

Coal Demand, Market Forces, and US Coal Mine Closures

Brett Jordan, Ian Lange, and Joshua Linn

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org



Coal Demand, Market Forces, and US Coal Mine Closures *

Brett Jordan[†]

University of Alaska Anchorage

Ian Lange[‡]

Colorado School of Mines

Joshua Linn[§]

University of Maryland and Resources for the Future

Abstract

Economic transitions have the potential to displace workers and cause social unrest. Coal mine closures in the eastern United States due to the changing electricity system and the resulting employment losses in rural areas have become salient issues for all levels of government. Previous research has not distinguished among the potential causes of recent mine closures, such as rising production costs and decreasing coal demand from the electricity sector. This analysis utilizes unique data on coal mine and power plant operation to estimate the impact of supply and demand factors on mine closure. We model closure as a function of expected profits, which allows us to compare the effects on mine closure of specific demand and supply shocks to expected mine profits. Our results suggest that each shock substantially affected coal mine employment. Increasing costs of producing Appalachian coal have had the largest impact on closures with lower natural gas prices and lower electricity demand each accounting for a substantial number of closures additionally.

JEL Classification: L51, L71, Q35, Q53

Keywords: Coal Mining; Firm Exit; Fuel Procurement

*This work is supported by NSF grant number 1559339 Energy Security Implications of Premature Coal Mine Closure. The authors are grateful for the helpful comments of Jurgen Brune, David Evans, Kenneth Gillingham, John Grubb, Ronald Shadbegian, and seminar participants at the Political Economy Research Center, Colorado School of Mines, Resources for the Future, and the Front Range Energy Camp. The authors wish to thank David Manning, Muhammad Abdullah Khawar, and Ensieh Shojaeddini for their excellent research assistance.

[†]bwjordan2@alaska.edu

[‡]*Corresponding author:* Division of Economics & Business, Colorado School of Mines, Golden, CO 80401, USA. ilange@mines.edu

[§]linn@umd.edu

1 Introduction

The US electricity sector is undergoing a transformation, driven by technological developments such as innovation in natural gas production and wind power, and greater use of energy efficiency. The transformation has dramatically reduced US consumption and production of coal—both relative to expectations in the mid-2000s and in absolute terms. Coal production declined 19 percent from 2015 to 2016 and by 38 percent since its peak in 2008 ((Energy Information Administration, 2017a), Table 6.1). Between 2008 and 2015, coal-mining employment declined by 23 percent (Mine Safety and Health Administration, 2016) and 31 percent of coal mines closed (The National Institute for Occupational Safety and Health, 2016), mostly in the eastern United States.

The result of the 2016 presidential election led to some federal policies and proposals that aim to reverse the downward trend of coal production and employment. One of those policies is a reversal of an Obama administration order to review the federal coal leasing program. In early 2016, the Obama administration placed a moratorium on new coal leases on federal lands, beginning a review of leasing policy that might have ultimately raised costs of producing coal from federal lands (Department of Interior, 2016). However, the plight of coal miners and coal-mining communities across the United States generated much discussion during the 2016 presidential election. The closure of a coal mine is a highly salient event with potential knock-on effects to local economies. Although coal mining accounts for a small share of total US employment, mining significantly contributes to local economies in many areas, such as West Virginia (Black et al., 2005). Like other resource dependent regions, Appalachia has exhibited depressed long run income growth, consistent with the so-called resource curse or drag (Douglas and Walker, 2016; Jacobsen and Parker, 2016).

The potential efficacy of coal policies is an open question. A number of recent studies (e.g., Department of Energy (2017), Linn and Muehlenbachs (2018), Cullen and Mansur (2017), Fell and Kaffine (forthcoming), and Holladay and LaRiviere (2017)) have concluded that natural gas prices and renewables substantially reduced coal consumption at US power

plants between 2008 and 2014. Houser et al. (2017) apportion reductions in coal consumption to electricity market and regulatory factors, while Linn and McCormack (2017) analyze the effects of market and regulatory forces on the retirements of coal-fired power plants. However, most of the employment losses in the coal sector have occurred among coal mines rather than power plants. Although Brinkman et al. (2017) examine the longer-term trends in coal mine closures, the literature has not attempted to explain the recent coal mine closures. Standard microeconomic theory would suggest that a decline in coal demand from the electricity sector would reduce mine profitability and cause closures, but other factors could contribute to closures, such as increased interregional competition¹, declines in economically recoverable reserves, or regulations that raise production costs. Disentangling the contributions of these factors to mine closure would be essential for understanding how hypothetical policies affecting coal demand and supply, such as those considered by the Obama and Trump administrations, would affect coal-mining closures and employment.

In this paper, using data from 2002 to 2012, we estimate the effects of changes in coal mine profitability on mine closures and mining employment. We isolate individual demand- and supply-side contributions to recent closures, comparing the importance of these factors in explaining recent closures. This study period was chosen based on availability of data, particularly confidential data from Form EIA-923. We focus on closure decisions among Appalachian coal mines, which account for most of the observed closures during this period. Appalachian mines are typically smaller and higher-cost mines relative to the western mines (Gillingham et al., 2016).

A mine's closure depends on expectations of future profits. Given the large numbers of mines, coal shippers, and coal-fired electricity plants in the United States, it is not computationally feasible to estimate and simulate a mine-level dynamic stochastic model that allows for competitive behavior by mines and shippers. Instead, we adapt and extend methods developed in the literature, which allows us to estimate a mine's expected profits and

¹The Clean Air Act of 1990 led most coal plants to alter their boilers in a way that allows for the burning of coal from different regions (Ellerman et al., 2000).

approximate equilibrium outcomes under imperfect competition. We estimate the expected discounted profits in each year from 2002 to 2012 for each Appalachian coal mine. Profits depend on coal demand from the electricity sector, coal attributes, transportation costs, and production costs. A mine's profits also depend on its market power, export demand, and potential competition from other mines. The estimation of expected profits takes advantage of unique data on marginal costs and coal prices, and a novel identification strategy for estimating demand for each mine's coal that accounts for imperfect competition among mines and the endogeneity of coal prices and demand. We adapt Wolak (2003) to approximate equilibrium coal shipping behavior given market power for shippers and other factors affecting the equilibrium, and we use a computational model of the electricity generation sector to construct plausibly exogenous measures of the future demand for a mine's coal. Johnson et al. (2017) use a computational model to predict the generation at individual coal-fired plants in a particular period, and we build on this approach by using a more accurate computational model and by estimating future coal demand, which is necessary to account for the forward-looking nature of closure decisions (see Linn and McCormack (2017) for details on the model).

Having estimated each mine's expected profits, we use a duration model to estimate the effect of profits on the decision to close the mine. We find that expected profits strongly predict closure, such that a 10 percent decrease in profits raises the closure probability by 0.2 percentage points above the baseline closure rate, which implies a 3 percent increase in the closure rate.

Using the estimated effects of expected profits on closure, we analyze scenarios to compare the relative effects of demand, supply, and policy factors in explaining coal mine closures between 2002 and 2012. Conceptually, each scenario includes a shock to a particular demand or supply-side determinant of profitability. We compute the effect of the shock on expected profits, and the effect of the resulting profit change on closure. For each scenario, we compare the effect of the shock on closures with the total number of closures caused by declining

profitability during the sample. That is, we net out a baseline closure rate that reflects factors not included in the profitability model, such as regulation of coal mines. This adjustment accounts for the high degree of job losses from mine closings that are typically observed in the industry (Dunne and Merrell, 2001), independently of the factors driving mine profitability that we include in our model.

On the demand side, given the recent findings in the literature on the effects of electricity demand and natural gas prices on domestic coal consumption, we focus on these two factors, which have jointly reduced coal consumption by about one-third (Linn and McCormack, 2017). Between 2006 and 2012, electricity consumption increased by an annual rate of 0.6 percent, whereas the Energy Information Administration had projected a 2 percent annual growth rate during this period (Energy Information Administration (2015b), Table 2.2). Between 2008 and 2012, real natural gas prices declined by 69 percent (Energy Information Administration, 2017c). Both declining natural gas prices and slower growth in electricity consumption reduced coal demand and mining profitability, raising closures. The reduction in natural gas prices explains 125 (30 percent) of the mines that closed between 2002 and 2012 as a result of declining profits. The reduction in electricity consumption explains 135 (33 percent). The analysis accounts for the fact that the shocks occurred at different times during our sample period. The fact that the electricity consumption shock occurred prior to the natural gas price shock partly explains the larger effect of the consumption shock.

We compare the demand-side explanations of closures with a supply-side explanation. Between 2002 and 2012, real per-ton extraction costs in Appalachia nearly doubled, and we show that declining worker productivity rather than rising wages explains most of the cost increase. Consequently, we define the supply-side shock to a mine as the decrease in a mine's productivity relative to the maximum productivity that we observe at the mine. Because of rising extraction costs and other factors, generally we observe declining productivity over time across mines in the data. The declining productivity caused 274 mine closures between 2002 and 2012, or 67 percent of the total closures in the sample caused by declining prof-

itability. The supply-side shock caused a roughly 20 percent decrease in US coal mining employment, and the two demand-side shocks each affected employment substantially.

The conclusion that declining mine productivity explains more closures than declining coal demand is perhaps surprising, given the focus of the literature and public debate on demand rather than supply side factors. However, this conclusion is consistent with the magnitudes of the shocks. During the sample period, declining productivity reduced annual operating profits three times as much as did lower natural gas prices or electricity consumption.

We note that the three counterfactuals in isolation cause closures exceeding the number of closures observed as a result of declining profits. This is because the demand and supply shocks affected the same mines, and many mines close in multiple counterfactuals. Nonetheless, the results suggest that the supply shock had a larger effect on mine closures than did each of the demand shocks in isolation. Although explaining the falling productivity lies outside the scope of this paper, we note that changes in permitting procedures or declining reserves could partly explain the declining productivity.

The analytical structure allows us to estimate the effects of potential policies on closures. As noted above, the Obama administration began a review of federal leasing policy, and some advocates, analysts, and policymakers have discussed raising the cost of producing coal on federal lands to reflect social costs of coal consumption, such as greenhouse gas emissions. To the extent that western and eastern coal mines compete with one another, raising costs of producing from federal lands increases profits at eastern mines and reduces the probability that those mines close. However, our coal mine profitability model implies that western and eastern coal are rather weak substitutes for one another. A hypothetical policy that raises costs of western mines and causes the 10 highest-cost mines to close would reduce closures by 84 mines, which is smaller than the estimated effects of the other supply and demand scenarios on closures.²

²Western production is heavily concentrated in two large operations near Gillette, Wyoming. Other mines in the region are one-tenth to one-hundredth as large as these two mines.

In addition to the novel empirical strategy for estimating mine closures and considering counterfactuals, our results contribute to the recent literature on coal mine policy. The Obama administration’s lease moratorium motivated a number of studies on coal mine behavior. Gillingham et al. (2016) discuss potential options for improving the coal mine leasing process and discuss the potential for substitution in consumption of federal and nonfederal coal. Gerarden et al. (2016) use a dynamic simulation model to show that increasing the costs of mining on federal lands has a small impact on nonfederal coal production. Our analytical approach and results differ from theirs. Whereas they use a forward-looking computational model and compare simulation results across hypothetical future policy scenarios, ours is the first to use observed closure decisions to estimate the effects of demand- and supply-side factors on closure. For computational reasons, their model aggregates across mines and power plants, while we perform a mine-level analysis. Their focus on hypothetical policies affecting coal on federal lands contrasts with our focus on the potential explanations for recent coal mine closures.

The rest of the paper is organized as follows. Section 2 discusses the coal mine and power plant data and their trends, Section 3 explains our estimation strategy, Section 4 provides our estimation results and counterfactual simulations, and Section 5 concludes.

2 Background

2.1 Data

The analysis uses three main data sources. The first is the Mine Safety and Health Administration (MSHA) Part-50 database. The data include the location, production, and employment for every US coal mine. The second is the Energy Information Administration (EIA) Form 423/923 data on coal purchases. We utilize the complete Form EIA-423/923 data that contain fuel expenditures and quantities for both regulated and deregulated power plants. We obtained the data for deregulated plants under a nondisclosure agreement with

the EIA. The data have been used in other recent papers, such as Linn et al. (2014) and Cicala (2015). The third dataset is the Environmental Protection Agency’s (EPA) Continuous Emissions Monitoring System (CEMS). For most fossil fuel-fired electricity generation units in the United States, the data include hourly fuel consumption, generation, and emissions. These primary datasets are supplemented with wage information from the Bureau of Labor Statistics’ (BLS) Quarterly Census of Employment and Wages (QCEW) and financial disclosure information (Form 10-K) from publicly traded firms.

2.2 Trends in Coal Consumption

Coal became the backbone of the electricity generation system in the United States following the Organization of the Petroleum Exporting Countries (OPEC) oil embargo in 1973-4. The federal government encouraged coal-fired power plant construction and encouraged production of coal as a substitute for oil. Between the mid-1970s and mid-2000s, the share of coal-fired generation in total electricity generation increased from about 44 to 50 percent ((Energy Information Administration, 2017b) Table 7.2A)). As a result, the amount of coal consumed by the electricity sector increased from 664 million tons in 1984 to 851 million tons in 2014 (Energy Information Administration, 1990, 2015a). Coal exports have fluctuated over time; for example, 81 million tons were exported in 1984, 58.5 million tons in 1999, and 97 million tons in 2014.³

Over time, the consumption of coal came to be almost exclusively by the electric power sector. According to the EIA, between 1988 and 2014, the amount of US coal consumed by the industrial, residential, and commercial sectors declined from 83 million tons ((Energy Information Administration, 1998), Tables 71 and 75) to 45 million tons (Energy Information Administration (2015a), Table 26). Another large consumer of coal is the coking/steel industry. This industry has also had a decline in coal consumption, from 41 million (Energy


³See Energy Information Administration (1994), Table 76; Energy Information Administration (2000), Table 76; and Energy Information Administration (2015a), Table 8. By comparison, coal imports were around 9 million tons in 2012 (Energy Information Administration, 2015a), and historically the United States has been a net exporter of coal.

Information Administration (1998), Table 73) to 21 million short tons (Energy Information Administration (2015a), Table 26) over the same time period. Coal that is used in coking has specific properties and generally is not used by power plants, though sometimes the two sectors compete for coal. Given the relatively small amount of coal consumption by nonelectric power sources, we do not model these sectors specifically, but we account for fluctuations in their consumption over time.

2.3 Trends in Coal Production

There are three major coal-producing basins in the United States: Appalachian, Interior, and Western. See Figure 1 for a map of the basins. Most coal production in the West is from federal lands, whereas most coal production in the East is from private lands. Compared with eastern coal, western coal tends to have lower heat content (the amount of heat released from burning a certain weight of coal). The lower heat content puts western coal at a disadvantage relative to eastern coal, but two factors favor western coal: lower production costs and lower sulfur content (environmental regulations introduce a shadow cost on sulfur emissions from power plants). Figure 2 shows that real per-ton production costs increased between 2002 and 2012 in all regions, and that the cost increase was greatest in Appalachia. Figure 3 shows that this increase in production cost is largely driven by changes in labor productivity. The solid line shows the wage rate (dollars per hour) divided by labor productivity (tons per hour) for each basin in each year, while the dotted line is the same ratio but with constant wages. The gap between the two curves represents the contribution of wage changes to the observed cost changes. As the figure reveals for Appalachian mines, declining productivity has been a much larger labor cost driver than wage increases, evidenced by the small gap between the two plotted series. The rising relative cost of producing Appalachian coal has likely contributed to the decline of Appalachian coal production relative to other regions, which Figure 4 shows. Regulation of sulfur dioxide emissions from the electricity sector, health and safety regulation (Gowrisankaran et al., 2015), and changes in permitting of new

Figure 1: US Coal Basins




Source: Energy Information Administration

mines also have contributed to this trend.

Because coal shipments play an important role in the empirical analysis, Figure 5 illustrates the geographic patterns of coal shipments. For the main coal-producing states, the figure provides a heat map of the shipments of coal to other states for the years 2002-12. The states with federal coal leases, primarily Wyoming, Utah, and Colorado, sell to a wide range of states. However, there are states, particularly in the Southeast, to which very little western coal is shipped. This is likely due to transportation costs, as coal is relatively bulky and transportation firms often have market power (Busse and Keohane, 2007). Some states, such as Ohio and Tennessee, receive coal from multiple regions. Other states, especially in the Southeast and Northeast, receive coal only from Appalachian mines. Additionally, Figure 5 shows that some states consume their own production exclusively, such as Texas and North Dakota. The fact that Western low cost coal does not reach all consumers could be due to transportation costs or to incomplete and sluggish pass-through of mining costs to delivered prices (Chu et al., 2017). Eyer and Kahn (2017) find that the amount of coal


Figure 2: Estimated Variable Cost by Basin, 2002-12



Note: Average variable costs from sample of 10-K filings for publicly traded firms. Coal basin definitions follow Stoker et al. (2005). APP: Appalachian, INT: Interior, WST: West, LIG: Lignite, PRB: Powder River Basin.


bought from in-state mines has been increasing. These patterns suggest that market power of shipping firms and coal mines may vary geographically. Our model of coal mine profitability, described in the next section, accounts for such potential variation.

Figure 3: Estimated Labor Cost by Basin, 2002-12



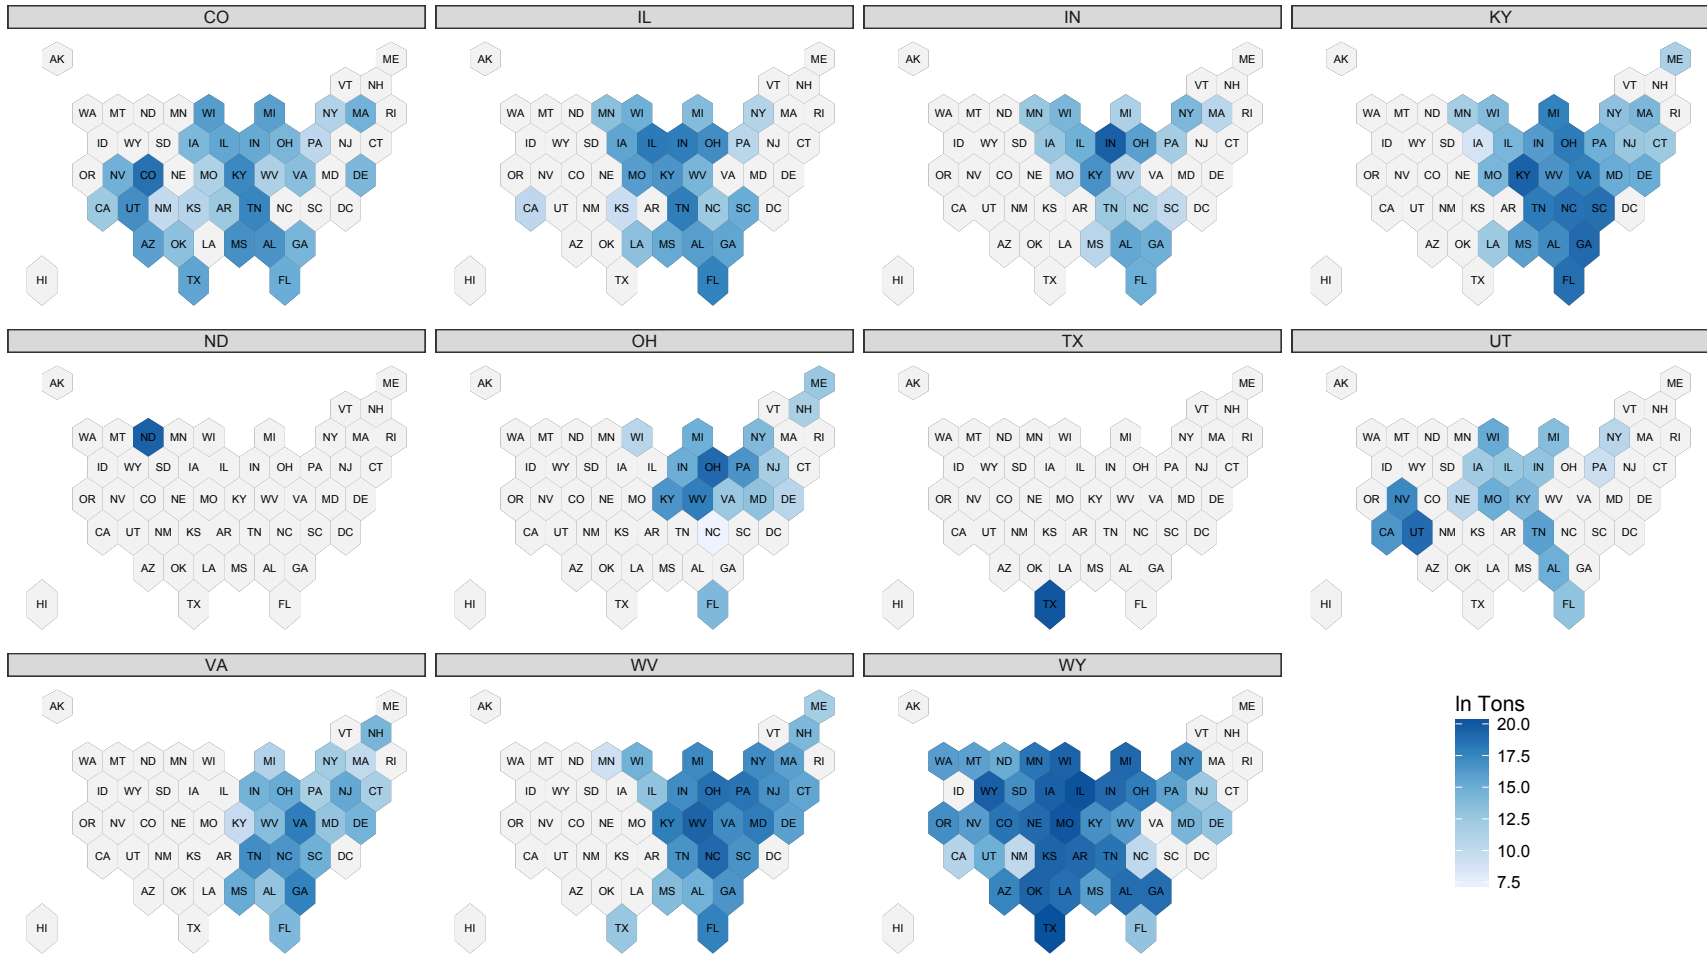
Note: Labor productivity data come from the Mine Safety and Health Administration (MSHA) Part-50, and wage data come from the BLS QCEW. Coal basin definitions follow Stoker et al. (2005). APP: Appalachian, INT: Interior, WST: West, LIG: Lignite, PRB: Powder River Basin.

Figure 4: Annual Coal Production by Major Basin, 2002-12



Note: Annual production constructed from Mine Safety and Health Administration (MSHA) Part-50. Coal basin definitions follow Stoker et al. (2005). APP: Appalachian, INT: Interior, WST: West, LIG: Lignite, PRB: Powder River Basin.

Figure 5: Deliveries by Selected Supply State, 2002-12



13

Note: Totals represent cumulative shipments (in logged tons) from 2002 to 2012. Each panel represents origin state; each stylized mapped state represents destination states. Shipment data from Form EIA-923.

2.4 Trends in Appalachian Mine Closures

The Appalachian region has been a major coal producer in the United States since coal became a widely used energy source. In the recent history of coal mining, interregional competition along with environmental regulation played a strong role in determining the profitability of Appalachian coal mines. Throughout the last 40 years, the number of mines operating in the Appalachian region has decreased. Among factors contributing to this decline has been the competition with coal from the Powder River Basin (PRB) in Wyoming. PRB coal has lower production costs than Appalachian coal and also lower sulfur content, as noted above. The competition arises from declining shipping costs as well as regulation of power plant sulfur dioxide emissions (Ellerman et al., 2000). As Table 1 reveals, the number of Appalachian mines operating has fallen by two-thirds in the last 40 years. The rate of mine closure has been relatively consistent, with a small respite in the 2000s likely due to the fact that many power plants had installed capital controls to reduce sulfur emissions. Coal prices started to decline in the late 1980s and remained low for much of the 1990s (Joskow, 1990; Bonskowski, 2000), which led to the consolidation of mines to favor larger mines and those with thicker seams of coal (Bonskowski, 2000). Over this period, average per-mine employment has increased, likely due to the movement toward larger mines.

Table 1: Average Yearly Appalachian Closures Observed by Decade

Decade	Avg Closed/Year	Avg Opening/Year	Avg Mines Operating	Avg Employee Count	Avg Closure Rate	Avg Opening Rate
1980s	1029.20	415	9217.30	32.30	0.11	0.04
1990s	498.70	208	5475.10	34.64	0.09	0.04
2000s	247.30	112	3674.80	39.33	0.07	0.03
2010s	241.25	74	2637.50	50.25	0.09	0.03

Source: Calculated by authors from MSHA Part 50 Data

2.5 Federal Coal Leasing Policy

In addition to promoting the construction of coal-fired power plants in the 1970s, the federal government encouraged coal production by reforming the process for leasing federal lands to mine coal. The Mineral Leasing Act of 1920 and the 1976 Federal Coal Leasing Amendments

Act (FCLAA) authorize leasing of federal lands for coal production. The FCLAA was passed after the OPEC oil embargo and instructs the Bureau of Land Management (BLM) to manage coal leases to ensure maximum economic recovery of coal and to ensure fair market value in the lease sales.

The main pathway firms have used to acquire leases is lease-by-application. After a firm submits an application to acquire a lease for a specific deposit, BLM reviews the application to determine fair market value of the deposit and to assess compliance with other regulations. If BLM approves the lease, it holds a public sealed-bid auction for the lease. The winner of the auction is given a lease term of 20 years with the option to renew every 10 years after the initial term⁴. BLM has generally renewed leases when leaseholders have extracted commercial quantities of coal.⁵

While the federal coal leasing program provides a large amount of revenue to the government, the program may not provide fair market value. Both the Government Accountability Office (GAO) and the Department of the Interior's Office of Inspector General (OIG) have issued reports raising concerns about the BLM not receiving fair market value for federal coal leases. Government Accountability Office (2013) reports that lease valuation techniques are not standardized across states and that 96 of the 107 leases analyzed had only one bidder. Office of Inspector General (2013) found cases where BLM state offices accepted bonus bids less than the statutory minimum, and reported that a large majority of lease modifications were sold at the minimum bid. Both reports argue that the price of coal exports should be taken into account in the determination of fair market value of coal deposits. Because of these and other concerns, the Obama administration put a moratorium on new federal

⁴Leaseholders make three payments to the BLM. The first is a bonus bid paid when the firm wins the lease auction. The second is an annual rent payment of \$3 per acre. The third payment is a statutory minimum 12.5 percent gross royalty on the value of coal produced (the statutory minimum is 8 percent for underground coal mines, but most federal coal leases are surface mines). In 2015, the bonus bids provided around \$450 million in revenue to the federal government and the royalty rates provided \$680 million (Office of Natural Resource Revenue , 2015).

⁵Leases can also be modified to add adjacent tracts, typically when the adjacent tract is unlikely to have been mined otherwise. The maximum size of a modification is 960 acres. Lease modifications are not generally auctioned but the BLM has to ensure fair market value is received for the tract (Bureau of Land Management, 2012).

Table 2: 2015 Federal Coal Leases

State Office ^a	Leases	Acres	% of Leased Acreage ^b
Colorado	54	88,471	18%
Eastern States	7	16,939	4%
Montana	49	48,408	10%
New Mexico	21	42,756	9%
Utah	73	85,037	18%
Wyoming	102	200,559	42%
Total	306	482,170	100%

Source: Data from Bureau of Land Management.

^a MT state office includes ND; NM state office includes OK

^b Individual percentages do not add to 100 because of rounding.

coal leases in January 2016 while the Department of Interior studied whether and how to reform the coal leasing program (leases in operation were unaffected). In March 2017, President Trump signed an executive order to lift the moratorium. The Trump administration has also sought policies that would increase coal demand, such as by weakening carbon emissions regulations of the electricity sector.

Because of the high quality (low sulfur, proximity to surface, and thick seams) of coal deposits in Wyoming, along with the potential cost advantage of federal leases, coal mining on federal land constitutes approximately 40 percent of total US coal production (Government Accountability Office, 2013). By comparison, federal lands accounted for 11 percent of total US production in 1980 and 30 percent in 1993.⁶ In 2015, there were 306 active federal coal leases and 482,000 acres under lease, with most of the leased acreage in western states (see Table 2). Currently, most production on federal land occurs in the PRB of Wyoming. In 2013, 96 percent of federal coal production came from five states — Colorado, Montana, New Mexico, Utah, and Wyoming — with approximately 86 percent from Montana and Wyoming (Krupnick et al., 2015).

⁶See Energy Information Administration (1990), Table FE1; and Energy Information Administration (1994), Table 12.

3 Empirical Strategy

The intuition underlying the empirical strategy is that a capacity-constrained mine with higher expected profits is less likely to close than a mine with lower expected profits. Based on this intuition, the empirical strategy consists of two stages. In the first stage, for each mine that operates in any year between 2002 and 2012, we estimate the expected profits of that mine, given expectations of demand for the mine’s coal from coal-fired power plants, other domestic industries (such as industrial boilers), and exports. The profits model is forward-looking, in that profits in a particular year equal the present discounted value of future profits over subsequent years.

In the second stage, we use a duration model to estimate the effect of expected profits on the probability a mine closes between one year and the next. The two-stage approach allows us to incorporate expectations of future market conditions, which would not be possible if we simply estimate the effect of contemporaneous market conditions on closure. We describe the estimation of the profit function and then the closure analysis.

3.1 Expected Mine Profits

3.1.1 A mine’s profit maximization

Here, we describe the basic structure of the model. We construct a partial equilibrium model, in which a mine’s profits depend on the demand for the mine’s coal as well as its production costs.

The market consists of mines, indexed by i ; power plants, indexed by j ; and shippers. All agents maximize their profits given the behavior of other agents in the market.

Consider a mine that is open in period t and is deciding whether to close between periods t and $t + 1$. The mine’s profits in any year $s > t$, $\pi_{it,s}$, depend on the price of its coal ($p_{it,s}$), its marginal costs ($m_{it,s}$), and the quantity produced ($q_{it,s}$).

The mine may have market power, because coal is a differentiated product and because

of transportation costs. The demand for a mine’s coal depends on the qualities of the coal, including ash content, heat content, and sulfur content. Transportation costs typically account for a large share of the delivered price of coal. Consequently, a coal mine located nearer to a coal-fired power plant can generally sell the coal to the power plant at a lower cost than a mine located farther from the same power plant (i.e., including transportation costs). The transportation costs create market power for the mine located closer to the plant.

Because of the market power, the mine chooses its production quantity mindful of the effect of this choice on the equilibrium price. Specifically, the mine chooses the profit-maximizing quantity in each year s :

$$\begin{aligned} \pi_{it;s} = \max(p_{it;s}(q_{it;s}) - m_{it;s}(q_{it;s}))q_{it;s} \\ \text{subject to } q_{it;s} < k_i \end{aligned} \tag{1}$$

where k_i is the maximum production level. We refer to the optimization in Eq. 1 as the profit function. Marginal costs are a nondecreasing function of quantity, and the production quantity is a decreasing function of the mine’s price.

Because we do not directly observe profits, we must compute profits from Eq. 1. This requires estimating both the demand curve and marginal cost function. Once we have estimated these equations, we can use projected market conditions for each year $s > t$ to compute expected profits in year s . Then, we discount future profits back to the present and compute the sum of expected discounted profits.

The demand curve is a function of the price of the mine’s coal and the demand from electricity generators net of the supply from other mines. Demand also comes from domestic consumers outside the electricity sector, such as industry, as well as exports. For convenience, we combine these two sources of demand in the variable $X_{s;t}$. That is, we allow demand for coal from mine i in year $s > t$ to depend on the price of the coal, residual demand from power plants, and shipments to industry and export markets:

$$\ln q_{is;t} = \alpha \ln p_{it;s} + \gamma \ln Q_{it;s} + \delta \ln X_{t;s} \quad (2)$$

subject to $q_{it;s} < k_i$ for each $s > t$

where $p_{it;s}$ is the price the mine receives (in dollars per Btu); $Q_{it;s}$ is residual demand for coal from US power plants (defined next); $X_{t;s}$ is demand for coal from the rest of the world, as well as from domestic consumers outside the electricity sector, as noted above; and α , γ , and δ are coefficients.

In principle, we could estimate residual demand by modeling a noncompetitive equilibrium in which each mine accounts for effects of its quantity choice on the quantity choices of the other mines. Such a model would incorporate each mine’s market power. However, such a model would be computationally infeasible given the large number of mines and the state space; existing models, such as that used by Gerarden et al. (2016), contain considerable aggregation across both mines and plants. This aggregation would not be compatible with the mine-level closure analysis that is the focus of this paper. Given this difficulty, we do not explicitly model simultaneous production decisions by all mines and consumption decisions by all electricity plants. Instead, we construct a proxy for the residual demand for the mine’s coal that incorporates information about coal demand from electricity generators, coal supply from other mines, and market power by mines.

More specifically, we define the residual demand as the potential amount of a mine’s coal that coal-fired power plants may purchase, given supply decisions of other mines. Intuitively, factors that can affect residual demand include coal demand from coal-fired power plants, supply from other coal mines, and sulfur dioxide regulation that affects the implicit price on the sulfur content of the coal. For example, a reduction in natural gas prices reduces overall demand for coal from coal-fired power plants, causing a mine’s residual demand curve to shift to the origin. Alternatively, the closure of a nearby mine increases a mine’s residual demand because the mine is competing with fewer other mines. As another example, tightening the national sulfur dioxide emissions cap can raise residual demand if the mine’s coal has a low

sulfur content.

We construct a proxy for residual demand based on estimated coal demand from the electricity sector and a reduced-form representation of market power in mining and shipping. For coal mine i , residual demand equals aggregate coal demand $Q_{t,s}$ net of the coal supplied by all other mines to meet aggregate demand. Aggregate demand is simply the sum of coal consumption across all power plants.

Turning to coal supply from other mines, the supply is a weighted sum of coal consumption by each coal-fired power plant, with weights equal to the share of a plant's consumption from a particular mine, s_{jk} , where j indexes plants and k indexes other mines. For mine i supply from all other mines is the double summation across all coal-fired power plants j , and all mines k other than mine i , of the product of the share and corresponding coal consumption: $\sum_{k \neq i} \sum_j s_{jk} c_{jt,s}$. Therefore, supply by other mines varies across mines and over time. The time-series variation arises from changes in coal demand from coal-fired power plants, interacting with the shares s_{jk} . Thus, residual demand is given by $Q_{it,s} = \sum_j c_{jt,s} - \sum_{k \neq i} \sum_j s_{jk} c_{jt,s}$.

Next, we discuss the interpretation of the residual demand. The share plays an important role in residual demand. It represents the typical shipping patterns between mines and plants, which are the outcomes of market power exercised by mines and shipping firms. In this way we capture in a reduced-form manner the effects of such market power on shipping patterns. That is, we do not model the market power of the market participants directly, but instead use the equilibrium outcomes of this market structure to construct proxies for residual demand.

To illustrate the sources of residual demand variation, consider coal mine i located close to a coal-fired power plant j . The plant purchases all of its coal from the nearby mine. In that case, an increase in generation by plant j causes a proportional increase in residual demand for mine i . As another example, suppose plant j purchases 75 percent of its coal from mine 1 and the remaining 25 percent of its coal from mine 2. In that case, an increase

in consumption by plant j would have a larger effect on residual demand for mine 1 than for mine 2.

We make a few notes about this approach to estimating the present discounted value of future profits. First, we abstract from the shipping market. Shippers may have market power that affects the output prices faced by mining firms (Busse and Keohane 2007). However, it is infeasible to dynamically model coal mine production, shipping, and electricity generation simultaneously at the level of disaggregation needed for the closure analysis. Given this difficulty, we capture the effects of market power of shippers, as well as transportation costs between mines and plants, in the consumption shares. That is, the market shares represent the equilibrium outcomes that result from the market power.

Second, we assume that each mine takes production decisions of other mines as exogenous. The fact that each firm in our data accounts for a small share of aggregate eastern production supports this assumption.

Finally, the model is deterministic. Mines are assumed to remain open if the present discounted value of future profits is positive. In a dynamic stochastic model, mines might remain open even with negative profits if the option value of remaining open is sufficiently high. However, as we show below, the closure regressions should not yield biased estimates as long as mine-specific shocks to the net present value of expected profits are uncorrelated with mine-specific shocks to the option value of remaining open (i.e., shocks to option value affect all mines proportionately in a county or year).

3.1.2 Estimating the demand for a mine's coal

A mine's profits depend on demand from coal plants, demand for exports, and marginal costs. Here, we first outline the estimation of demand for the mine's coal, demand for exports, and a mine's marginal costs. Demand for the mine's coal is a function of the mine's coal price and residual demand after accounting for supply from other mines. Both price and residual demand may be endogenous and we instrument for price using cost shifters, and

we instrument for residual demand using consumption by individual electricity generators predicted from a simulation model combined with a reduced-form representation of shipping markets.

Coal exports depend on global coal prices and domestic consumption. Because domestic consumption may be endogenous, we use the same simulation model to predict domestic consumption as we use for the demand instrument. Finally, marginal costs depend on wages and production. Over the next three subsections, we describe the estimation in more detail.

Next, we provide the estimation details, beginning with the demand equation. The estimating equation for a coal mine’s demand is adapted from Eq. 2. We assume that the same demand equation holds in any year $s > t$, and use t to index years in the estimating equation:

$$\ln q_{it} = \alpha \ln p_{it} + \gamma \ln Q_{it} + \tau_t + \epsilon_{it} \quad (3)$$

In Eq. 3, year fixed effects (τ_t) account for aggregate unobserved shocks to coal production. The error term accounts for measurement error in the dependent variable and the influence of unobserved factors on production. The year fixed effects are co-linear with national exports in Eq. 2, so we do not include exports in the estimating equation. Instead, we estimate export demand separately as we explain below. In the demand equation, each mine’s price and quantity data come from bilateral monthly deliveries from origin mine to destination plant reported to EIA (for certain years we observe the origin county and not the mine, and use county-level prices). The data include confidential data on plant fuel purchases that EIA collects in Form EIA-423/923.

A complication with the EIA data is that the 2002-8 prices reflect only county-level delivered prices (inclusive of the free on board [FOB] mine price plus transportation cost); after 2008, Form EIA-923 prices distinguish between FOB and delivered prices as well as providing a mine identifier. The more complete, post-2008 data are used to estimate transportation costs for each coal-producing county. The average fraction of delivered costs attributable to

transportation is then used to discount the estimated delivered prices and infer FOB prices for 2002-8. This approach is similar to the methodology in Preonas (2017), which shows that there are important cross-sectional differences in transportation costs.

We expect the price coefficient (α) to be negative. The coefficient captures the underlying market structure. For example, if a mine's production has a large effect on the equilibrium price, the price coefficient would be smaller than if the production has a small effect on the price (i.e., closer to perfect competition). We expect the coefficient on residual demand (γ) to be positive.

Both the price and residual demand may be endogenous. Regarding price, for example, a mine with unusually productive workers may be able to sell coal at a lower price than an otherwise identical mine. Failing to measure worker productivity would yield biased estimates of the price coefficient in the equation. A mine's price may also be measured with error. To address this endogeneity and measurement error, we instrument for price using cost shifters: the average productivity at the county level. Productivity is calculated as the number of employee hours required to produce one ton of coal reported by MSHA. Productivity in a given county depends largely on average seam thickness, which in turn is based on historical operations and geologic factors. For robustness, we have also tested alternative instruments, yielding largely similar results (available from the authors by request). We also control for the cost of burning sulfur, which depends on coal sulfur content and current sulfur emissions spot prices.

Recall that residual demand is the demand for a mine's coal from power plants, after accounting for supply from other mines; we expect residual demand to positively affect a mine's production. Residual demand depends on consumption from coal-fired power plants and production from other mines, both of which may be endogenous. Each power plant's consumption of coal, c , is endogenous to the price of coal and may be correlated with other factors that depend on the mine. Moreover, observed production at a mine is mechanically related to observed residual demand as we have defined it, because the mine's production

is used to compute residual demand. Therefore, any measurement error in the mine’s production would bias the residual demand coefficient. Given these issues, rather than using observed consumption to compute aggregate demand, we use the consumption predicted by an hourly unit commitment-style model that covers most US fossil fuel-fired plants (including nearly all coal-fired plants). The appendix outlines the structure of the model and primary data sources, and Linn and McCormack (2017) provide further detail. Here, we summarize the features of the model that are most relevant to our analysis.

In the model, fossil fuel-fired plants are dispatched on an hourly basis over an entire year, taking fuel prices and aggregate demand as exogenous (markets correspond to the three major US interconnections). For each year t , we simulate the model using observed fuel prices and aggregate demand in that year.

The model is based on an economic dispatch model, in which units are dispatched each hour according to their marginal costs, so that low-cost units operate when demand is low. As demand increases, higher-cost units turn on to equate supply and demand. The model departs from a conventional economic dispatch model by approximating constraints on starting up and shutting down units; including minimum generation levels; accounting for transmission congestion in a reduced-form manner similar to Davis and Hausman (2016); and including uncertainty in aggregate demand and unit availability. Linn and McCormack (2017) show that the model reproduces observed unit-level outcomes accurately, and more accurately than a standard dispatch model that does not include these features.

The simulations yield coal consumption by coal-fired power plant and year t . Therefore, fuel prices, electricity demand, and other factors create variation in potential demand over time for a particular coal-fired plant. We compute aggregate demand by year by summing consumption across plants.

We also address the potential endogeneity of coal shipments between mines and plants. Specifically, the shares, s_{ij} , reflect the outcomes of any market power that the mines, shipping companies, or electricity plants may have. That is, we do not model the market power of

the market participants directly, but instead use the equilibrium outcomes of this market structure. Importantly, the shares are computed in the first year of the sample and do not vary over time. This reduces concerns that the shares may respond endogenously to unobserved factors at mines, shipping companies, or power plants. To the extent that market power and transportation costs are persistent over time, one expects the initial consumption shares to be highly correlated with the actual consumption shares during the estimation sample. Using the presample shares introduces measurement error, however, and we explain below how the simulations are constructed to yield unbiased estimates of counterfactual profits. Note further that using the initial shares implies that we do not include mines that enter during the sample; as we argue below, it is unlikely that accounting for entry in the estimation or simulations would substantially affect our results.

In short, rather than introducing bias by using observed coal plant consumption and mine-plant shares, we use the electricity sector model to simulate $c_{js;t}$, and we compute s_{ij} from the data but do not allow the shares to vary during the sample period. We treat as exogenous to coal mines the inputs to the electricity sector simulations, which include natural gas prices and electricity demand.

3.1.3 Estimating equilibrium exports

Recall that exports enter the mine's demand equation, Eq. 2. In the demand estimation, the year fixed effects absorb exports. In principle, we could insert observed US exports in the profit calculation, but doing so would prevent us from adjusting US exports in the counterfactuals considered below. For example, if natural gas prices had been lower than observed levels, we would expect US consumption of coal to be lower than observed. Exports might increase if mines find it more profitable to produce for international rather than domestic consumption. Consequently, we estimate the relationship between equilibrium US exports and US coal consumption, controlling for international prices. We use the following estimating equation:

$$\ln X_t = \lambda_1 \ln P_t + \lambda_2 \ln Q_t + \lambda_3 t + \lambda_4 t^2 + \epsilon_{it} \quad (4)$$

where X_t equals national exports in year t , P_t is the average price of exported coal from the United States, Q_t is US consumption of coal, t and t^2 are linear and quadratic time trends, ϵ_{it} is an error term, and the λ s are parameters to be estimated. Data for export quantities and export prices are from EIA. The export data are available for each port and year, but the port-level data are very noisy. Consequently, we aggregate across ports to compute national exports. The quadratic time trend accounts for nonlinear and time-varying shocks to export demand or supply. One expects US exports to increase with international prices, because higher global prices make it more profitable to produce for international markets rather than the US market. Exports should decrease with US coal demand because, all else equal, higher US demand raises the profitability of producing for domestic consumption rather than for international markets, reducing exports. For these reasons, we expect λ_1 to be positive and λ_2 to be negative.

In Eq. 4, we assume that international coal prices are exogenous to exports. This assumption is supported by the fact that aggregate US exports of steam coal typically account for less than 5 percent of global shipments of steam coal. Consequently, mines and exporters likely take export prices as exogenous. Note that the quadratic time trend controls for the effects of the global macroeconomy on coal demand.

On the other hand, US coal consumption is clearly endogenous to US exports, as unobserved factors (e.g., regulation that raises the cost of coal production) may affect both US coal consumption and equilibrium exports. Given this consideration, rather than the observed consumption, we use the predicted US consumption from the simulation model described above. This variable should be exogenous to exports for similar reasons as for the demand equation discussed above.

3.1.4 Estimating a mine’s marginal costs

When marginal costs are not observed, it is common in the empirical industrial organization literature to estimate marginal costs from the first order condition for price in a firm’s profit maximization problem (e.g., Berry et al. 1995). In our case, although marginal costs are generally not publicly available at the mine level, we do observe cost-per-ton data reported in corporate filings from a selection of large, publicly traded mining companies. We use the information given in the firm’s 10-K filings to estimate a marginal cost function that allows us to impute marginal costs for each mine in the data. Summary statistics are presented in Appendix D, Table 8.⁷ The table shows that the mines included in the cost estimation sample are larger, on average, than the mines in the full MSHA sample. The cost sample data appear to cover the range of production values, however, as the standard deviation is large and the largest mine appears in both datasets. The average wage and start year are similar across the two datasets, as is the proportion of underground mines (relative to surface operations). The mines with the cost data are more dispersed geographically, while the majority of mines (in terms of number of unique mines, regardless of scale) are located in the Appalachian region.

Other important cost drivers, such as the production level (q_{it}), hourly wages, mine type (surface or underground), and region (basin), are reported to the MSHA. We estimate the marginal cost function by merging the two sources of data, and regressing costs per ton on mine characteristics, as specified:

$$\begin{aligned} \ln m_{it} = & (\beta_1 \ln q_{it} + \beta_2 \ln q_{it}^2 + \beta_3 wage_{it}) * Region_i * Type_i \\ & + \beta_4 Company_{it} + \beta_5 Year_t * Region_i + \beta_6 Vintage_i * Region_i + \eta_{it} \end{aligned} \quad (5)$$

where the vector β is to be estimated. $Wage$ is the county-level wage rate for mining industries; $Type$ is an indicator variable denoting whether the mine is surface, underground,

⁷The companies included in the cost dataset are Arch Coal, Alliance Resource Partners, Consol Energy, Peabody Energy, and Westmoreland Coal. This set of companies provides a diverse sample of regions and mine types, but omits smaller, privately owned mines.

or both; *Region* is an indicator variable for the region or basin in which the mine operates (values are Appalachian, Interior, Lignite, Powder River Basin, or Western); *Company* is a company fixed effect controlling for management efficiency effects; *Year* is a time fixed effect; *Vintage* is the year the mine opened; and η is an error term. Note that the fact that wages are computed at the county rather than mine level reduces concerns that the estimates may be biased by mine-specific productivity shocks or bargaining power in labor markets. *Vintage* plays an important role because it allows for the possibility that production costs increase as firms mine low-cost coal before high-cost coal. As in Eq. 3, wage data come from the BLS QCEW. Regional definitions are adopted from Stoker et al. (2005). After estimating this equation, we assume that the coefficients are the same for mines that do not belong to the estimation for sample as for mines that do belong to the sample. This assumption allows us to extrapolate the estimates and predict marginal costs for all mines in the MSHA data. The estimates from estimation of Eq. 5 are presented in Table 9 in Appendix D.

3.2 Profitability and Mine Closure

The key independent variable in the closure regression is the present discounted value of expected profits by mine and year. To estimate the present discounted value of profits, we simulate the electricity sector model, which we used to construct the residual demand instrument, over all years $s > t$ using EIA projections of fuel prices and aggregate demand. For example, for $t = 2001$, we use the EIA projections from the 2001 Annual Energy Outlook for $s = 2001$ through 2020. Henceforth, we refer to s as the projection year. For any particular hour in year s , the aggregate demand equals the demand in the same hour in year t , multiplied by the EIA projected annual growth rate between years t and s . We use these simulations to compute residual demand for each mine i , year t , and projection year s .

Profits for each mine, year, and projection year are calculated by using the fitted values of Eqs. 3, 4, and 5 into Eq. 1. More specifically, we invert Eq. 3 to express the mine's price as a function of quantity, net aggregate demand, and other factors. We used observed global

coal prices and aggregate domestic consumption to compute exports.

Equation 2, the coal demand equation, includes the parameter δ , which is the mine's share of exports. We estimate this parameter by starting with the difference between a mine's annual reported MSHA production and reported Form EIA-923 sales. This difference is our initial guess of the mine's exports, but it may include measurement error because we do not directly observe the mine's exports. To attempt to reduce the measurement error, we account for a mine's proximity to US ports from which coal is exported. Specifically, we observe exports from each port, and for each mine we construct a weighted average of US exports using the inverse of the distance between the mine and port as weights. For example, a mine located close to a major eastern port has a higher value of the export index than a mine located farther from the port. We regress the mine's estimated exports on this index, and we use the fitted values from this prediction to estimate the mine's share in national exports. The predicted shares are consistent with expectations; for example, East Coast mines located close to major export ports have larger predicted shares than mines located further inland. For further discussion of the modeling of coal imports and exports see Appendix C.

Using the estimated parameters and observed exogenous inputs, we maximize profits by mine, year, and projection year, and compute the sum of discounted profits by mine and year. We use a real discount rate of 10 %, which represents a typical cost of capital for mines during the sample.

We use the estimated present discounted value of profits by mine and year to estimate the effect of profitability on closure. Given fixed annual production capacity, expected market conditions, and reserve-dependent costs, increased profitability should extend the life of the mine, reducing the probability that it closes in a given year.

A survival regression is employed to empirically model the relationship between profitability and a mine's closure decision. Specifically, the analysis utilizes a Cox proportional hazard model that takes the form

$$P[Close_{it}|TimeSinceOpen, X_{it}] = h_0(t)exp(BX_{it}) \quad (6)$$

where the probability that mine i closes in year t depends on the duration the mine has been open and a set of covariates X_{it} , the most important of which is the mine's profitability. Our estimation excludes mine-year observations where the expected profit or loss is more than \$1 billion; this cutoff causes us to drop 0.8 percent of the sample and reduces the effects of outliers on the coefficient estimates. Some of the reported specifications of Eq. 6 include county fixed effects in X_{it} . The unspecified baseline hazard function, $h_0(t)$, represents the underlying probability of closure. This term plays an important role in the interpretation of the estimation results and counterfactuals, as it represents an estimate of the baseline level of closure that arises from factors not included in the independent variables in X_{it} . For example, the term would include closures caused by nonlinear effects of mine age on marginal costs or by changes in shipping costs (recall that the consumption shares in the residual demand do not vary over time). As we noted in the Introduction and show below in our data, the typical closure rate for coal mines is rather high.

Note that the key independent variable is the sum of profits in years $s > t$, discounting profits in each year s back to year t . The present discounted value of a mine's profits in year t could be positive even if profits in particular years $v > t$ are negative. We expect that in such cases the mine remains open because it is the discounted value of profits driving the closure decision, and not the profits in any particular year.

Information on the characteristics of mines in our sample, both those that closed and those that remained open, is given in Table 3. The table reveals that there are many fewer PRB coal mines than Appalachian mines. This is because coal in the PRB is located close to the surface with thick seams, making it easier to capitalize on scale economies. Eq. 6 is estimated for Appalachian mines only; the analysis includes all mine closures observed between 2002 and 2012.

The data for mine closure comes from the MSHA Part 50 data. The data include a

status code for each mine that states whether the mine is in operation, temporarily closed, or permanently abandoned. We specify a mine as closed in the year that the status is recorded as permanently abandoned or abandoned and sealed. In principle, one could define closure in other ways, such as based on a mine’s employment. We provide evidence that this specification of closure is correct in Appendix B by showing that production and employment are close to zero after a mine closes. The vast majority of mines fall out of the sample after their observed closure date, which provides additional confirmation of the validity of this measure.

Table 3: Summary Statistics for Mines in the Sample

	By End of Sample		Overall
	Closes	Remains open	
Mean Production (tons/qtr)	65,071	328,831	198,819
Average Workers per Mine	28.4	63.9	45.9
Average Profit	303,334	458,418	376,338
% Surface Mines	46.9	61.8	54.3
Number of Mines	2,220	1,200	3,420
Appalachian Mines	2,086	1,049	3,135
PRB Mines	4	17	21
Observations	27,933	28,582	56,515

Note: Mine summary statistics constructed from MSHA Part 50. Profits estimated using method described in Section 3.1.2.

Before presenting the estimation results, we discuss briefly the implications of using baseline estimated profits rather than actual (i.e., unobserved) profits in Eq. 6. Using estimated profits results in measurement error and a different coefficient than if we could hypothetically observe profits and include them in the estimation. However, the counterfactuals considered below involve changes to estimated profits. We use Eq. 6 to estimate the effects of the counterfactual profits on closures and compare with predicted closures using the baseline estimated profits. This approach should yield unbiased estimates of the baseline and counterfactual closures.

4 Results

4.1 Estimation of Demand, Export, and Marginal Cost Equations

Profits are calculated for each mine based on its profit maximization problem from Eq. 1. The demand, exports, and marginal cost equations (Eqs. 3, 4, and 5) are inputs into the profit maximization problem, and we first report results from these equations. Table 4 shows the results from the estimation of Eq. 3, and Figure 6 in the main text and Table 9 in Appendix D shows the result of estimation of Eq. 5. In Table 4 the coefficients on price and residual demand both have the expected sign. The ordinary least squares estimate of the price coefficient is likely to be biased for reasons discussed above. The price coefficient is negative in the specification that instruments for price. The table reports the first-stage F-statistic, which reduces concerns about weak instruments bias. Table 10 in Appendix D shows the first stage regression results, where the coefficient on the instrument has the expected sign and is statistically significant. In Table 4 the coefficient on residual demand is positive, as expected, and is statistically significant. The regression also includes the interaction of sulfur content with the acid rain permit price (and exogenous demand shifter). The negative sign is consistent with expectations, and suggests that an increase in the sulfur price has a larger negative effect on production from mines with higher sulfur content than other mines.

Figure 6 compares predicted and observed marginal costs to assess the fit of the marginal cost estimation (Table 9 in Appendix D reports the underlying coefficient estimates). The figure shows that our predictions of marginal costs are quite close to reported costs.

Table 5 shows the results of estimating Eq. 4. The main coefficient of interest is the coefficient on domestic consumption, because domestic consumption varies across the counterfactuals discussed below. The coefficient is precisely estimated and fairly stable across the specifications shown in the table, which differ by whether we include a quadratic or cubic time trend, take first differences, or include quarter fixed effects. The coefficient on

the export price varies across specifications, in both sign and magnitude, but this coefficient does not play a direct role in the counterfactuals.

We use the cost, profit-maximizing quantity, export elasticities, and corresponding price to compute each mine's profits by year and projection year, over an 18-year horizon. We discount future profits using a 10% discount rate, yielding present discounted profits by mine and year.

The resulting profits approximate unobserved profits over the period of study. To assess the validity of the profit estimates, for the median Appalachian mine, Figure 7 compares the estimated profits with observed variables that may be correlated with unobserved profits, at least over certain sub-periods of the data. The upper panel plots the FOB price of coal sold under contract (rather than on the spot market) and the estimated production cost per ton of coal mined and sold. We use the contract price in this analysis because approximately 78 percent of coal sold during the period of study was sold under contract; this fraction trended upward over the study period. The FOB price is calculated by discounting the estimated transaction price by the average share of transportation costs in total costs for coal leaving the corresponding county, yielding an estimated mine-mouth price. The distinction between contract and spot prices is important in this context because, although contract and spot prices were generally highly correlated with one another through 2008, contract prices did not experience the same dramatic collapse following the 2008 financial crisis that spot prices did. The contract price rises steeply in the first half of the sample period before leveling off. Figure 7 also shows that production costs rose steadily over the sample, with a pronounced upturn in 2012.

The middle panel plots current annual profits (in 2012 US\$100,000) and expected discounted profits (in 2012 US\$1,000,000); note the order-of-magnitude difference in the series. The two profit measures are highly correlated with one another, and both peak in 2008 and drop sharply after 2010.

One expects the present discounted value of profits to be highly correlated with stock

prices. The bottom panel plots the relative change in stock prices over the period for four large coal-mining firms and the average of these changes. While the profit estimates and stock prices follow a similar trend pre-2008, they diverge post-2008. Two reasons are likely responsible. First, in 2011 most publicly traded coal-mining companies used debt to purchase metallurgical coal deposits at the peak of a metallurgical coal boom. As the price of metallurgical coal fell after 2011, those metallurgical coal deposits failed to bring in revenues to cover the debt obligations (Roberts, 2016). Second, many of these companies also own Australian coal-mining assets that sell coal into the Pacific market. The Pacific coal market has not seen a reduction in price similar to that in the United States and this has helped keep stock prices higher than if their coal mines were all located in the United States. Thus, we do not observe a perfect correlation between estimated profits and stock prices, and the correlation is notably weaker after 2008. After 2008, profits decline sharply because of the decrease in coal demand caused by declining natural gas prices and electricity demand growth.

Table 4: Demand Function Estimation

	<i>OLS</i>	<i>Price IV</i>
	(1)	(2)
ln Price	0.467** (0.221)	-1.116* (0.580)
ln Residual Demand	0.740*** (0.056)	0.712*** (0.066)
Sulfur Cost (\$/ton)	-0.00003 (0.0001)	-0.0002 (0.0001)
Observations	1,266	1,266
Adjusted R ²	0.642	0.577
First Stage F-stat		311.83

Note:

*p<0.1; **p<0.05; ***p<0.01
Errors clustered at the county level. 159 clusters.

Figure 6: Fitted vs. Actual Costs





Figure 7: Prices, Costs, and Profitability Estimates in the Baseline Results for Median Appalachian Mine



Top panel: Median free on board (FOB) price of transactions under contract and the estimated production costs for Appalachian mines. Free on board prices are calculated by discounting the reported delivered price by average transportation costs from the Appalachian basin.

Middle panel: Annual, contemporaneous profits (in hundreds of thousands of dollars) and expected discounted profits (in millions of dollars) for the median Appalachian mine as calculated by this study.

Bottom panel: Annual average stock prices for four large, public US coal firms and their combined average price over the study period.

Table 5: Export Elasticity Results

Outcome: ln Exports	(1)	(2)	(3)	(4)
ln Export Price	0.44*** (0.12)	0.10 (0.19)	-0.12 (0.31)	0.48*** (0.12)
ln Sim. Domestic Cons.	-0.53** (0.24)	-0.42* (0.25)	-0.43*** (0.17)	-0.47** (0.23)
Trend Poly Order	2	3	0	2
Quarter Fixed Effects	No	No	No	Yes
First Difference	No	No	Yes	No
Observations	56	56	56	56

Note: *p<0.1; **p<0.05; ***p<0.01

4.2 Closure Results

Results of the estimation of Eq. 6, which links estimated profits and closures, can be found in Table 6.⁸ Columns (1) through (2) show estimates of the hazard model, and columns (3) and (4) from a linear probability model for comparison. All columns include profits and profits squared while columns (2) and (4) adds county fixed effects to the model. In column (1), the estimates imply an elasticity of closures to profits of -0.32, which is reported at the bottom of the table. Column (2) reports a slightly larger elasticity, and suggests that the county fixed effects control for unobserved variation in closure rates across counties. Because of that, we use column (2) as the preferred model. For comparison with the preferred model in column (2), we also report results assuming a non-linear relationship between profits and closure in columns (3) and (4). In Figure 8 we plot predicted closure rates resulting from the hazard model and the actual closure rates, by mine opening year (vintage). The model fits the data reasonable well across the most frequently observed mine ages (5-35 years). However, the model under-predicts closure early in a mine's life, and over-predicts closure for the few mines between 35-40 years old (albeit with large measured error).

⁸Additionally, regressions including the amount of economically recoverable reserves at the mine were undertaken while one of the authors had a nondisclosure agreement for Form EIA-7A. One reason a mine could close is that the deposit no longer has economically recoverable reserves and our estimated profits do not include information about estimated reserves. In other words, our marginal cost estimates depend on the mine's age, but declining reserves could cause costs to increase nonlinearly with age, which would affect closure independently of estimated profits. The coefficient on recoverable reserves had the expected sign (negative) while the magnitude of the modeled profit coefficient was not different in economic significance. Unfortunately, we are not able to include reserves in the mine closure regressions reported in the text.

Table 6: Mine Closure Results

Model	Hazard (1)	Hazard (2)	OLS (3)	OLS (4)
Modeled Profit (billion \$)	-1.154*** (0.333)	-1.925*** (0.351)	-0.0569*** (0.0186)	-0.112*** (0.0197)
Modeled Profit Squared	-2.998*** (0.770)	-2.759*** (0.752)	-0.117*** (0.0290)	-0.0764*** (0.0295)
County Fixed Effect	No	Yes	No	Yes
Closure Rate/ Profit Elasticity	-0.32	-0.35	NA	NA
Observations	10,679	10,679	10,679	10,679
Number of Mines	2,528	2,528	2,528	2,528
Observed Closures	941	941	941	941
Adjusted R ²			0.006	0.022

Note: *p<0.1; **p<0.05; ***p<0.01


Mine closure as a function of modeled profts. Estimation using Cox proportional hazards model shown in columns (1) and (2) and OLS linear probability model shown in columns (3) and (4).

4.3 Effects of Market Conditions and Policy on Coal Mine Closure

We assess the impact of four counterfactual demand and supply scenarios on the profitability and closure of Appalachian coal mines, as compared with the closures predicted by our baseline model.

The first scenario considers the reduction in electricity consumption after 2005, relative to expectations in 2005. In the 2005 Annual Energy Outlook (AEO), the EIA projected approximately 2 percent annual consumption growth over the subsequent two decades. Actual consumption growth was close to zero between 2005 and 2015, due to the 2008-09 economic recession and other factors, such as the adoption of energy efficiency. The counterfactual uses projected consumption from the AEO 2005 rather than from the AEO's 2006-12, which are used in the baseline. Therefore, the counterfactual consumption growth is higher than that used in the baseline. The higher consumption growth in the counterfactual should increase discounted profits, reducing counterfactual closure relative to the baseline. Comparing the baseline and counterfactual identifies the effects of lower-than-expected electricity consumption growth on mine closure.

Figure 8: Estimated vs. Observed Closure Rates, by Mine Age



Note: Plotted line shows the average observed closure rate by mine age. Plotted points show the estimated closure rate modeled from the hazard model shown in column (2) of Table 6. Vertical lines represent the 95% confidence interval for the estimated mean closure rate.

The second scenario considers the decrease in natural gas prices that occurred after 2008. Coal and natural gas are substitutes as fuels for electricity generation. Numerous studies document the short-run and long-run reduction in coal consumption caused by the decline in natural gas prices after 2008 (e.g., Linn and Muehlenbachs (2018), Cullen and Mansur (2017), Fell and Kaffine (forthcoming), and Holladay and LaRiviere (2017)). However, there has been limited empirical evidence linking low gas prices to Appalachian coal production; this paper provides the first substantiation of such an effect. To model this effect, rather than using the projected fuel prices from the AEO in each year between 2009 and 2012, we instead use the AEO 2008 to project prices in each year after 2008. Effectively, this scenario supposes that projections of fuel prices after 2008 were the same as projections made in 2008. For example, when we estimate baseline profits in year 2009 and projection year 2012, we use the AEO 2009 for the projected prices in 2012. In the counterfactual, we use the (higher) AEO 2008 projected prices for 2012. The higher natural gas prices in the counterfactual are expected to raise coal demand and increase mine profits, reducing closure.

The third scenario considers the dramatic increase in Appalachian mining cost between

2002 and 2012. As described in Section 2.3, productivity declined substantially in the Appalachian region over the study period. We expect such declines to result in higher costs, reduce profits, and increase closure. In this counterfactual, we assume that productivity for each mine was equal to the maximum observed level for the mine’s entire life. For most mines, the maximum occurs near the beginning of the sample. Comparing the baseline and counterfactual predicted closures shows the effects of the observed productivity decline on mine closure.⁹

The fourth scenario addresses the impact of federal coal leasing policy on Appalachian coal mine profits and closure.¹⁰ Because most production of federally leased coal comes from the PRB region, the scenario focuses on this area. We use the model to predict the effects of a hypothetical policy on mine profits and closure, due to the absence of variation in federal policy during the sample period. Because the PRB mines operate on very thin margins (less than \$1/ton), small increases in costs could cause the mines to become unprofitable and close or not renew their leases. In this scenario, we consider a hypothetical policy that raises PRB production costs sufficiently to close the 10 PRB mines that have the highest costs (these have costs above \$12 per ton).

The three residual demand scenarios (high gas prices, high electricity consumption, lower competing supply) are constructed by manipulating the residual demand variable in the demand equation, Q in Eq. 3. For the gas price and electricity demand scenarios, the electricity model is re-simulated using the counterfactual consumption or gas prices, resulting in counterfactual electricity consumption by coal-fired plant and projection year. These counterfactual consumption levels are used to compute counterfactual residual demand. The natural gas price scenario has a large effect on an individual mine if a large share of that mine’s coal is sold to coal-fired electricity generators that are in close competition with gas-fired generators.

⁹Alternatively, we could set counterfactual productivity for each mine equal to its value in 2002 for the same mine. We note that the results are qualitatively similar if we take that approach instead.

¹⁰For this scenario we assume that Interior basin mines do not alter production. While this is a strong assumption it helps highlight how much competition there is between the Western and Appalachian basins.

For the PRB leasing scenario, residual demand is recalculated assuming that $s_{ij} = 0$ for all i, j associated with the PRB mines selected for closure. In other words, if the highest-cost PRB mine is mine $i = 1$, closing this mine yields $s_{1j} = 0$ for all j . Closing $i = 1$ and $i = 2$ yields $s_{1j} = 0$, and $s_{2j} = 0$, and so on. We then rescale the shares so that they sum to one for each plant, which means implicitly that when a PRB mine closes coal plants increase consumption from other mines proportionately to their pre-closure consumption from those mines. This calculation raises residual demand for all Appalachian mines, relative to the baseline. This effect is relatively large for mines that sell to coal-fired plants that purchase a large share of coal from PRB mines in the baseline. In the counterfactual, those plants have to purchase from Appalachian mines instead, increasing residual demand, and by a different amount for each mine. However, by construction, all mines experience some increase in their residual demand regardless of whether they supply directly to affected plants as overall market supply has contracted.

Because the counterfactuals affect different components of discounted profits, the results depend on different aspects of the estimation. For example, the consumption, natural gas, and PRB counterfactuals affect residual demand, because of which hypothetically changing the residual demand coefficient in Eq. 3 would affect the results of these scenarios but not the cost scenario. In contrast, the third scenario is constructed by using the counterfactual costs in the profit calculations.

Figure 9 presents the results of the three counterfactual scenarios and the baseline. The upper panel compares coal prices under the baseline and four counterfactuals, while the lower panel presents total discounted profits (the figure also presents a PRB scenario, which we describe below). The electricity consumption and gas price scenarios have noticeable impacts on prices and profits. Both the consumption and gas price scenarios raise coal prices above the baseline levels, reflecting the higher coal demand. By the end of the sample period, the gas price counterfactual has a larger effect on prices than does the consumption counterfactual. Both counterfactuals cause profits to be higher than baseline profits, particularly


by the end of the sample. The productivity scenario reduces coal prices because the higher productivity increases production. The supply scenario has a larger effect on profits than the two demand scenarios (except at the very end of the sample).

Next, we calculate the number of mines that would have remained in operation in each scenario. Because we observe employment at each mine, we use the avoided closures in each scenario to estimate avoided layoffs. Table 7 presents results for the four counterfactual scenarios. We use the closure model fitted to the baseline profits to estimate baseline closures and employment losses. We use counterfactual profits and the estimated coefficients to estimate counterfactual closures and employment losses, and then compare these levels with the baseline levels for the entire 2002-12 period. By construction, the baseline and consumption counterfactual (Dmd 05) levels equal one another for the years 2002-05, and the baseline and gas counterfactual (Gas 08) levels equal one another for the years 2002-08. The baseline and productivity counterfactual (Prod Max) levels can differ from one another in each year. Because of this setup, the cost counterfactual could have larger effects than the other counterfactuals because it occurred earlier in time. Nonetheless, we compare the counterfactuals over the entire 2002-12 period because we are interested in the effects of each shock on closures and employment over the entire period, not in any specific year.

In the first row of Table 7, we present the estimated change in the hazard rate for the average mine in each scenario. Multiplying the estimated hazard reduction to a mine's life yields the number of additional years a mine would operate under each scenario. The table presents average number of additional years across mines. The scenarios add between 0.27 and 2.17 years of mine life. Adding these years to the observed closure date, we are able to determine whether the additional profitability would have been enough to defer closure to some date after the end of our study period, 2012. We estimate that our scenarios would have deferred closure for between 125 mines (in the consumption scenario) to 274 mines (in the maximum productivity scenario).

For reference, we observe the closure of 1,699 mines between 2002 and 2012. It is impor-

Figure 9: Appalachian Mines' Profitability and Trends by Scenario and Year



Top panel: Median FOB contract prices for Appalachian mines by scenario. Base is the baseline result, Prod. Max assumes mines operated at their observed maximum productivity for the entire period, Dmd 05 assumes demand had grown at the rate projected by EIA's AEO 2005, Gas 08 assumes that the AEO 2008 natural gas price forecast had been correct (essentially eliminating natural gas price declines from the "fracking" boom), and PRB assumes the 10 highest-cost PRB mines had closed in 2002.

Bottom panel: Median estimated annual expectations for discounted future profitability of Appalachian mines by scenarios described above.

tant to note that many of these mines may close for reasons other than the factors included in estimated profits, such as the depletion of reserves, regulatory compliance, or physical damage to the mine. From our survival model, it turns out that 1,287 mines closed independently of changes in profits during our sample period. While this number may seem high, it is consistent with typical closure rates observed prior to our sample. Therefore, it is most relevant to compare the number of avoided closures in the counterfactual scenarios with the 412 mines that closed due to declining profits in our sample period. Based on this comparison, each of the demand counterfactuals (in isolation) caused roughly one-third of the mines to close, and the supply shock caused the closure of two-thirds of the mines.

Furthermore, with the observed average employment levels at these mines, we are able to determine how many layoffs might have been deferred or avoided in each counterfactual scenario. Depending on scenario, this figure ranges between 4,551 and 15,023. By comparison, there were about 75,000 coal mine employees in 2002 in the United States. Note that these estimates do not incorporate mine openings (which might increase in the counterfactuals) or shifting of workers from one mine to another; we discuss the issue of openings below.

Table 7: Implied Avoided Mine Closure and Layoffs by Scenario

	Prod. Max	Gas 08	Dmd 05	PRB 10
Decrease in Hazard	0.15	0.06	0.07	0.02
Avg Additional Years	3.2	0.4	0.7	0.4
Closures Avoided	352	142	157	91
Layoffs Avoided ¹	15,023	6,294	6,883	4,551

¹Layoffs not reflective of overall employment as they do not reflect hiring
Note: Prod. Max assumes mines operated at their observed maximum productivity for the entire period, Dmd 05 assumes demand had grown at the rate projected by EIA’s AEO 2005, Gas 08 assumes that the AEO 2008 natural gas price forecast had been correct (essentially eliminating natural gas price declines from the “fracking” boom), and PRB assumes the 10 highest-cost PRB mines had closed in 2002.

Figure 9 and Table 7 show that the PRB scenario has a smaller effect on coal prices, mine profits, closures, and employment than the other scenarios. Therefore, although a hypothetical policy that raises production costs on federal lands and leads to the closure of

a substantial number PRB mines would benefit Appalachian mines, the effect is relatively small. This result reflects the fact that PRB and Appalachian coal are not close substitutes for one another,¹¹ given transportation costs and differences in characteristics.

We note that we do not model coal mine openings in the counterfactuals. Although we do not model mine openings explicitly, the information in Table 1 provides a rough approximation of how many openings our model may be missing. Throughout the 2000s, when coal prices were relatively high, the Appalachian region saw around 110 openings per year. From 2010 to 2012, there were only 75 openings per year. Thus around 150 mines would have been added to our sample of around 2,800 mines operating from 2010-2012. These mines would have taken some of the increased residual demand in each scenario that we now apportion to the existing mines, which would reduce the profits shown in Figure 9. However, the change in openings would be an order of magnitude smaller than the change in closings reported above, suggesting that the profitability shocks primarily affect closings rather than openings. Moreover, even if the counterfactuals did affect openings, they would affect total coal mine employment only to the extent that labor productivity differs between opening and closing mines.


Similarly, we do not explicitly model the impact of natural gas prices on coal demand that occurs through the change in transportation costs as discussed in Preonas (2017). He shows that increases in the price of natural gas would lead to increases in coal prices that are not as large as would have been seen with a competitive transportation sector. Since this effect is not modeled, our gas price counterfactual profit number is larger than it would be had the transportation sector been modeled explicitly. However, the size of the impact given in Preonas (2017) implies that the change in profits would be relatively small.

The above scenarios provide insight to the relative importance of electricity consumption, natural gas prices, mine worker productivity, and PRB leasing in explaining Appalachian mine closures. However, the mine closures that appear to be unexplained by expected profits

¹¹More precisely, we observe relatively few power plants that purchase substantively from both PRB and Appalachian mines in 2002, the year consumption shares (s_{jk}) are determined.

in Table 7 naturally leads to the question, how high would profits need to be in order to avoid closure of a larger number of mines? Using the parameter estimates fitted from Eq. 6, we arbitrary increase each mine's observed profits from 0 to \$10 million by \$0.1 million steps. From the computed change in each mine's chance of survival, we determine whether the given profit increase would be enough to defer closure to some time outside of our sample period. The profits required for a given number of mines to avoid closure in this way are plotted in Figure 10. For reference, the number of mines avoiding closure in the maximum productive scenario is highlighted on the figure. An important takeaway of Figure 10 is that the profit-closure response becomes relatively inelastic after approximately 500 mines have avoided closure, each having experienced around a \$7.5 million increase in discounted profits. To induce additional reductions in closures, more substantial increases in profits would be required. This result implies that constructing counterfactuals that include multiple shocks (such as natural gas prices and consumption), would not substantially increase the total number of closures. For this reason, we construct the counterfactuals by considering each shock one at a time, rather than combining the shocks.

Figure 10: Increase in Profits Required to Avoid Mine Closures



5 Conclusion

The fate of coal-mining communities continues to receive attention from politicians and the popular press. Although there has been heated debate at all levels of government on policies affecting coal mining, little is known about the recent causes of coal mine closures or about the effects of hypothetical future policies on the sector.

This paper models a coal mine's closure decision as a function of its expected profitability. This approach allows us to simulate how those profits and closure decisions would change under a number of counterfactual scenarios. We find that rising production costs explain about two-thirds of the observed coal mine closures caused by declining profits between 2002 and 2012. Natural gas prices and reduced electricity consumption independently explain about one-third of the closures.

While this paper has addressed four potential causes for the decline in Appalachian

coal mining, our modeling framework is appropriate for the evaluation of other policies and market forces that may have affected this industry and region. These include the labor to capital substitution that has occurred within and between coal basins (especially since 1990), emissions regulations directed at coal-fired plants, and environmental regulation directed at coal-mining operations, particularly restrictions around eastern US surface mining (so-called mountain top removal).

Another important area of future research is better understanding the environmental, economic, sociological, and health impacts from the continued decline of an industry that has defined life in rural Appalachia for a century.

References

- Dan Black, Terra McKinnish, and Seth Sanders. The economic impact of the coal boom and bust. *The Economic Journal*, 115(503):449–476, 2005.
- Richard Bonskowski. The u.s. coal industry in the 1990s: Low prices and record production. Energy Information Administration, 2000.
- Jacob Brinkman, Jay Schafler, and Catherine Johnson. A study of us coal mine closures since 1994. *Mining Engineering*, 69(10):18–37, 2017.
- Bureau of Land Management. Blm coal resources: A brief overview. Technical report, 2012.
- James Bushnell. Oligopoly equilibria in electricity contract markets. *Journal of Regulatory Economics*, 32:225–245, 2008.
- Meghan R. Busse and Nathaniel O. Keohane. Market effects of environmental regulation: Coal, railroads, and the 1990 clean air act. *Rand Journal of Economics*, 38(4):1159–1179, 2007.
- Anya Castillo and Joshua Linn. Incentives of carbon dioxide regulation for investment in low-carbon electricity technologies in texas. *Energy Policy*, 39(3):1831–1844, 2011.
- Yin Chu, Scott Holladay, and Jacob LaRiviere. Pass-through from fossil fuel market prices to procurement costs of the u.s. power producers. *Journal of Industrial Economics*, 65: 842–871, 2017.
- Joseph Cullen and Erin Mansur. Inferring carbon abatement costs in electricity markets: A revealed preference approach using the shale revolution. *American Economic Journal: Economic Policy*, 9(3):106–133, 2017.
- Lucas Davis and Catherine Hausman. Market impacts of a nuclear power plant closure. *American Economic Journal: Applied Economics*, 8(2):92–122, 2016.

Department of Energy. Staff report to the secretary on electricity markets and reliability, 2017. Accessed on: 2018-1-10.

Department of Interior. Secretary Jewell launches comprehensive review of federal coal program, 2016. URL <https://www.doi.gov/pressreleases/secretary-jewell-launches-comprehensive-review-federal-coal-program>. Accessed on: 2018-1-10.

Stratford Douglas and Anne Walker. Coal mining and the resource curse in the eastern United States. *Journal of Regional Science*, 57(4):568–590, 2016.

Timothy Dunne and David Merrell. Gross employment flow in U.S. coal mining. *Economics Letters*, 71(2):217–224, 2001.

A. D. Ellerman, P. L. Joskow, R. Schmalensee, J.-P. Montero, and E. M. Bailey. *Markets for Clean Air: the U.S. Acid Rain Program*. Cambridge University Press, New York, 2000.

Energy Information Administration. Coal production 1989. Technical report, 1990.

Energy Information Administration. Coal industry annual 1993. Technical report, 1994.

Energy Information Administration. Coal industry annual 1997. Technical report, 1998.

Energy Information Administration. Coal industry annual 1999. Technical report, 2000.

Energy Information Administration. 2014 annual coal report. Technical report, 2015a.

Energy Information Administration. Electric power annual 2013, 2015b. URL <https://www.eia.gov/electricity/annual/archive/03482013.pdf>. Accessed on: 2017-12-28, Table 2.2.

Energy Information Administration. May 2017 monthly energy review - table 6.1 coal overview, 2017a. URL <https://www.eia.gov/totalenergy/data/browser/index.php?tbl=T06.01#/?f=A>. Accessed on: 2017-06-27.

- Energy Information Administration. May 2017 monthly energy review - table 7.2a electricity net generation: Total (all sectors), 2017b. URL <https://www.eia.gov/totalenergy/data/browser/index.php?tbl=T07.02A#/?f=A>. Accessed on: 2017-06-27.
- Energy Information Administration. Henry hub natural gas spot price, 2017c. URL <http://tonto.eia.gov/dnav/ng/hist/rngwhhda.htm>. Accessed on: 2017-07-02.
- Jonathan Eyer and Matthew Kahn. Prolonging coal's sunset: The causes and consequences of local protectionism for a declining polluting industry. National Bureau of Economic Research Working Paper 23190, 2017.
- Harrison Fell and Daniel Kaffine. The fall of coal: Joint impacts of fuel prices and renewables on generation and implications for policy. *American Economic Journal: Economic Policy*, forthcoming.
- Todd Gerarden, W. Spencer Reeder, and James Stock. Federal coal program reform, the clean power plan, and the interaction of upstream and downstream climate policies. National Bureau of Economic Research Working Paper 22214, 2016.
- Kenneth Gillingham, James Bushnell, Meredith Fowlie, Michael Greenstone, Charles Kolstad, Alan Krupnick, Adele Morris, Richard Schmalensee, and James Stock. Reforming the u.s. coal leasing program. *Science*, 354(6316):1096–1098, 2016.
- Government Accountability Office. Coal leasing. Technical report, 2013.
- Gautam Gowrisankaran, Charles He, Eric Lutz, and Jefferey Burgess. Productivity, safety, and regulation in underground coal mining: Evidence from disasters and fatalities. National Bureau of Economic Research Working Paper 21129, 2015.
- Scott Holladay and Jacob LaRiviere. The impact of cheap natural gas on marginal emissions from electricity generation and implications for energy policy. *Journal of Environmental Economics and Management*, 85:205–227, 2017.

- Trevor Houser, Jason Bordoff, and Peter Marsters. Can coal make a comeback? Columbia University Center on Global Energy Policy, 2017.
- Grant Jacobsen and Dominic Parker. The economic aftermath of resource booms: Evidence from the boomtowns in the american west. *The Economic Journal*, 126(593):1092–1128, 2016.
- Reid Johnson, Jacob LaRiviere, and Hendrik Wolff. Fracking, coal, and air quality. IZA Discussion Paper 10170, 2017.
- Paul Joskow. The performance of long-term contracts: Further evidence from coal markets. *Rand Journal of Economics*, 21(2):251–274, 1990.
- Alan Krupnick, Joel Darmstadter, Nathan Richardson, and Katrina McLaughlin. Putting a carbon charge on federal coal: Legal and economic issues. Resources for the Future Discussion Paper 15-13, 2015.
- Joshua Linn and Kristen McCormack. The roles of energy markets and environmental regulation in reducing coal-fired plants profits and electricity sector emissions. RFF Report, 2017.
- Joshua Linn and Lucija Muehlenbachs. The heterogeneous impacts of low natural gas prices on consumers and the environment. *Journal of Environmental Economics and Management*, 89:1–28, 2018.
- Mine Safety and Health Administration. Coal employment by state and county cy09to15, 2016. URL https://www.msha.gov/sites/default/files/Data_Reports/Charts/Coal_Employment_by_State_and_County_CY09to15.pdf. Accessed on: 2017-06-27.
- Office of Inspector General . Coal management program. Technical report, 2013.
- Office of Natural Resource Revenue . Statistical information. Technical report, 2015.

Louis Preonas. Market power in coal shipping and implications for u.s. climate policy. Energy Institute at Hass Working Paper 285, 2017.

David Roberts. The u.s. coal industry is falling apart. here's the surprising reason why, 2016. URL <https://www.vox.com/2016/2/22/11090878/us-coal-industry-falling-apart>.

Thomas Stoker, Ernst Berndt, A. Denny Ellerman, and Susanne Schennach. Panel data analysis of u.s. coal productivity. *Journal of Econometrics*, 127(2):131–164, 2005.

The National Institute for Occupational Safety and Health. Number of active coal mines by year, 2006-2015, 2016. URL <https://www.cdc.gov/niosh/mining/UserFiles/statistics/15g04aac.svg>. Accessed on: 2017-07-02.

Frank Wolak. Measuring unilateral market power in wholesale electricity markets: The california market, 1998-2000. *American Economic Review*, 93(2):425–430, 2003.

A Overview of Hourly Generator Operation Model

This appendix outlines the model used to estimate coal consumption by coal-fired power plant, year, and projection year. We describe the model and then the data. Linn and McCormack (2017) provide additional details about the model.

A.1 Market Equilibrium

The model contains three markets for each year, t , and projection year, v , with the markets corresponding to the three major US interconnections (East, West, and Texas). Hourly electricity demand is exogenous and perfectly inelastic with respect to the equilibrium electricity price. Generation from nonfossil generators, such as nuclear and wind-powered, is also exogenous.

There is a set of fossil fuel-fired units in each market that have already been constructed, and for which capital costs are sunk. Each unit takes fuel prices, electricity prices, and generation from other units as exogenous, and chooses a production level to maximize profits. At the beginning of each year and projection year, the unit exits the market if discounted future profits are negative (in the model, profits are deterministic).

There is also a set of potential entering units. For these units, construction costs have not yet been paid. These units enter if discounted future profits exceed the construction costs.

Each market clears each hour of the year and projection year, and the equilibrium is determined differently in each of the three types of hours. The first type of hour is a peak hour, which is the hour with the maximum electricity demand for the corresponding day and market. Near-peak hours occur within several hours of the peak hour, and on the same day. Other hours include all hours that are neither peak nor near-peak hours.

For peak hours, there is a system operator who solicits bids from the units to supply electricity during that hour. In its bid, each unit specifies a generation level (typically, its

rated generation capacity) and a minimum price above which the unit will produce that amount of electricity. The operator stacks the bids in order of increasing price bids and accepts bids so that total generation from accepted bids equals electricity demand.

For near-peak hours, all units that operate in the peak hour must generate an amount of electricity no less than a unit-specific minimum generation level. Units that do not operate in the peak hour cannot operate in the near-peak hours. This structure prevents units from turning on and off many times per day. Units that operate in the near-peak period submit bids to the system operator that include generation levels and minimum prices. The operator stacks the bids by increasing prices, and accepts bids so that generation equals demand. For near-peak hours, total generation includes the generation levels from the accepted bids, plus the minimum generation levels of the units whose bids are not accepted. The equilibrium price is equal to the highest bid among the accepted bids. Note that some units operate at their minimum generation levels even if their marginal costs are greater than the price; they do so because the resulting losses are more than offset from the profits earned during the peak hour. This market structure thus approximates behavior in a dynamic model that includes fixed costs of starting up or shutting down and unit-level constraints on changing generation levels across hours (e.g.,(Bushnell, 2008; Castillo and Linn, 2011)).

During the other hours, equilibrium is determined according to simple economic dispatch. Because a unit's operation in other hours does not affect its profits in near-peak or peak hours, in other hours all units submit minimum price bids equal to their marginal costs. The operator stacks bids by increasing minimum prices, and accepts bids to equate total generation of accepted bids and demand. The equilibrium price is equal to the marginal costs of the unit with the highest costs among those with accepted bids.

Profits for each unit are computed for each year and simulation year, and are discounted to the present year using a 10 percent discount rate. Units with negative profits exit, and the model is resimulated using the resulting smaller set of units. The procedure is iterated until all units have non-negative profits. After the model is solved, we compute coal consumption

by plant, year, and projection year.

A.2 Data and Parameter Estimation

The fuel prices are from EIA projections in the AEO from various years. For example, for the year 2001, we use projections from years 2001 through 2020 from the AEO 2001.

For electricity demand, we begin by computing total observed fossil generation by year, hour, and market from CEMS. We then use the corresponding AEO to scale up hourly generation by the projected electricity consumption growth. For example, for year 2001 and projection year 2002, we use the hourly CEMS generation from 2001, which we scale up by the projected consumption growth between 2001 and 2002 from the AEO 2001. Implicitly, we assume that consumption grows proportionately across hours in a year.

The set of existing fossil units (i.e., those that have sunk construction costs) includes all units in CEMS that operate in the corresponding year. Potential entrants include all units that are observed to enter between the year and 2015, which is the most recent year of CEMS data. Potential entrants also include all units in the Form EIA-860 data that are listed as proposed units, meaning that they have entered permitting and construction phases sometime after 2001.

Each unit's fuel costs are computed from the fuel prices and its observed heat rate (the ratio of average heat input to generation observed in CEMS). Nonfuel costs and minimum generation levels are estimated as described in Linn and McCormack (2017). The solution algorithm and the process for determining which units exit and which enter are also described in Linn and McCormack (2017).

B Mine Closure Date Checks

The definition of a mine closure is clear in practice but can be difficult to ascertain in the data. Our measure of mine closure is the permanent closure of an operating mine.¹² To ensure that we are correctly specifying a mine closure, we look at reported employment and production in the years leading up to and after our defined year of closure. Figures 11 and 12 show box plots of the ratio of a given year's employment or production relative to the maximum amount ever reported for the mine in the data. The bold line reports the median, and the top and bottom of the box report the 25th and 75th percentiles while the ends of the whiskers are the 5th and 95th percentiles and the dots reveal any outliers. Figure 12 shows that production levels drop to zero on the year of closure, with a few mines reporting some production the year after closure, likely the result of coal already mined but not yet processed for delivery. Figure 11 shows an initial drop in employment levels when a mine closes, then a rebound 6 to 10 years later. This residual may be due to remediation of the mine site that is required by both federal and state laws, but these employment/production data are reported for only a handful of mines after reported closure occurs, indicating that the vast majority of operations simply leave the sample.

¹²Specifically, mine closure occurs on the date MSHA classifies the mine as either abandoned or abandoned and sealed.

Figure 11: Ratio of Employment to Maximum Employment





Figure 12: Ratio of Production to Maximum Production



C Imports and Exports

Coal imports and exports are an additional consideration in the modeling approach. While coal trade is only a small portion of the market as a whole, it may have especially large impacts on a disaggregate basis for particular mines or plants. While foreign mines are not considered in the closure regression models, their supply to particular plants is readily incorporated into the modeling framework as consuming coal plants note the country of origin for purchases when these purchases are not made from domestic mines. Exports are somewhat more difficult. The EIA reports exports at an aggregated port level (the point of export) rather than mine level. We model these ports in much the same way as we model plants (but taking the given export quantities as exogenous rather than simulating their demand). To determine the share that each mine contributes to a given port's exports (and therefore the implied export demand for the particular mine), we exploit several characteristics. We utilize the geographic distance from each mine to each port, the export quantities at each port, and the quantity of coal that is produced at a given mine, but not purchased by a plant (i.e., missing from the Form EIA-923 transaction data). While the difference between observed production in the MSHA data and the purchased coal by plants as recorded by Form EIA-923 is unlikely to directly measure export quantities (as this number captures stockpiling and related production activities), we model this difference as a function of distance from port to mine and the export quantiles of ports in order to create a more credible measure of coal export demand at a mine level. We expect mines located closer to ports with large export volumes to have a greater share of produced but unsold coal represented in export quantities. The exports regression takes the specific form in Eq. 7 and is estimated using OLS.

$$\ln(q_{it}^{Produced} - q_{it}^{Sold}) = \ln\left(\sum_j \frac{Exports_{jt}}{Distance_{ij}}\right) + YEAR_t + MONTH_t + \epsilon_{it} \quad (7)$$

The logged difference between each mine i 's production (measured by MSHA) and pur-

chases by plants (measured by EIA) is regressed on the logged, inverse distance-weighted quantities of exports from each port j and time effects. The fitted values from Eq. 7 are used to calculate the corresponding share of each port's exports to each mine's production, s_{ij} as above. Export demand from each mine is calculated from these share values and the observed level of exports from each port.

D Additional Tables

Table 8 provides information comparing the data used for the marginal cost estimation and the larger MSHA dataset for which these parameter estimates were fit. The companies included in the cost dataset are Arch Coal, Alliance Resource Partners, Consol Energy, Peabody Energy, and Westmoreland Coal. This set of companies provides a diverse sample of regions and mine types, but omits smaller, privately owned mines. Table 8 shows that the mines included in the cost estimation sample are larger, on average, than the mines in the full MSHA sample. The cost sample data appear to cover the range of production values, however, as the standard deviation is large and the largest mine appears in both datasets. The average wage and start year are similar across the two datasets, as is the proportion of underground mines (relative to surface operations). The mines with the cost data are more dispersed geographically, while the majority of mines (in terms of number of unique mines, regardless of scale) are located in the Appalachian region.

Results from estimation of Eq. 5 are presented in Table 9.

The first stage results from estimation of Eq. 2 are presented in Table 10.

Table 8: Summary Statistics of Corporate Filings (Cost) Sample Compared with MSHA mines

Variable	Sample	N	Mean	St. Dev.	Min	Max
Annual Production ('000 tons)	Cost	839	10.25	19.81	0.00	116.20
	MSHA	62,207	0.74	4.16	0.00	116.23
County Wages (All Industries, \$/week)	Cost	882	693.83	147.56	395	1,144
	MSHA	62,090	660.02	127.74	364	1,266
Mine Start Year	Cost	446	1979	17.05	1925	2012
	MSHA	62,177	1984	17.92	1950	2012
(% Appalachian Mines)	Cost	362	41%			
	MSHA	54,635	88%			
(% Underground Mines)	Cost	434	49%			
	MSHA	26,039	42%			

Table 9: Marginal Cost Equation Estimates

	<i>Dependent variable: Cost Per Ton</i>
	Coefficient (Std Error)
q	0.058 (0.527)
q ²	-0.009 (0.634)
wages	0.0004 (0.0004)
Mine Type:S/U	0.190 (0.229)
Mine Type:U	0.101 (0.216)
Company:Consol	-0.0001 (0.016)
Company:Peabody	-0.199*** (0.029)
Company:Westmoreland	0.042 (0.050)
Vintage Year	-0.0002 (0.0003)
q x Region:INT	0.018 (0.529)
q x Region:LIG	-0.058 (0.528)
q x Region:PRB	-0.056 (0.527)
q x Region:WST	-0.057 (0.527)
q ² x Region:INT	0.004 (0.634)
q ² x Region:LIG	0.009 (0.634)
q ² x Region:PRB	0.009 (0.634)
q ² x Region:WST	0.009 (0.634)
wages x Region:INT	-0.001** (0.001)
wages x Region:LIG	-0.0004 (0.0005)
wages x Region:PRB	-0.001 (0.0005)
wages x Region:WST	-0.001 (0.0005)
q x Mine Type:S/U	-0.058 (0.529)
qx Mine Type:U	-0.071 (0.528)
q ² x Mine Type:S/U	0.010 (0.634)
q ² x Mine Type:U	0.011 (0.634)
wages x Mine Type:S/U	-0.0004 (0.0004)
wages x Mine Type:U	-0.0002 (0.0004)
Region:INT x Mine Type:S/U	0.230 (4.638)
Region:INT x Mine Type:U	-0.591** (0.289)
Region:WST x Mine Type:U	1.341*** (0.386)
Region:INT x Vintage Year	0.0002 (0.001)
Region:LIG x Vintage Year	0.0002 (0.002)
Region:PRB x Vintage Year	-0.001 (0.004)
Region:WST x Vintage Year	0.0001 (0.002)
q x Region:INT x Mine Type:U	-0.002 (0.529)
q x Region:WST x Mine Type:U	0.154 (0.847)
q ² x Region:INT x Mine Type:U	-0.006 (0.634)
q ² x Region:WST x Mine Type:U	-0.058 (0.779)
wages x Region:INT x Mine Type:U	0.001* (0.001)
wages x Region:WST x Mine Type:U	-0.0002 (0.001)
Constant	3.335*** (0.604)
Year Fixed Effects	Yes
Region Fixed Effects	Yes
Year x Region	Yes
Observations	399

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 10: First-stage Estimates of Demand Regression

	ln Price
Productivity ('000 hours/ton)	-0.020*** (0.001)
ln Residual Demand	0.003 (0.004)
Sulfur Cost	-0.0002*** (0.00002)
Constant	4.984*** (0.070)
Year Fixed Effects	Yes
Observations	1,266
Adjusted R ²	0.612

Note: *p<0.1; **p<0.05; ***p<0.01

Excluded instrument is a supply-side shifter: county-average mining labor productivity. Residual demand and sulfur cost are exogenous demand shifters.