

The Local Economic and Welfare Consequences of Demand Shocks for Coal Country

Daniel Kraynak

Working Paper 25-03
October, 2025

The Local Economic and Welfare Consequences of Demand Shocks for Coal Country

Daniel Kraynak*

Abstract

This paper estimates the welfare costs of declining coal demand from the power sector on coal mining regions of the US. Using an instrumental variable derived from a stylized model of the electricity sector, I estimate that coal producers shed jobs and wages primarily in coal mining and adjacent industries. In-migration, home values, and public education expenditures also decline. Applied in a spatial equilibrium framework, my estimates imply about \$0.85 billion in costs to coal country residents resulting from a net decline of \$8.03 billion in thermal coal production value from 2007-2017.

JEL Codes: H72, J23, Q32, Q40, L71

Keywords: Coal; Coal Country; Labor markets; Distributional impacts

*US Environmental Protection Agency, National Center for Environmental Economics. Correspondence email: kraynak.daniel@epa.gov. Much of this project was developed as a graduate student in the Department of Agricultural and Resource Economics, University of Maryland, College Park - for that I am especially thankful to Roberton Williams III, Joshua Linn, Louis Preonas, and James Archsmith. Many additional thanks to Pete Maniloff, Wade Davis, Ann Wolverton, participants of the AERE 2024 Summer Conference, NAREA 2024 Workshop, Camp Resources XXVIII, the Colorado University Workshop in Environmental and Resource Economics and several anonymous referees for many useful comments. The views expressed in this paper are those of the author, and do not necessarily represent those of the US EPA. No official or unofficial agency endorsement should be inferred.

1. Introduction

From 2005 to 2020, coal-fired electricity generation in the US declined dramatically, and domestic coal mining employment fell by approximately 25,000 jobs (35%).¹ The expected continuation of such a decline has spurred significant interest in policies to limit the welfare impact of structural changes in electricity markets on the coal producing communities that have supplied them for over 100 years (Morris, 2016). By one estimate, existing policy initiatives have already committed some \$410 million dollars to applicant counties in coal producing regions (Shelton et al., 2021).

The impacts of declining coal demand on coal country are of interest to economists for several reasons. First, they are a significant demand-side shock to labor markets in producing regions, which relates to a vast literature that studies the economic consequences of trade and natural resource shocks (Autor, Dorn, and Hanson, 2013; Allcott and Keninston, 2018). Second, because coal is an emissions-intensive fuel, economic impacts on coal producers have become a prominent distributional consequence of various environmental policies designed to limit air pollution. Third, in the US and elsewhere, some of the policy response is likely to take the form of place-based policies, such as geographically-targeted grants (Furnaro et al., 2021). The success of these policies ultimately depends on whether they address the real economic impacts caused by a declining coal industry.

Identifying the effects of demand shocks on coal producers is challenging because coal producing regions in the US differ from the rest of the country in important ways. This difficulty stems from the fact that many mining regions in the US were settled from 1870 to 1930 when coal was the dominant source of fuel in the US and coal mining employment reached its peak (Matheis, 2016)². These regions developed economies closely tied to the coal industry and have endured associated cycles in population, employment, wages, and poverty, despite coal mining's long-term decline in employment at the national level (Betz et

¹Data is from the Federal Reserve Bank of St. Louis. Difference calculated Jan. 2005 to Jan. 2020.

²As documented in Matheis (2016), US coal mining employment during this period reached its maximum of more than six hundred thousand employees, declining precipitously after the 1940s. Nationally, this trend continued largely uninterrupted through the present, with the exception of a short-lived employment bump in the 1980s.

al., 2015; Deaton and Niman, 2012). As a result, it is difficult to construct counterfactuals for economic indicators in these regions, and correspondingly to estimate the impacts of changes in the coal industry on local economic indicators in coal country.

To overcome these challenges, this paper leverages downstream variation from the power sector. Combining this credible source of identifying variation with a rich set of outcomes, this paper aims to provide a more complete picture of the local labor market impacts of shocks to coal country than previous literature. Additionally, I incorporate several of the estimates into a measure of economic welfare, opening the door to a comparison of costs borne in coal producing areas and environmental benefits from reduced coal combustion, such as those estimated by Johnsen, LaRiviere and Wolff (2019). To isolate aggregate shocks to coal demand from electricity markets that are exogenous to labor markets in mining counties, I predict power plant dispatch by constructing regional electricity models using hourly generation data from thermal power plants.³ I then map changes in electricity markets up the supply chain to coal producers, using comprehensive data on contracts between large coal-fired power plants and mines. The motivation for the empirical strategy is a comparison of producing counties with greater or lesser exposure to recent declines in coal demand from the electricity sector, eliminating the need to compare historical producers with the rest of the country.

The model outputs are applied as an instrumental variable to estimate the effect of production declines at mines on a broad set of county-level outcomes in a long-difference, within-coal-country specification using data from 2007-2017. I find that a typical decline in coal production value per capita at the county level reduces mining employment and wages and causes modest employment and wage spillovers across industries, on the order of 30% of the initial impact on the mining sector. Substantial effects occur in Appalachia and in the Illinois Basin, and point estimates are similar for Western coal producers (though these estimates are more imprecise). On the 5-10 year time horizon studied, I find evidence consistent with the idea that many out-of-work coal miners find new jobs and unemployment effects on local areas are small. I also provide causal estimates of the effect of demand

³I employ the EPA's Continuous Emissions Monitoring Systems (CEMS) data and model static, least-cost dispatch of generators following a large economic literature dating at least to Borenstein et al. (2002).

shocks on several knock-on effects previously undocumented in the literature. I find that local school district expenditures and state aid to general purpose governments decline, and that poverty increases. Motivated by estimated declines in population and in-migration, I summarize the causal estimates in the context of a theoretical model of location choice. My estimates of household-level declines in economic welfare average around 5% of initial wages per household, and imply a total welfare impact on the order of \$0.85 billion resulting from an \$8.03 billion decline in thermal coal production (10-15% of baseline production value).

The primary contribution of the paper is a set of credible estimates of the impacts of a recent contraction in the coal industry on coal country's residents.⁴ My results confirm some of the analysis in Weber (2020), which finds that coal producing regions have experienced declines in aggregate earnings and employment, while contrasting its finding of nonexistent local spillovers across industries.⁵ My results on public goods also complement those of Metcalf and Wang (2019), which finds little evidence for the hypothesis that coal mining layoffs lead to opioid usage. This paper also offers causal estimates of the phenomena suggested in Morris, Kaufman, and Doshi (2020), which provides descriptive evidence suggesting that there have been declines in local public revenue related to declines in coal production.

Second, by exploring a broader set of outcomes than previous studies, estimating capitalization effects, and explicitly estimating economic welfare changes using a model of spatial equilibrium, this paper contributes a more complete picture of the effect of demand shocks on coal country's residents which complements recent working papers exploring specific channels of economic impacts and exploring policy options (Blonz, Troland, and Tran, 2023; Hanson, 2023; Colmer et al., 2024). In doing so, it creates a connection to literature that has explored broad effects of economic shocks on labor markets in other contexts (Feler and Senses, 2017;

⁴In using policy-relevant variation from the electricity sector, this paper also connects labor markets in coal country to a sizeable environmental and industrial organization literature studying the effects of environmental regulations and the natural gas boom on electricity markets (e.g., Johnsen, Lariviere, Wolff, 2019; Linn and McCormack, 2019; Davis, Holladay and Sims, 2021; Watson, Lange, Linn, 2023).

⁵My findings align more closely with literature measuring the effects of demand and price shocks in the 1970s and '80s in Appalachia, which demonstrates that the economies of coal producing areas follow trends in coal demand, with employment and wage spillovers mostly to the locally traded goods sector. The same literature also documents the effects that booms and busts have on migration and educational decision-making, suggesting that non-wage channels are also important in the transmission of demand shocks to the local economy (Black et al., 2005 a,b).

Bartik, Currie, Greenstone and Knittel, 2019), as well as literature examining place-based policy (Busso, Gregory and Kline, 2013; Kline and Moretti, 2014).

The next section (Section 2) motivates the rest of the paper with background on coal markets and the data available to perform the analysis. Section 3 outlines the econometric strategy. Section 4 describes construction of the instrumental variable in detail. Sections 5 and 6 present 2SLS estimates of the effects of declining coal production on a variety of outcomes. Section 7 ties together the causal estimates of the paper and estimates changes in economic welfare. Section 8 concludes. Additional details on the data, supplementary empirical analysis, and the welfare calculations can be found in the Appendix.

2. Background and Data

From 2002 to 2007, the share of coal-fired electricity generation in the US was relatively constant.⁶ Over the next 10 years it fell to nearly half of its previous market share, reducing demand for coal from mines (Figure 1). Both during and prior to the period of declining coal demand from the power sector, coal deliveries to power plants tracked closely with total production reported from mines as demonstrated by Figure 2. The figure shows mine-reported production from the Mine Safety and Health Administration (MSHA) and Energy Information Administration (EIA) Form 7A alongside plant-reported coal deliveries from a combination of EIA 423 and EIA 923. Both data series are reported in short tons and benchmarked against EIA published aggregates. All of the data sources show a significant decline from the early and mid-2000s to the later 2010s.⁷ Coal mining employment, also reported in the MSHA data and EIA form 7A, follows a similar pattern over time as mine production and deliveries to power plants (Figure 3). These simple figures are highly suggestive that labor demand and coal production at mines are responsive to coal consumption by power plants, which links aggregate power plant demand to local labor markets.

⁶Cicala (2015) shows that a flurry of state-level electricity market restructuring halted in the wake of the California electricity crisis of 2000-2001, suggesting that state-level regulation was also relatively stable during this time.

⁷Over the longer run, 90-95% of thermal coal produced in the US ends up at a US power plant. However, some care is required in interpreting Figure 2, because coal is stored relatively cheaply at both the minemouth and onsite at power plants, so not all production and delivery is contemporaneous even at the annual level.

The core empirical task of the paper is to credibly estimate how the trends outlined above impact labor markets in coal country. I draw measures of local labor market outcomes from several publicly available sources of data. Data on wages, employment, and unemployment come from the Bureau of Labor Statistics' Quarterly Census of Employment and Wages and Local Area Unemployment Statistics. I also examine effects on local public revenues using data from the Census of Governments. A full Census of Governments only occurs in years ending in 2 and 7, so I construct a harmonized panel of 5-year differences at the county level to maintain a consistent estimation strategy throughout. Finally, I present evidence of effects on population, migration, home values, and local public goods. These measures come from a variety of sources which are outlined in detail in the Data Appendix.⁸

The estimation sample includes all counties with a coal mine recorded in the MSHA data as of 2000, 255 counties. One dominant producer and extreme outlier in terms of productivity (Campbell County, Wyoming of the Powder River Basin) is omitted from all estimates, as it accounts for around 35% of coal production by short tons for most years in the sample - a figure which remains stable from 2007-2017. Summary statistics for the estimation sample in 2007 are shown in Table 1. On average, coal counties are relatively rural, with an average population of around 62,000. Though the share of coal miners in the workforce averaged around 3%, the maximum in the sample is more than 40%, and coal counties produced an average of \$111.5 million of thermal coal in 2007 (valuation is discussed below). Notably, the average mining sector job in coal country had about 20% higher wages than the average across all jobs. Appendix Table A1 shows statistics on variables relevant to the study but not covered in Table 1.

One of the key contributions of the paper is a new estimation strategy using instrumental variation from the power sector, described in further detail in the following sections. Here, I briefly outline the data sources used to instrument for coal demand from the electricity sector. Data on electricity generation by fossil fuel plants comes from the Environmental Protection Agency's Continuous Emissions Monitoring Systems (EPA CEMS) data, which contains hourly gross generation and emissions for most fossil-powered units in the US. I use

⁸Population data comes from the US Census and migration from the IRS county-to-county migration files. Several indicators (such as educational attainment and median home values) are drawn from the ACS.

the EPA’s eGRID data as my source of metadata for CEMS units. The EIA 423 and EIA 923 data for power plants contains information on coal characteristics and origin reported by power plants. These data are used to infer a long-run average heat content of coal from each coal producing county, and to link between coal mines and the coal-fired power plants they supply. I value the coal by applying state-level fuel prices from the EIA’s State Energy Data System (SEDS) to the estimated Btus of coal supplied by the mines.⁹ Production at the mine-mouth and employment of coal miners is reported by the MSHA and EIA Form 7.

3. Empirical Strategy

The empirical strategy is designed to estimate the effect of recent declines in coal production by mines on coal producing communities. Using nearly all¹⁰ of the 255 counties with a coal mine present in MSHA data since 2000, I start by estimating the coefficients of models with the form of Equation (1). Counties are indexed by c and years are indexed by t . The difference operator Δ indicates a 5-year difference between time t and $t - 5$.¹¹

$$\frac{\Delta Y_{ct}}{Pop_{c,t-5}} = \beta_1 \left[\frac{-\Delta D_{ct}}{Pop_{c,t-5}} \right] + X'_{c,t-5} \beta_2 + r'_c \gamma_t + \epsilon_{ct} \quad (1)$$

The variable ΔY_{ct} is a county-level change in local outcomes such as employment in different industries, and is normalized by start-of-period ($t - 5$) population estimates from the US Census. The key independent variable is ΔD_{ct} , the change in total value of coal produced at mines in county c , similarly normalized. This variable captures both declining production value (for which miners might be laid off, work fewer hours, or face wage cuts) and

⁹While the natural physical unit of measurement in the EPA CEMS and EIA 423 and 923 data is Btus of coal, the MSHA/EIA supply-side data is measured in short tons, meaning that conversion to dollar amounts for continuity with the coal demand measure is not trivial. To do so, I use the EIA 423 and 923 data to infer a long-run average Btu content of coal produced in each county. Where sufficient data is not available, I impute the average of the five nearest neighbor counties (based on centroid distance), allowing me to use state-level prices from the SEDS database to convert to dollar amounts.

¹⁰As mentioned above, Campbell County, WY is excluded from all estimation due to its status as a very large producer and an extreme outlier. Estimation results are sensitive to its inclusion.

¹¹I use 5-year differences ending in years 2 and 7 due to the data restrictions imposed by the Census of Governments. They also align relatively conveniently with trends in coal demand - relative stability from 2002-2007, declining dispatch but with some modestly increasing production from 2007-2012, and rapid declines from 2012-2017.

mine closures (where the number of miners goes to 0). State-year variation in coal prices is also reflected. Production values are smoothed with two lags at the start of each period (e.g., $t = 2007$ is averaged with 2006 and 2005) to reduce measurement error due to coal storage both at mines and plants, as well as frictions that might slow labor market adjustment. The coefficient of interest is β_1 , which measures the relationship between changes in coal production value per capita and county outcomes.

I use Equation (1) to estimate labor market outcomes for which the object of interest is an aggregate change within the county, such as the total number of jobs lost or the total decrease in wage earnings by sector. The model is estimated using two stacked five-year differences, 2007-2012 and 2012-2017, which breaks the ten-year period of rapidly declining coal dispatch into an earlier period with small impacts (in which coal production and employment actually increased in many counties) and a later period in which production and employment in mining rapidly declined. A vector of time fixed effects (γ_t) varies at the regional level via interaction with a dummy vector indicating each county's region (r_c), to control for transportation cost differences across regions and other economic and cultural differences that might independently drive trends in local outcomes, with coal producing regions adapted from Stoker et al. (2005). A county-level map of the regions is available in the Appendix.

The unit of analysis is the county, but estimation is structured around the year 2000 commuting zone to acknowledge the important economic interactions between counties within a commuting zone (Fowler, Rhubart, and Jensen, 2016). Standard errors are clustered at the commuting zone level, of which there are 111 in the sample. The $X_{c,t-5}$ vector contains several relevant controls measured at the start-of-period ($t - 5$) computed at the commuting zone level (including non-producing counties) that might independently impact mining employment, wages, and local labor market conditions more generally. Specifically, I include the baseline share of mining employment in the commuting zone, the share of manufacturing employment, share of employment in the transportation and utility sectors, and share of the population with a bachelor's degree. Additionally, I include the baseline value of oil and gas production calculated using Enverus and EIA data (aggregated to the commuting zone level) to mitigate concerns that co-location of oil and gas production with coal reserves influences

the estimates.

I also use Equation (2), which shares almost the exact same structure as Equation (1) with one key difference - the dependent variable is typically a log change, and any normalization by population or employment is *contemporaneous*. While this specification has the disadvantage of omitting zero-valued outcomes, it has the advantage of directly estimating an approximate proportional change in the outcome variable. This is useful to estimate changes in quantities that have most meaning with contemporaneous normalizations, such as average wages (per job) or average public expenditures per capita.

$$\Delta \log(Y_{ct}) = \beta_1 \left[\frac{-\Delta D_{ct}}{Pop_{c,t-5}} \right] + X'_{c,t-5} \beta_2 + r'_c \gamma_t + \epsilon_{ct} \quad (2)$$

Coefficients β_1 from Equation (1) and (2) might have a causal interpretation if one is willing to assume that contemporaneous changes in coal production value are exogenous or as-good-as random with respect to pre-existing labor market trends in the county of production. It is not difficult to imagine channels through which this assumption is potentially violated. Some mines may face a secular decline in productivity that also drives lower production and lower wages for miners. Some commuting zones may have better amenities or tighter labor markets that affect the ability of mines to compete on output price and decrease their production. Put another way, if local economic conditions impact mine output, then ordinary least squares estimates based on Equations (1) and (2) suffer from reverse causality. To overcome these issues, I instrument for changes in coal production value using predicted demand from a model of the electricity sector mapped back to coal counties, described in greater detail in the next section.

4. Construction of the Coal Demand Instrument

The instrumental variables strategy of the paper identifies the proportion of declining coal production at mines that is driven by factors reducing generation by coal-fired power plants in electricity markets (most importantly, declining natural gas prices). To parsimoniously summarize the influence of these factors, I predict yearly variation in aggregate coal

consumption by power plants using simplified models of electricity generation at the regional level covering the continental US. Long run changes in coal deliveries to plants are mapped to coal producing counties using market shares calculated during the period of relative stability from 2002-2007. The instrument follows a shift-share logic, where market shares are held fixed and are multiplied by exogenous variations in demand, yielding an instrument for coal demand from each coal mining county in the US. However, the structure of the instrument differs substantially from the canonical setting using industry shares and growth rates (Bartik, 1991), or recent applications using mining employment shares and growth rates (Weber, 2020). Instead, the shares reflect the historical importance of supply counties to electricity generation regions, and the shocks rely on regional variations in aggregate coal demand. This section describes construction of the instrument in more detail - the dispatch models, the shares, and the formal definition of the instrument. The final subsection discusses identifying assumptions.

a. Dispatch Models

The dispatch models are least-cost economic dispatch models. Generating units are dispatched according to their costs, a very simple approximation of actual electricity market operation in the US.¹² Low-cost units operate when total load (electricity usage) is low and higher-cost units operate only at high levels of load. The model follows the approach of Johnsen, LaRiviere and Wolff (2019) or JLW, where supply curves are constructed for each electricity load region and year and dispatch is run hourly using observed state-level fuel prices and region-level demand. Inputs are the EPA's CEMS data and the EIA's state-level fuel prices (from SEDS). Similar to JLW, the primary purpose of the electricity sector modelling in this paper is to summarize relevant market variations to be used in causal estimation, rather than pursuing a highly detailed or accurate representation of electricity market operation. Minimal adjustments are made to the CEMS data so that the sources of variation are relatively transparent.

¹²This is sometimes achieved through a coordinated market mechanism as in the ISO/RTO regions, and sometimes by regulated utility companies managing wholesale generation.

The main departure of this paper’s dispatch model from that of JLW is the use of the EPA’s eGRID Subregions (polygons depicted in Figure 4) rather than North American Electric Reliability Corporation (NERC) regions. EPA eGRID subregions divide the continental US into 20 electricity load regions within which electricity flows relatively freely, rather than 8 as with NERC. This choice was motivated to balance the goal of choosing the smallest electricity market definition possible (increasing predictive power of the instrument) with ensuring that each coal producing county only enjoyed a small market share in terms of fuel delivered to market participants (strengthening the argument for exogeneity).¹³ Market shares of mining counties and the instrument’s predictive power are discussed in the following subsections.

Within a market and year, a model run consists of 4 main steps. First, using only CEMS data, an average heat rate for each thermal generating unit is calculated by averaging heat input in millions of British thermal units (mmBtu) per megawatt-hour (MWh) of electricity generation. Second, unit capacity is calculated using the observed maximums of generation on days in which the unit produced for more than 1.5 hours.¹⁴ Third, a supply curve is constructed by multiplying unit heat rates by state-level fuel input prices, yielding an estimated marginal cost for each unit. Ranking units by marginal cost and calculating a running sum of observed capacities yields the supply curve or “merit order”. Fourth, for each hour of the year, the supply curve is intersected with total observed generation by fossil fuel plants in the load region, and the lowest cost plants required to meet this demand are “dispatched” for the full hour.¹⁵

To illustrate, Figures 5a and 5b show an example of the supply curve modelled for a market in the southern US (SRMV, a subregion of SERC in the Southeastern US) in years

¹³In reality, neither NERC regions nor EPA eGRID Subregions are the actual boundaries at which electricity dispatch occurs, which are known as balancing authority areas (BAAs). In most cases BAAs are contained within eGRID subregions, and electricity is traded across BAAs, to a lesser extent across EPA subregions, and even less across NERC regions.

¹⁴This procedure follows prior literature and limits measurement error from unobservable plant changes such as maintenance and retrofits.

¹⁵An important model simplification is implicit here. This procedure approximates electricity supply assuming that the universe of non-CEMS generators are price insensitive. This is not always, but often the case (e.g., most renewables are non-dispatchable, nuclear plants face very costly constraints in altering output).

2007 and 2017. The gray shape behind each estimated supply curve shows the (arbitrarily rescaled) density of hourly fossil load in the region. In other words, the height of the gray curve is proportional to the number of hours that total regional load was equal to the corresponding value on the x-axis. Generating units on in the left part of the figure are frequently dispatched in the model, while generating units to the far right are predicted to have little or no generation. In 2007, the lowest cost units were all coal-fired power plants. By 2017 there was a high degree of overlap in the marginal costs of gas units and coal units. Coal and oil-fired units shifted substantially rightwards in the figure. Overall, there is a noticeable downward shift in the cost curve for each quantity of generation, due almost entirely to declining natural gas prices which shifted gas plants lower in the merit order.

The dispatch model yields a value of predicted generation for nearly the universe of coal-fired generators in the US, derived from the interaction of generation technologies, market-level load, and state-level fuel prices. Figure 6 plots predicted versus actual generation for the same region-years whose supply curves are depicted in Figures 5a and 5b. The model has significant predictive power, but fails to capture many realistic features of grid operation, including dynamic and congestion constraints that shift the true merit order. Additional prediction error is induced by measurement error from gaps in CEMS reporting requirements, such as those highlighted in Cicala (2022).¹⁶

The prediction outputs from the dispatch models are the exogenous variations that form the core of the estimation strategy. Each coal plant’s predicted generation is converted to a predicted coal demand (via the estimated heat rate). Annual region-level coal demand is calculated by summing across all coal plants in the region, which forms the “shift” component of the instrumental variable.

b. Mapping from Load Regions to Coal Counties

To map changes in aggregate demand for coal derived from the electricity sector model to mining counties, I use producing-county-to-load-region shares that do not vary during

¹⁶For instance, in a data appendix Cicala (2022) presents evidence that for early CEMS years (1999-2012), combined cycle units may not report generation from the “second” or heat recovery cycle. In the present context, these units would have substantially overstated heat rates, and consequently overstated costs.

the estimation period following Watson, Lange, and Linn (2023). For each load region b , I compute the share of coal Btus delivered by each supply county c over the entire period from 2002-2007.¹⁷ Figure 7 shows a strip plot that displays the entire distribution of market shares separated by eGRID Subregion. No coal producing county in the estimation sample enjoys a market share in the period from 2002 to 2007 of more than 40% in terms of Btus of coal delivered from county to load region and only 3 counties enjoy a market share of more than 30% in any load region.¹⁸ This fact supports the assumption that EPA’s eGRID Subregions are sufficiently large as a load region definition such that no one coal producing county can exert a dominant influence on load region dispatch, which lends credibility to the use of load region shocks as exogenous variation.

Predicted demand for coal deliveries from county c is derived by mapping yearly variation in region-level coal demand during the years of falling dispatch back to previous suppliers. Mathematically, predicted demand for coal from county c in year t , \widetilde{D}_{ct} is given by

$$\widetilde{D}_{ct} = \sum_b \left[\frac{Deliv_{cb0}}{Deliv_{b0}} \right] \widetilde{Demand}_{bt} \quad (3)$$

for $t \in \{2007, \dots, 2017\}$. Reading Equation (3) from right to left, for each electricity load region b , annual modelled coal demand \widetilde{Demand}_{bt} is aggregated from the dispatch model in years $t \in \{2007, \dots, 2017\}$. These load region demands are mapped back to supply county c via its share of total deliveries to load region b , $\frac{Deliv_{cb0}}{Deliv_{b0}}$, where $t = 0$ denotes the entire period from 2002-2007. Summing across shocks from the 20 load regions yields the expected demand for coal from county c . All market share calculations are performed in Btus and later converted to dollars using the EIA SEDS data.

¹⁷This calculation relies on EIA 423 and 923 data, which has a plant reporting threshold of 50 or more MW of nameplate capacity, but still contains about 95% of all US coal deliveries (Figure 2).

¹⁸Campbell County, Wyoming is the only exception to this condition, and is excluded from all calculations and estimation other than previous descriptive analysis of trends in national production.

c. Testing Instrument Validity

Valid IV estimation of the coefficient β_1 in Equations (1) and (2) requires that the predicted demand instrument is highly correlated with county-level production (relevance) and operates only on local labor market outcomes through its direct, contemporaneous impact on the coal industry (exclusion restriction). This subsection investigates these conditions.

The predicted demand instrument is highly correlated with county level production across different specifications and is robust to controls. Figure 8 plots the IV (5-year changes in predicted coal demand value per capita) against observed changes in coal production value in USD per capita. Each county has two observations corresponding to the periods 2007-2012 and 2012-2017. Figure 8a differentiates points by region suggesting that it is highly (but imperfectly) predictive across all coal regions, while Figure 8b differentiates points by period showing that the majority of declines in demand and production occurred in the latter period from 2012-2017. First stage estimates are available in Appendix B. An important conclusion from this exercise is that the dispatch models, despite being fairly simple and coarsely aggregated, have substantial predictive power. If measurement errors in the dispatch model procedure were very important, then the instrument would likely be weak, which turns out not be the case in practice.

One advantage of this paper’s approach is that estimation follows a difference-in-differences format, with clear pre- and post-treatment periods. In light of recent literature formalizing identification in shift-share settings, failing to reject parallel pre-trends in this paper’s context would support the exogeneity of the instrument in terms of “shares” or the exposure measure. So following the recommendation of Goldsmith-Pinkham et al. (2020) and the example of Autor, Dorn, and Hanson (2013), I conduct two tests of the parallel pre-trends assumption analogous to conventional difference-in-differences with continuous treatment intensity.¹⁹

¹⁹Because inputs to the dispatch models are region-level electricity load, technology mixes, and state-level fuel prices, it may also be plausible to view exogeneity in terms of the shocks under the assumption that those inputs are sufficiently “large” relative to mining counties to be treated as exogenous (Borusyak, Hull, and Jaravel, 2022). Unfortunately, with only 20 regional shocks, it is difficult to assess whether these shocks can plausibly be considered conditionally uncorrelated with quasi-random assignment to units based on their observed distribution. For this reason I rely on tests of the parallel-trends assumption in terms of the overall

Table 2 begins the pre-trends test, showing reduced-form estimates of β_1 from Equation (4) below against employment in different sectors in different time periods. Equation (4) follows the same structure as Equation (1) of the main text, except outcomes are regressed directly against the instrument and commuting zone controls are omitted to explore the unconditional relationship of the instrument to important outcomes. The key regressor ($\Delta\widetilde{D}_{ct}$ per capita) is standardized to have mean 0 and standard deviation 1, so that coefficients measure the effect of a one standard deviation decline in predicted coal demand per capita on local employment across sectors.

$$\frac{\Delta Y_{ct}}{Pop_{c,t-5}} = \beta_1 \left[\frac{-\Delta\widetilde{D}_{ct}}{Pop_{c,t-5}} \right] + r'_c \gamma_t + \epsilon_{ct} \quad (4)$$

Panels A-C of Table 2 show reduced-form estimates of β_1 from Equation (4) by time period over the period of declining coal demand (2007-2017). Panel A shows the results of estimating with both of the later periods in the stacked difference model. The coefficient of -0.53 in column (1) of panel A means that a one standard deviation exogenous decline in coal demand is associated a 0.53 percentage point decrease in coal miners per capita in the county. All other coefficients in the table are interpreted similarly. Panels B and C of Table 2 estimate Equation (4) separately for the periods 2012 – 2017 (panel B) and 2007 – 2012 (panel C), where $\Delta\widetilde{D}_{ct}$ is the contemporaneous value of demand declines from power plants. The results in panel B are similar to those in panel A, suggesting that the dominant effect of demand declines occurred in this later half-decade. In panel C the results remain negative but become statistically weak.

Panel D of Table 2 serves as a placebo exercise of the reduced form specification. In this panel, past changes in employment are regressed against the average of future changes in predicted demand for coal from the electricity sector. This is done to alleviate concerns that the estimates capture a long-run secular decline in either productivity or general labor market conditions, rather than the contemporaneous effect of declining demand from coal power plants. The key right-hand-side regressor is the average value of long-run demand declines

instrument.

in the half-decades from 2007-2012 and 2012-2017. The results show no evidence of negative correlation between future demand declines and past (2002-2007) trends in employment across sectors. In fact, the correlations are relatively small and positive, which is evidence against such reverse causality and suggests that counties that fared worse from 2007-2017 fared slightly better from 2002-2007.²⁰

The results in panel D of Table 2 show that the measure of predicted coal demand used as an instrumental variable is uncorrelated with pre-existing trends in local labor markets, which supports a valid exclusion restriction. That is, no effect is found prior to the demand shocks that the IV is meant to measure. While this is comforting for direct estimates on employment and wages, it is possible that supply counties that would later suffer demand declines have other pre-existing trends that separate them from other coal counties.

To address the concern that the instrument may capture spurious trends in locations that would later suffer demand shocks, Figures 9a and 9b show coefficients from regressions similar to panel D of Table 2 that test this assumption for several outcomes relevant later in the paper.²¹ That is, Figure 9 shows regressions of past trends (2002-2007) against future declines in coal demand. While technically Figure 9a does not show strong statistical evidence of correlation between future demand declines and trends in total employment and wages, labor force participation, home values, migration, or local government revenue from 2002-2007, some of the point estimates are large and 95% confidence intervals are wide (a coefficient of 0.1 indicates a 10% difference). Comfortingly, Figure 9b suggests that adjusting for start-of-period commuting zone controls does a reasonable job of adjusting for pre-period differences in these variables, moving the largest coefficients close to zero. Eight of 10 coefficients are statistically indistinguishable from zero, bolstering the argument for a conditional-on-observables interpretation of the instrumental variables results in the rest of the paper. The marginally significant coefficients indicate that trends in unemployment have been somewhat

²⁰In this context, panel B also alleviates statistical power concerns that might arise in the falsification exercise - there is indeed sufficient variation to identify effects even in a 1-period estimation, though the main results of the paper use stacked differences to incorporate more data and increase precision.

²¹The specification for these regressions is a reduced-form version of Equation (2) given below:

$$\Delta \log(Y_{ct}) = \beta_1 \left[\frac{-\Delta \widetilde{D}_{ct}}{Pop_{c,t-5}} \right] + r'_c \gamma_t + \epsilon_{ct}$$

lower and government revenues higher for counties that later experienced negative demand shocks.

5. The Impact of Demand Declines on Employment and Wages

Table 3 presents instrumental variables estimates of β_1 from Equation (1) where the dependent variable is employment or total earnings per capita in coal mining or other sectors within the county according to BLS definitions. The key independent variable ΔD_{ct} per capita is the 5-year change in the total value of coal produced by mines per capita according to MSHA and EIA data, a supply-side measure of production. It is instrumented with $\widetilde{\Delta D}_{ct}$, changes in predicted demand for coal by power plants per capita as defined in Equation (3). For ease of interpretation, total production value per capita is standardized to have mean 0 and standard deviation 1. The coefficients measure the effect of a one standard deviation decline in production value per capita, or \$5,127 per capita (Table 1b) on local outcomes.

The first two columns in panel A of Table 3 show that a one standard deviation decline in coal production per capita is expected to reduce the number of miners per capita by 0.94 percentage points as measured by the MSHA, or 1.15 percentage points by the BLS measure of mining employment. The MSHA data is filtered to include only employees classified as miners at coal mines and so is likely a more conservative measure. Dividing Column (2) of panel A by Column (3) of panel A suggests that there are indeed employment spillover effects on the order of 37% according to BLS employment categories (69% according to MSHA employment).²² Columns (4) and (5) show negative but statistically insignificant point estimates on manufacturing employment and employment in non-tradeable (or locally consumed) sectors. Column (6) clarifies that most of the effects on the non-tradeable goods sectors appear to be dominated by sectors directly in business with coal mining such as utilities and transportation. These findings are most consistent with the idea that coal mining mostly affects the local economy through a direct business-to-business demand channel.

Panel B tells a similar story, where in column (2) the results imply that a decrease of

²²The latter figure is close to the central local employment multiplier proposed in Osman and Kemeny (2022), but lower than initial estimates proposed by Moretti (2010).

\$5,127 per capita of coal production decreased total mining wages by \$1,043 per capita and total wages by \$1,338, an estimated spillover of roughly 28%. Column (4) shows insignificant results across panels A and B on manufacturing employment and wages, while non-tradeables show a decline of \$214, about 65% of which is driven by adjacent industries, in particular those those that directly transport or use coal.²³

a. Benchmarking the Employment Estimates

One way to understand the economic importance of the employment effects is to benchmark them using the overall decline in coal production and employment during the sample period. Put another way, one can ask: How much of the observed decline in mining employment do the estimates explain? An important assumption simplifies this calculation - that the estimates measure absolute changes in employment (not solely relative changes), so that a decrease in the number of miners in a county means a similar decrease in the total number of miners nationally. This assumption is supported by the fact that the number of miners nationally was indeed rapidly declining as demand from the power sector declined (Figure 3).

The simplest way to perform this exercise is to read the coefficient in panel A of Table 3 as the effect of a \$5,127 exogenous decline in coal production value on mining employment. This is also an appropriate interpretation given that the normalization by population is lagged on both sides of the equation. Normalizing the total reduction in thermal coal production value from coal country between 2007 and 2017 (8.03 billion dollars) by \$5,127 and multiplying by the estimated effect size (0.0094) yields an estimate of about 14,722 miners, or approximately 83% of the observed decline (17,678 miners).

Aside from the fact that meaningful proportions of US-produced coal are exported and used for metallurgical purposes, it is intuitively reasonable that demand declines from the power sector may not explain all of recent employment declines at coal mines. Secular

²³Figure 10a displays these coefficients visually. Figure 10b shows complementary evidence - estimates of changes in log wages per job estimated with Equation (2). The patterns are similar across figures, but suggest that the decline in total wage earnings by miners is a result of fewer workers rather than lower wages per worker, while in other sectors it appears to be both fewer workers and lower wages per job.

productivity declines at mines, compliance or other cost pressures, and substitution of labor for capital in the industry more broadly may also be behind these trends. Nonetheless, this back-of-the-envelope calculation suggests that declining power plant demand is likely the dominant factor behind recent declines in coal mining employment. This fact in-and-of-itself suggests that market forces, regulatory policies, and their interactions which advantage some power generation technologies over others, exert a substantial economic impact on upstream fuel suppliers.²⁴

6. Impacts Beyond Employment and Wages

This section looks at outcomes beyond direct labor market impacts in the coal mining sector and other industries. The main reason to extend the analysis in this way is that estimated effects on wages and employment may not adequately characterize even the economic effects of demand shocks on coal communities. For instance, out-of-work coal miners have several margins on which they can respond to a job loss, such as remaining unemployed, leaving the labor force through retirement, or migrating. Moreover, indirect evidence on the price of locally consumed goods such as housing can provide evidence on the magnitude of the demand spillovers caused by lower employment. Finally, the effect of large demand shocks can have ambiguous effects on local amenities and publicly-provided services. Reduced tax revenue may cause declines in public service levels, but outmigration can simultaneously reduce congestion pressures on local public goods.

a. Population and the Labor Force

Table 4 shows the results of IV estimation using the full 2-period differenced model with population (from the Census) and migration (inferred by the IRS from tax receipts) as the outcome variables. Panel A presents estimates from Equation (1) with the dependent variable normalized by lagged population, and panel B from Equation (2), with the dependent variable log-differenced. Coefficients report the estimated effect of an exogenous one

²⁴For a detailed examination of the role of environmental policy as it relates to coal dispatch and retirement, see Linn and McCormack (2019), among others.

standard deviation decline in coal production value per capita in the county. Estimates in columns (1) and (2) suggest that reduced coal production may accelerate population decline, though perhaps in an unexpected way. In column (1) of panel B, total population declined by approximately 0.9% following a decline in production value relative to a counterfactual in which the decline did not occur. In columns (2) and (3) of panel B, IRS county-to-county migration files show a 5.4% decrease in in-migration for each standard deviation decline in coal production value per capita, with a very small (statistically zero) estimated increase in outmigration. The statistical strength of these results are robust to changing the specification as shown in panel A, where a one standard deviation decline in coal production value decreases in-migration by 0.19 persons per capita and population growth by 0.87 percent.

Table 5 presents further evidence on the activities of the working age population in a similar format to Table 4, with estimates from Equation (1) in panel A and from Equation (2) in panel B. Columns (1) and (2) reproduce results on employment with slightly different aggregation, with panel B demonstrating that mining employment decreased 13.5% while nonmining employment decreased only 1.7%. Columns (3) and (4) show small and insignificant effects of declines in coal production value on both the number of individuals not in the labor force (NILF) and the number of unemployed individuals.²⁵ Both estimates show small but noisily estimated increases. Combined with the results of Table 4, these estimates suggest that laid-off coal miners may re-integrate into the labor force quickly. Column (5) provides some evidence on the economic meaning of shocks to residents of coal country, suggesting that the number of poor individuals increased by an average of 3.6% in response to a one standard deviation decline in coal production value per capita in the county.

b. Public Finance and Public Goods

Table 6 examines local public finances, split into general purpose (GP) expenditures and revenues in panel A and school district (SD) expenditures and revenues in panel B. General purpose governments include counties as well as municipalities and township governments

²⁵I estimate the number not in the labor force (NILF) by subtracting the BLS labor force estimates from census estimates of population from ages 15-64 in each county.

with corresponding boundaries according to the Census of Governments. Panel A shows that expenditures and revenues collected by local general purpose governments show statistically insignificant changes, while state support shows a striking 20% decrease. This is likely due to the distribution of severance tax revenues collected at the state level, which are generally allocated by formula. Thus as production declines, state support declines. This may disadvantage local governments during downturns, because state aid is reduced when its marginal value may be relatively high. On the other hand, general purpose governments appear positioned to cushion this blow to total revenue, perhaps by raising taxes, consistent with the positive point estimate in column (3). Total school district revenues and expenditures on the other hand declined by a substantial 4.3% and 4.0%, respectively. This result is supported by the even larger proportional decrease estimated in local revenues from property taxes per standard deviation decline in coal production value per capita, because school district revenue is generally closely tied to property taxes.

Table 7 explores effects on two public goods commonly studied in the economics literature for which data is available, crime rates and education. Columns (1) and (2) show no measureable effect on either violent or property crime rates per 1000 residents. Both point estimates are negative, but have large standard errors. Column (3) suggests that an exogenous decrease in coal production value per capita may substantially increase teacher to student ratios, but the estimate is too noisy to make a definitive statement about the result, while column (4) suggests that the number of teachers declines by 2.2%, consistent with lower school funding. The overall message from Table 7 is that traditional measures of effects on public goods or “local amenities” show relatively weak evidence of positive or negative effects over the time horizon studied, aside from potential impacts on school funding.

c. Rents and Home Values

Table 8 shows estimates of effects on the real estate market. Columns (1) and (2) use data from the ACS on median rents and median home values. Column (1) shows a marginally significant decline in median rents of 2.1%, while column (2) shows a stronger and larger decline in home values of 3.2%. Column (3) reports changes in average dividend and interest

income reported by the IRS at the county level as a proxy for homeowners' non-wage income, which ultimately does not show strong evidence of declining income derived from asset ownership in coal counties. Nonetheless, results in Table 8 support a substantial demand-driven decline in the desirability of real estate in coal country and a drop in overall home prices.²⁶

d. Regional Heterogeneity

Appendix Table C1 examines effect heterogeneity at the regional level with subsample analysis. Each row contains four coefficients with estimates of the effect of a one standard deviation decline in coal production value per capita on changes in the outcome listed on the left hand side of the table. The table is split into three main categories - Employment and Wages, Real Estate and Migration, and Public Revenue. Because splitting the sample dramatically reduces the number of clusters for each regression, standard errors are estimated via clustered pairs bootstrap with 250 repetitions (re-sampling commuting zones).

Several patterns emerge from the labor market results in panel A. First, there is evidence of substantial employment and wage effects in the coal industry in Appalachia, the Illinois basin, and the West (although estimates for Western counties are very imprecise). The largest estimates (and most precisely estimated) are in Appalachia. Results in panels B and C are generally imprecise. Point estimates suggest home value declines in Appalachia, Illinois, and the West, and declining in-migration in Appalachia, the Interior, and the West. In panel C, school district expenditure declines appear in Appalachia, the Interior, and the West, with the largest (in percentage terms) occurring in the West. Overall, the subsample estimates suggest Appalachian, Illinois basin, Interior region, and Western producers may each have felt the effects of declining demand from the power sector to varying degrees, although region-specific conclusions must be treated with care due to the imprecision of the regional estimates.

²⁶I also estimated home value impacts using the Federal Housing Finance Agency's county-level Home Price Index (FHFA HPI). Estimated price declines were somewhat larger than those presented in Table 8, but these data suffer from low coverage of low-population counties suggesting a biased sample in the coal country context.

e. Spatial Spillovers

Appendix Table D1 examines spatial spillovers by estimating the impact of commuting zone level production declines on adjacent non-producing counties. That is, non-producing counties that share a commuting zone with coal producers constitute the estimation sample, and their outcomes are modelled as a function of declining coal production value per capita within the commuting zone. The results share a similar format to Table 3 (omitting the first column), with panel A focusing on employment and panel B on wages.²⁷ The results show three important patterns. First, the overall impact on employment and wages in adjacent counties reported in column (2) is modest and statistically insignificant. Second, results in both panels suggest modest declines in the non-tradeable sector (but not in transport or utilities), consistent with a demand-side spillover effect. Third, part of the total impact comes from a significant positive impact on mining and extractive industry employment. While this estimate reflects sector-wide mining employment (of various commodities) in adjacent counties, the average MSHA-defined coal mining employment among these counties in 2007 was in fact zero. A possible interpretation of this result is that coal miners shifted directly into the oil and gas industry given the frequent co-location of these two resources (especially in Appalachia and the West).

Two further points are worth emphasizing. First, this exercise underscores the importance of controlling for oil and gas production in the main specification. Second, the possibility that coal workers shifted locally to another fossil fuel extraction industry suggests that previous estimates in the paper may be conservative relative to a situation where this mechanism does not operate.²⁸

²⁷Effect sizes in the table are not precisely comparable to estimates in Table 3 because the spillover estimates are conditioned on state- rather than region-year fixed effects (spillover counties are not easily assigned to coal producing regions).

²⁸Put another way, even if they are accurate for the historical period studied here, this paper's marginal effects may be a conservative estimate of the negative economic impacts of future demand shocks on coal country if future shocks are not accompanied by a similar expansion in other job opportunities for miners.

7. Changes in Economic Welfare

This paper estimates the effect of demand-driven declines in the produced value of coal on a variety of outcomes in coal communities. The primary effects are declines in employment and wages, a decline in home values and local public revenues and expenditures, and declining in-migration. One way to understand these impacts is to consider their incidence across groups. On average, in-county wage earners and homeowners with children bear the incidence of the three largest estimated effects: a decline in wages, a decline in home equity, and a decline in education spending per capita.²⁹ Workers who commute out of the county, non-resident landlords, and renters may avoid some of the incidence of demand shocks. While informative, this exercise does not quite reach a statement about the economic welfare of coal country's residents without further structure from a model. This section offers several calculations to tie together the key estimates in the paper.

a. Capitalization

Building on the foundation of a Roback (1982) style spatial equilibrium model with inelastic housing supply and elastic migration, aggregate changes in home values are often applied as a measure of economic welfare, because they reflect changes in individuals' willingness-to-pay to reside (more specifically, own land) in a location. A ballpark estimate of the aggregate effect on home values can be constructed using the average median home value (across counties, Appendix Table A1) as a conservative estimate of the mean home value in coal country, combined with the national \$8.03 billion dollar decline in coal production, similar to the benchmarking exercise in Section 5. Normalizing the production decline by population, scaling to the effect of one standard deviation, and multiplying by the total number of households (identified by IRS tax returns also in Table A1) yields an estimate of about \$2.66 billion in home value declines, or about 33% of the size of the initial shock.

²⁹Using the results presented in Figure 10, a one standard deviation decrease in coal production per capita is predicted to reduce average wages by 4.2%, while column (2) of Table 8 indicated that a one standard deviation production decline per capita is capitalized into a 3.2% decline in median home values. Column (1) of Table 6 shows a 4.0% decline in school district expenditures per standard deviation decline in coal production value per capita.

This approach has some shortcomings - for instance, some of the decline in housing costs represents a transfer from property owners to renters.

b. Welfare Estimates based on Spatial Equilibrium

This section presents alternative welfare estimates from a more flexible model of mobile households in the rental market. Homogeneous agents in the model value real wages and local public goods or amenities. Migration responses by the agents in response to a location-specific shock reveal valuations of marginal changes in the wage and amenity characteristics of the location. Guided by expressions from the model, this section combines this paper’s causal estimates with parameters from the literature to estimate the economic welfare impacts of demand shocks that combine wage impacts with money-metric valuations of other impacts. While this approach does not directly estimate total welfare, it offers a defensible estimate by monetizing key causal effects, which will be made more precise throughout the section. Because the model and approach are similar to Bartik, Greenstone, Currie, Knittel (2019), which analyzes the willingness-to-pay of local communities for hydraulic fracturing, the model structure is introduced here and a detailed derivation is provided in Appendix E.

The model economy consists of many homogeneous individuals $1, \dots, i, \dots, N$ who value wages, rental rates, and public goods and receive an individual-region specific preference shock. Individuals choose a region c to maximize their indirect utility U_{ic} given by Equation (E1), where w_c , r_c , and G_c are the wage rate, cost of housing, and a public good in location c . The parameter σ_c scales the preference shock distribution according to other features that might vary by location.

$$\ln U_{ic} = \alpha \ln w_c + \beta \ln r_c + \gamma \ln G_c + \sigma_c \epsilon_{ic} \quad (5)$$

Focusing on the case with just two regions c and k implies regional populations N_c and $N_k = N - N_c$. Wages, rents, and public goods are modelled each as a function of a latent index θ that represents global demand for a good produced only in region c . Combined with distributional assumptions on the preference shocks, this basic structure leads to a simple

relationship between individual welfare and migration equilibria in the model.

Assuming region c is small relative to region k , shocking θ , aggregating over individuals, and applying empirical analogues leads to Equation (E7), which is the estimating equation for the wage-equivalent welfare impacts. Terms on the right-hand side of Equation (E7) are all observable or empirically estimated. They include the population of coal country and this paper’s estimates of changes in wages, housing costs, and public education expenditures with respect to demand shocks. The parameters $\hat{\alpha}$, $\hat{\beta}$, and $\hat{\gamma}$ by are migration elasticities in terms of wages, housing costs, and local amenities (including public education expenditures) from Diamond (2016).³⁰ The baseline average wage $\overline{w_{c0}}$ monetizes the utility impacts following Bartik et al. (2019).

$$\frac{1}{\hat{\alpha}} \frac{dU}{d\theta} = N_c * \left(\frac{d \ln w_c}{d\theta} + \frac{\hat{\beta}}{\hat{\alpha}} \frac{d \ln r_c}{d\theta} + \frac{\hat{\gamma}}{\hat{\alpha}} \frac{d \ln G_c}{d\theta} \right) * \overline{w_{c0}} \quad (6)$$

Intuitively, normalizing $\hat{\beta}$ and $\hat{\gamma}$ by $\hat{\alpha}$ implies that housing cost changes are valued at approximately 60% of wage changes and amenity changes are valued at roughly 20%. In practice this approach allows for a coherent, combined valuation of wage and nonwage (housing and public expenditure on education) as revealed by the empirical estimates presented in the paper.

This section’s estimates provide a model-based welfare interpretation of the causal effects of wage, housing, and public education expenditure changes. However, a key limitation arises from the partial equilibrium nature of the approach. Specifically, G_c should generally be a vector of local amenities, with education being just one element. Monetizing only the important causal estimates of the paper implicitly assumes no significant changes in other unobserved amenities, or that such changes offset each other. In other words, these estimates are conditional on the assumption that only changes in wages, housing, and public education expenditures affect welfare, with other amenities held constant or netting to zero. The reasons for taking this approach are discussed further below, after presenting the results.

³⁰These parameters are taken from panel A of Table 5 of Diamond (2016), and reflect migration elasticities estimated under the least restrictive possible model. Non college-worker parameter estimates on wages, housing costs, and Diamond’s amenities index are (3.261, -2.944, 0.771) and college-worker parameters are (4.976, -2.159, 0.638).

Table 9 presents the results of applying the parameters described above to Equation (E7) to estimate a household average impact in coal country and to estimate aggregate welfare effects on all of coal country. Each panel of Table 9 contains two sets of estimates that explore the sensitivity of the estimates to measuring housing costs either with rental prices or home values. The two panels use different sets of migration elasticities, one for college-educated workers and one for non-college workers, which are treated as fixed population parameters. Standard errors account for joint uncertainty in the empirical estimates of the paper and are obtained via 1000 pairs bootstrap repetitions (re-sampling commuting zones) of the entire welfare estimation procedure, first re-estimating coefficients and then applying Equation (E7).

Column (1) monetizes the welfare effect of the shock estimated throughout the paper - a one standard deviation or \$5,127 decline in coal production value per capita. Central estimated effects on average wages (4.2%), home values (3.2%), and public education expenditures (4.0%) are drawn from Figure 10, Table 8, and Table 6, and re-scaled to the household level where appropriate. Column (2) aggregates the per capita effects using the total decline in coal production per capita and total population of coal country, similar to the employment effect benchmarks in Section 5. Column (3) follows Bartik et al (2019). and scales the estimate in column (2) assuming the observed changes in wages, housing costs, and public good expenditures are constant and remain the same in perpetuity using a 5% discount rate.

Using the parameters for non-college workers in panel A, the estimates of changes in economic welfare in coal country in column (1) range from a decline of \$801 per capita to \$1153 per capita. Aggregating across all residents of coal country leads to a range of \$0.57 billion to \$0.83 billion in impacts. Assuming permanent declines in wages, rental rates, and school expenditures yields a present value impact of \$11.5 to \$16.6 billion. Larger declines in housing costs offsets some of the estimated welfare losses (a transfer to renters), while for smaller decreases in housing costs the combined effect of wage declines and reductions in public expenditures make mobile residents worse off. The first row of panel A and third row of panel B roughly coincide across parameters, suggesting a total aggregate decline in

economic welfare of \$0.85 billion.

How can we benchmark the magnitudes in Table 9? By the valuation techniques described in the Background and Data section, I estimate that US coal country produced \$28.3 billion dollars of thermal coal in 2007 before reaching a maximum value in the sample of more than \$34 billion dollars of thermal coal (in 2008), which dropped to \$20.3 billion in 2017. Compared to the decline from 2007 to 2017, the effects in column (2) range from about 7.1% of the value of produced coal at the smallest to 12.1% at the largest. Another useful set of benchmarks are the local impacts of fracking estimated in Bartik, Currie, Greenstone, and Knittel (2019). My monetized impacts of \$11.5 billion to \$19.4 billion in coal country are about one fifth of their range of welfare impacts across major fracking counties of \$57 billion to \$85 billion. So by these benchmarks at least, the estimates seem reasonable. These multi-year estimates (with assumed permanence) are small relative to annual local air quality benefits of \$17 billion estimated in Johnsen, Lariviere, and Wolff (2019).

The model is stylized and requires two related caveats. First, the structural parameters taken from Diamond (2016) reflect average migration and utility valuations for the US workforce from earlier decades, based on national data. These parameters may not fully capture the specific context of coal country, where migration responses or utility valuations could differ qualitatively. Second, unobserved changes in amenities, which are not modeled here, could push the welfare estimates in either direction. The results should therefore be interpreted as representing the causal effects of the variables analyzed in this paper. For example, it is possible that negative coal demand shocks could lead to improvements in congestion or environmental amenities, but evidence for such spillovers remains weak. Metcalf and Wang (2018) found no significant association between mining layoffs and opioid use, and Tables 6 and 7 show weak or no evidence of broader effects on public spending or crime rates.³¹

The value of the calculations in Table 9 is that they offer a coherent summary of changes in economic welfare derived from exogenous changes to the coal sector, grounded by empirical

³¹In theory, the full vector of amenity changes G_c could be inferred by using population shifts from Table 4. While the population estimates suggest a negative overall welfare change, the response is relatively small compared to the impacts on wages and home values. This raises the distinct possibility that limited mobility or strong local ties, as modeled in Zabek (2024), could be driving the muted response, rather than unobserved improvements in amenity values.

estimates. The migration elasticities from the literature are important adjustments, and suggest that wages are the main factor driving welfare changes. This seems sensible in the context of coal country - housing prices and rents are already low relative to the rest of the US (Appendix Table A1), and it is likely that congestion effects in local education (and other local public goods) are low, while non-college employment opportunities are relatively highly valued.

8. Conclusion

Many economic changes and policies that are welfare-enhancing in aggregate or over the long run carry acute costs to certain groups or regions. This paper contributes to the study of local economic shocks by examining US coal communities in a down-decade of coal demand driven by competition from alternative generation fuels and environmental policy. I find evidence of substantial labor market reallocation, including job loss and foregone wages that spill over across sectors, migration responses, and knock-on effects to home values and public revenues.

Instrumenting for declines in production value at mines with predicted changes in coal plant dispatch from a stylized regional model of the electricity sector, I find that a one standard deviation decline in coal production value per capita at the county level decreases non-mining employment by 1.7% and average in-county wages per job by 4.2%. Still, my estimates suggest that mining employment and wages make up the bulk of the direct labor market effects, with total effects 30-60% larger. Most spillovers appear to occur in directly connected industries such as transportation. For the same shock, I estimate that state support for general purpose government declines, and public school revenues and expenditures fall (4.3 and 4.0% respectively) driven by a significant decline in home values (3.2% at the median). Correspondingly, in-migration and medium-run trends in population are significantly impacted by declining coal production. Interpreted as capitalization effects, the home value estimates imply a total decline of around \$2.6 billion in individuals' valuation of life in coal country. Alternatively, my estimates of wage, rent, and public school expenditure

changes applied in a theoretical model of spatial equilibrium imply that mobile households showed declining valuations for living in coal country of about \$0.85 billion in the aggregate.

These results suggest the effect of declines in coal demand from the power sector on producing areas has thus far been meaningfully negative, but small relative to the local impacts of fracking. Negative multipliers on mining jobs and wages appear to be moderate and consistent with literature exploring previous cycles in coal demand. I do not find strong evidence of high or persistent unemployment, increasing crime, rapid outmigration, or declining expenditures by general purpose governments, despite reduced support from the state. On the other hand, it seems beyond doubt that coal demand shocks reduce the overall welfare of coal country's residents as evidenced by reduced in-migration, declining home values and school district expenditures on public education, and increasing poverty rates.

9. References

- Allcott, Hunt, and Daniel Keniston.** 2018. “Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America.” *The Review of Economic Studies*, 85(2): 695–731.
- Autor, David H, David Dorn, and Gordon H Hanson.** 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review*, 103(6): 2121–2168.
- Bartik, Alexander W, Janet Currie, Michael Greenstone, and Christopher R Knittel.** 2019. “The Local Economic and Welfare Consequences of Hydraulic Fracturing.” *American Economic Journal: Applied Economics*, 11(4): 105–55.
- Bartik, Timothy J.** 1991. *Who benefits from state and local economic development policies?* WE Upjohn Institute for Employment Research.
- Betz, Michael R, Mark D Partridge, Michael Farren, and Linda Lobao.** 2015. “Coal Mining, Economic Development, and the Natural Resources Curse.” *Energy Economics*, 50: 105–116.
- Black, Dan A, Terra G McKinnish, and Seth G Sanders.** 2005*a*. “Tight Labor Markets and the Demand for Education: Evidence from the Coal Boom and Bust.” *ILR Review*, 59(1): 3–16.
- Black, Dan, Terra McKinnish, and Seth Sanders.** 2005*b*. “The Economic Impact of the Coal Boom and Bust.” *The Economic Journal*, 115(503): 449–476.
- Blonz, Josh, Brigitte Roth Tran, and Erin E Troland.** 2023. “The Canary in the Coal Decline: Appalachian Household Finance and the Transition from Fossil Fuels.” *NBER Working Paper Series*, Working Paper No. 31072.
- Borenstein, Severin, James B Bushnell, and Frank A Wolak.** 2002. “Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market.” *American Economic Review*, 92(5): 1376–1405.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel.** 2022. “Quasi-Experimental Shift-Share Research Designs.” *The Review of Economic Studies*, 89(1): 181–213.
- Busso, Matias, Jesse Gregory, and Patrick Kline.** 2013. “Assessing the Incidence and Efficiency of a Prominent Place Based Policy.” *American Economic Review*, 103(2): 897–947.
- Cicala, Steve.** 2015. “When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation.” *American Economic Review*, 105(1): 411–444.
- Cicala, Steve.** 2022. “Imperfect Markets Versus Imperfect Regulation in US Electricity Generation.” *American Economic Review*, 112(2): 409–441.
- Colmer, Jonathan, Eleanor Krause, Eva Lyubich, and John Voorheis.** 2024. “Transitional Costs and the Decline in Coal: Worker-Level Evidence.” Centre for Economic Performance, LSE.

- Davis, Rebecca J, J Scott Holladay, and Charles Sims.** 2022. “Coal-Fired Power Plant Retirements in the United States.” *Environmental and Energy Policy and the Economy*, 3(1): 4–36.
- Deaton, B James, and Ekaterina Niman.** 2012. “An Empirical Examination of the Relationship Between Mining Employment and Poverty in the Appalachian Region.” *Applied Economics*, 44(3): 303–312.
- Diamond, Rebecca.** 2016. “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000.” *American Economic Review*, 106(3): 479–524.
- Feler, Leo, and Mine Z Senses.** 2017. “Trade Shocks and the Provision of Local Public Goods.” *American Economic Journal: Economic Policy*, 9(4): 101–143.
- Fowler, Christopher S, Danielle C Rhubart, and Leif Jensen.** 2016. “Reassessing and Revising Commuting Zones for 2010: History, Assessment, and Updates for US ‘Labor-Sheds’ 1990–2010.” *Population Research and Policy Review*, 35: 263–286.
- Furnaro, Andrea, Philipp Herpich, Hanna Brauers, Pao-Yu Oei, Claudia Kemfert, and Wesley Look.** 2021. “German Just Transition: A Review of Public Policies to Assist German Coal Communities in Transition.” *Resources for the Future Report*, 21–13.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift.** 2020. “Bartik Instruments: What, When, Why, and How.” *American Economic Review*, 110(8): 2586–2624.
- Hanson, Gordon.** 2023. “Local Labor Market Impacts of the Energy Transition: Prospects and Policies.” *HKS Faculty Research Working Paper Series*.
- Johnsen, Reid, Jacob LaRiviere, and Hendrik Wolff.** 2019. “Fracking, Coal, and Air Quality.” *Journal of the Association of Environmental and Resource Economists*, 6(5): 1001–1037.
- Kline, Patrick, and Enrico Moretti.** 2014. “People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs.” *Annu. Rev. Econ.*, 6(1): 629–662.
- Linn, Joshua, and Kristen McCormack.** 2019. “The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions.” *The RAND Journal of Economics*, 50(4): 733–767.
- Look, Wesley, Joey James, Evan Fedorko, Sophie Pesek, Dan Mazzone, Aurora Barone, and Rebecca Shelton.** 2022. “POWER for Transition: Investment in Coal Communities Through the Partnerships for Opportunity and Workforce and Economic Revitalization (POWER) Initiative.” *Resources for the Future*.
- Manson, Steven M.** 2020. “IPUMS National Historical Geographic Information System: Version 17.0.”
- Matheis, Mike.** 2016. “Local Economic Impacts of Coal Mining in the United States 1870 to 1970.” *The Journal of Economic History*, 76(4): 1152–1181.
- Metcalf, Gilbert E, and Qitong Wang.** 2019. “Abandoned by Coal, Swallowed by Opioids?” *NBER Working Paper Series*, Working Paper No. 26551.

- Moretti, Enrico.** 2010. “Local Multipliers.” *American Economic Review*, 100(2): 373–377.
- Morris, Adele.** 2016. “Build a Better Future for Coal Workers and their Communities.” *Climate and Energy Economics Discussion Paper, Brookings Institution, Washington DC.*
- Morris, Adele C, Noah Kaufman, and Siddhi Doshi.** 2021. “Revenue at Risk in Coal-Reliant Counties.” *Environmental and Energy Policy and the Economy*, 2(1): 83–116.
- Osman, Taner, and Tom Kemeny.** 2022. “Local Job Multipliers Revisited.” *Journal of Regional Science*, 62(1): 150–170.
- Roback, Jennifer.** 1982. “Wages, Rents, and the Quality of Life.” *Journal of Political Economy*, 90(6): 1257–1278.
- Roemer, Kelli, Daniel Raimi, Rebecca Glaser, et al.** 2021. “Coal Communities in Transition: A Case Study of Colstrip, Montana.” *Washington, DC: Resources for the Future Report.*
- Stoker, Thomas M, Ernst R Berndt, A Denny Ellerman, and Susanne M Schennach.** 2005. “Panel Data Analysis of US Coal Productivity.” *Journal of Econometrics*, 127(2): 131–164.
- Tolbert, Charles M, and Molly Sizer.** 1996. “US Commuting Zones and Labor Market Areas: A 1990 Update.” United States Department of Agriculture, Economic Research Service.
- Watson, Brett, Ian Lange, and Joshua Linn.** 2023. “Coal Demand, Market Forces, and US Coal Mine Closures.” *Economic Inquiry*, 61(1): 35–57.
- Weber, Jeremy G.** 2020. “How Should We Think About Environmental Policy and Jobs? An Analogy with Trade Policy and an Illustration from US Coal Mining.” *Review of Environmental Economics and Policy.*
- Zabek, Mike.** 2024. “Local Ties in Spatial Equilibrium.” *American Economic Journal: Macroeconomics*, 16(2): 287–317.

10. Figures

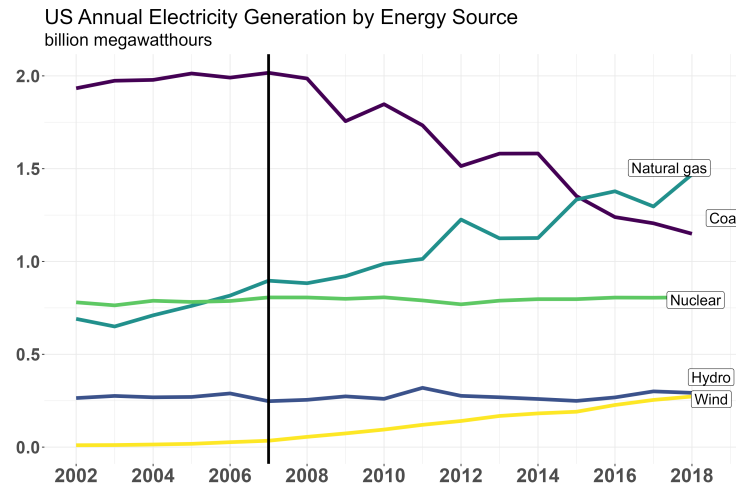


Figure 1: **Aggregate Fossil Fuel Generation.** US fossil fuel generation by source since 2002, (author reproduction) from EIA published aggregates. Coal generation was the major source of electricity in the US, but began to decline rapidly in 2008, and was eventually overtaken by natural gas generation.

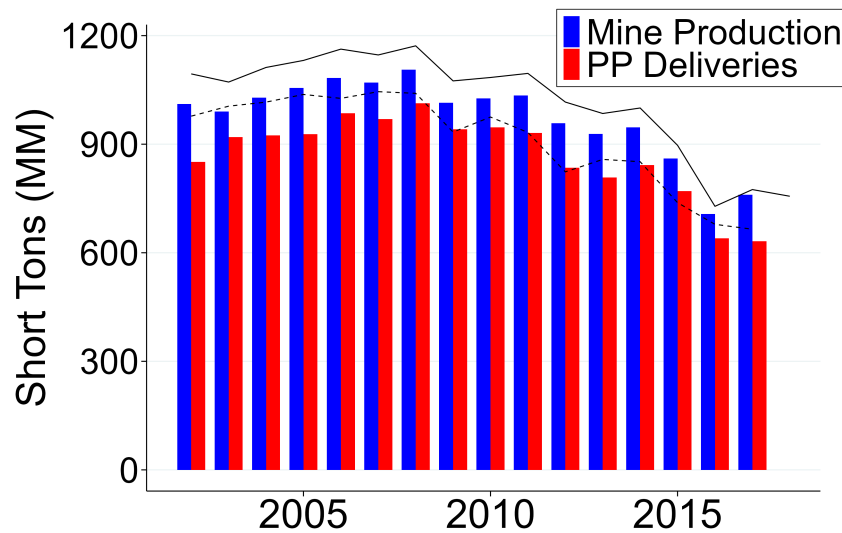


Figure 2: **Coal Production and Power Plant Deliveries.** Blue bars are derived from MSHA and EIA form 7A. Red bars are derived from EIA 423 and 923. The solid and dotted lines are EIA published aggregates from EIA's annual coal reports, which generally exceed but appear to match up well with the aggregated microdata.




Figure 3: **US Coal Mining Employment.** Coal mining employment (total miners) reported in MSHA/EIA-7A represented by vertical bars and FRED data series (CEU1021210001) by the dotted line. The definition of mining employment in the MSHA data is more restrictive so the reported totals are lower, but the trends match closely.

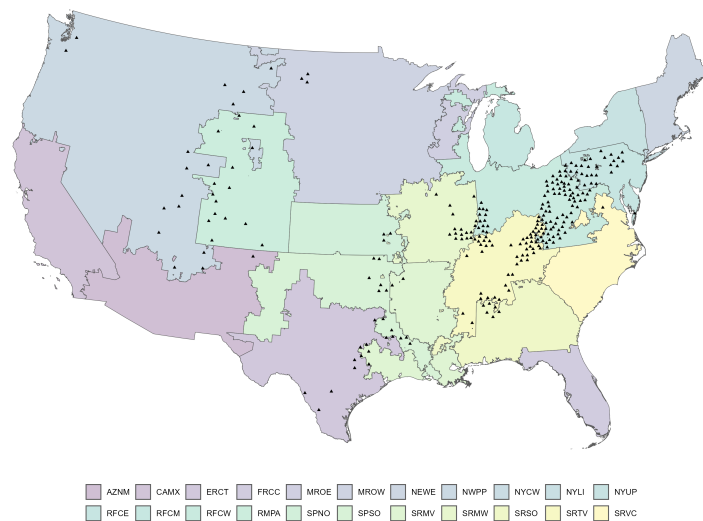
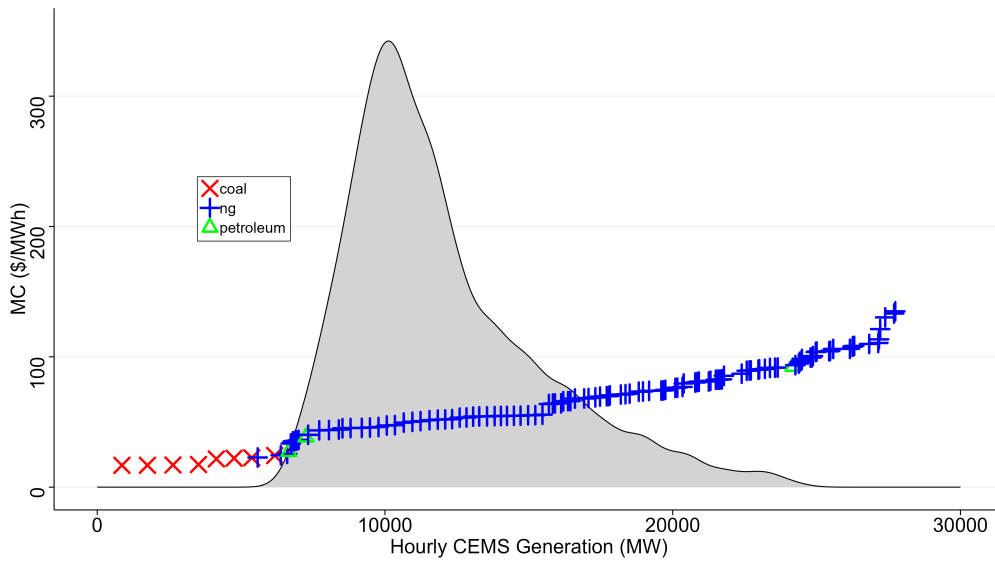



Figure 4: **Subregions and Coal County Centroids.** This figure shows approximate boundaries of EPA eGRID Subregions as colored polygons and the centroids of coal-supplying counties as black triangles. eGRID Subregions are an intermediate division of electricity zones between NERC (North American Electric Reliability Corporation) regions and BAAs (balancing area authorities).



(a) SRMV region, 2007



(b) SRMV region, 2017

Figure 5: **Example Supply Stacks.** These figures show two examples of estimated electricity load region supply curves for the SRMV eGRID subregion. The grey shape in each figure shows the distribution of total fossil fuel load from CEMS plants at the hourly level. Marginal costs for CEMS plants are estimated as the unit-level average heat rate times fuel prices drawn from the EIA SEDS database. Panel A has one very high marginal cost plant removed from the figure to maintain readability.




Figure 6: **Dispatch Predictions.** This figure shows predicted generation from a least-cost dispatch model of the SRMV load region against actual observed gross generation in 2007 and 2017. Predicted values are derived by constructing the supply curves depicted in the previous figures and “intersecting” them with observed load from CEMS plants in each hour (modelling hourly demand as perfectly inelastic). All units with costs below the intersection of supply and demand are dispatched for the full hour. Summing over all hours in the year yields annual predicted generation for each unit as a function of state-level fuel prices, the region’s fuel mix, and total fossil load.


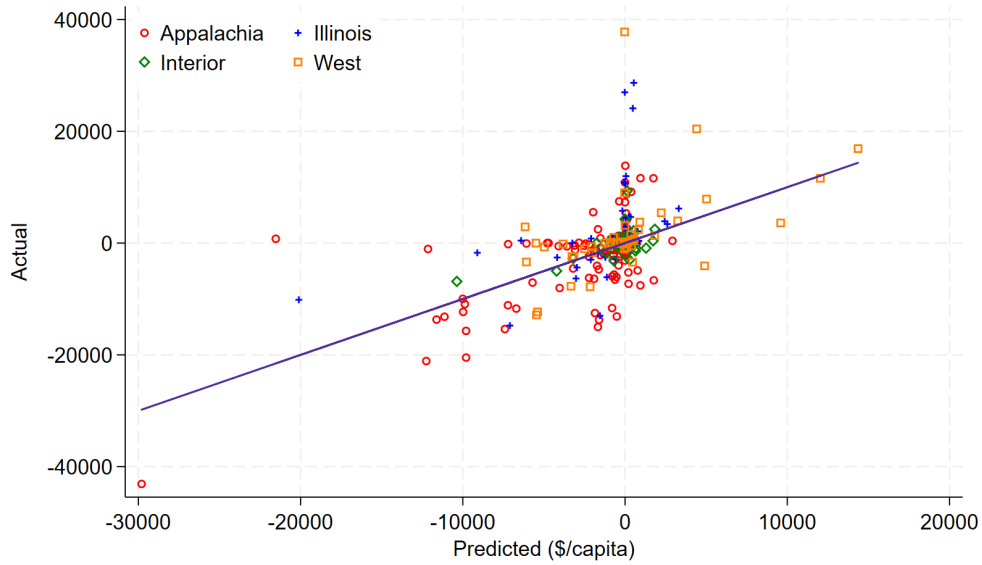
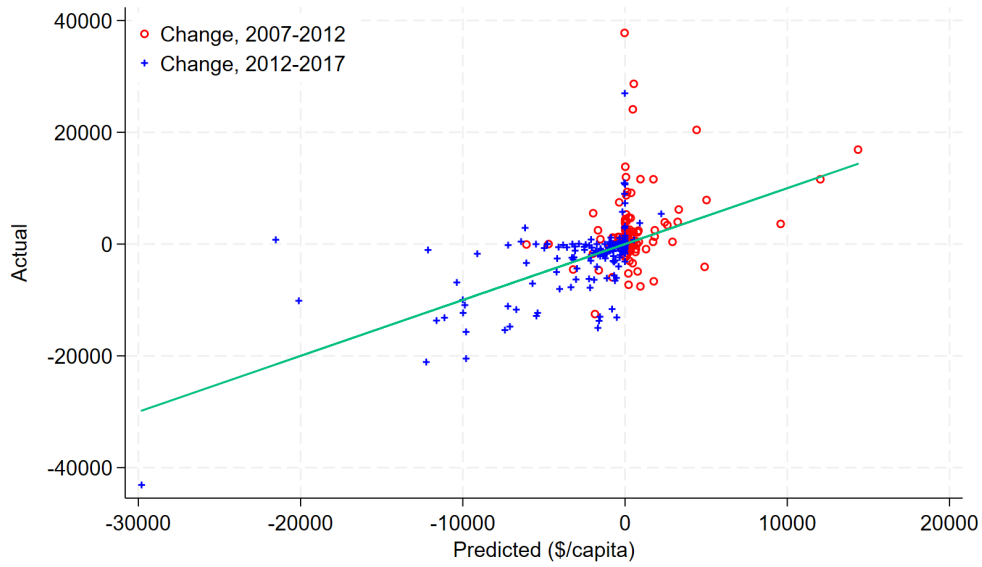


Figure 7: **Market Shares.** Figure shows a strip plot of the share of Btus of coal delivered by each coal producing county to each eGRID Subregion in the continental US over the period from 2002-2007. Each row on the y-axis represents an eGRID Subregion and the x-axis measures the share of fuel deliveries. Campbell County, WY is excluded from all calculations.

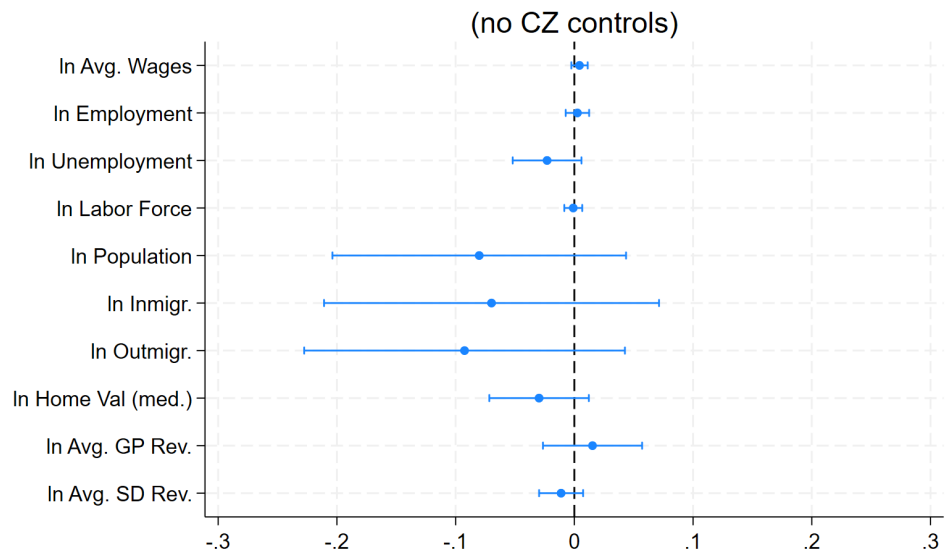


(a) By region

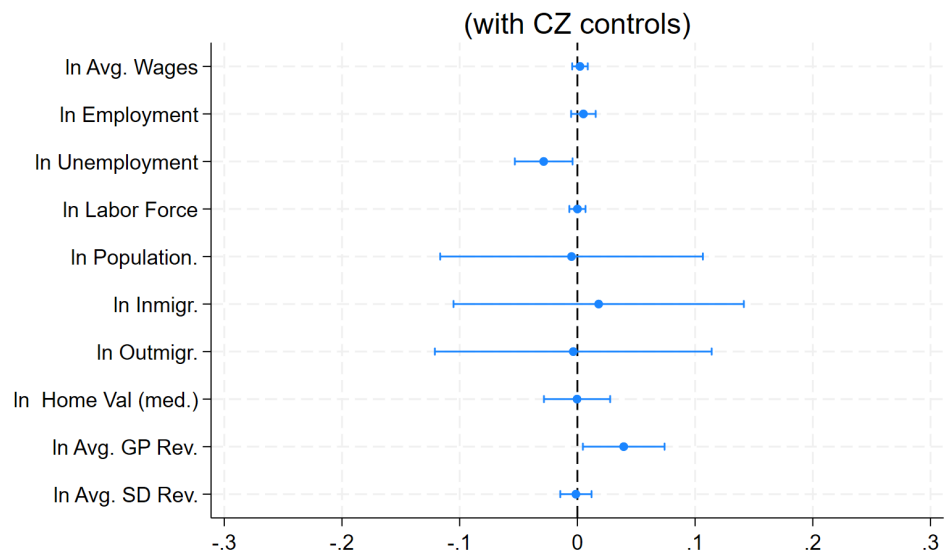


(b) By period

Figure 8: **Predicted Coal Demand vs. Actual Coal Production.** These figures show the first stage relationship between the predicted coal demand instrument and observed coal production value in dollars. See Section 4 for details.

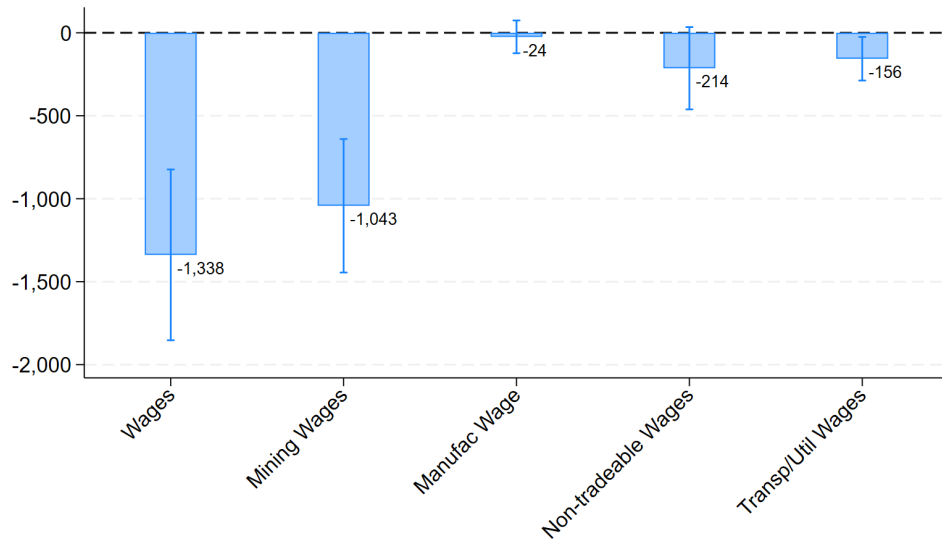


(a) No adjustment for commuting zone characteristics




(b) Adjusted for commuting zone characteristics

Figure 9: **Reduced-Form placebo regressions, 2002-2007.** See Section 4 for details. GP Rev. = General Purpose Revenue (county, municipality, township). SD = School District Revenue.



(a) Total wages per capita



(b) Log of wages per job

Figure 10: **Wage Effects**. See Section 5 for details.

11. Tables

Table 1: Summary Statistics

	(1)	(2)	(3)	(4)	(5)
	N	mean	sd	min	max
Production (MM short tons)	254	2.625	4.957	0	34.47
Production (MM Btus)	254	56.94	107.3	0	804.3
Production value (MM USD)	254	111.5	208.9	0	1,721
Coal mining emp. share	254	0.0305	0.0614	0	0.430
Mining emp. share	254	0.0553	0.0708	0	0.452
Avg. Mining Wage (000s)	240	46.28	15.42	17.53	92.92
Wage per job (000s)	254	38.03	6.393	26.83	66.44
AGI per HH (000s)	254	45.34	11.39	24.56	92.23
Population (000s)	254	62.73	150.6	1.865	1,848
Labor Force (000s)	254	30.77	82.82	1.131	1,067
Total Employment (000s)	254	27.18	90.15	0.660	1,173
Unemployed (000s)	254	1,416	3,011	48	34,024
GP Exp. per capita (000s)	254	1.718	1.117	0.145	7.989
SD Exp. per capita (000s)	254	1.629	0.730	0	4.297

(a) Summary statistics for coal counties in 2007. Campbell County, Wyoming excluded.

	p5	p25	p50	p75	p95	iqr	sd
Appalachia	-11,122	-600	-10	16	2,015	616	4,669
Illinois	-6,102	-193	0	1,408	10,918	1,601	6,315
Interior	-3,022	-625	-4	48	2,211	673	1,958
West	-7,719	-496	0	744	11,593	1,239	6,713
Total	-7,559	-527	0	122	5,307	648	5,127

(b) Distribution of $\Delta D_{ct}/Pop_{c,t-5}$, change in production value per capita by mines in county c . iqr = inter-quartile range. sd = standard deviation.

Table 2: Reduced-Form Effects of Demand Declines by Period
 Dependent variable: % point change in emp/capita over 5-year difference

	(1)	(2)	(3)	(4)	(5)
	Coal Miners	Mining emp. (BLS)	Tot. Emp.	Manufacturing	Non-Tradeables
<u>Panel A: 2007-2017</u>					
Demand Decline (Δ s.d.)	-0.53*** (0.12)	-0.63*** (0.12)	-0.93*** (0.16)	-0.01 (0.03)	-0.19* (0.09)
Obs	508	508	508	508	508
<u>Panel B: 2012-2017</u>					
Demand Decline (Δ s.d.)	-0.60*** (0.13)	-0.67*** (0.11)	-0.97*** (0.18)	0.01 (0.03)	-0.22* (0.10)
Obs	254	254	254	254	254
<u>Panel C: 2007-2012</u>					
Demand Decline (Δ s.d.)	-0.14 (0.13)	-0.46 (0.38)	-0.69* (0.29)	-0.11 (0.06)	-0.07 (0.29)
Obs	254	254	254	254	254
<u>Panel D: 2002-2007 (placebo)</u>					
Demand Decline (Δ s.d.)	0.02 (0.08)	0.05 (0.10)	0.17 (0.17)	0.02 (0.07)	0.08 (0.14)
Obs	254	254	254	254	254
Region-time FE	yes	yes	yes	yes	yes
CZ controls	no	no	no	no	no

Notes: This table reports reduced form regressions of 5-year changes in employment per capita on the instrumental variable $\widehat{\Delta D_{ct}}/Pop_{c,t-5}$. Outcome variables come from the MSHA (column 1) and the BLS (all other columns). Each coefficient comes from a separate regression. Panel A shows coefficients estimated using Equation (4) with two periods of data in the stacked difference model. Panels B and C estimate one-period regressions of contemporaneous changes in employment against changes in the instrument. Panel D regresses changes from 2002-2007 against the average of future values of the instrument. Non-tradeables are defined as all employment except mining and agriculture, manufacturing, and government. The instrument is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 3: Results on Employment and Wages
 Dependent Variables: (A) % point change in emp/capita, (B) wages/capita

	(1)	(2)	(3)	(4)	(5)	(6)
<u>Panel A. Employment</u>						
	Coal miners (MSHA)	Mining	Total	Manufacturing	Non-Tradeable	Util./Transp.
Coal Value (Δ s.d.)	-0.94*** (0.12)	-1.15*** (0.22)	-1.59*** (0.30)	-0.07 (0.07)	-0.22 (0.22)	-0.17*** (0.06)
F excl.	21.28	21.28	21.28	21.28	21.28	21.28
Obs	508	508	508	508	508	508
<u>Panel B. Wages</u>						
		Mining	Total	Manufacturing	Non-Tradeable	Util./Transp.
Coal Value (Δ s.d.)		-1,043*** (205)	-1,338*** (263)	-24 (50)	-214* (127)	-156** (67)
F excl.		21.28	21.28	21.28	21.28	21.28
Obs		508	508	508	508	508
Region-time FE	yes	yes	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in employment and wages per capita on the treatment variable $\Delta D_{ct}/Pop_{c,t-5}$, coal production value changes per capita. Outcome variables come from the MSHA (column 1, wages not reported) and the BLS (all other columns). Each coefficient comes from a separate regression. Panels A and B both show coefficients estimated using Equation (1) with two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the "Empirical Strategy" section. Non-tradeables is defined as all employment except mining and agriculture, manufacturing, and government. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 4: Results on Population and Migration

	(1)	(2)	(3)
<u>Panel A: ($\Delta Y_{ct}/Pop_{c,t-5}$)</u>			
	Population	Inmigr.	Outmigr.
Coal Value (Δ s.d.)	-0.87*** (0.22)	-0.19** (0.09)	0.01 (0.07)
F excl.	21.28	21.28	21.28
Obs	508	508	508
<u>Panel B: ($\Delta \ln Y_{ct}$)</u>			
	Population	Inmigr.	Outmigr.
Coal Value (Δ s.d.)	-0.009*** (0.002)	-0.054*** (0.019)	0.003 (0.013)
F excl.	21.28	21.28	21.28
Obs	508	508	508
Region-time FE	yes	yes	yes
CZ controls	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in population and migration on the treatment variable $\Delta D_{ct}/Pop_{c,t-5}$, coal production value changes per capita. Outcome variables come from the US Census (column 1) and the IRS (columns 2 and 3). Each coefficient comes from a separate regression. Panel A shows coefficients estimated using Equation (1) while panel B shows coefficients estimated using Equation (2) with log differenced outcomes. Both panels use two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the "Empirical Strategy" section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 5: Results on Labor Force Components

	(1)	(2)	(3)	(4)	(5)
Panel A: ($\Delta Y_{ct}/Pop_{c,t-5}$)	Mining	NonMining	NILF	Unemployed	Poor
Coal Value (Δ s.d.)	-1.15*** (0.22)	-0.44* (0.25)	0.33 (0.38)	0.04 (0.07)	0.66** (0.32)
F excl.	21.28	21.28	21.28	21.28	21.28
Obs	508	508	508	508	508
Panel B: ($\Delta \ln Y_{ct}$)	Mining	NonMining	NILF	Unemployed	Poor
Coal Value (Δ s.d.)	-0.135*** (0.040)	-0.017** (0.009)	0.004 (0.023)	0.018 (0.025)	0.036** (0.016)
F excl.	17.62	21.28	21.00	21.28	21.28
Obs	455	508	506	508	508
Region-time FE	yes	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in employment, labor force participation, and poverty on the treatment variable $\Delta D_{ct}/Pop_{c,t-5}$, coal production value changes per capita. Outcome variables come from the BLS (columns 1-4) and the Census' SAIPE (column 5). Each coefficient comes from a separate regression. Panel A shows coefficients estimated using Equation (1) while panel B shows coefficients estimated using Equation (2) with log differenced outcomes. Both panels use two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the "Empirical Strategy" section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 6: Results on Public Finances

	(1)	(2)	(3)	(4)
Panel A: General Purpose	GP Expend.	GP Rev.	GP Rev. (local)	GP Rev. (state)
Coal Value (Δ s.d.)	0.015 (0.036)	-0.007 (0.035)	0.027 (0.044)	-0.200*** (0.045)
F excl.	21.30	21.30	21.30	21.30
Obs	507	507	507	507
Panel B: School District	SD Expend.	SD Rev.	SD Rev. (local)	SD Rev. (state)
Coal Value (Δ s.d.)	-0.040** (0.017)	-0.043*** (0.013)	-0.075*** (0.026)	-0.032 (0.021)
F excl.	20.38	20.38	20.38	20.38
Obs	469	469	469	469
Region-time FE	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in per-capita local public finance on the treatment variable ΔD_{ct} , coal production value changes per capita. Outcome variables come from the Census of Governments. Each coefficient comes from a separate regression. Panels A and B both show coefficients estimated using Equation (2) with log differenced outcomes. Both panels use two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the "Empirical Strategy" section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 7: Results on Local Public Goods

	(1)	(2)	(3)	(4)
<u>Panel A:</u>	Crime Rate (viol.)	Crime Rate (prop.)	Students/Teachers	Teachers
Coal Value (Δ s.d.)	-0.011 (0.030)	-0.035 (0.043)	0.039 (0.026)	-0.022*** (0.005)
F excl.	21.28	21.28	17.56	16.29
Obs	508	508	468	462
Region-time FE	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in measures of public goods on the treatment variable ΔD_{ct} , coal production value changes per capita. Outcome variables come from the FBI's Uniform Crime Reporting Data (columns 1 and 2) and the National Center for Education Statistics (columns 3 and 4). Each coefficient comes from a separate regression. Coefficients are estimated using Equation (2) with log differenced outcomes and two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the "Empirical Strategy" section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 8: Results on Rents and Home Values

	(1)	(2)	(3)
<u>Panel A:</u>	Med rent	Med Home Val	Dividends (IRS)
Coal Value (Δ s.d.)	-0.021* (0.011)	-0.032*** (0.012)	-0.006 (0.025)
F excl.	21.28	21.28	21.28
Obs	508	508	508
Region-time FE	yes	yes	yes
CZ controls	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in measures of housing costs and asset appreciation on the treatment variable ΔD_{ct} , coal production value changes per capita. Each coefficient comes from a separate regression. Coefficients are estimated using Equation (2) with log differenced outcomes and two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the "Empirical Strategy" section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 9: Welfare Estimates

	(1)	(2)	(3)
	Household (\$/cap)	Aggregate (\$ B)	Aggregate LR (\$ B)
Panel A: Non-College parameters			
Δ housing costs 2.1%	-1153.15 (545.64)	-0.83 (0.39)	-16.55 (7.83)
Δ housing costs 3.2%	-800.97 (646.73)	-0.57 (0.46)	-11.50 (9.28)
Panel B: College parameters			
Δ housing costs 2.1%	(-1352.20 (406.84)	(-0.97 (0.29)	(-19.41 (5.84)
Δ housing costs 3.2%	(-1182.94 (477.32)	(-0.85 (0.34)	(-16.98 (6.85)

Notes: The table reports estimates of the effect of demand declines on the economic welfare of residents of coal country. Parameters are drawn from Diamond (2016), those in Panel A correspond to Non-College-educated workers those in Panel B to College-educated workers. Column (1) reports estimates the per household impact of a one standard deviation decline at the county level, while column (2) aggregates across households by scaling the effect to the \$8.03 billion decline in coal production value observed over sample and the 7,300,375 households in coal producing counties inferred from IRS tax data. Column (3) inflates column (2) under the assumption that the observed changes in wages, housing costs, and public spending on education are one-time and permanent following Bartik et al. (2019) assuming a 5% discount rate. The rows of the panels explore the sensitivity of the estimates to different effect sizes on housing costs. The calculations are converted to dollars using a mean household wage and salary of \$35,958 inferred from BLS and IRS data. Standard errors are reported in parenthesis and are obtained by bootstrap accounting for the joint distribution of uncertainty in effects on wages, housing costs, and school expenditures per capita.

Appendix (For Online Publication)

Appendix A: Supplemental Descriptive Figures and Tables




Figure A1: Coal Supply Regions from Stoker et al. (2005). Boundaries of coal supply regions drawn at the county level.

Table A1: Supplemental Summary Statistics

	(1)	(2)	(3)	(4)	(5)
	N	mean	sd	min	max
Below poverty line (000s)	254	9.019	16.58	0.175	180.2
Median Home Value (000s)	254	116.0	58.94	36.06	493.1
Median Rent	254	626.3	126.5	289.6	1,242
Annual Inmigr.	254	2,212	5,416	66	76,116
Annual Outmigr.	254	2,213	5,661	86	78,034
Teachers	247	635.8	1,268	17.20	12,654
Students	254	9,648	20,903	221	251,213
Students/Teachers	254	15.16	2.112	9.600	22.20
GP Rev per capita (000s)	254	1.946	1.302	0.302	6.764
SD Rev per capita (000s)	254	1.898	0.869	0	5.298
IRS exemptions (000s)	254	57.47	137.2	1.862	1,673
IRS Returns (000s)	254	28.74	71.68	0.857	873.5

Additional summary statistics for coal counties in 2007. Campbell County, Wyoming excluded.

Appendix B: First Stage Results

Table B1: First Stage Estimates

	(1)	(2)
<u>Panel A:</u>		
Predicted (\$/capita)	0.945*** (0.171)	0.865*** (0.184)
F	15.56	28.27
R2	0.354	0.380
Obs	508	508
Region-time FE	yes	yes
CZ controls	no	yes

Notes: This table reports first stage estimates of the relationship between the predicted coal demand instrument and observed coal production value in dollars. See section 4 for details.

Appendix C: Supplemental Results - Regional Heterogeneity

Table C1 examines effect heterogeneity at the regional level. Each row contains 4 coefficients with estimates of the effect of a one standard deviation decline in coal production value per capita on changes in the outcome listed on the left hand side of the table. The table is split into three main categories - Employment and Wages, Real Estate and Migration, and Public Revenue. In general, Appalachia appears have the largest effect sizes across most categories (and the most precise estimates due to sample size). See Section 6 for further discussion of these results.

Table C1: Heterogeneity Across Regions

	(1)	(2)	(3)	(4)
	East/Appalachia	Illinois	Interior	West
<u>Panel A: Employment and Wages</u>				
miners per cap	-1.13*** (0.08)	-0.94*** (0.16)	-0.23 (16.35)	-0.31 (0.52)
mining wages per cap	-1,139.18*** (327.83)	-489.88 (896.23)	540.46 (8,131.96)	-941.03 (15,459.65)
ln(employment)	-0.07*** (0.02)	-0.02 (0.02)	0.07 (1.81)	-0.04 (0.06)
ln(avg wage)	-0.06** (0.03)	-0.01 (0.10)	0.04 (71.28)	-0.03 (0.28)
<u>Panel B: Real Estate and Migration</u>				
ln (med. home value)	-0.021 (0.016)	-0.035 (0.028)	0.091 (9.038)	-0.095 (0.268)
ln (med. rent)	-0.021 (0.024)	-0.023 (0.026)	-0.006 (0.149)	-0.039 (0.227)
ln (inmigr.)	-0.062 (0.047)	0.046 (0.086)	-0.232 (1.283)	-0.070 (1.606)
ln (outmigr.)	0.009 (0.022)	0.007 (0.500)	-0.165 (1.293)	-0.004 (0.736)
<u>Panel C: Public Revenue</u>				
ln (avg GP exp.)	0.014 (0.082)	0.021 (0.120)	-0.058 (54.600)	0.007 (2.022)
ln (avg SD exp.)	-0.039 (0.042)	0.049 (0.180)	-0.027 (65.456)	-0.103 (6.275)
No. Counties	149	41	31	33
Time FE	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes

Notes: This table reports region-level 2SLS regressions of 5-year changes in labor market outcomes, home values, migration, and local public finance on the treatment variable $\Delta D_{ct}/Pop_{c,t-5}$, coal production value changes per capita. Coefficients in the first two rows are estimated using Equation (1) and all other rows using Equation (2) with log differenced outcomes. All rows use two periods of data in the stacked difference model. Time FE is an indicator for the later period. Control variables are described in the "Empirical Strategy" section. Treatment is standardized to have mean 0 and standard deviation 1. SE are estimated via clustered pairs bootstrap with 250 repetitions (re-sampling commuting zones).

Appendix D: Supplemental Results - Spillovers

Table D1 examines spatial spillovers by estimating the impact of commuting zone level production declines on adjacent non-producing counties. That is, non-producing counties that share a commuting zone with coal producers constitute the estimation sample, and their outcomes are modelled as a function of declining coal production value per capita within the commuting zone. Table D1 results share a similar format to Table 3 (omitting the first column), with panel A focusing on employment and panel B on wages. See Section 6 for a discussion of the results.

Table D1: Spillover Employment and Wage Regressions

	(1)	(2)	(3)	(4)	(5)
Panel A: Employment	Mining (BLS)	Total	Manufac	Nontradeable	Utility/Transp
CZ Coal Value (Δ s.d.)	0.48***	0.18	0.01	-0.50	0.03
	(0.13)	(0.28)	(0.05)	(0.43)	(0.09)
Obs	638	638	638	638	638
Panel B: Wages	Mining (BLS)	Total	Manufac	Nontradeable	Utility/Transp
CZ Coal Value (Δ s.d.)	208.06***	139.00	7.85	-114.75	54.56
	(76)	(219)	(21)	(196)	(63)
Obs	638	638	638	638	638
Region-time FE	yes	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in measures of employment and wages on the treatment variable $\Delta D_{ct}/Pop_{c,t-5}$, coal production value changes per capita. Each coefficient comes from a separate regression. Coefficients are estimated using Equation (1) with per-capita outcomes and two periods of data in the stacked difference model. All regressions are conditional on **state-time** fixed effects and control variables (described in the "Empirical Strategy" section). Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Appendix E: Welfare Estimates based on Spatial Equilibrium

The model economy consists of many homogeneous individuals $1, \dots, i, \dots, N$ who value wages, rental rates, and public goods and receive an individual-region specific preference shock. Individuals choose a region c to maximize their indirect utility U_{ic} given by Equation (E1), where w_c , r_c , and G_c are the wage rate, cost of housing, and a public good in location c . The parameter σ_c scales the preference shock distribution according to other non-modelled features affecting mobility that might vary by location, such as local ties.

$$\ln U_{ic} = \alpha \ln w_c + \beta \ln r_c + \gamma \ln G_c + \sigma_c \epsilon_{ic} \quad (\text{E1})$$

Conditional on location choice, individuals inelastically supply one unit of labor. Consider, in the two region case, the comparison that individual i makes between region c and another location k . The individual chooses region c if Inequality (E2) holds.

$$\alpha[\ln w_c - \ln w_k] + \beta[\ln r_c - \ln r_k] + \gamma[\ln G_c - \ln G_k] + [\sigma_c \epsilon_{ic} - \sigma_k \epsilon_{ik}] > 0 \quad (\text{E2})$$

Intuitively, Inequality (E2) holds when an individual's idiosyncratic preference for region c exceeds the utility premium the individual would receive from living in k . If the differences $\sigma_c \epsilon_{ic} - \sigma_k \epsilon_{ik}$ are draws from an exponential distribution with shape parameter $\sigma^{-1} = \frac{\sigma_k}{\sigma_c}$, then the log share of individuals residing in region c is approximated by Equation (E3).

$$\sigma \ln\left(\frac{N_c}{N}\right) = \alpha[\ln w_c - \ln w_k] + \beta[\ln r_c - \ln r_k] + \gamma[\ln G_c - \ln G_k] \quad (\text{E3})$$

An elastic housing supply function given by Equation (E4) closes the model and shrinks the range of parameter values that yield corner solutions.¹ Equilibrium is achieved when all agents choose their preferred location, and the housing market clears.

$$\ln r_j = \eta_{1j} + \eta_{2j} N_j, \quad j \in \{c, k\} \quad (\text{E4})$$

To reach a usable expression for changes in willingness to pay for the bundle of characteristics of region c , assume that region c is small relative to location k ² and take the derivative of Equation (E3) with respect to $\ln w_c$. This leads to Equation (E5), in which α in the utility model is proportional to the elasticity of (N_c/N) with respect to the wage. Analogously, β and γ are also proportional to the elasticities of population share in region c to rents and local amenities in c .

$$\frac{d \ln(N_c/N)}{d \ln w_c} = \alpha / \sigma \quad (\text{E5})$$

¹Endogenizing rents creates a feedback between population shares and home prices captured by the semi-elasticity η_{2j} from Equation (E4) that pushes the desirability of locations towards equality.

²The assumption of the "smallness" of location c follows Bartik et al (2019). and implies no general equilibrium effects - e.g., marginal changes in prices in region c do not affect prices in k . It can be thought of as letting k represent "everywhere else" and greatly simplifies the relationship between the migration function and the utility function.

Next, let wages, rents, and public good expenditure each be a function of a latent index θ that may represent global demand for a good produced only in region c . Taking the derivative of U_{ic} for a resident of location c with respect to θ and dividing by α yields Equation (E6), a money-metric expression for the change in indirect welfare from residing in location c (equivalently, the change in willingness to pay for changes in the bundle of characteristics associated with location c).

$$\frac{1}{\alpha} \frac{d \ln U_{ic}}{d \theta} = \frac{d \ln w_c}{d \theta} + \frac{\beta}{\alpha} \frac{d \ln r_c}{d \theta} + \frac{\gamma}{\alpha} \frac{d \ln G_c}{d \theta} \quad (\text{E6})$$

The value of Expressions (E5) and (E6) is that they relate elasticities of migration with respect to the characteristics of region c to a utility function in a simple economic model of location decisions. For a marginal change in the latent index θ , migration decisions reveal effects on individuals' willingness to pay for the bundle of wages and amenities they enjoy in location c , and hence reveal changes in the economic welfare of its residents. Appealing to the envelope theorem argument in Busso, Gregory, and Kline (2013) suggests that to assess the aggregate effect of a change in θ on the model economy, we need only sum Expression (E6) over individuals in region c . Intuitively, this is because individuals just on the margin between locations are indifferent between alternatives, implying that the relevant welfare effects of the shock to θ are only the price changes it induces, almost as though the population were immobile.

I combine my estimates of changes in wages, housing costs, public good expenditures with parameters from Diamond (2016), who estimated 10-year migration elasticities in terms of wages, housing costs, and local amenities (including public education expenditures), to approximate economic welfare changes from demand shocks to coal country. I estimate the total wage-equivalent welfare impact of observed declines in coal production value using the expression given by Equation (E7),

$$\frac{1}{\hat{\alpha}} \frac{dU}{d\theta} = N_c * \left(\frac{d \ln w_c}{d \theta} + \frac{\hat{\beta}}{\hat{\alpha}} \frac{d \ln r_c}{d \theta} + \frac{\hat{\gamma}}{\hat{\alpha}} \frac{d \ln G_c}{d \theta} \right) * \overline{w_{c0}} \quad (\text{E7})$$

where $\overline{w_{c0}}$ is the average baseline wage, which monetizes the proportional changes implied by Equation (E6) following Bartik et al. (2019). The set of estimated parameter values are drawn from Diamond (2016), one set of migration elasticities estimated for college-educated workers and one for noncollege workers. An unweighted average of the parameter vector is $(\alpha, \beta, \gamma) = (4.1185, -2.5515, 0.7045)$.³ Normalizing β and γ by α implies that housing cost changes are valued at approximately 60% of wage changes and amenity changes are valued at roughly 20%. In practice this approach allows me to create a coherent, combined valuation of the three largest components of welfare changes revealed by the causal estimates - the effect on wages, housing prices, and public expenditures on education. See section 7 for further discussion and presentation of results.

³These parameters are taken from Panel A of Table 5 of Diamond (2016), and reflect migration elasticities estimated under the least restrictive possible model. Noncollege-worker parameter estimates on wages, housing costs, and Diamond's amenities index are (3.261, -2.944, 0.771) and college-worker parameters are (4.976, -2.159, 0.638). Note here that I interpret the estimates as implicitly including the σ term.

Appendix F: Data Appendix

This section describes the data sources and cleaning in more detail.

i. Energy Market Data

Electricity market data comes from the EPA’s Continuous Emissions Monitoring Systems data. This contains hourly gross generation and emissions data for many of the fossil fuel power plants in the US. I downloaded these data directly from the EPA’s FTP site. Due to availability issues with plant metadata, I used EPA’s eGRID data to infer unit subregion and primary fuel types each year, filling gaps forward for years where eGRID is not available.

Data on the **supply side of the coal market** comes from the Mine Safety and Health Administration (MSHA) and EIA Form 7A, which includes production in short tons and average employment for mines across the US. **Demand side data on the coal market** is derived from a combination of EIA (joint with FERC) 423 and EIA 923 data.

Combining the demand and supply side of the coal market, **fuel prices** are taken from the EIA’s State Energy Data System (SEDS).

ii. Outcomes

Wages, employment, unemployment, and labor force participation all come from the BLS. I corroborated income effects with data from the BEA and IRS, but ultimately cut those estimates from the paper.

Public finance data comes from the Census of Governments which I obtained via the Government Finance Database.

Population data comes from various US Census datasets, most importantly the Inter-censal population estimates (both total and broken down by age and sex). **In-migration and outmigration** are obtained from the IRS’s county-to-county migration files, which they infer from changes in tax filings.

Median home values, rents, and educational attainment at the county level are derived from the American Community Survey (ACS) and collected via IPUMS NHGIS. To build the long differences, I use the decennial census in 2000 as an estimate of $t_1 = 2002$, and other values are inferred using the 5-year files (2005-2009, 2010-2014, 2015-2019). I considered other outcomes from this Census/ACS combination, but ultimately ruled them out because they did not appear to harmonize well.

I also estimated home value effects using the Federal Housing Finance Agency’s (FHFA) Home Price Index (HPI) published at the county level. I found larger negative impacts than those estimated using the ACS survey data, but ultimately cut these results from the paper because of a substantial number of missing observations (low coverage of the indices for low-population counties).

Crime rates come from the FBI’s Uniform Crime Reporting Data and **student and teacher counts** come from the National Center for Education Statistics.

iii. Other controls

Oil and gas production value come from Enverus (formerly DrillingInfo) through a subscription. These data provide gas and oil production at the well level for most areas of the country, which I aggregated up to the county-year level and valued using the West Texas Intermediate (WTI) and Henry Hub prices available from the EIA.