents the formal estimates of the effects of resource booms on manufacturing employment. Each of the five columns parallels the specification in Table 4; column 2 again contains the estimates of τ from Equation (6). A resource boom is associated with an *increase* in manufacturing sector growth in treatment counties relative to control. The coefficients are relatively robust, ranging from 0.44 percentage points to 0.57 percentage points. Using the "high treatment" group increases the coefficient further, to 1.09 log points. This further suggests that it is expansion and contraction of the resource sector, not some unobserved correlated factor, which drives these results.

6.2.1 Manufacturing Sub-Sectors

Of course, the manufacturing sector is not monolithic. Manufacturing firms that are upstream or downstream of the oil and gas sector should grow during resource booms. If these linked firms are more likely to be located in treatment counties due to lower transport costs, this could cause the overall manufacturing sector to grow in treatment counties even if non-linked firms contract. As shown in the model, increases in local labor costs may be less problematic for firms that sell non-tradable goods into local markets, since competitors also face this input cost increase. On the other hand, such cost increases will be especially harmful for firms that sell tradable goods. Finally, firms in labor intensive sectors should be more strongly affected by an increase in labor costs.

Column 1 of Table 9 tests these predictions by presenting estimates of Equation (6) using county-level data on manufacturing employment from the Census of Manufactures (for 1972) and the Longitudinal Business Database (for 1976-2010). The first set of rows uses the change in the natural log of county total manufacturing employment. The second set of rows uses the change in natural log of total county-level employment in non-linked manufacturing establishments. The third set of rows analyzes non-linked highly-tradable establishments. The fourth set examines non-linked establishments in "labor intensive" two-digit SIC codes, by which we mean industries where total salaries divided by value added is larger than 0.2. Despite the fact that in theory, these sub-sectors should be more likely to contract during resource booms, there is no evidence of Dutch Disease. Indeed, the estimated τ coefficients are not statistically different in these more vulnerable sectors from the estimate for the manufacturing sector as a whole. Although Census non-disclosure rules make it difficult to include many robustness checks in working papers, the qualitative results are robust to a variety of different specifications and configurations of control variables.

Column 2 of Table 9 presents estimates of Equation (6) for the analogous sub-sectors, using county-level data on total value of sales revenues from manufacturing establishments from the Census of Manufactures. In this regression the difference operator takes five-year differences instead of one year differences, as the CM data are only observed quinquennially. Thus, the $\widehat{\tau} = 0.09$ in the first row implies that a resource boom (bust) over the previous five-year period is associated with a nine-percent increase (decrease) in sectoral output. This corresponds to a 1.8 percent annual growth rate. For all of the four samples in column 2, resource booms are associated with large increases in manufacturing output.

6.3 Manufacturing Productivity Effects

One reason why resource booms could cause the manufacturing sector to grow is if they caused productivity to increase. Table 10 presents tests of this using Equation (9). In all three columns, establishments are observed every five years, meaning that the coefficient estimates reflect changes over five year periods, not one-year periods. The first column uses the natural log of value added per worker, while the second column considers TFP. Value added per worker increases (decreases) by about one percentage point in treatment counties during boom (bust) five-year periods.

For the first three sets of rows, there is no statistically significant association between TFP and resource booms. Furthermore, the standard errors allow us to rule out that TFP increases by as much as value added per worker, suggesting that firms increase capital and/or materials inputs during the boom instead of increasing TFP. Non-linked labor intensive industries have the largest point estimate increase in value added per worker, and their TFP also does appear to increase (decrease) during booms (busts).

Column 3 of Table 10 tests whether resource booms allow manufacturing plants in treatment counties to increase output prices. There is no statistically significant association between resource booms and output prices for any of the four samples we analyzed, and the standard errors in the first two rows are tight enough to rule out price changes of more than 0.75 percentage points per five year period, or about 0.15 percentage points per year. These results suggest that effects on physical productivity or prices are not sufficient to explain the increase in manufacturing employment and revenues.

7 Conclusion

The rise in oil and gas prices and drilling activity in the past decade has caused economists and policymakers to again consider whether natural resource production benefits producer economies or whether there is a "Natural Resource Curse." In this paper, we use confidential establishment-level data from the US Census of Manufactures and Longitudinal Business Database to estimate the effects of expansions and contractions of the oil and gas sector on growth since the early 1970s. There are three key results. First, oil and gas booms increase growth rates in producer counties by 60 to 80 percent relative to non-producer counties, and a necessary condition for Dutch Disease is satisfied: wages increase by 0.3 to 0.5 percentage points per year during a boom. Second, however, manufacturing growth is positively associated with natural resource booms: manufacturing employment and output both rise. Third, there is little evidence that oil and gas booms significantly increase productivity.

In making sense of these results, recall that our sample includes the 18 percent of national manufacturing employment that is in rural counties. We hypothesize that these establishments sell at least some of their output locally instead of in national or international markets, and they thus increase quantities produced in response to the positive local demand shock. If this is true, this provides a compelling counterexample to the traditional Dutch Disease story that resource booms cause manufacturing to contract. While resource booms might likely cause producers of traded goods to contract, our results suggest that many manufacturing plants in rural America's resource-intensive counties are producing for local markets as well.

References

[1] Aragon, Fernando, and Juan Pablo Rud (2011). "Natural Resources and Local Communities: Evidence from a Peruvian Gold Mine." Working Paper, Simon Fraser University (April).

[2] Atkin, David (2012). "Endogenous Skill Acquisition and Export Manufacturing in Mexico." Working Paper, Yale University (July).

[3] Bartik, Timothy (1991). Who Benefits from State and Local Economic Development Policies? Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

[4] Black, Dan, Terra McKinnish, and Seth Sanders (2005a). "The Economic Impact of the Coal Boom and Bust." *The Economic Journal*, Vol. 115, No. 503 (April), pages 449-476.

[5] Black, Dan, Terra McKinnish, and Seth Sanders (2005b). "Tight Labor Markets and the Demand for Education: Evidence from the Coal Boom and Bust." *Industrial and Labor Relations Review*, Vol. 59, No. 1 (October), pages 3-16.

[6] Blanchard, Olivier, and Lawrence Katz (1992). "Regional Evolutions." Brookings Papers on Economic Activity, Vol. 1992, No. 1, pages 1-75.

[7] Bleakley, Hoyt, and Jeffrey Lin (2012). "Portage and Path Dependence." *Quarterly Journal of Economics*, Vol. 127, No. 2, pages 587-644.

[8] Brown, Chip (2013). "North Dakota Went Boom." *The New York Times Magazine*, January 31, 2013.

[9] Brunnschweiler, Christa, and Erwin Bulte (2008). "The Resource Curse Revisited and Revised: A Tale of Paradoxes and Red Herrings." *Journal of Environmental Economics and Management*, Vol. 55, No. 3, pages 248-264.

[10] Bureau of Economic Analysis (2013. "Gross-Domestic-Product-(GDP)-by-Industry Data." http://www.bea.gov/industry/gdpbyind_data.htm

[11] Busse, Meghan and Nathaniel Keohane (2007). "Market Effects of Environmental Regulation: Coal, Railroads, and the 1990 Clean Air Act." *RAND Journal of Economics*, Vol. 38, No. 4, pages 1159-1179.

[12] Busso, Matias, Jesse Gregory, and Patrick Kline (2012). " Assessing the Incidence and Efficiency of a Prominent Place-Based Policy." Forthcoming, *American Economic Review*.

[13] Carrington, William (1996). "The Alaskan Labor Market during the Pipeline Era." Journal of Political Economy, Vol. 104, No. 1 (February), pages 186-218.

[14] Caselli, Francesco, and Guy Michaels (2013). "Do Oil Windfalls Improve Living Standards? Evidence from Brazil." *American Economic Journal: Applied Economics*, Vol. 5, No. 1 (January), pages 208-238.

[15] Collier, Paul, and Benedikt Goderis (2009). "Commodity Prices, Growth, and the Natural Resource Curse: Reconciling a Conundrum." Working Paper, Oxford University (December).

[16] David, Paul, and Gavin Wright (1997). "Increasing Returns and the Genesis of American Resource Abundance." *Industrial and Corporate Change*, Vol. 6, No. 2, pages 203–45.

[17] Domenech, Jordi (2008). "Mineral Resource Abundance and Regional Growth in Spain, 1860-2000. *Journal of International Development*, Vol. 20, pages 1122-1135.

[18] Dube, Oeindrila, and Juan Vargas (2012). "Commodity Price Shocks and Civil Conflict: Evidence from Colombia." Working Paper, New York University (July).

[19] Ellison, Glenn, and Edward Glaeser (1997). "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach." *Journal of Political Economy*, Vol. 105, No 5, pages 889-927.

[20] Ellison, Glenn, Edward Glaeser, and William Kerr (2010). "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns." *American Economic Review*, Vol. 100, No. 3 (June), pages 1195-1213.

[21] Emery, Herb, Ana Ferrer, and David Green (2011). "Long Term Consequences of Natural Resource Booms for Human Capital Accumulation." Canadian Labour Market and Skills Research Network Working Paper No. 74 (February).

[22] Frankel, Jeffrey (2010). "The Natural Resource Curse: A Survey." Working Paper, Harvard Kennedy School (February).

[23] Glaeser, Edward, and Joshua Gottlieb (2009). "The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States." *Journal of Economic Literature*, Vol. 47, No. 4, pages 983-1028.

[24] Glaeser, Edward, Sari Pekkala Kerr, and William Kerr (2012). "Entrepreneurship and Urban Growth: An Empirical Assessment with Historical Mines." Working Paper, Harvard University (July).

[25] Greenstone, Michael, Richard Hornbeck, and Enrico Moretti (2010). "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings." *Journal of Political Economy*, Vol. 118, No. 3, pages 536-598.

[26] Harding, Torfinn, and Anthony Venables (2013). "The Implications of Natural Resource Exports for Non-Resource Trade." OxCarre Research Paper 103 (January).

[27] Hooker, Mark, and Michael Knetter (2001). "Measuring the Economic Effects of Military Base Closures." *Economic Inquiry*, Vol. 39, No. 4 (October) pages 583-598.

[28] Ismail, Kareem (2010). "The Structural Manifestation of the 'Dutch Disease': The Case of Oil Exporting Countries." IMF Working Paper 10/103 (April).

[29] Kline, Patrick (2008). "Understanding Sectoral Labor Market Dynamics: An Equilibrium Analysis of the Oil and Gas Field Services Industry." Working Paper, UC Berkeley (November).

[30] Kline, Patrick, and Enrico Moretti (2012). "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority." Working Paper, University of California at Berkeley (September).

[31] Krugman, Paul (1987). "The Narrow Moving Band, the Dutch Disease, and the Competitive Consequences of Mrs. Thatcher." *Journal of Development Economics*, Vol. 27, pages 41-55.

[32] Krugman, Paul (1991a). Geography and Trade. Cambridge, MA, and London: MIT Press; Louvain, Belgium: Louvain University Press.

[33] Krugman, Paul (1991b). "Increasing Returns and Economic Geography." *Journal of Political Economy*, Vol. 99, No. 3, pages 483-499.

[34] Madrian, Brigitte, and Lars Lefgren (1999). "A Note on Longitudinally Matching Current Population Survey (CPS) Respondents." NBER Technical Working Paper 247 (November).

[35] Manzano, Osmel, and Roberto Rigobon (2001). "Resource Curse or Debt Overhang?" NBER Working Paper 8390 (July).

[36] Matsuyama, Kiminori (1992). "Agricultural Productivity, Comparative Advantage, and Economic Growth." *Journal of Economic Theory*, Vol. 58, pages 317-334.

[37] Mehlum, Halvor, Karl Moene, and Ragnar Torvik (2006). "Cursed by Resources or Institutions?" *The World Economy*, pages 1117-1131.

[38] Michaels, Guy (2010). "The Long Term Consequences of Resource-Based Specialisation." *The Economic Journal*, Vol. 121 (March), pages 31-57.

[39] Mine Safety and Health Administration (2013). "Mine Data Retrieval System." http://www.msha.gov/drs/drshome.htm

[40] Monteiro, Joana, and Claudio Ferraz (2012). "Does Oil Make Leaders Unaccountable? Evidence from Brazil's Offshore Oil Boom." Working Paper, Pontificia Universidade Catolica do Rio de Janeiro (May).

[41] Moretti, Enrico (2004). "Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions." *American Economic Review*, Vol. 94, No. 3 (June), pages 656-690.

[42] Notowidigdo, Matthew (2011). "The Incidence of Local Labor Demand Shocks." NBER Working Paper No. 17167 (June).

[43] Oreopoulos, Philip, Till von Wachter, and Andrew Heisz (2012). "The Short- and Long-Term Career Effects of Graduating in a Recession." Working Paper, Columbia University.

[44] Oyer, Paul (2008). "The Making of an Investment Banker: Stock Market Shocks, Career Choice, and Lifetime Income." *The Journal of Finance*, Vol. 63, No. 6 (December), pages 2601-2628.

[45] Peretto, Pietro (2012). "Resource Abundance, Growth, and Welfare: A Schumpeterian Perspective." *Journal of Development Economics*, Vol. 97, pages 142-155.

[46] Robins, James M., Andrea Rotnitzky, and Lue Ping Zhao (1995). "Analysis of Semiparametric Regression Models for Repeated Outcomes in the Presence of Missing Data." *Journal of the American Statistical Association*, Vol. 90, No. 429, pages 106-121.

[47] Romer, Paul (1986). "Increasing Returns and Long-Run Growth." *Journal of Political Economy*, Vol. 94, No. 5 (October), pages 1002-1037.

[48] Sachs, Jeffrey, and Andrew Warner (1995). "Natural Resource Abundance and Economic Growth." NBER Working Paper 5398 (December).

[49] Sarafinelli, Michel (2012). "Good Firms, Worker Flows, and Productivity." Working Paper, UC Berkeley (November).

[50] Torvik, Ragnar (2001). "Learning by Doing and the Dutch Disease." *European Economic Review*, Vol. 45, pages 285-306.

[51] U.S. Energy Information Administration (2012). "OPEC Revenues Fact Sheet." http://www.eia.gov/cabs/OPEC_Revenues/Factsheet.html

[52] U.S. Energy Information Administration (2013). "Detailed Oil and Gas Field Maps." http://www.eia.gov/pub/oil_gas/natural_gas/analysis_publications/maps/maps.htm

[53] U.S. Geological Survey (2013). "National Oil and Gas Assessment." http://energy.usgs.gov/OilGas/AssessmentsData/NationalOilGasAssessment.aspx.

[54] van der Ploeg, Frederick (2011). "Natural Resources: Curse or Blessing?" Journal of Economic Literature, Vol. 49, No. 2, pages 366-420.

[55] van Wijnbergen, Sweder (1984). "The 'Dutch Disease': A Disease After All?" *The Economic Journal*, Vol. 94, No. 373 (March), pages 41-55.

[56] Wright, Gavin, and Jesse Czelusta (2007). "Resource-Based Growth Past and Present." In Lederman, Daniel, and William Maloney, Eds., Natural Resources: Neither Curse Nor Destiny. Stanford, CA: Stanford University Press.

Tables

Table 1: Descriptive Statistics

Variable	N	Mean	SD	Treatment Mean	Control Mean
1969 Employment	2427	8085	8136	7907	8225
1969 Retail Employment	2312	1123	1251	1204	1239
1969 Mfg Employment	1898	1785	2266	1616	1911
1969 Earnings ($million)	2427	262	300	257	265
1969 Earnings/Worker ($000)	2427	30.4	6.2	30.5	30.4
1969 Population	2427	20,400	19,255	20,512	20,312
1969 Density (People/Sq. Mile)	2427	30.6	24.4	30.0	31.0
Average Oil and Gas Output ($000s/1969 Inhabitant)	2427	5.49	39.7	12.5	0
Δln(Employment)	101,934	0.0126	0.0416	0.0120	0.0130
Δln(Retail Employment)	97,104	0.0041	0.0878	0.0023	0.0054
Δln(Mfg Employment)	79,716	0.0014	0.1387	0.0024	0.0007
Δln(Earnings)	101,914	0.0169	0.1295	0.0162	0.0175
Δln(Earnings/Worker)	101,914	0.0043	0.1201	0.0042	0.0045
Δln(Population)	101,934	0.0069	0.0223	0.0057	0.0079
Oil and Gas Output ($million/year)	104,225	43.4	211	99.0	0.0

Table 2: Average Shipment Distances

NAICS	Description	Average Miles	1(Highly Tradable)
311	Food manufacturing	289	
312	Beverage and tobacco product manufacturing	217	
313	Textile mills	798	Yes
314	Textile product mills	881	Yes
315	Apparel manufacturing	1072	Yes
316	Leather and allied product manufacturing	1097	Yes
321	Wood product manufacturing	329	
322	Paper manufacturing	554	Yes
323	Printing and related support activities	675	Yes
324	Petroleum and coal products manufacturing	146	
325	Chemical manufacturing	853	Yes
326	Plastics and rubber products manufacturing	711	Yes
327	Nonmetallic mineral product manufacturing	185	
331	Primary metal manufacturing	558	Yes
332	Fabricated metal product manufacturing	610	Yes
333	Machinery manufacturing	884	Yes
334	Computer and electronic product manufacturing	1176	Yes
335	Electrical equipment, appliance, and component manufacturing	913	Yes
336	Transportation equipment manufacturing	752	Yes
337	Furniture and related product manufacturing	711	Yes
339	Miscellaneous manufacturing	972	Yes

Notes: Based on 2007 Commodity Flow Survey.

Table 3A: Upstream Manufacturing Industries

SIC	Description	Percent
2899	Chemicals and Chemical Preparations, NEC	5.1
2992	Lubricating Oils and Greases	2.1
324	Hydraulic Cement	6.3
3491	Industrial Valves	3.0
3492	Fluid Power Valves and Hose Fittings	3.0
3494	Valves and Pipe Fittings, NEC	3.0
3498	Fabricated Pipe and Pipe Fittings	3.0
3533	Oil and Gas Field Machinery and Equipment	7.0
3561	Pumps and Pumping Equipment	1.2
3563	Air and Gas Compressors	1.2
3566	Speed Changers, Industrial High-Speed Drives, and Gears	1.4
3568	Mechanical Power Transmission Equipment, NEC	1.4
3621	Motors and Generators	1.4

Notes: We define an industry as being upstream of the oil and gas sector if more than one percent of its output is sold to the oil and gas industry. This table presents all upstream industries and their oil and gas output shares. Based on 1987 Bureau of Economic Analysis input-output tables.

Table 3B: Downstream Manufacturing Industries

SIC	Description	Percent
281	Industrial Inorganic Chemicals	2.1
2865	Cyclic Organic Crudes and Intermediates	2.1
2869	Industrial Organic Chemicals, NEC (aliphatics)	2.1
2873	Nitrogenous Fertilizers	8.1
2874	Phosphatic Fertilizers	8.1
2895	Carbon Black	6.2
291	Petroleum Refining	68.8
2999	Products of Petroleum and Coal, NEC	31.5

Notes: We define an industry as being upstream of the oil and gas sector if more than one percent of its inputs are purchased from the oil and gas industry. This table presents all downstream industries and their oil and gas input cost shares. Based on 1987 Bureau of Economic Analysis input-output tables.

Table 4: Aggregate Effects

Outcome	1	2	3	4	5
Employment	0.0077	0.0076	0.0132	0.0075	0.0064
	(0.0005)***	(0.0005)***	(0.0008)***	(0.0005)***	(0.0005)***
Earnings	0.0131	0.0110	0.0181	0.0109	0.0114
	(0.0009)***	(0.0010)***	(0.0017)***	(0.0010)***	(0.0010)***
Population	0.0038	0.0037	0.0061	0.0036	0.0032
	(0.0003)***	(0.0003)***	(0.0005)***	(0.0003)***	(0.0003)***
Wage	0.0054	0.0034	0.0049	0.0034	0.0050
	(0.0006)***	(0.0008)***	(0.0013)***	(0.0008)***	(0.0008)***
Mfg Wage	0.0067	0.0051	0.0095	0.0051	0.0048
	(0.0007)***	(0.0008)***	(0.0014)***	(0.0008)***	(0.0009)***
Year Controls	Yes	Yes	Yes	Yes	Yes
Division-Yr Controls		Yes	Yes	Yes	Yes
County Fixed Effects		Yes	Yes	Yes	Yes
THigh instead of T			Yes		
Common Support				Yes	Yes
Inv. Prob. Weights					Yes
Observations	101,934	101,934	79,422	99,456	99,456
Fixed Effect Groups	-	2,427	1,891	2,368	2,368

Notes: *, **, ***: Statistically different from zero with 90, 95, and 99 percent certainty, respectively. Standard errors are robust and clustered by county.

Table 5: Aggregate Effects: Geographic Spillovers

Outcome	1	2	3	4	5
Employment	0.0086	0.0076	0.0052	0.0003	0.0073
	(0.0005)***	(0.0005)***	(0.0006)***	(0.0005)	(0.0007)***
Earnings	0.0144	0.0110	0.0065	0.0002	0.0097
	(0.0009)***	(0.0010)***	(0.0013)***	(0.0010)	(0.0014)***
Population	0.0047	0.0037	0.0027	-0.0007	0.0031
	(0.0003)***	(0.0003)***	(0.0003)***	(0.0003)***	(0.0004)***
Wage	0.0059	0.0034	0.0012	-0.0001	0.0024
	(0.0007)***	(0.0008)***	(0.0010)	(0.0008)	(0.0011)**
Mfg Wage	0.0072	0.0051	0.0043	-0.0045	0.0027
	(0.0007)***	(0.0008)***	(0.0011)***	(0.0010)***	(0.0012)**
Year Controls	Yes	Yes	Yes	Yes	Yes
Division-Yr Controls		Yes	Yes	Yes	Yes
County Fixed Effects	Yes	Yes	Yes	Yes	Yes
State x Year Controls			Yes		
C Counties in T States				Yes	
Omit C in T States					Yes
Observations	101,934	101,934	101,934	57,204	67,956
Fixed Effect Groups	2427	2427	2427	1362	1618

Notes: *, **, ***: Statistically different from zero with 90, 95, and 99 percent certainty, respectively. Standard errors are robust and clustered by county.

Table 6: CPS Wage Regressions

Sector	1	2	3	4	5	6
All						
τ	0.0096	0.0060	0.0110	0.0048	0.0040	0.0025
$SE(\tau)$	(0.0012)***	(0.0019)***	(0.0029)***	(0.0014)***	(0.0011)***	(0.0025)
N	3,735,973	3,735,973	1,407,414	1,032,542	1,037,410	336,583
Manufacturing						
τ	0.0081	0.0006	0.0132	0.0025	0.0029	0.0011
$SE(\tau)$	(0.0015)***	(0.0024)	(0.0026)***	(0.0019)	(0.0028)	(0.0039)
N	608,927	608,927	285,041	178,277	178,781	73,722
Non-Linked Mfg						
τ	0.0073	-0.0003	0.0105	0.0010	0.0045	-0.0012
$SE(\tau)$	(0.0015)***	(0.0026)	(0.0029)***	(0.0025)	(0.0034)	(0.0049)
N	539,785	539,785	255,254	156,959	157,415	65,524
Year Controls	Yes	Yes	Yes	Yes	Yes	Yes
Division x Year Controls		Yes	Yes		Yes	Yes
County Fixed Effects				Yes	Yes	Yes
Age, Education, Race Controls	Yes	Yes	Yes			
Though 1990 Only			Yes			Yes

Notes: *, **, ***: Statistically different from zero with 90, 95, and 99 percent certainty, respectively. Standard errors are robust and clustered by state.

Table 7: Manufacturing Effects

Outcome	1	2	3	4	5
Mfg Employment	0.0057	0.0046	0.0109	0.0044	0.0046
	(0.0014)***	(0.0016)***	(0.0025)***	(0.0016)***	(0.0018)***
Year Controls	Yes	Yes	Yes	Yes	Yes
Division-Yr Controls		Yes	Yes	Yes	Yes
County Fixed Effects		Yes	Yes	Yes	Yes
THigh instead of T			Yes		
Common Support				Yes	Yes
Inv. Prob. Weights					Yes
Observations	79,716	79,716	60,816	78,288	78,288
Fixed Effect Groups	-	1898	1448	1864	1864

Notes: *, **, ***: Statistically different from zero with 90, 95, and 99 percent certainty, respectively. Standard errors are robust and clustered by county.

Table 8: Manufacturing Effects: Geographic Spillovers

Outcome	1	2	3	4	5
Mfg Employment	0.0055	0.0046	0.0040	0.0003	0.0036
	(0.0014)***	(0.0016)***	(0.0020)**	(0.0019)	(0.0023)
Year Controls	Yes	Yes	Yes	Yes	Yes
Division-Yr Controls		Yes	Yes	Yes	Yes
County Fixed Effects	Yes	Yes	Yes	Yes	Yes
State x Year Controls			Yes		
C Counties in T States				Yes	
Omit C in T States					Yes
Observations	79,716	79,716	79,716	45,864	54,096
Fixed Effect Groups	1,898	1,898	1,898	1,092	1,288

Notes: *, **, ***: Statistically different from zero with 90, 95, and 99 percent certainty, respectively. Standard errors are robust and clustered by county.

Table 9: Manufacturing Subsectors

Outcome:	ln(Mfg Employment)	ln(Revenues)
Sample	1	2
All Mfg		
τ	0.007	0.09
$SE(\tau)$	(0.002)***	(0.02)***
N	82,000	22,000
Non-Linked		
τ	0.006	0.07
$SE(\tau)$	(0.002)***	(0.02)***
N	82,000	22,000
Non-Linked Highly Tradable		
τ	0.011	0.12
$SE(\tau)$	(0.003)***	(0.03)***
N	82,000	22,000
Non-Linked Labor Intensive		
τ	0.011	0.10
$SE(\tau)$	(0.003)***	(0.02)***
N	82,000	22,000
Division x Year Controls	Yes	Yes
County Fixed Effects	Yes	Yes

Notes: *, **, ***: Statistically different from zero with 90, 95, and 99 percent certainty, respectively. Standard errors are robust and clustered by county.

Table 10: Productivity Effects

Outcome:	ln(VA/Worker)	ln(TFP)	ln(Price)
Sample	1	2	3
All Mfg			
τ	0.010	0.004	-0.001
$SE(\tau)$	(0.004)**	(0.003)	(0.004)
N	426,000	202,000	126,000
Non-Linked			
τ	0.010	0.004	-0.003
$SE(\tau)$	(0.004)**	(0.003)	(0.004)
N	415,000	196,000	122,000
Non-Linked Highly Tradable			
τ	0.010	0.000	0.008
$SE(\tau)$	(0.005)**	(0.004)	(0.010)
N	226,000	115,000	31,000
Non-Linked Labor Intensive			
τ	0.014	0.008	0.009
$SE(\tau)$	(0.004)***	(0.004)**	(0.009)
N	210,000	100,000	31,000
Division x Year Controls	Yes	Yes	Yes
Industry x Year Controls	Yes	Yes	Yes
Firm Fixed Effects	Yes	Yes	Yes

Notes: *, **, ***: Statistically different from zero with 90, 95, and 99 percent certainty, respectively. Standard errors are robust and clustered by county. The sample for column 1 is all establishments in the Census of Manufactures from 1963-2007. The sample for column 2 is all establishments in the CM for 1972-2002. The sample for column 3 is all products with valid prices in the CM from 1963-2007.

Figures

Figure 1: North Dakota Case Study

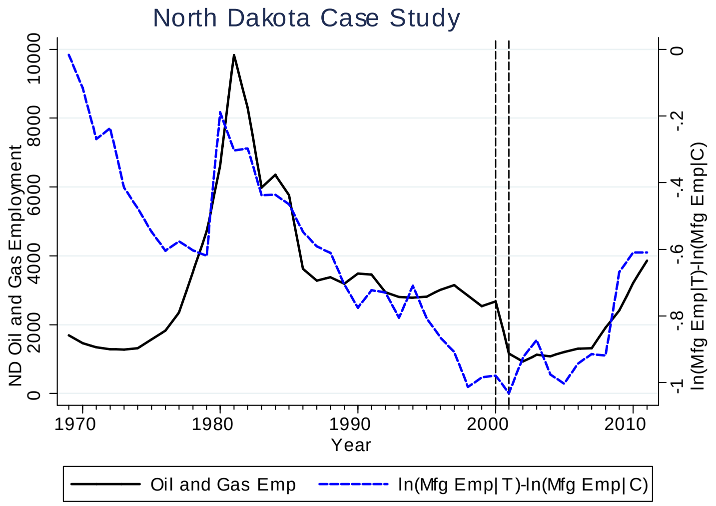

Notes: The solid black line shows total statewide employment in the oil and gas sector. The dotted blue line shows the unconditional difference in mean ln(Manufacturing Employment) between "treatment" counties that produce oil or gas and "control" counties that produce no natural resources. The vertical black lines highlight the change from SIC to NAICS classification systems between 2000 and 2001, which artificially re-classifies some jobs out of the oil and gas and manufacturing sectors.

Figure 2: Real Oil and Gas Prices

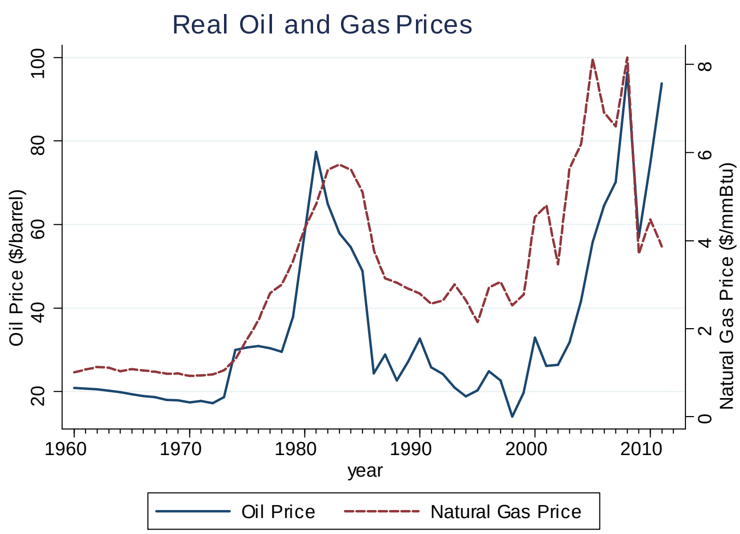

Notes: Prices are in real 2010 dollars.

Figure 3: Oil and Gas Production and Rig Counts

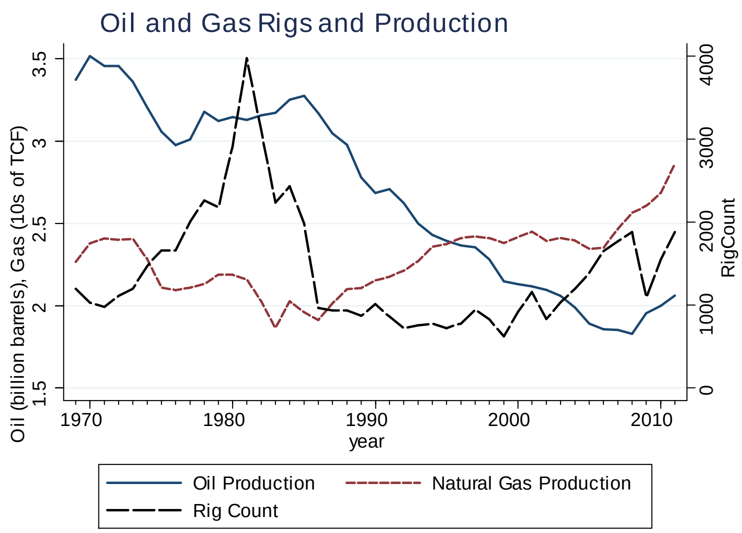

Figure 4: National-Level Manufacturing and Oil and Gas Employment

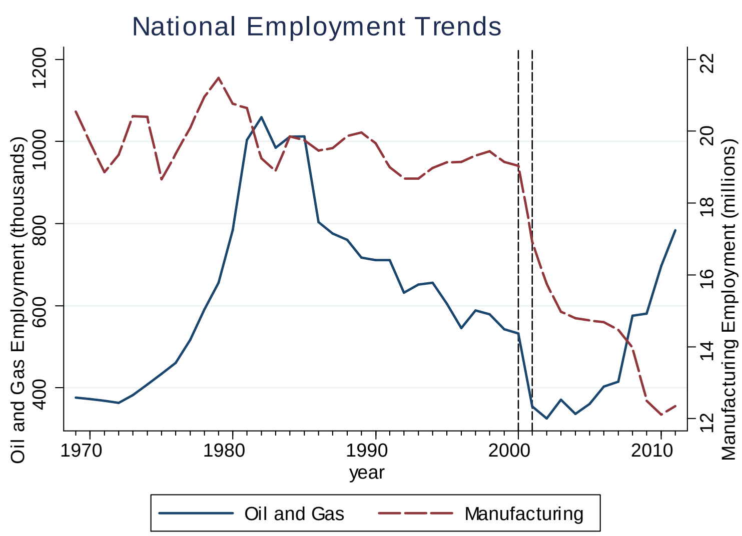

Figure 5: County Resource Intensity

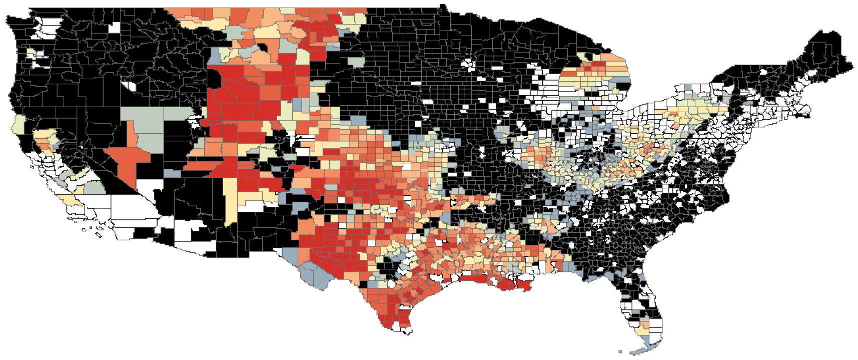

Notes: Counties in white have 1969 population density>100 people per square mile or 1969 population > 250,000. Counties in black have no oil, gas, or coal production between 1969 and 2011. Colors from blue to red indicate increasing intensity of oil and gas production over the 1969-2011 period, per 1969 population.

Figure 6: Treatment and Control Counties

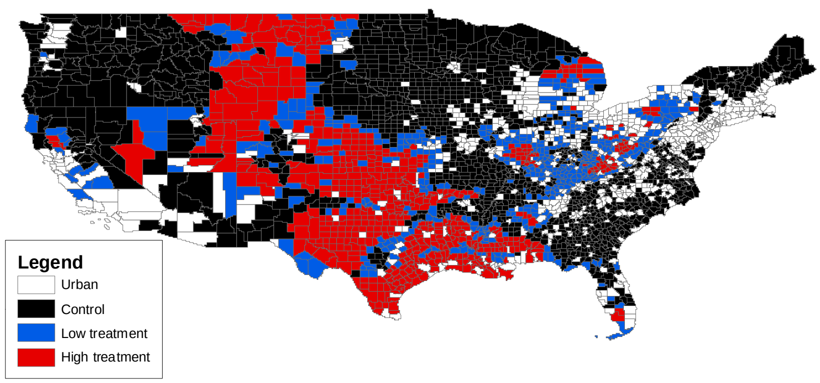

Notes: Counties in white have 1969 population density>100 people per square mile or 1969 population > 250,000. Counties in black have no oil, gas, or coal production between 1969 and 2011. Treatment counties have positive oil or gas production between 1969 and 2011, and "High Treatment" counties average more than $1000 in oil and gas production over that period per 1969 population.

Figures 7A-7E: Effects on Employment, Earnings, Population, Wages, and Manufacturing Wages

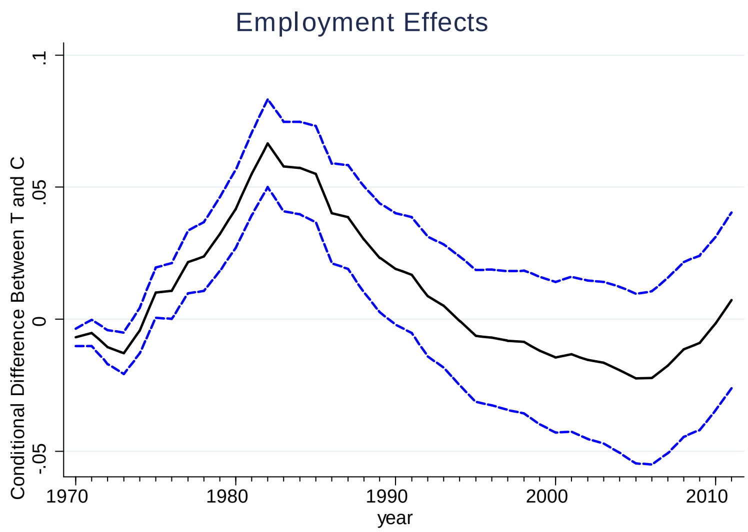

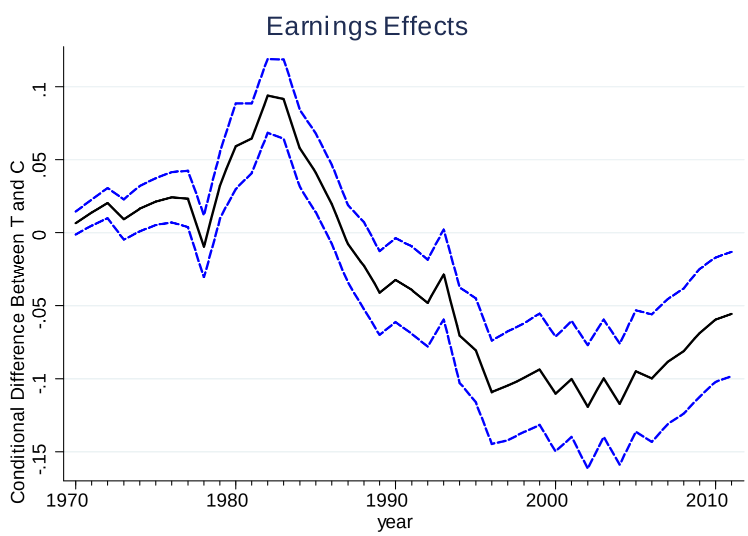

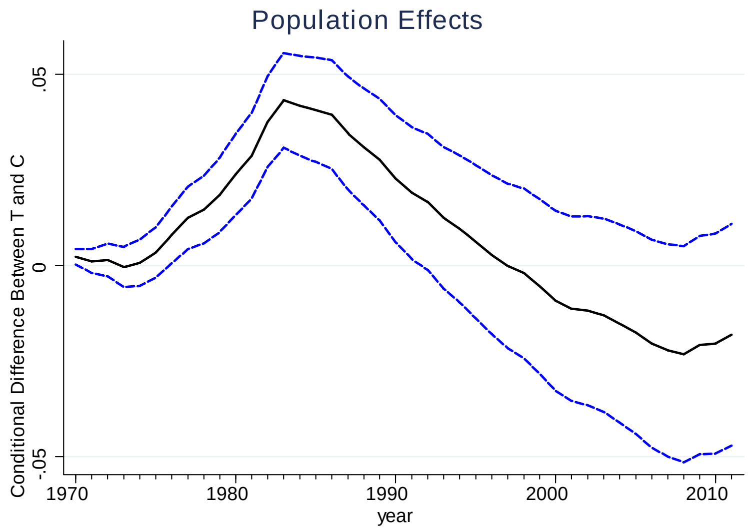
Population Effects
Conditional Difference Between T and C
.05
0
-.05
1970
1980
1990
2000
2010
year

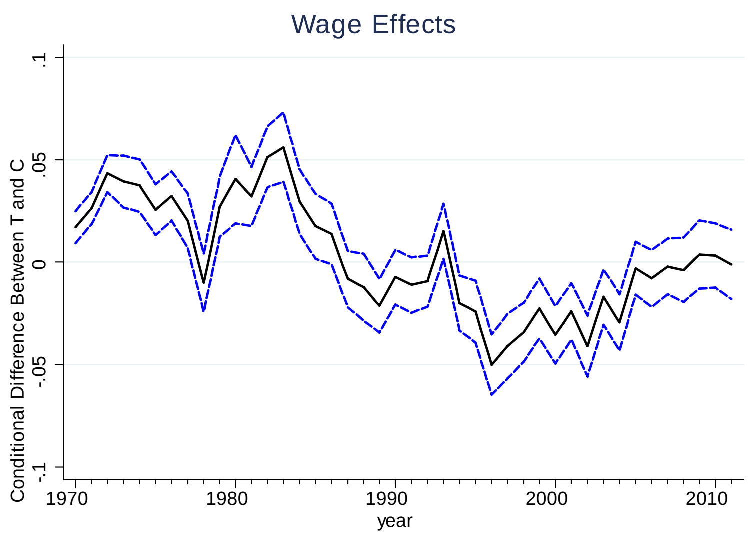
Wage Effects
Conditional Difference Between T and C
.1
.05
0
-.05
-.1
1970
1980
1990
2000
2010
year

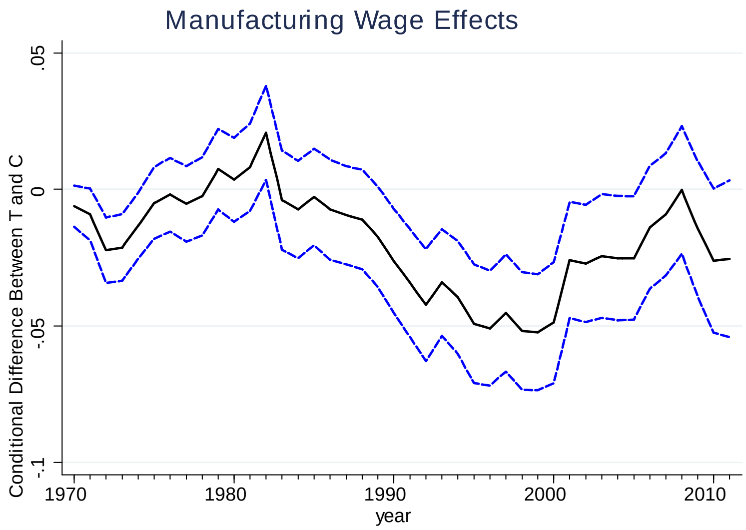

Notes: These figures present the coefficients and 90 percent confidence intervals from estimating Equation (5), with different aggregate outcome variables.

Figure 8: Aggregate Trends

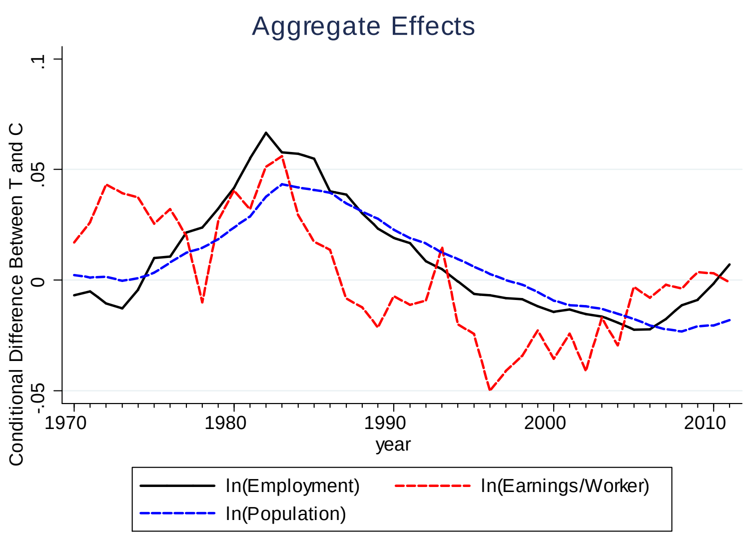

Notes: This shows the regression coefficients from Equation (5) for county aggregate employment, population, and earnings per worker.

Figure 9A-9C: CPS Wage Coefficients

All Sectors

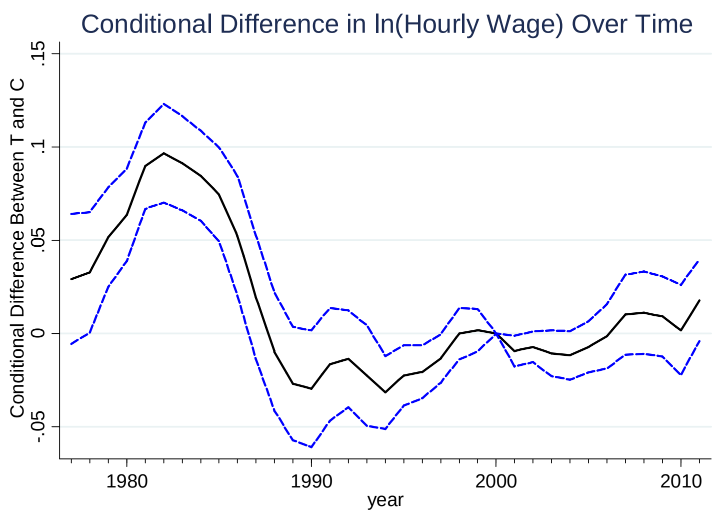

Manufacturing Wages

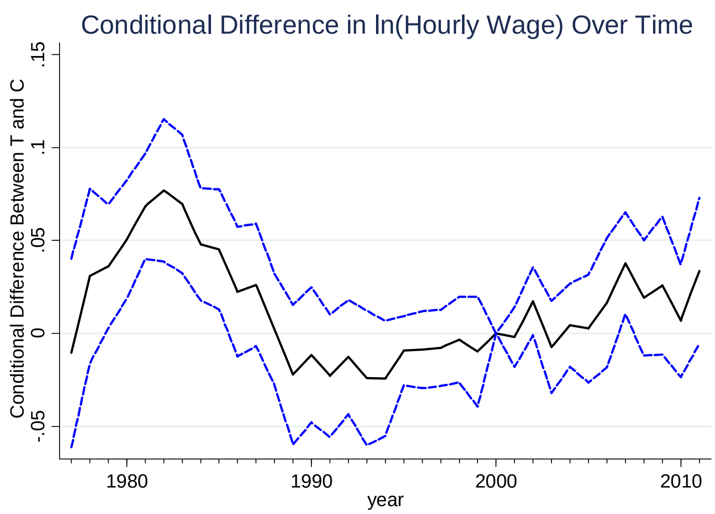

Non-Linked Manufacturing Wages

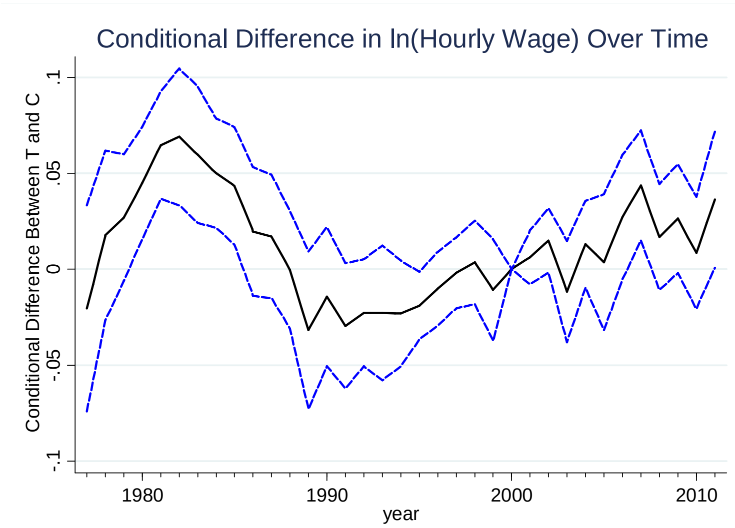

Notes: Figures 8A-8C present the coefficients and 90 percent confidence intervals from estimating Equation (5), with samples of all workers, manufacturing workers, and non-linked manufacturing workers, respectively.

Figure 10: Manufacturing Employment Effects

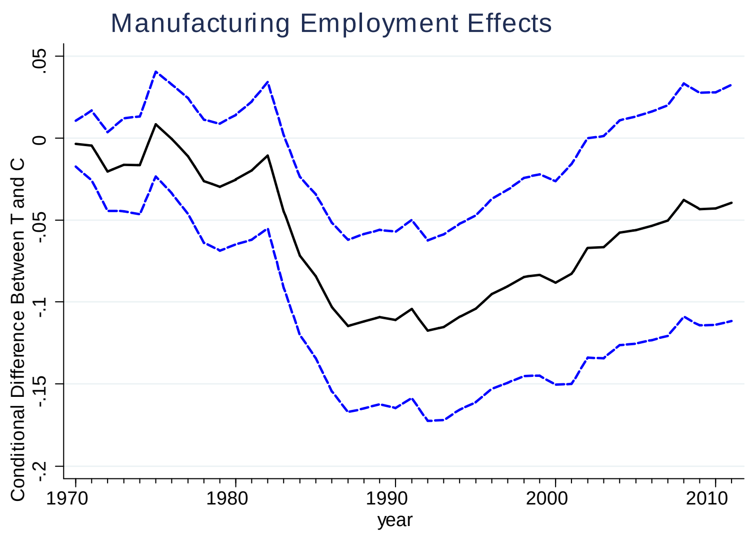

Notes: This figure shows the coefficients and 90 percent confidence intervals from estimating Equation (5) with log of manufacturing employment as the dependent variable.

Figure 11: Annual Manufacturing Effects vs. Oil and Gas Sector Growth

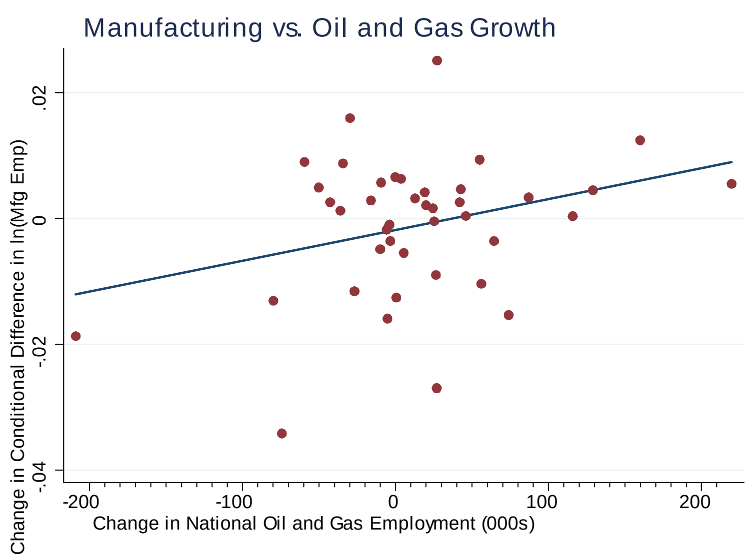

Notes: This figure plots the estimated change in natural log of manufacturing employment for treatment compared to control counties for each year of the sample against the change in national oil and gas sector employment.

Appendix: For Online Publication

The Local Economic Effects of Natural Resource Booms in Modern America

Hunt Allcott and Daniel Keniston

Appendix

Appendix Table 1: State-Level Sources of Oil and Gas Production Data

State	Fuel	Title	Source	Years
IL	Gas	Natural Gas Production in Illinois	Bryan Huff, Illinois State Geological Survey	1973-1992
IL	Oil	Historic County Production in Illinois	Bryan Huff, Illinois State Geological Survey	1932-2011
IN	Gas	Petroleum Data Management System	http://igs.indiana.edu/PDMS/WellSearch.cfm	1863-2011
IN	Oil	Petroleum Data Management System	http://igs.indiana.edu/PDMS/Fields.cfm	1953-2011
KS	Oil, Gas	County Production	http://www.kgs.ku.edu/PRS/petro/interactive.html	1960-2011
KY	Oil	Oil and Gas Production	http://kgs.uky.edu/kgsmap/OGProdPlot/OGProduction.asp	1883-2011
KY	Gas	Oil and Gas Production	http://kgs.uky.edu/kgsmap/OGProdPlot/OGProduction.asp	1986-2011
LA	Oil, Gas	Crude and Natural Gas Production by Parish	Sharron Allement, Louisiana Office of Conservation	1965-1977
MI	Oil, Gas	Michigan's Oil and Gas Fields, 1965-1982	http://www.michigan.gov/deq	
MT	Oil, Gas	Annual Reviews for the Years 1965-1985	http://bogc.dnrc.mt.gov/annualreview/	1965-1985
NV	Oil, Gas	Historical Production	Lowell Taylor, Nevada Division of Minerals	1954-2011
NY	Oil, Gas	New York Natural Gas and Oil Production	http://www.dec.ny.gov/energy/1601.html	1967-2011

Notes: This details additional state-level sources of oil and gas production data that are used to augment the DrillingInfo database.