

Sources of Partisan Change: Evidence from the Shale Gas Shock in American Coal Country*

Alexander F. Gazmararian[†]

September 22, 2023

Abstract

What explains the shift to Republicans in places that historically voted for Democrats? This paper tests a new explanation for part of this reversal. The shale gas revolution displaced coal, which intensified the salience of national environmental regulations and increased support for Republican presidential candidates. Analysis of presidential elections from 1972 to 2020 with a difference-in-differences design finds that the shale gas shock increased Republican vote share by 4.9 percentage points. Leveraging geospatial data, media analysis, and interviews, I show that voters blamed environmental regulations for their community's decline and that the backlash was more likely to occur where the shale shock was least visible. The attribution of blame for economic dislocation helps to explain electoral behavior in places "left behind," and sheds light on political responses to climate policy.

*Princeton University's Institutional Review Board approved the semi-structured interview protocols (IRB 13849) and the surveys (IRB 13942, IRB 14813). Appendix P affirms how the study conforms to the APSA Principles and Guidance for Human Subjects Research. The author gratefully acknowledges the financial support of Princeton University's Niehaus Center for Globalization and Governance. Thanks to Chris Achen, Christian Baehr, Fiona Bare, Jonathan Bendor, James Bisbee, Steven Callander, Michael Cerny, Amanda Clayton, Lawrence Goulder, Gavin Medina-Hall, Vincent Heddesheimer, John Kastellec, John Londregan, Robert Keohane, David Konisky, Haillie Na-Kyung Lee, Daniel Lyng, Nolan McCarty, Tali Mendelberg, Helen Milner, Jonathan Mummolo, Torsten Persson, Markus Prior, Kris Ramsay, Dani Reiter, Michael Ross, Robert Staiger, Dustin Tingley, Sam van Noort, Erik Voeten, James Vreeland, Hye Young You, Noah Zucker, and seminar audiences at Princeton University, Emory University, Stanford's Graduate School of Business and Doerr School of Sustainability, the Graduate Climate Lab, and the Political Economy of Climate and the Environment (PECE) Mini-Conference for helpful feedback.

[†]Ph.D. Candidate, Department of Politics, Princeton University. Email: afg2@princeton.edu. Web: alexgaz.org

The American political landscape has undergone a remarkable transformation. Voters in places that traditionally supported Democrats, like coal-producing areas, have swung far in the opposite direction. In 2000, coal mining counties were 2.5% more Democratic on average than the rest of the nation, casting their ballots for Al Gore, a Democrat who promised to tackle climate change. Since then, coal mines have shuttered. In 2020, coal mining counties were 10.3% more Republican on average than the rest of the country, voting for Donald Trump, a Republican who pledged to roll back environmental regulations. These changes in coal country are part of a broader rightward shift where people in places that historically backed Democrats have increasingly voted for Republicans. What explains this partisan reversal?

One prominent explanation is that working-class voters “left behind” by economic trends are responsible for this partisan reversal. Dislocation from automation, offshoring, and trade have deteriorated social and economic conditions in communities (Broz, Frieden, and Weymouth 2021) and reduced upward mobility (Ballard-Rosa, Goldstein, and Rudra 2022), which has sparked a backlash to elites thought responsible and amplified inter-group conflict to the benefit of far-right politicians (Autor et al. 2020; Baccini and Weymouth 2021; Ballard-Rosa, Jensen, and Scheve 2021; Walter 2021).

However, others argue that rather than responding to deteriorated economic circumstances, voters have turned to the Republican Party because of its appeal to “cultural wedge issues” (Frank 2007). The growth of these social divisions may be the product of long-term cultural changes (Norris and Inglehart 2019), as well as more recent developments like the election of the first Black president and an increase in immigration, which have threatened the status of dominant groups (Abrajano and Hajnal 2015; Sides, Tesler, and Vavreck 2018; Reny, Collingwood, and Valenzuela 2019; Mutz 2018).

Yet, a theoretical and empirical challenge that confronts both the left behind thesis and its detractors is the question of how voters attribute blame for economic changes (Cramer 2016; Healy and Malhotra 2013). Ambiguity around attribution is one reason for a mix of

findings regarding whether citizens support left- or right-wing parties in response to economic disruption (Margalit 2019). Indeed, causal judgments of responsibility are a crucial link between personal experiences and politics (Iyengar 1989). For example, automation is responsible for many manufacturing job losses, yet impacted individuals have largely blamed trade policy and immigrants (Wu 2022; Mutz 2021).

The primary contribution of this paper is to propose and test a novel explanation for part of the partisan reversal in a context that also allows for an exploration of how voters attribute blame for economic disruption. I argue that an underappreciated global technological shock has had outsized political consequences: the shale gas revolution. Unlike studies of places that benefited from shale gas (e.g., Sances and You 2022; Cooper, Kim, and Urpelainen 2018), I focus on the economic losers. The cheap supply of gas unlocked by hydraulic fracturing displaced coal-fired power generation (Acemoglu et al. 2023; Holladay and LaRiviere 2017) and caused job losses across the coal industry (Coglianese, Gerarden, and Stock 2020). Following this decline, coal country—places that produce coal—flipped from its historical support of Democratic presidential candidates to Republicans.

I argue that this electoral change is because the closure of coal mines increased the salience of the Democratic Party’s “issue ownership” of the environment (Egan 2013), a connection that the Republican Party politicized. However, this is a misattribution of the primary cause of the industry’s decline. The best estimates show that the shale shock accounts for 92% of the total decrease in coal production, not environmental regulations (Coglianese, Gerarden, and Stock 2020). So why did voters blame policy instead of markets?

Building on studies about the political effects of local context (e.g., Hopkins 2010; Huckfeldt and Sprague 1995), I theorize that the attribution of blame for an economic outcome depends on the visibility of its source. Visibility refers to the extent to which an event’s cause is directly observable, akin to Arnold’s (1990) concept of traceability. Since it is costly to gather information (Downs 1957), people will gravitate to easily identifiable sources. Information from one’s immediate surroundings is more accessible, relevant, and salient

(Ansolabehere, Meredith, and Snowberg 2014; Bisbee 2019; Mondak, Mutz, and Huckfeldt 1996), so it should receive greater weight when attributing blame. The target of blame may not always be the actual cause because what one sees may be influenced by elites or a partial picture. However, local information should lead to more accurate attributions when the true source of an economic shock is directly observable.

In this paper’s context, the argument about visibility indicates that the inability of voters to observe fuel-switching from coal to gas far down the supply chain in electricity markets created latitude for the Republican Party and interest groups to blame environmental regulations. The decline of unions that serve as information brokers and the broader community’s weaker understanding of the industry’s economics contributed to this misattribution. As a result, voters left the party that they (correctly) believed pursued a policy they (incorrectly) believed caused their community job losses.

Evaluating the effect of the shale shock on voting in coal country and the moderating effect of visibility presents two empirical challenges. First, there is an endogeneity problem where communities that extract coal may have pre-existing political differences that influence vote choice and the decision to mine for coal. Failing to account for endogeneity would produce biased causal estimates. Second, there is a measurement challenge in quantifying the visibility of the shale gas shock, given the lack of individual-level data.

To overcome the endogeneity challenge, I leveraged the timing of the shale gas boom that coal country—and even financial markets—did not anticipate (Holladay and LaRiviere 2017). I employed a difference-in-differences design that matched coal counties with their neighbors (Imai and Ratkovic 2014; Ho et al. 2007). With this setup, I estimated the effect of the shale gas shock on Republican presidential vote share in affected counties from 1972 to 2020 using the fixed effects counterfactual estimator (FEct) (Liu, Wang, and Xu 2022). The analysis focuses on presidential elections because the Democratic Party’s issue ownership of the environment is most salient at the national level, whereas local Democrats can strategically break from the party line on coal (Canes-Wrone, Brady, and Cogan 2002).

I find that the shale shock increased two-party Republican presidential vote share in coal counties by 4.9 percentage points on average relative to matched neighboring counties. The partisan reversal takes place in electorally consequential states like Pennsylvania and is notable given coal country’s historical support for Democrats. Using demographic data, I also constructed bounds on the bias from compositional changes, showing that the partisan reversal is largely due to voter conversion and mobilization.

I conducted three analyses to evaluate whether these results are spurious with the global financial crisis (GFC) or racial backlash. First, falsification tests examine the effect of the shale shock in less-exposed anthracite coal counties; iron, silver, and sandstone mining counties; and counties that experienced non-coal layoffs. Second, empirical models with time-varying shale shocks pinpoint when communities experienced economic dislocation. Third, a sensitivity analysis shows there would have to be an unobserved confounder with more than 1,500% the correlation of the county share of white residents with the treatment and the outcome to alter the finding (Cinelli and Hazlett 2020).

Testing the argument about misattribution is difficult because there is no survey data with sufficient geographic and temporal coverage that also measures blame for economic shocks. To overcome this challenge, I leveraged administrative data on the location of all new gas-fired power plants from 1990 to 2020. I developed a measure of shale shock visibility by calculating the distance between a county’s center and new gas power plants. In counties where the shale shock is less visible, I find there is a larger shift to the Republican Party, even when controlling for new jobs created by power plants and balancing on covariates that predict a county’s distance from new gas plants. A sensitivity analysis indicates there would have to be an unobserved confounder at least 43% as large as the correlation of a county’s poverty level with the distance to a new gas plant and the outcome to change this conclusion.

As evidence that voters blamed environmental regulations, I supplemented the quantitative analysis with 60 semi-structured interviews with coal miners, community members, and politicians in Southwest Pennsylvania. This investigation provides suggestive evidence that

voters were worried about the Democratic Party’s support for environmental regulations. Crucially, the interviews show that voters did not blame the shale shock nor the Democratic Party’s lack of response to protect coal communities in decline.

I also coded local news coverage of coal mine closures in Southwest Pennsylvania and find that environmental regulations received the most blame. A parallel analysis of cable news closed-caption transcripts shows that the national media, especially Fox News, was more likely to discuss regulations than the shale gas shock as the cause of coal’s decline.

To explain why a county-wide shift in voting occurred despite the small share of direct coal employment, I build on studies about community-centric preferences in coal country (Gaikwad, Genovese, and Tingley 2022; Gazmararian and Tingley 2023). I argue that social ties led people to interpret job losses through a collective lens. To assess this claim, I conducted original in-the-field surveys in Southwest Pennsylvania during 2021 and 2022 to measure social network ties to the fossil fuel industry ($N = 606$). I expended considerable resources to access the hard-to-reach population by conducting convenience samples at county fairs. I find that for the modal respondent, all ten of the closest individuals in her network work in fossil fuels. While likely an exaggeration, this perception indicates how people might become concerned about regulations even if they did not lose their jobs.

The paper makes three contributions that the conclusion elaborates upon. First, I propose and test a new explanation for part of the partisan reversal in the US, which nuances our understanding of the rightward electoral shift. Second, I provide evidence demonstrating the plausibility of how the visibility of an economic shock’s source affects vote choice. The argument about visibility moves beyond debates over voter competence (e.g., Achen and Bartels 2016) to, instead, differentiate the conditions under which voters’ behavior approximates a boundedly rational model of decision-making. Lastly, I provide credible estimates of the electoral effects of energy transitions, which emphasize the importance of “just” transition measures as a part of climate policy (Bolet, Green, and Gonzalez-Eguino 2023; Gaikwad, Genovese, and Tingley 2022; Kono 2020).

1 Global Technological Disruption and Partisan Reversal

1.1 Economic Impact of the Shale Gas Revolution on Coal Communities

The shale gas boom was a global technological shock resulting from the widespread application of hydraulic fracturing and horizontal drilling, initially developed in the 20th century but not deployed at scale until the 2000s. As gas became cheaper than coal, electric utilities began replacing coal-fired power plants with gas. The left panel of Figure 1 depicts the change in coal and gas production, and the right panel shows the corresponding fuel-switching in the electricity sector. Although coal industry employment has been on a slow decline due to mechanization, the shale shock caused considerable economic dislocation in mining communities (DOE 2017; Coglianese, Gerarden, and Stock 2020). Between 2008 and 2020, County Business Patterns (CBP) data show a 38% decline in coal industry employment.




Figure 1: **Shale Shock Began in 2008.** The left panel shows normalized coal and gas production in the US. The right panel shows net electricity generation. Source: EIA

Coal's decline had community-wide consequences, which are essential for understanding the political responses of people not directly employed in the industry. First, decreases in coal production undermined the provision of local public goods, such as tax revenue for public schools, infrastructure, and emergency services, so residents felt the impacts (Gazmararian and Tingley 2023, Appendix O). Second, the industry has social and cultural relevance that resonated with residents, instilling a community-orientation in their political preferences

(Gaikwad, Genovese, and Tingley 2022).

The social ties that residents have to people in mining are one mechanism facilitating this community orientation in coal country. As evidence, I conducted in-the-field surveys, which show how residents are embedded in social networks where their friends and family work in fossil fuels. This helps to explain why coal’s decline could cause county-wide shifts in voting despite the small share directly employed (Appendix M.3).

Determining whether voters attributed the right cause for the industry’s decline requires a precise assessment of the reason for coal’s fall. A recent credible econometric evaluation that accounts for the relative contribution of nine factors, including falling gas prices, environmental regulations, renewable energy incentives, and overall electricity demand, concluded that the declining price of gas accounts for 92% of the decrease in coal production (Coglianese, Gerarden, and Stock 2020). Appendix A.1 reviews econometric studies, which concur that the shale shock is the primary cause of coal’s decline.

Despite this decline, the federal policy response has been minimal compared to countries like Germany (Sheldon, Junankar, and Pontello de Rosa 2018). The Obama administration’s POWER Initiative provided grants (Appendix H.2), but it began years after coal started declining and was not fully enacted (Lawhorn 2022, 3). One assessment determined that the federal government’s approach to the coal transition has not provided “the certainty and adequate support needed” (Roemer and Haggerty 2021, 5).

Local responses to the decline of coal have also struggled or have been nonexistent. Representative surveys of local policymakers in fossil fuel communities document the lack of capacity as a formidable barrier (Gazmararian and Tingley 2023). In a review of local strategies to cope with coal power plant closures in the West, Haggerty et al. (2018, 1) concluded that there has been an “absence of effective strategies...” There are economic constraints that inhibit local responses, such as the difficulty of attracting new employers to rural areas that used to rely on a single industry (Interviews A1, A11, A18).

1.2 Left-Behind and Cross-Pressured Voters

The shale gas shock occurred in a political context where coal-producing counties were historically more likely than the rest of the country to vote for Democratic presidential candidates. However, after the shale shock, Figure 2 shows how coal counties flipped to Republicans on average, some by double digits. This trend intensified in 2012 when coal job losses grew and exceeded the national rightward shift observed at the same time.

This historical advantage for the Democrats in coal country was due partly to union ties (Dark 1999). Other issues also drew voters to the Democratic Party since rates of unionization vary (Gaventa 1982), and there were also some regional differences in baseline support. This average Democratic advantage created the conditions for a partisan reversal.

Democratic voters in these areas were cross-pressured. The Democratic Party, in addition to defending labor, also backed environmental regulations that threatened coal (Mildenberger 2020). The Democratic Party has long “owned” environmental issues, being seen as more competent and consistent, while the Republican Party has been more pro-business and anti-regulatory (Egan 2013). Appendix L presents an analysis of party platforms to demonstrate this partisan divide. As early as 1988, the Democratic platform included a call to address climate change, while the Republican platform promised to defend the coal industry. To the extent that Democratic voters were aware, they faced a dilemma: supporting a party that protects workers’ rights could jeopardize their industry’s viability.

I argue that the decline of coal due to the shale shock increased Republican support through two mechanisms. First, the decline of coal raised the salience of the Democratic Party’s environmental issue ownership since voters increasingly saw a conflict between their economic self-interest and partisan preference. Second, the Republican Party also sought to politicize the issue by blaming environmental regulations, which also mobilized new voters.

While coal has undergone busts before, I hypothesize that this did not precipitate a partisan reversal because labor issues used to be of greater salience. During the 1980s bust, miners cared deeply about Reagan-era policies weakening mine safety. For example, the back




Figure 2: **Average Republican Presidential Vote Share, County-Level.** $N = 40,424$. The plot compares all counties. Subsequent analysis matches coal counties with neighbors.

cover of the May 1985 issue of the UMW Journal, referencing Reagan’s policies, said in large font: “Will You Be Alive To Vote On Election Day? Vote As If Your Life Depended On It Because It Does.” As systemic evidence of labor’s salience, I collected microfilm archives of the titles and tables of contents for UMW Journal editions from 1975–1985. I find using a structural topic model that safety is the most frequently mentioned topic (Appendix B.2). While labor issues have not disappeared, mines have become relatively safer as measured by fatalities (Figure B1). Thus, when the shale shock happened, there was more room for cross-pressured voters to be converted and unaligned voters to be mobilized.

1.3 Visibility and Attribution of Economic Shocks

Why did voters blame policy for coal’s decline instead of market forces? To answer this question, I build on studies on political context (e.g., Hopkins 2010; Huckfeldt and Sprague 1995) to develop an argument that the visibility of an economic outcome’s cause shapes voters’ (mis)attributions. When the cause is directly observable, voters should be more likely to make accurate attributions. This idea of visibility relates to Arnold’s (1990) argument about the traceability of policy costs, which posits that voters are more likely to hold politicians

accountable for laws with highly visible costs since it is easier to connect lawmakers' actions with their economic fortunes.

I assume that people are not maximizers, where they seek out all information before making decisions. Information is costly to obtain due to the effort required to search, process, and evaluate claims about economic outcomes (Downs 1957; Huckfeldt and Sprague 1995), so voters will often satisfice by gathering enough data to make a sufficient decision (Simon 1955). This leads to reliance on accessible messages. In an information-scarce environment, there is room for politicians and societal groups to frame the cause of economic outcomes in a way favorable to their interests. When voters rely on misleading messages, they may appear to behave in non-rational ways (e.g., Achen and Bartels 2016), such as when citizens blame foreigners for jobs lost to robots (Mutz 2021; Wu 2022).

However, I argue that when there is visible information about the cause of economic outcomes, voters should be able to make more accurate attributions, thus appearing to behave in a boundedly rational manner. Politicians should have less latitude to blame an erroneous source since people can see the reason for their economic misfortune. This argument builds on studies about how the local economy affects electoral behavior due to its accessibility, relevance, and salience (Bisbee 2019; Mondak, Mutz, and Huckfeldt 1996; Ansolabehere, Meredith, and Snowberg 2014; Reeves and Gimpel 2012; Larsen et al. 2019). However, my argument differs from many of these studies in that local conditions are not a heuristic for national conditions but function as information about local economic outcomes.

The logic of my argument relates to theories about how political context shapes evaluations of national issues. For example, Hopkins (2010) shows how immigration influxes, made salient by the national media, cause people to become more concerned about immigration. Relatedly, cross-national research suggests that institutions influence the “clarity of responsibility” for economic performance (Powell and Whitten 1993). An empirical advance in this paper is to overcome measurement challenges in capturing local visibility. A theoretical advance is to complement Hopkins’s (2010) contention that national frames politicize local

changes. My argument suggests that the *absence* of local information can be politicized while its *presence* can serve as a ground-up constraint on elite messages.

Local information is not a panacea. For one, politicians and the media can politicize certain experiences while deemphasizing others (Mutz 1994). Local conditions could lead voters astray by making it easier for politicians and interest groups to blame erroneous but visible causes of economic shocks. Visibility should lead to more accurate attributions when the true cause of an economic outcome is directly observable.

1.4 (Mis)Attribution of the Shale Shock

I argue that the low visibility of coal-to-gas switching in electricity markets made it harder for voters to understand the cause of coal’s decline. This created an opportunity for Republicans and interest groups to blame environmental regulations, which was in their self-interest to win votes and avoid regulations. In response, Democratic politicians tried to communicate how the shale shock was responsible for coal’s decline. Diana DeGette, a senior Democratic representative, said in a high-profile hearing, “the reality [is] that as natural gas becomes cheaper than coal and more and more other utilities and others transfer to natural gas, it is the invisible hand of the free market” (Appendix A.2).

When judging these competing messages, many coal country residents could not directly observe coal-to-gas switching, so the framing that regulations were responsible was likely compelling. My interviews suggest that people could see how the coal industry in their area was subject to environmental rules, though ones implemented several administrations prior. One individual described how a coal-fired power plant alongside a well-traveled road had to install large scrubbers, a decision stemming from the 1990 CAA Amendments, not the cause of coal’s decline three decades later (Interview A11).

While miners usually have a good sense of the industry’s business model and have access to unions as information brokers, the decline in unionization among coal miners, as documented in Appendix H.1, meant fewer workers had access to quality information. Instead,

they increasingly relied on messaging from mine owners, which emphasized regulatory costs (Interview B26, D1, D4).

Even in the mines that remained unionized, contemporary accounts point to a growing fissure between members and leadership. When explaining why the UMW did not endorse a Democratic presidential candidate in 2012, which it had done in every election since 1972, a union representative described the ground-up pressure: “We’re a very Democratic union and we try to listen to the rank and file. They’ve sent a clear message that they’re not supportive of the environmental rules that are being put in place.” This is despite the same representative recognizing that the EPA rules were enacted under previous administrations and that the shale gas boom was responsible for coal’s decline. However, he said, “a lot of our members don’t realize that” (Harder 2012). This lack of awareness is consistent with my interviews of workers in non-unionized (B14, B18) and unionized mines (B31, B32, C11).

Community members not employed in the mines likely had an even greater lack of understanding, which made them more susceptible to messages blaming regulations. While residents care about issues seen to affect their community, absent directly observable changes, people might not be equipped with the information necessary to render accurate attributions for coal’s decline.

Thus, in places where coal-to-gas switching was more visible, voters should have been less persuaded by the claim that regulations were at fault. The lack of visibility helps to explain the political response of voters in coal country. Had they identified cheap gas prices as the cause, they may have instead supported restrictions on hydraulic fracturing.

2 Research Design

The primary hypothesis is that counties with greater economic dependence on coal should become more likely to vote for Republican presidential candidates after the shale shock. Further, this rightward shift should be more pronounced in places with less visible coal-to-gas switching in electricity markets. The main test of the first hypothesis leverages the

unanticipated timing of the shale gas boom and uses a difference-in-differences model that matches coal counties to adjacent counties.

The analysis focuses on coal counties throughout the US because the mechanism behind the shale shock’s effect should apply across all coal-producing areas. While there are geographic differences in rates of unionization, mode of mining, the sulfur content of coal, and historical level of production (Roemer and Haggerty 2021), the empirical models account for time-invariant factors by estimating within-unit changes in voting and for time-varying factors with controls.

2.1 Outcome: Two-Party Republican Presidential Vote Share

Since the theorized reason for the partisan reversal is the misattribution of *national* Democratic policies, the analysis focuses on presidential elections. It is at the national level where there is a clear environmental divide between the parties (Appendix L). By contrast, there is a well-documented phenomenon where lawmakers strategically diverge from the national party’s line to stay in step with their constituents (Canes-Wrone, Brady, and Cogan 2002). Local Democrats in coal-producing areas often break from the national party on environmental issues, so it is not possible to make predictions about non-presidential races without data on candidate platforms. For example, Joe Manchin, the Democratic Senator from West Virginia, is more pro-coal than his party, so he may not have been blamed for the industry’s decline. Likewise, Democratic Pennsylvania state representative, Pam Snyder from Greene County, vocally supports the coal industry, so she would have also escaped blame. Democratic *presidential* candidates, however, take on the mantle of their party’s platform.

Data on elections come from Leip (2020). The analysis includes contests from 1972 through 2020 due to the reshuffling of party allegiances in the post-racial realignment period.¹ Counties represent the lowest level of aggregation that matches the employment data and

¹The analyses using a control for oil and gas employment begin in 1976 due to data availability.

are of theoretical relevance due to the community-wide effects of coal’s decline.

Since the empirical strategy examines election outcomes, it cannot distinguish whether the reason for the electoral change is voters blaming the Democratic Party for causing the decline of coal, their failure to respond with transitional support, or both. However, the interviews discussed below indicate that the miners, residents, and leaders of a major coal community blamed Democrats for causing the decline, not their failure to respond.

2.2 Shale Shock Timing

Studies of the shale gas boom identify 2008 as a structural break in gas prices, which is the threshold this paper’s analyses employ since it also matches the closest election (Acemoglu et al. 2023). While the recession may have contributed to the decline in gas prices due to the fall in consumption, the fracking boom kept these prices low, resulting in an entirely new price regime (Holladay and LaRiviere 2017). Qualitative accounts validate the 2008 date. Energy historian Yergin (2020, 11) notes, “2008 was the moment when the bell rang. That year, US natural gas output went up instead of down, as had been the general expectation.”

The shock was so sudden that financial markets did not anticipate it. Future prices exceeded spot prices before the shale boom, indicating that investors incorrectly expected constrained supply (Holladay and LaRiviere 2017). Even the EIA—the entity responsible for energy forecasts—did not predict it. One assessment concluded the shale gas boom “was hardly, or not at all, predicted or projected by forecasters, analysts, or industry experts even a year or two before its emergence” (Reed et al. 2019, 1). Lastly, the interviews below indicate that coal country residents did not anticipate the shale shock.

The primary threat to inference with a temporal break is spuriousness. There could be another event in 2008, like the financial crisis or Obama’s election, that is actually responsible for the partisan reversal. When presenting results, I rule out spuriousness with falsification tests and by estimating time-varying shale shocks.

2.3 Causal Identification Strategy

The primary research design is a difference-in-differences (DD) model comparing coal counties that were economically vulnerable to shale shock with matched adjacent counties that were not. Treated counties are those that had more than 1% of local employment in the coal industry during 2005-2007, the years before the shale shock (Eckert et al. 2020). Theory, data, and validation exercises inform this treatment categorization. First, the community-wide effects of coal’s decline imply that even small levels of employment should have wider political effects. Second, an analysis of the distribution of the data indicates that 1% is an important cutpoint that meaningfully distinguishes counties with coal from those without. Third, a validation exercise indicates that the treatment definition exhibits a strong, positive correlation with county GDP from coal. Lastly, the results are robust to using a continuous measure of employment in an event study design and when defining the treatment in terms of the size of layoffs (Appendix C.3).

The main assumption for causal inference is that absent the shale gas shock, exposed coal counties would have followed the same average trend in Republican presidential vote share as the control group. When presenting results, I provide rigorous tests of this assumption. To enhance the comparability of the counties, I also matched units using covariate balancing propensity scores (CBPS) (Imai and Ratkovic 2014). The pool of possible matches included counties within two degrees of adjacency to a treated unit. The model to estimate CBPS used county-level covariates that may predict coal mining activity: race, ethnicity, foreign-born, education, income per capita, poverty, rurality, population, age, and female workforce participation.² I selected matches with replacement between treatment and control counties using these propensity scores (Ho et al. 2007). Appendix D.1 reports how this procedure improves balance, and shows the location of relevant counties.

Since the analysis focuses on adjacent counties, spillover effects are possible if the control group also experienced some of coal’s decline because of economic interdependencies such as

²Appendix C.1 describes the operationalization of these covariates.

commuting workers. However, to the extent that there are negative economic spillovers, this would introduce bias against the hypothesis because control counties would also be facing similar economic pressures theorized to cause increased Republican support.

Another potential concern is that there is variation in local responses to coal’s decline. However, the discussion above indicates that local strategies to cope with coal’s decline have been absent or largely ineffective. Even if the minimal local responses softened the economic toll of coal’s decline, this would introduce bias against the hypothesis because it would mean there would be counties that would have shifted farther to the Republican Party absent local initiatives.

The estimand is the average treatment effect on the treated units (ATT). To estimate the counterfactual Republican vote share had a coal county not experienced the shale shock, $Y_{it}(0)$, I employed the FEct estimator, which is more reliable than the standard two-way fixed effects model when there are heterogeneous treatment effects or unobserved time-varying confounders (Liu, Wang, and Xu 2022). FEct takes treated observations as missing and imputes their potential outcomes.

$$Y_{it}(0) = \mathbf{X}_{it}^\top \beta + \alpha_i + \eta_t + \epsilon_{it}$$

In this model, \mathbf{X}_{it} is a time-varying covariate for hydraulic fracturing employment,³ α_i is a county fixed effect, η_t is an election fixed effect, and ϵ is the error term.

3 Shale Shock Effect on Republican Presidential Vote Share

Figure 3 plots the results from estimating the effect of the shale shock on two-party Republican presidential vote share in coal country. In 2008 at the shale shock’s start, coal counties began to diverge from their matched neighbors. In 2012 as the employment effects of coal-to-gas switching were more intensely felt, vote share for the Republican party strengthened and persisted into subsequent elections. The substantive magnitude of the shale shock is

³Robustness checks include a range of covariates (Table D5).

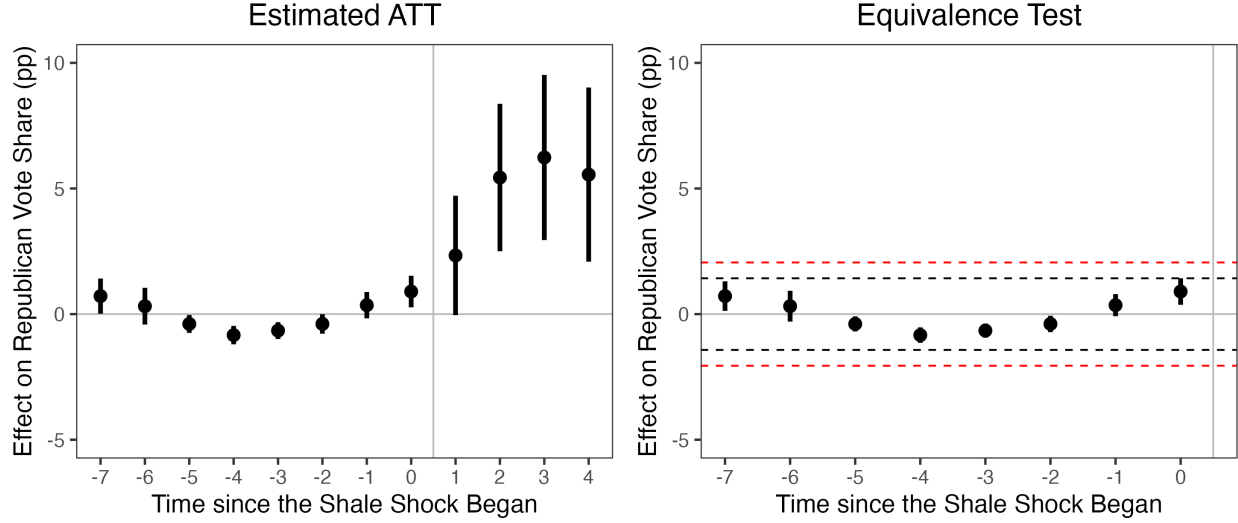


Figure 3: **Effect of the Shale Shock on Republican Presidential Vote Share.** The left plot shows the dynamic treatment effects estimates for elections before and after the shale shock. Bars denote 95% confidence intervals from 2,000 block bootstrap replications clustered by county. Model controls for hydraulic fracturing employment. The right plot shows the pretreatment average prediction errors and their 90% confidence intervals. The red dashed lines denote the equivalence range set at $[-0.36\hat{\sigma}, 0.36\hat{\sigma}]$. The black dashed lines mark the minimum range.

appreciable: there is a 4.9 percentage point increase in Republican presidential vote share in coal counties compared to their matched neighbors in the post shale shock period compared to the average pre-shock period (Table D3). This estimate is similar in magnitude to one study, developed simultaneously with mine, that assessed the electoral effect of coal layoffs in Appalachia between 2000 and 2016 (Egli, Schmid, and Schmidt 2022). Notably, this is not just a Trump story. The partisan reversal became most pronounced in the 2012 election.

A visual inspection of the dynamic treatment effects plot indicates that in the 3-5 elections before shale shock, there appears to be a slightly negative trend, while in the period immediately before 2008, there seems to be a slightly positive trend. I used systematic statistical tests to investigate whether these visual perturbations represent violations of the parallel trends assumption.

First, I conducted an equivalence test as recommended by Hartman and Hidalgo (2018). Whereas a standard placebo test evaluates the null hypothesis of no pretrends violation, which can suffer from limited power, the equivalence approach flips the null so the starting

assumption is there is a pretrends violation of a certain magnitude. Rejecting this flipped null hypothesis implies that there is a high probability that the placebo ATT falls within a pre-specified narrow range, indicating trends are parallel. Equivalence holds only when tests for all pretreatment periods produce significant results, a conservative standard. The p -value for the equivalence test is less than 0.001, which indicates a high confidence that equivalence holds. Visually, the dashed black line in the right panel of Figure 3 depicts the smallest symmetric bound within which the null of inequivalence can be rejected, which is well within the 90% confidence intervals for the equivalence range depicted by the dash red line.

Second, I conducted placebo tests proposed by Liu, Wang, and Xu (2022) that use the three periods before the shale shock begins. Table D3 contains the results, which show no statistically detectable violation of the parallel trends assumption. These diagnostic tests indicate that the pretrends are statistically parallel.

These results are robust to alternative model specifications, sensitivity analyses, and separate causal identification strategies. The results hold when using the matrix completion (MC) estimator (Appendix D.3). Sensitivity analysis using a separate specification, with a full set of controls to permit benchmarking, indicates that there would have to be a confounder orthogonal to the covariates with over 1,500% the strength of the correlation between the share of white residents, coal mining, and the outcome, to change the conclusions (Appendix D.4). I also implemented two separate research designs and reach consistent results. The first employed a shift-share model to capture vulnerability to the shale shock (Appendix G). The second used a flexible event study with staggered treatment onset operationalized by different shale shock timing for each state (Appendix E.1).

Since the analysis examines county-level vote share, an inferential challenge is determining how much of the partisan reversal is due to voter conversion and mobilization as theorized versus compositional changes from migration and mortality. Appendix I contains results from two tests demonstrating the robustness of the findings to compositional changes.

First, I constructed bounds on bias from compositional changes using a simple model of voting and population data. Second, I estimated two-period DiD models with progressively more proximate pre-treatment periods that minimize the possible size of compositional bias. These analyses indicate that the partisan reversal is largely due to voter conversion and mobilization.

I also assessed whether the electoral effects of the shale shock are more muted in communities that received federal assistance to help with coal’s decline (e.g., Bolet, Green, and Gonzalez-Eguino 2023). I collected new data on the universe of federal grants for the POWER Initiative and find that the partisan reversal occurs regardless of whether a county received federal assistance. There is suggestive evidence that counties receiving assistance shifted more to the Republican Party, but this is likely because of a selection effect where these places are the most economically distressed (Appendix H.2).

In addition, I examined whether the shale shock had different effects in counties with higher levels of unionization. The long-run decline of unions, as depicted in Figure H4, may have made it easier for a partisan reversal to occur since they could not act as effective information brokers. When estimating treatment effect heterogeneity, I detected that necessary statistical assumptions are unmet, so I refrain from drawing inferences from these estimates. The main results also obtain when including unionization as a control (Appendix H.1.4).

3.1 Alternative Explanations

Since the analysis used 2008 as a structural break, the results could be spurious with the financial crisis or Obama’s election. The GFC alternative explanation would require that coal counties were differentially exposed to the financial crisis and their exposure predicted vote choice. However, there is inconclusive evidence that the GFC impacted voting (Hall, Yoder, and Karandikar 2021). Still, I investigated this possibility by controlling for the 2006 pre-crisis county-level income-to-debt ratio, a validated measure of GFC exposure (Mian and Sufi 2011), and the results hold (Table D5).

For racial backlash to Obama to explain the reversal, there would have to be an unobserved social factor that renders counties susceptible to racial backlash that is present in the treatment group but not the matched control. While coal-producing areas tend to be more rural and white, the matching procedure balances counties along these dimensions. That said, racial backlash may still be at play in *both* coal and neighboring counties. However, the shale shock explains why there is a larger shift to Republicans in impacted counties.

Alternative explanations based on spuriousness rest on their coincidence with 2008 as the start of the shale shock. While 2008 marks a structural break in national gas prices, there is spatial and temporal heterogeneity that I can leverage to attend to concerns about using a single temporal break. First, I estimated state-specific shale shocks that account for staggered treatment onset (Appendix E.1). Second, I constructed a time-varying shale shock measure with the ratio of coal-to-gas prices (Appendix E.2). Results hold across these specifications.

Lastly, the shale shock could have affected voting in counties that benefited from hydraulic fracturing if new jobs offset coal’s decline, reducing regulations’ economic salience or changing migration incentives. This is less concerning since it implies that the model would underestimate the shale shock’s effect. Alternatively, if the shale shock created demand for Republican presidential candidates who promised not to regulate the gas industry (Cooper, Kim, and Urpelainen 2018; Sances and You 2022), the results would be partly spurious in counties with coal and shale deposits. To address this possibility, the models include time-varying controls for hydraulic fracturing employment.

3.2 Falsification Tests of the Shale Shock Effect on Coal Country

To increase confidence that the effect of the shale shock transmits through its economic impact on coal counties, I conducted falsification tests that evaluate the effect of the shale shock on other mining communities, including coal mining areas with varying economic exposure to coal-gas switching in electricity markets. Appendix J contains extensive details

regarding these tests, including a discussion of covariate balance and pretrends.

First, if the shale shock is driving the electoral backlash, the partisan reversal should only take place in counties that mine coal displaced by gas in electricity markets. To test this claim, I compared counties that mine thermal coal displaced by coal-to-gas switching with counties mining metallurgical anthracite coal less exposed to electricity markets. The results show that it is the economically vulnerable thermal coal mining counties that exhibit a partisan reversal, which indicates that it is the economic changes in coal mining communities driving the electoral change (Appendix J.4).

Second, I examined the effect of the shale shock in non-coal mining communities. The logic is that if a latent characteristic of mining communities was activated post-2008, the shale shock should also cause partisan reversal in these areas. But, to the contrary, the shale shock does not affect Republican vote share in these non-coal mining counties.

Third, I attempted to falsify the alternative explanation that the results are simply the consequence of layoffs for reasons independent of coal. For example, it could be that economic distress activates cultural grievances or leads to retrospective voting, which pushed voters to the right. This falsification test examined the effect of having more than a standard deviation increase in layoffs during an election year. However, there is no effect of non-coal layoffs on Republican vote share after the shale shock.

These falsification tests help rule out the possibility that the results are spurious with a post-2008 event, like the GFC or Obama’s election, or political responses to economic distress generally. There is a unique shale shock effect on coal communities that operates through an economic channel.

4 Evidence of (Mis)attribution

4.1 Geospatial Measurement of Coal-to-Gas Switching Visibility

I argue that the visibility of coal-to-gas switching in electricity markets influences attribution, which moderates the effect of the shale gas shock on the partisan reversal in coal country.

An observable implication of this theory is that in places where it is more apparent that gas is displacing coal, the effect of the shale gas boom on Republican vote share should attenuate since voters are more likely to make the correct attribution of blame.

Measuring visibility and attribution presents an exceptional challenge because there are no panel surveys with questions that probe a respondent’s understanding of the cause of coal’s decline in their communities. Yet, it is possible to test the plausibility of this hypothesis if there exists a reasonable measure of visibility that can be used to see if voter behavior in the aggregate is consistent with the theory.

To overcome this measurement challenge, I collected plant-level data from the EIA 860 form on the location of new gas-fired power plants. This administrative data has complete geographic and temporal coverage from 1992-2020, which is necessary to study electoral behavior across space and time. For each year, I calculated the distance from the center of each county to the nearest newly constructed gas-fired power plant (Appendix K). Voters in counties farther away from these changes in electricity markets should be more likely to misattribute the cause of coal’s decline and, thus, vote for the Republican candidate.

While this measure captures the objective visibility of the shale shock across space and time, a limitation is that it does not reveal the information that voters observed. Such a mismatch could occur if county boundaries do not correspond with the relevant political context in one’s mind (Wong et al. 2012). However, if there is bias from the measure, it likely would cut against the hypothesis because voters would be unaware of the objective information that the measure indicates is present.

I also conducted a small survey to assess whether coal country residents in communities with gas plants are aware of their presence. An overwhelming majority (81%) in a Southwest Pennsylvania county report passing the gas power plant at least once a year. I also find in an exploration of local news articles and press releases accompanying new gas power plants in Pennsylvania that the shale shock is often referenced as the cause (Appendix K.9). This awareness is unsurprising because rural areas have few industries, so new construction

projects are notable and accompanied by groundbreaking ceremonies covered in the local press. It is furthermore consistent with the evidence cited above about how people are aware of their economic surroundings, such as changes in immigration (Hopkins 2010), unemployment (Ansolabehere, Meredith, and Snowberg 2014), fuel prices and foreclosures (Reeves and Gimpel 2012), housing prices (Larsen et al. 2019), and trade shocks (Bisbee 2019).

The main causal identification challenge is that distance to new gas power plants may be endogenous to political and economic factors that correlate with vote choice. I address endogeneity by estimating CBPS weights that maximize covariate balance with respect to the interaction of the coal treatment indicator and the moderator, so the assignment of the treatment and moderate are statistically independent of potential outcomes (Appendix K).⁴ To assess the plausibility of this assumption, I conducted a sensitivity analysis that finds an extreme confounder 43% as strong as the correlation of poverty—the strongest predictor of distance to a gas plant—and the interaction of the moderator and the treatment would be unlikely to change the results (Appendix K.6).

A threat to inference would be if the economic gains from constructing a new gas power plant led counties to be less likely to vote for the Republican Party. To account for this possibility, the models control for power plant employment which should also be positively correlated with other economic benefits (Appendix K.7). This means the main estimate should be interpreted as the moderating effect of distance to new gas power plants *net* local economic changes, which helps to isolate the role of visibility (Appendix K.8).

Figure 4 shows the effect of the shale shock on Republican presidential vote share in coal counties, conditional on their within-state change in distance to a new gas-fired power plant. The estimates are from a linear regression model with election fixed effects, state fixed effects, time-varying controls for power plant and hydraulic fracturing employment, balancing weights, and heteroskedasticity robust standard errors clustered by county. As the visibility of coal-gas switching decreases (counties are farther away from new gas-fired

⁴Results also robust when using covariates instead of balancing weights (Appendix K).




Figure 4: **Moderating Effect of More Visible Coal-to-Gas Switching.** The plot shows the average marginal effect of being a coal county before and after the shale shock, conditional on the within-state change in distance from a new gas-fired power plant. The logged distance to a new gas power plant is standardized, so negative (positive) x -axis values indicate that a county is closer (farther). The model includes state and year fixed effects, time-varying controls for hydraulic fracturing and power industry employment, and balancing weights. HC1 standard errors are clustered by county (Table K1). Shaded bands denote 95% confidence intervals.

power plants), coal county residents become more likely to vote for Republicans after the shale shock.⁵ This finding is consistent with the claim that people in places with greater visibility of coal-to-gas switching are more likely to draw the correct attribution for coal's decline.

4.2 Interviews

I conducted interviews to identify whether people on the ground misattributed environmental regulations as the cause of coal's decline. The interviews took place in Southwest Pennsylvania in three waves during the summers of 2021, 2022, and 2023. Participants included coal miners, gas drillers, local politicians, bureaucrats, union organizers, and environmental activists. I spoke to 60 subjects over online video conferences and in person. Appendix M.2 contains details on sampling. With repeated visits to the same field site, I was able to reinterview some participants, which provided deeper insights having had established trust in

⁵Appendix K.5 checks interaction effect assumptions (Hainmueller, Mummolo, and Xu 2019).

the community.

Southwest Pennsylvania is useful to study because it is a historical and contemporary coal-producing region. Pennsylvania is also a consequential state in presidential elections. Compared to other coal counties in the treatment group, the study site was more dependent on coal mining and Democratic prior to the shale shock. The study site has comparable socio-demographic characteristics in relation to other coal mining counties. Considerable hydraulic fracturing takes place at the study site, but it amounts to no more than 5% of local employment on average, whereas coal accounts for 17%. There is low visibility of coal-to-gas switching in electricity markets at the study site because there were no new gas power plants constructed there after the shale gas boom (Appendix M.1). Studies of Western coal communities raise similar themes as in my interviews (Cha 2020), which suggests the findings should travel.

The main semi-structured question to probe respondents' thinking was an iteration of, "This area used to vote for Democrats. What do you think caused that to change?" A limitation of this question is that it asks for a retrospective evaluation that may be less accurate due to the limitations of human memory. Still, given the lack of contemporaneous interview evidence, it is the best option to learn about the considerations on voters' minds.

The potential for social desirability bias may make it harder to detect if the race or gender of presidential candidates motivated vote choice, an alternative explanation mentioned before. However, social desirability should not lead people to emphasize environmental regulations when they could instead mention other popular issues like abortion or guns. To address the potential for social desirability, I listened for coded language, such as concern about likability, which might be deployed instead of explicit gendered or racial claims.

The following quotes provide rich, suggestive evidence of a perception that the Democrat's environmental regulations were responsible for coal's decline, despite the shale shock being at fault. These interviews took place in a community where coal-to-gas was not visible, so there should be a misattribution of responsibility.

In response to a question about why his county flipped from blue to red, a long-time resident and self-identified Democrat pointed to environmental policies: “You know Obama did a lot when it came to coal, and, you know, the attack on coal...You know, you have people coming off that Obama administration era that are ticked off with the Democrats; they’re already ready to rally. And...Trump at the federal level coming in, and it just, it created this synergism that just killed the Democratic Party” (B20). The idea of an attack—or “war”—on coal has its roots in GOP and interest group messaging that aimed to blame the Obama administration’s environmental regulations (Revesz and Lienke 2016).

This response blaming environmental regulations is not isolated. Another resident said, “...that push toward the Republican side is [because of] the perceived attack on coal. And that has done huge damage in our area. I really don’t think they [Republicans] zero in on much of anything else other than, you know, attack on livelihoods. You can drive down the road when it’s election time, and you’ll see, ‘save the coal jobs’ [on signs], and then whoever’s sponsoring it and who they’re voting for” (B24). When asked about the biggest changes in the industry, a coal miner said, “regulation and downsizing...regulations continue to get tougher.” He did not mention the shale shock (B26).

I also asked local politicians about the partisan reversal. Learning from the perspective of local leaders is insightful because they have a synoptic understanding of the issues that move people. The views of local Democratic politicians are also helpful for mitigating social desirability bias since these individuals would be characterizing the behavior of their political opponents. In response to a question about the reason for the local rise of the Republican Party, a local Democratic leader said, “I give the man [Trump] credit. They drank the Kool-Aid. He backed off on some of the environment stuff...They bought into Mr. Trump and save the coal industry” (Interview A18). Another Democratic politician remarked, “[Trump’s message to bring back coal] resonated off the charts...” (B1). According to this framing, coal had to be brought back by undoing the supposed regulations that caused its decline.

The interviews also suggest that the primary source of frustration with the Democratic

Party is for *causing* the decline of the coal industry rather than a failure to *respond* with transitional policies. In listening to coal miners, community residents, and elected officials, not once did I hear someone say they wished there had been a larger federal government response to their economic dislocation. In terms of local responses, one community leader, when asked about efforts for economic revitalization, said, “There was there was limited interest, I think, from the local government of doing those things” (A1). The focus was on blame for the cause of the problem rather than the response to the decline.

4.3 News Coverage and Presidential Debates

I conducted two additional analyses to assess further the claim that coal country residents blamed regulations for their industry’s decline. First, I collected all articles covering the coal industry from two local newspapers in Southwest Pennsylvania and coded the reasons attributed as the cause of coal’s decline (Appendix N). The top reason is environmental regulations. Then, I analyzed closed-caption transcripts from Fox News and CNN to see how cable news discussed coal relative to environmental regulations and fracking. Both channels, but especially Fox News, were more likely to discuss the decline of coal in the context of regulations rather than fracking. When fracking was mentioned, it was generally in the context of whether the Democratic presidential candidate would ban it. Conservative-slanted media, in particular, connected coal’s decline to environmental regulations.

Second, I analyzed the transcripts of recent presidential debates (Appendix L.4). Searching for references to the coal industry, I find that the Republican presidential candidates pinned the blame for coal’s decline on the Democrats’ support for environmental regulations. While not systematic, the focus on coal and environmental regulations during limited debate time indicates that political elites thought it would resonate with some voters.

5 Concluding Discussion

This paper proposed and tested a new argument for part of the rightward shift in the US: the shale gas shock. This global technological disruption precipitated a partisan reversal in coal

country. The decline of coal increased the salience of the Democratic Party’s environmental issue ownership, which increased support for Republican presidential candidates in impacted counties where the source of the shock was least visible. While the partisan reversal in coal country is part of a broader rightward shift, these electoral changes took place in politically consequential states like Pennsylvania that influence presidential election outcomes.

The study’s primary limitation is the absence of individual-level panel data, which would be helpful to characterize the extent of conversion versus mobilization and measure individual attributions of responsibility for coal’s decline. The need for more panel data results from sparse survey coverage in coal-producing counties and the lack of appropriate questions measuring attributions for economic outcomes. Still, it is encouraging that others with fine-grained administrative data show partisan conversion during the same period in coal states like Pennsylvania (Hill, Hopkins, and Huber 2021), and the analysis of compositional changes indicates the findings are robust.

The paper’s argument about the effects of economic shocks on voting should apply in other cases where an industry’s decline increases the salience of a pre-existing issue cleavage. For example, the decline of manufacturing due to trade and automation has had political effects (Autor et al. 2020). However, unlike the Democrats’ consistent environmental issue ownership, the party’s position on trade has shifted as its political base evolved. Issue positions are important, but perhaps even more so is the attribution of blame.

In this respect, this paper’s argument about visibility should generalize beyond the shale gas shock. For economic disruptions from trade, offshoring, and automation, there is also political contestation over the attribution of blame (Mutz 2021; Wu 2022). Future research should construct measures of the visibility of these shocks, such as individual exposure to automation in the workplace and daily lives or whether workers had to train their foreign replacements before offshoring.

The paper makes three contributions across scholarly literatures. First, my new explanation for part of the partisan reversal nuances our understanding of the rightward shift

in working-class communities that traditionally supported Democrats. While a burgeoning literature emphasizes racial backlash (Abrajano and Hajnal 2015; Sides, Tesler, and Vavreck 2018; Mutz 2018; Reny, Collingwood, and Valenzuela 2019) and globalization as a cause of the populist right’s rise (Autor et al. 2020; Baccini and Weymouth 2021; Milner 2021; Walter 2021), this paper shows how the rightward shift in the US is multi-faceted: the confluence of racial backlash, deindustrialization, trade shocks, and the decline of coal, as I argue, led to an increase in Republican vote share, but for distinct reasons.

Second, I build on studies about political context (e.g., Hopkins 2010; Huckfeldt and Sprague 1995) to develop an argument about how visibility can explain when voters make more accurate attributions of responsibility for economic outcomes. This theory moves beyond skepticism about voter competence (e.g., Achen and Bartels 2016) to, instead, differentiate the conditions under which voters’ behavior approximates a bounded rationality model of decision-making and under which they behave in apparently non-rational ways. This argument complements other critiques of “blind retrospection” (Ashworth, Bueno de Mesquita, and Friedenbergh 2018), which call for a better understanding of what voters learn from external events. Visibility provides one explanation for what voters find informative and how they act under information constraints. In parallel research, I am further developing the theoretical basis for the role of visibility and gathering individual-level data, which is necessary to more confidently assess its political effects. As a first cut, this paper hopes to have demonstrated the plausibility of the claim that the visibility of an economic shock’s cause influences attribution and vote choice.

Lastly, I executed a credible research design that provides estimates of the potential electoral backlash from the clean energy transition. The shale shock is a case of unmanaged decline, where the government did not provide meaningful transitional assistance. Future research should explore whether reducing transition costs could quell electoral backlash, as some work has begun to examine (Kono 2020; Bolet, Green, and Gonzalez-Eguino 2023). There has been mixed success with compensation for trade losses (Broz, Frieden, and Wey-

mouth 2021). The lack of credible commitments to assist workers and communities contributes to the challenge (Gazmararian and Tingley 2023). The political success of the clean energy transition may depend on whether governments can lessen the costs of climate action, which could otherwise spur backlash that rolls back reforms.

References

- Abrajano, M., and Z. L. Hajnal. 2015. *White Backlash*. Princeton University Press.
- Acemoglu, D., et al. 2023. *Climate Change, Directed Innovation, and Energy Transition: The Long-run Consequences of the Shale Gas Revolution*. NBER.
- Achen, C., and L. Bartels. 2016. *Democracy for Realists*. Princeton University Press.
- Ansolabehere, S., M. Meredith, and E. Snowberg. 2014. “Mecro-Economic Voting: Local Information and Micro-Perceptions of the Macro-Economy.” *Economics & Politics* 26.
- Arnold, R. D. 1990. *The Logic of Congressional Action*. Yale University Press.
- Ashworth, S., E. Bueno de Mesquita, and A. Friedenberg. 2018. “Learning about Voter Rationality.” *American Journal of Political Science* 62 (1): 37–54.
- Autor, D., et al. 2020. “Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure.” *American Economic Review* 110 (10): 3139–3183.
- Baccini, L., and S. Weymouth. 2021. “Gone For Good: Deindustrialization, White Voter Backlash, and US Presidential Voting.” *American Political Science Review* 115 (2).
- Ballard-Rosa, C., J. Goldstein, and N. Rudra. 2022. “Trade As Villain: The Fading American Dream and Declining Support for Globalization.” *Journal of Politics* Forthcoming.
- Ballard-Rosa, C., A. Jensen, and K. Scheve. 2021. “Economic Decline, Social Identity, and Authoritarian Values in the United States.” *International Studies Quarterly*, sqab027.

- Bisbee, J. 2019. “What You See out Your Back Door: How Political Beliefs Respond to Local Trade Shocks.” Unpublished Manuscript, 2019.
- Bolet, D., F. Green, and M. Gonzalez-Eguino. 2023. “How to Get Coal Country to Vote for Climate Policy: The Effect of a ‘Just Transition Agreement’ on Spanish Election Results.” Preprint, accessed April 17, 2023. <https://doi.org/10.2139/ssrn.4394195>.
- Broz, J. L., J. Frieden, and S. Weymouth. 2021. “Populism in Place: The Economic Geography of the Globalization Backlash.” *International Organization*, 1–31.
- Canes-Wrone, B., D. W. Brady, and J. F. Cogan. 2002. “Out of Step, Out of Office.” *American Political Science Review* 96 (1): 127–140.
- Cha, J. M. 2020. “A Just Transition for Whom?” *Energy Research & Social Science* 69:101657.
- Cinelli, C., and C. Hazlett. 2020. “Making Sense of Sensitivity: Extending Omitted Variable Bias.” *Journal of the Royal Statistical Society Series B* 82 (1): 39–67.
- Coglianesi, J., T. Gerarden, and J. Stock. 2020. “The Effects of Fuel Prices, Environmental Regulations, and Other Factors on US Coal Production, 2008-2016.” *Energy Journal*.
- Cooper, J., S. E. Kim, and J. Urpelainen. 2018. “The Broad Impact of a Narrow Conflict: How Natural Resource Windfalls Shape Policy and Politics.” *Journal of Politics* 80 (2).
- Cramer, K. J. 2016. *The Politics of Resentment*. University of Chicago Press.
- Dark, T. 1999. *The Unions and the Democrats*. Cornell University Press.
- DOE. 2017. *Staff Report to the Secretary on Electricity Markets and Reliability*.
- Downs, A. 1957. *An Economic Theory of Democracy*. Harper and Row.
- Eckert, F., et al. 2020. *Imputing Missing Values in the US Census Bureau’s County Business Patterns*. w26632. National Bureau of Economic Research.
- Egan, P. J. 2013. *Partisan Priorities*. Cambridge University Press.

- Egli, F., N. Schmid, and T. S. Schmidt. 2022. “Backlash to Fossil Fuel Phase-Outs.” *Environmental Research Letters* 17 (9): 094002.
- Frank, T. 2007. *What’s the Matter with Kansas?* Picador.
- Gaikwad, N., F. Genovese, and D. Tingley. 2022. “Creating Climate Coalitions: Mass Preferences for Compensating Vulnerability in the World’s Two Largest Democracies.” *American Political Science Review* 116 (4): 1165–1183.
- Gaventa, J. 1982. *Power and Powerlessness*. University of Illinois Press.
- Gazmararian, A. F., and D. Tingley. 2023. *Uncertain Futures: How to Unlock the Climate Impasse*. Cambridge University Press.
- Haggerty, J. H., et al. 2018. “Planning for the Local Impacts of Coal Facility Closure: Emerging Strategies in the U.S. West.” *Resources Policy* 57:69–80.
- Hainmueller, J., J. Mummolo, and Y. Xu. 2019. “How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice.” *Political Analysis* 27 (2): 163–192.
- Hall, A., J. Yoder, and N. Karandikar. 2021. “Economic Distress and Voting: Evidence from the Subprime Mortgage Crisis.” *Political Science Research and Methods* 9 (2): 327–344.
- Harder, A. 2012. “Coal Miners’ Union Sits Out Presidential Race.” *National Journal*, 2012.
- Hartman, E., and F. D. Hidalgo. 2018. “An Equivalence Approach to Balance and Placebo Tests.” *American Journal of Political Science* 62 (4): 1000–1013.
- Healy, A., and N. Malhotra. 2013. “Retrospective Voting Reconsidered.” *Annual Review of Political Science* 16 (1): 285–306.
- Hill, S., D. Hopkins, and G. Huber. 2021. “Not by Turnout Alone: Measuring the Sources of Electoral Change, 2012 to 2016.” *Science Advances* 7 (17): eabe3272.

- Ho, D. E., et al. 2007. “Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference.” *Political Analysis* 15 (3): 199–236.
- Holladay, S., and J. LaRiviere. 2017. “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.
- Hopkins, D. J. 2010. “Politicized Places: Explaining Where and When Immigrants Provoke Local Opposition.” *American Political Science Review* 104 (1): 40–60.
- Huckfeldt, R. R., and J. Sprague. 1995. *Citizens, Politics and Social Communication: Information and Influence in an Election Campaign*. Cambridge University Press, 1995.
- Imai, K., and M. Ratkovic. 2014. “Covariate Balancing Propensity Score.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76 (1): 243–263.
- Iyengar, S. 1989. “How Citizens Think about National Issues: A Matter of Responsibility.” *American Journal of Political Science* 33 (4): 878–900.
- Kono, D. Y. 2020. “Compensating for the Climate: Unemployment Insurance and Climate Change Votes.” *Political Studies* 68 (1): 167–186.
- Larsen, M. V., et al. 2019. “When Do Citizens Respond Politically to the Local Economy?” *American Political Science Review* 113 (2): 499–516.
- Lawhorn, J. 2022. *The POWER Initiative*. Congressional Research Service, 2022.
- Leip, D. 2020. *Dave Leip’s Atlas of U.S. Presidential Elections*.
- Liu, L., Y. Wang, and Y. Xu. 2022. “A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data.” *American Journal of Political Science* Forthcoming.
- Margalit, Y. 2019. “Political Responses to Economic Shocks.” *Annual Review of Political Science* 22 (1): 277–295.

- Mian, A., and A. Sufi. 2011. "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis." *American Economic Review* 101, no. 5 (2011): 2132–2156.
- Mildenberger, M. 2020. *Carbon Captured: How Business and Labor Control Climate Politics*. MIT Press.
- Milner, H. V. 2021. "Voting for Populism in Europe: Globalization, Technological Change, and the Extreme Right." *Comparative Political Studies* 54 (13): 2286–2320.
- Mondak, J., D. C. Mutz, and R. Huckfeldt. 1996. "Persuasion in Context: The Multilevel Structure of Economic Evaluations." In *Political Persuasion and Attitude Change*, edited by D. Mutz, P. Sniderman, and R. Brody, 249–266. U. Michigan Press.
- Mutz, D. 1994. "Contextualizing Personal Experience: The Role of Mass Media." *Journal of Politics* 56 (3): 689–714.
- Mutz, D. C. 2018. "Status Threat, Not Economic Hardship, Explains the 2016 Presidential Vote." *Proceedings of the National Academy of Sciences* 115 (19): E4330–E4339.
- . 2021. "(Mis)Attributing the Causes of American Job Loss." *Public Opinion Quarterly* 85, no. 1 (2021): 101–122.
- Norris, P., and R. Inglehart. 2019. *Cultural Backlash*. Cambridge University Press.
- Powell, G. B., and G. D. Whitten. 1993. "A Cross-National Analysis of Economic Voting." *American Journal of Political Science* 37 (2): 391–414.
- Reed, A., et al. 2019. "Interrogating Uncertainty in Energy Forecasts: The Case of the Shale Gas Boom." *Energy Transitions* 3, no. 1 (2019): 1–11.
- Reeves, A., and J. G. Gimpel. 2012. "Ecologies of Unease: Geographic Context and National Economic Evaluations." *Political Behavior* 34 (3): 507–534.

- Reny, T., L. Collingwood, and A. Valenzuela. 2019. "Vote Switching in the 2016 Election: How Racial and Immigration Attitudes, Not Economics, Explain Shifts in White Voting." *Public Opinion Quarterly* 83, no. 1 (2019): 91–113.
- Revesz, R. L., and J. Lienke. 2016. *Struggling for Air: Power Plants and the "War on Coal"*. Oxford University Press.
- Roemer, K. F., and J. H. Haggerty. 2021. "Coal Communities and the U.S. Energy Transition: A Policy Corridors Assessment." *Energy Policy* 151:112112.
- Sances, M. W., and H. Y. You. 2022. "Voters and Donors: The Unequal Political Consequences of Fracking." *Journal of Politics* 84 (3): 1667–1682.
- Sheldon, P., R. Junankar, and A. Pontello de Rosa. 2018. *The Ruhr or Appalachia?* University of New South Wales.
- Sides, J., M. Tesler, and L. Vavreck. 2018. *Identity Crisis*. Princeton University Press.
- Simon, H. A. 1955. "A Behavioral Model of Rational Choice." *The Quarterly Journal of Economics* 69 (1): 99–118.
- Walter, S. 2021. "The Backlash Against Globalization." *Annual Review of Political Science* 24 (1): 421–442.
- Wong, C., et al. 2012. "Bringing the Person Back In: Boundaries, Perceptions, and the Measurement of Racial Context." *The Journal of Politics* 74 (4): 1153–1170.
- Wu, N. 2022. "Misattributed Blame? Attitudes toward Globalization in the Age of Automation." *Political Science Research and Methods* 10 (3): 470–487.
- Yergin, D. 2020. *The New Map*. Penguin Press.

Online Appendix: Sources of Partisan Change

A	Shale Shock	2
A.1	Primary Cause of Coal’s Decline	2
A.2	Contemporaneous Awareness	4
A.2.1	Government Officials	4
A.2.2	News Media	4
B	Mine Safety Salience	6
B.1	Mine Fatalities	6
B.2	Estimating Topics in UMW Journal	7
C	Measurement	8
C.1	County-Level Controls	8
C.2	Summary Statistics	10
C.3	Treatment Definition	11
C.3.1	Rationale for 1% Cutoff	11
C.3.2	Robustness to Alternative Treatments	12
D	DiD Research Design	13
D.1	Matching Diagnostics	13
D.2	FEct Regression Results	16
D.3	Matrix Completion Estimator	17
D.4	Sensitivity Analysis	19
E	Time-Varying Shale Shocks	23
E.1	State-Specific Shale Shocks	23
E.1.1	Estimating State-Specific Shale Shocks	23
E.1.2	Structural Break Estimation Procedure	26
E.1.3	Treatment History	28
E.1.4	State-Specific Shale Shocks Results	28
E.2	Coal-to-Gas Price Ratio	30
E.2.1	Measurement	30
E.2.2	Summary Statistics	30
E.2.3	Time-Varying Shale Shock Results	30
F	Employment Loss as Treatment	34
F.1	Matching Diagnostics	34
F.2	Treatment History	35
F.3	Employment Loss Results	36

G	Shift-Share Research Design	38
G.1	Causal Inference Assumptions	38
G.2	Shift-Share Results	39
H	Heterogeneous Treatment Effects	41
H.1	Unions	41
H.1.1	Mine-Level Unionization Rates	41
H.1.2	Results Controlling for Unionization	42
H.1.3	National Trends in Coal Unionization	44
H.1.4	Treatment Heterogeneity by Union Strength	44
H.2	Transitional Assistance	46
H.2.1	Measurement	46
H.2.2	Treatment Heterogeneity by Transitional Assistance	47
I	Robustness to Compositional Change	49
I.1	Study Demonstrating Conversion in Coal States During Study Period	49
I.2	Reducing Bias through Estimation Strategy	49
I.3	Bounding Compositional Bias Using Counterfactuals	50
I.3.1	Simple Model of Voting and Population Change	50
I.3.2	Model Operationalization	51
I.3.3	Counterfactual Estimates	52
I.4	Reducing Compositional Bias with More Proximate Pre-Treatment Windows	53
I.5	Migration Data	55
J	Falsification Tests	57
J.1	Non-Coal Mining Communities	57
J.2	Non-Coal Layoffs	58
J.3	Falsification Test Results	59
J.4	Anthracite Coal Falsification Test	62
J.4.1	Empirical Validation of Anthracite Resilience to Shale Shock	62
J.4.2	Anthracite Falsification Test Results	65
K	Moderating Effect of Visibility	66
K.1	Data and Measurement	66
K.2	Covariate Balance	68
K.3	Estimating Equation	69
K.4	Visibility Regression Results	69
K.5	Interaction Effect Diagnostics	71
K.5.1	Linearity Assumption Diagnostic Tests	71
K.5.2	Flexible Estimation Strategies	72
K.6	Sensitivity Analysis	73
K.7	New Gas Plants Do Not Offset Tax Revenue Loss	75
K.8	Post-Treatment Bias	76
K.8.1	DAGs Assessing Post-Treatment Bias	76
K.8.2	Sequential-g Estimation	77

K.9	Microfoundational Evidence	80
K.9.1	Survey	80
K.9.2	News	80
L	Long-Standing National Partisan Environmental Issue Cleavage	83
L.1	Studies on US Environmental Policy Polarization	83
L.2	Analysis of Party Platforms	83
L.3	Climate Policy Came on National Agenda in 1980s/90s	85
L.4	Politicization of Environmental Regulations in Presidential Debates	86
M	Fieldwork	87
M.1	Study Site Representativeness	87
M.2	Interviews	89
M.3	Surveys	92
N	News Media	96
N.1	Local News Media Coverage	96
N.2	National News	98
O	Community-Oriented Preferences	102
P	Research Ethics	103

A Shale Shock

A.1 Primary Cause of Coal’s Decline

This appendix reviews findings from econometric studies that evaluate the effect of the shale gas boom on coal production. These studies substantiate the claim that the shale gas boom is the primary reason for the coal industry’s decline, which implies that placing blame on environmental regulations is a misattribution.

- Coglianese, Gerarden, and Stock (2020) conclude that 92% of the reduction in coal production between 2008 and 2016 is due to declining gas prices relative to coal. Their main methodological approach uses monthly state-level data to decompose the decline of coal’s share of electricity generation into the effect of nine different factors. To measure exposure to environmental regulations, they construct a time series at the state level that measures whether individual CAA regulations affect coal generators in a given state.
- Brehm (2019) conclude that “[f]alling US natural gas prices caused more gas-fired (and less coal-fired) electricity generation...[and] 65–85% of new gas-fired plants were constructed because of lower gas prices.”
- Though not their primary focus, Davis, Holladay, and Sims (2021, 15) find that “low natural gas prices increase the probability of [coal plant] retirement.”
- Knittel, Metaxoglou, and Trindade (2016, 242) make a summary judgment that “[t]he coal-to-gas switching observed in electricity generation has been attributed to a large extent to the introduction of hydraulic fracturing, or fracking, which has led to a dramatic reduction in the price of natural gas..” They further write, “The dramatic increase in fracking activity led to an unprecedented decrease in the price of natural gas, changing fundamentally the landscape of the US electric power sector via coal-to-gas switching” (246).
- Mohlin et al. (2019) examine carbon emissions and conclude that gas substitution for coal accounts for 60% of the emissions decline from 2005 to 2015. While an indication of fuel-switching, emissions differ from the outcome of interest: actual coal production.
- In a review of the literature on the shale shock, Black et al. (2021, 313) write, “Coal-fired electricity generation declined by more than 60% between 2007 and 2019 (authors’ calculations using US EIA data), which was driven by the increasing competitiveness of natural gas-fired electricity generation...”
- Watson, Lange, and Linn (2023, 37) use data from 2002 to 2012 and estimate that “[b]y 2012, the effect of low gas prices reduced average mine life by 5.7 years, while lower-than-expected electricity demand reduced mine life by 4.3 years.” While they identify declining mine productivity as an important factor, they note that “[s]ince 2012, which is the end of our study period, natural gas prices have remained low, and industry-wide productivity increased because of the continuing closure of lower-productivity mines. Thus, while declining productivity between 2002 and 2012 substantially reduced mine

lifetimes, after around 2010, demand-side factors (particularly low gas prices) became more important.”

- Linn and McCormack (2019, 20) examine 2005 to 2015 with a new electricity sector model and conclude that “the market shocks explain...99 percent of the decline in coal-fired plant profits...This importance of the natural gas price shock is consistent with the empirical literature.” Profits are a distinct outcome but should be correlated with coal production.
- Linn and Muehlenbachs (2018, 2) examine 2001 to 2012 with plant-level coal and natural gas prices to estimate elasticities that imply “roughly a 1:1 shift from coal to gas-fired generation from a natural gas price decrease.”
- Joskow (2013) describes how “[c]heap natural gas is in turn leading to higher capacity factors for existing natural gas electric generating plants, reducing the dispatch of coal plants...,” but does not quantify the contribution of gas to this change.
- DOE (2017) examines electricity markets from 2002 to 2017 and finds that the relatively low natural gas prices are the biggest driver of recent coal plant retirements. The Trump administration DOE report concluded, “[t]he biggest contributor to coal and nuclear plant retirements has been the advantaged economics of natural gas-fired generation” (13).

A.2 Contemporaneous Awareness

There is documentary evidence from news articles and Congressional hearings that policy-makers were aware that the shale gas revolution was the reason for coal’s decline while it was happening. Thus, it is reasonable to be puzzled about why people did not recognize the shale shock as the cause of their community’s decline while it happened.

A.2.1 Government Officials

For example, on 29 October 2013, the US Congress held a hearing entitled, “EPA’s Regulatory Threat to Affordable, Reliable Energy: The Perspective of Coal Communities.” Republican Tim Murphy from Pennsylvania opened the hearing by preempting an argument he knew would come from the Democratic committee members. He said,

We will hear from some who say coal plants are closing because natural gas is cheaper. Not true. They are closing because the EPA refuses to work out solutions that help coal move forward to be even cleaner than it already is (p. 3)

Indeed, the Democratic committee member, Diana DeGette from Colorado, opened her statement by saying,

We do also need to talk about the real reality that as natural gas becomes cheaper than coal and more and more other utilities and others transfer to natural gas, it is the invisible hand of the free market. Utilities are moving to natural gas because it makes business sense. So we do need to talk about that, and as we think about what is happening with the loss of jobs in coal country, we need to think about the inevitable hand of the free market and what we do about that. (p. 5)

A.2.2 News Media

On 27 February 2011, an article in the *Pittsburgh Post-Gazette* entitled “A Boom Without A Bust?; For Now, The Future Is Bright for Shale Economics.” The article reported how “natural gas also remains cheap enough to be an attractive option to coal. Natural gas is cleaner burning than coal, too, one of the reasons that more than two full years have passed since ground was broken on a new coal-fired power plant in the United States.”

On 11 January 2013, *NPR*’s flagship show, *All Thing’s Considered*, ran a segment called, “Coal Loses Crown As King Of Power Generation.” The story described how for one utility, “within a few years only a third of the company’s power plants will run on coal. The company has already built three new natural gas plants.” The article went on to note how, “After decades in which coal was king of electricity generation, natural gas is making a bid for the title.” The journalists clearly noted the shale gas boom as the cause: “In board rooms across the country, electric companies are deciding that many coal plants, especially small, older ones, just don’t make economic sense any more. One factor is the expectation that low prices for natural gas will continue because of the shale gas boom across the country.”

An op-ed written on 3 July 2012 in *Charleston Gazette* — a local West Virginia newspaper — also placed clear blame on the shale shock, indicating contemporaneous awareness of the

shifts in electricity markets. Note that this article was written before the general election. The article entitled, "Hardship; More coal layoffs," said:

Much of coal's retreat can be blamed on an abundance of cheap natural gas loosed by horizontal drilling and "fracking" in the deep Marcellus Shale. Gas that cost \$10 per million Btu during the Bush administration has dropped to \$2.50. Coal can't compete with such low-cost fuel. Many plants are converting to gas.


...

Some politicians and coal executives blame federal pollution controls - which they call a "war on coal" - for the mining downturn. But that's merely a small part of the picture. Sen. Jay Rockefeller wisely told coal leaders to "accept reality" and try to solve their difficulty without accusing Washington.

B Mine Safety Salience

B.1 Mine Fatalities


Figure B1: Coal Miner Fatalities per 10,000 Workers, 1972-2020



Notes: Data from MSHA. Estimates are adjusted by the number of workers in a given year. Office workers were included starting in 1973. Blue line represents the LOESS fit. Shaded bands are 95% confidence intervals.

B.2 Estimating Topics in UMW Journal

Figure B2: Estimated Topic Proportions in the UMW Journal, 1975-1985



Notes: Estimates from using a structural topic (STM) model (Roberts et al. 2014). STM estimates include a covariate for the year of publication. Editions of the UMW Journal run from August 1975 to August 1985. The author accessed microfilms archives of the UMWJ and hand-recorded the title of each edition as well as the table of contents where possible. The number of topics is chosen by benchmarking results between 5, 7, and 10 topics.

C Measurement

C.1 County-Level Controls

County-level covariates used for matching and in empirical models come from the 2000 Decennial Census. Employing measures from 2000 minimizes bias from post-treatment controls. For example, the shale shock likely created migratory pressures that altered the age, gender, and racial distribution of the population. Including time-varying measures of these factors after treatment would introduce a subtle form of selection bias. The following list describes the operationalization of each measurement. Each measure is standardized within-state by subtracting the state mean from a county’s observed value and dividing by the state deviation.⁶

- Race: Percent of white residents.
- Ethnicity: Percent of Hispanic/Latino residents.
- Foreign-born: Percent of foreign-born residents.

The race, ethnicity, and foreign-born covariates measure racial homogeneity that may be predictive of racial backlash (e.g., Abrajano and Hajnal 2015; Sides, Tesler, and Vavreck 2018; Mutz 2018; Jardina 2019), as well as political context like immigration (e.g., Hopkins 2010).

- Education: Percent of residents over the age of 25 with a bachelor’s degree or higher. This variable accounts for the “diploma gap,” which Sides, Tesler, and Vavreck (2018) emphasize as a determinant of racialized party realignment.
- Income per capita: County income per capita in 1999 dollars (logged).
- Poverty: Percent of the county residents who live below the federal poverty line.

Income and poverty (as well as education) are correlated with class, which is thought to characterize the population that moves from blue to red (e.g., Frank 2007; Morgan and Lee 2017; Morgan 2018).

- Rurality: Percent of the county that is non-agricultural rural.
- Population: County population (logged) accounts for the size of the labor market. Rurality and population account for the propensity to mine for coal, which is less likely to occur in densely populated urban areas. Previous studies also identify rural resentment as one factor driving support for the Republican Party (e.g., Cramer 2016).
- Age: Percent of the county under 40 years of age, which accounts for age cohort effects and labor market composition.

⁶This standardization isolates the relevant variation because the models that include these covariates have state fixed effects (Mummolo and Peterson 2018). These pre-treatment covariates are not appropriate to include in the models that estimate within-county changes in Republican voting since the covariates are perfectly colinear with the county fixed effects and would drop out in estimation.

- Female workforce participation: Percentage of female residents in the labor force. On the relevance of female workforce participation and far-right backlash, see Clark, Khoban, and Zucker (2022).

C.2 Summary Statistics

Table C1: Summary Statistics

	Coal			Matching Pool			Rest of Country		
	Mean	SD	NA	Mean	SD	NA	Mean	SD	NA
Pretreatment:									
White	0.90	0.00	0	0.87	0.00	0	0.84	0.00	7
Hispanic	0.03	0.00	0	0.06	0.00	0	0.06	0.00	7
Foreign-born	0.01	0.00	0	0.03	0.00	0	0.04	0.00	7
College	0.08	0.00	0	0.10	0.00	0	0.11	0.00	7
Income per capita (log)	9.62	0.00	0	9.70	0.00	0	9.77	0.00	7
Poverty	0.18	0.00	0	0.15	0.00	0	0.13	0.00	7
Rural	0.65	0.00	0	0.57	0.00	0	0.54	0.00	7
Population (log)	10.21	0.00	0	10.12	0.00	0	10.27	0.00	7
Under 40 years	0.54	0.00	0	0.54	0.00	0	0.54	0.00	7
Debt-to-income	1.37	0.00	0	1.74	0.00	0	1.77	0.00	15
Female workforce	0.21	0.00	0	0.23	0.00	0	0.24	0.00	7
Time-Varying:									
Gas plant distance (1992–2020)	0.20	0.99	0	0.06	1.02	0	-0.03	0.99	5
Hydraulic fracturing employment	0.85	2.37	115	0.58	2.50	657	0.44	2.43	2333
Coal employment	11.12	13.64	115	0.46	2.63	657	0.02	0.46	2333
Power employment	0.74	2.91	115	0.18	1.09	657	0.09	0.52	2333
Two-party Republican vote share	57.62	15.22	0	60.42	13.99	0	58.28	14.18	1
Observations:									
Counties	116	NA	NA	658	NA	NA	2338	NA	NA
<i>N</i>	1505	NA	NA	8552	NA	NA	30368	NA	NA

Notes: Coal counties are those with greater than 1% of local employment in coal during 2005–2007. The matching pool consists of counties within two degrees of adjacency of a coal county. The rest of the country is a residual category that includes all counties not in the first two. Pre-treatment socio-demographic measures come from the 2000 U.S. Census. Data on the debt-to-income ratio are from 2006 and are based on Mian and Sufi (2011). County-level employment data come from the CBP reports that have had missing values imputed according to Eckert et al. (2020). Election data come from Leip (2020).

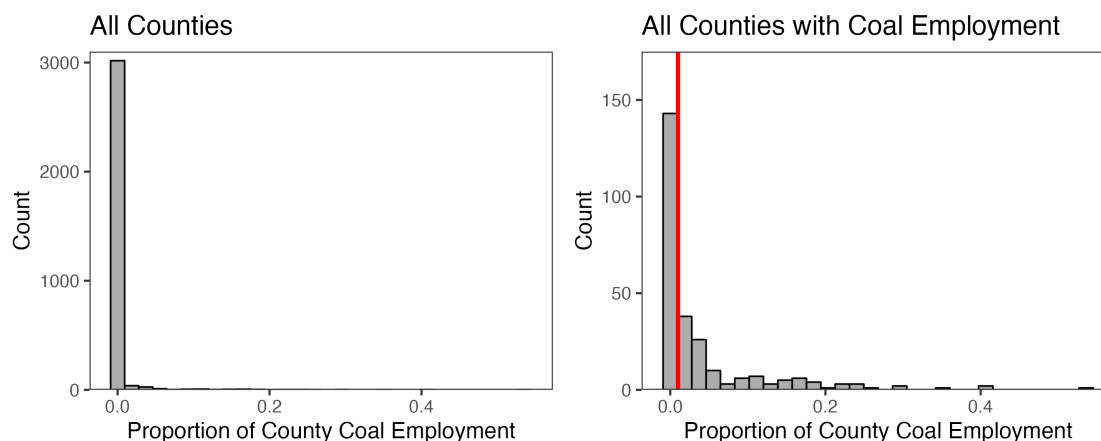
C.3 Treatment Definition

C.3.1 Rationale for 1% Cutoff

The threshold for treatment group membership is having more than 1% of county employment in the coal industry in any of the three years prior to the shale shock. Theory and data guide this choice. In terms of theory, the population of interest is counties where there was an active coal mining prior to the shale shock. It does not have to be a tremendous amount because of how the community impacts of coal's decline are theorized to shape preferences. However, it is reasonable to suspect that these social forces are less likely to be influential in vote choice in a place where only 0.2% of the county, for instance, works in coal.

Of course, the selection of any threshold involves some arbitrariness. To make this choice less ad hoc, I let the distribution of the data guide the decision. Figure C1 shows the distribution of the county-level average proportion of coal employment as a share of overall county employment from 2005-2007. The left plot shows that across much of the country, the modal county has no coal employment. The right figure shows that among counties with any coal employment, the modal county has less than 1% of employment in coal. In counties with less than 1% of local coal employment, the average share of local coal employment is 0.2% (median 0.1%). If these counties were included in the treatment group, they would receive disproportionate weight, despite being a poor match to the theory. By contrast, in counties defined as part of the treatment group with more than 1% of local coal employment, the average share of local coal employment is 9% (median 4%). This exploration of the raw data led to the selection of 1% as the threshold.

Figure C1: Distribution of County-Level Proportion of Coal Employment



Notes: Employment data from CBP. Coal industry employment as a proportion of county employment is averaged over 2005-2007. The left panel shows the distribution of coal employment in all counties, while the right panel shows the distribution in counties with any coal employment. The red line denotes the 1% threshold employed to define treatment group status.

As an additional test of construct validity, I estimate the correlation of the treatment indicator with the share of county GDP from coal extraction. GDP data come from the Bureau of Economic Analysis in the Department of Commerce. To match the employment data, which are averaged in the years prior to the shale shock, I also averaged the county

GDP data in the same years, so the regression analyzes a common cross-section. Table C2 presents the results from regressing the share of county GDP from coal extraction on different operationalizations of the treatment. Model 1 shows that the treatment has a strong, positive correlation with county GDP from coal extraction; going from under 1% local coal employment to 1% or more correlates with having 20.2% more local GDP from coal extraction. Model 2 shows that the coefficient is larger if we were to define the treatment threshold at 10%, but that this would not dramatically improve model fit compared to the original treatment definition. Model 3 shows that using the continuous share of coal employment (standardized for interpretation) does not have as large of a correlation with county GDP compared to our treatment definition in Model 1. In all, this correlative analysis lends support to using the 1% of coal employment as the threshold for treatment assignment.

Table C2: Linear regression of county GDP from coal extraction on the treatment and alternative coal employment thresholds

	(1)	(2)	(3)
Intercept	0.05*** (0.00)	0.05*** (0.00)	0.06*** (0.00)
Above 1% Coal Employment (=1)	0.20*** (0.02)		
Above 10% Coal Employment (=1)		0.37*** (0.03)	
Coal Employment			0.04*** (0.00)
N	3084	3084	3084
Adjusted R^2	0.070	0.081	0.094

Notes: HC2 standard errors clustered by county. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

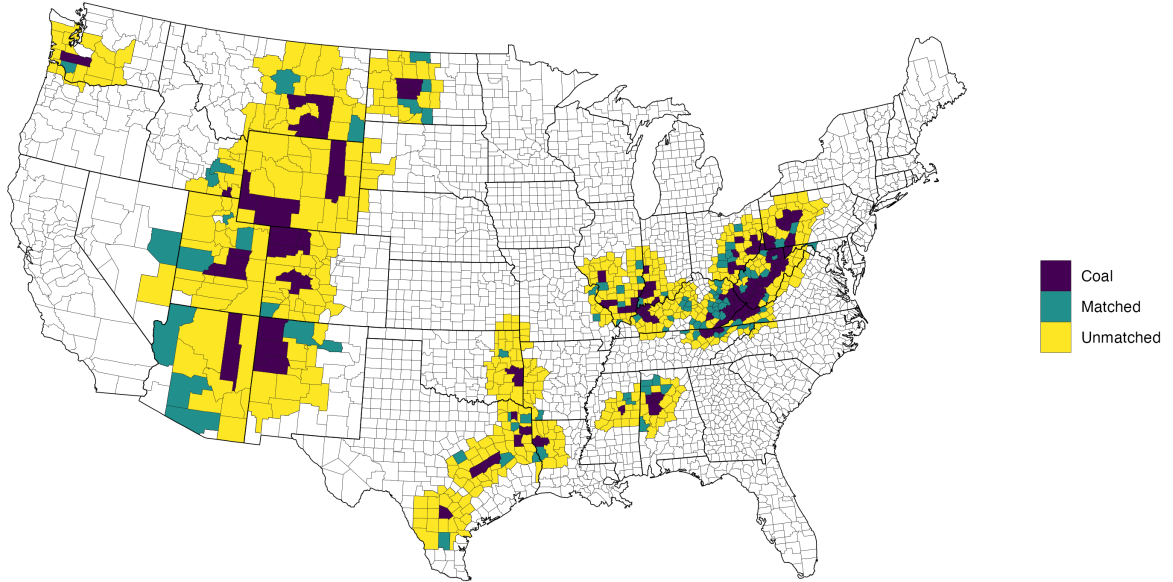
C.3.2 Robustness to Alternative Treatments

The results are robust to using a continuous measure (Figure J1; Table J2). The advantage of a continuous measure is that it avoids the potential limitations of reliance on a threshold to define treatment status. Furthermore, the results are also robust when defining the treatment in terms of actual layoff experience (Appendix F).

D DiD Research Design

D.1 Matching Diagnostics

Figure D1: Matched Treatment and Control Counties



Notes: Treatment counties are those with more than 1% of local employment in the coal industry from 2005-2007. The pool of possible control counties contains those with two degrees of adjacency from the treatment group.

Table D1: Covariate Balance Before and After Weighting

	Mean		Standardized Mean	
	Control	Treated	Control	Treated
Original:				
White	0.87	0.90	6.22	6.42
Hispanic	0.06	0.03	0.50	0.25
Foreign-born	0.03	0.01	0.86	0.48
College	0.10	0.08	2.03	1.58
Income per capita (log)	9.70	9.62	49.07	48.65
Poverty	0.15	0.18	2.30	2.74
Rural	0.58	0.65	2.23	2.53
Population (log)	10.12	10.21	8.30	8.38
Under 40 years	0.54	0.54	9.56	9.52
Female workforce	0.23	0.21	7.36	6.72
Balanced:				
White	0.90	0.90	6.42	6.42
Hispanic	0.03	0.03	0.25	0.25
Foreign-born	0.01	0.01	0.48	0.48
College	0.08	0.08	1.58	1.58
Income per capita (log)	9.62	9.62	48.65	48.65
Poverty	0.18	0.18	2.74	2.74
Rural	0.65	0.65	2.53	2.53
Population (log)	10.21	10.21	8.38	8.38
Under 40 years	0.54	0.54	9.52	9.52
Female workforce	0.21	0.21	6.72	6.72


Notes: Treatment is whether a county had coal employment between 2005-2007. Figure D1 plots the spatial distribution of treatment counties and the population of possible matches. Covariate balancing propensity score estimated (Imai and Ratkovic 2014). Total of 775 counties, including treated units and the matching pool of possible control units. Covariate data from the 2000 U.S. Census. Coal employment data from CBP.

Table D2: Covariate Balance After Nearest Neighbor Matching

	Control (N=97)		Treatment (N=116)		Diff.	<i>p</i>
	Mean	SD	Mean	SD		
White	0.90	0.13	0.90	0.15	0.00	0.90
Hispanic	0.05	0.15	0.03	0.07	-0.02	0.24
Foreign-born	0.02	0.02	0.01	0.02	0.00	0.27
College	0.08	0.03	0.08	0.04	0.00	0.72
Income per capita (log)	9.61	0.19	9.62	0.17	0.01	0.82
Poverty	0.18	0.08	0.18	0.07	0.00	0.73
Rural	0.67	0.26	0.65	0.21	-0.01	0.68
Population (log)	10.03	1.22	10.21	0.79	0.18	0.21
Under 40 years	0.54	0.05	0.54	0.05	0.00	0.84
Female workforce	0.21	0.03	0.21	0.03	0.00	0.56

Notes: Treatment is whether a county had coal employment between 2005-2007. Figure D1 plots the spatial distribution of treatment counties and the population of possible matches. Covariate balancing propensity score estimated (Imai and Ratkovic 2014). Nearest-neighbor matching using the method in Ho et al. (2007). Covariate data from the 2000 U.S. Census. Coal employment data from CBP.

Figure D2: Two-Party Republican Vote Share Pretrends for Matched Counties



Notes: Treatment counties are defined as those with more than 1% of local employment in the coal industry during 2005-2007. Control counties are those matched with coal counties using nearest-neighbor matching. Election data from Leip (2020).

D.2 FEct Regression Results

Table D3: Shale Shock Effect on Two-Party Republican Presidential Vote Share

	Estimate	S.E.	CI _{2.5%}	CI _{97.5%}	<i>p</i> -value
ATT:					
Observations equally weighted	4.89	1.47	2.00	7.78	0.00
Units equally weighted	5.06	1.47	2.16	7.95	0.00
Covariates:					
Hydraulic Fracturing Employment	0.02	0.14	-0.26	0.29	0.91
Placebo Tests:					
-2 to 0 election interval	1.00	0.76	-0.48	2.48	0.19


Notes: Presidential elections from 1972–2020 in the matched sample of 213 counties ($N = 2,766$). Standard errors and confidence intervals constructed using 2,000 block bootstrap replications clustered at the unit level.

D.3 Matrix Completion Estimator

This appendix reports the results of using the matrix completion (MC) estimator for the counterfactual outcomes (Athey et al. 2021; Liu, Wang, and Xu 2022). This estimator works by constructing a lower rank approximation of the outcome data matrix using information from untreated observations to account for potential time-varying confounders. The advantage of this estimator over the FEct is that it provides more leverage to account for potential time-varying confounding.

Figure D3 shows the dynamic treatment effect estimates from the MC estimator of the shale shock on the two-party Republican vote share in exposed coal counties. The results are consistent in magnitude and precision as those presented in the main text. That is, the shale shock caused Republican presidential vote share to increase by over 4 percentage points in exposed coal counties (Table D4). As before, the equivalence test indicates that equivalence in the pretrends holds with high confidence.

Figure D3: MC Estimates of Dynamic Treatment Effect of the Shale Shock on Two-Party Republican Vote Share



Notes: The left plot shows the dynamic treatment effects estimates using the FEct estimator with the matrix completion algorithm. The MC method was selected using cross-validated comparison with an interactive fixed effects (IFE) model. k -fold cross-validation is employed to select the hyper-parameters in the matrix completion algorithm. Treated counties are those with coal employment three years before the 2008 shale shock, which are matched to adjacent control counties using socio-economic covariates ($N = 213$). The model includes county and election fixed effects and controls for hydraulic fracturing employment. The bars denote 95% confidence intervals from 2,000 block bootstrap replications clustered at the unit level. The right plot shows the pretreatment average prediction errors and their 90% confidence intervals. The red dashed lines denote the equivalence range set at $[-0.36\hat{\sigma}, 0.36\hat{\sigma}]$ as proposed by Hartman and Hidalgo (2018), whereas the black dashed lines mark the minimum range. The p -value indicates equivalence holds with high confidence.

Table D4: Matrix Completion Estimate of the Shale Shock Effect on Two-Party Republican Presidential Vote Share

	Estimate	S.E.	CI _{2.5%}	CI _{97.5%}	<i>p</i> -value
ATT:					
Observations equally weighted	4.76	1.39	2.04	7.48	0.00
Units equally weighted	NA	1.38	NA	NA	NA
Covariates:					
Hydraulic Fracturing Employment	0.00	0.13	-0.25	0.26	0.97
Placebo Tests:					
-2 to 0 election interval	1.23	0.70	-0.14	2.59	0.08

Notes: The matrix completion method was selected using cross-validated comparison with an interactive fixed effects (IFE) model. Cross-validation is employed to select the hyper-parameters in the matrix completion algorithm. Presidential elections from 1972–2020 in the matched sample of 213 counties ($N = 2,766$). Standard errors and confidence intervals constructed using 2,000 block bootstrap replications clustered at the unit level.

D.4 Sensitivity Analysis

A sensitivity analysis supplies an estimate of how strong an unobserved confounder—one that is orthogonal to the covariates in the model—would have to be to bring the ATT estimate to 0 (a bias of 100% of the original estimate). I use the `sensemakr` package, which implements the methodology proposed by Cinelli and Hazlett (2020). Since the interpretation of sensitivity analysis requires benchmark covariates and the only covariate in the main model, hydraulic fracturing employment, has a weak correlation with the outcome, I estimate a linear regression model with a fuller set of pre-treatment covariates. While this requires substituting county fixed effects (that are otherwise co-linear with the pre-treatment covariates) with state fixed effects, this should make the sensitivity analysis even more conservative since there is likely to be even greater unobserved confounding than in the primary analysis. Specifically, I regress two-party Republican presidential vote share (y_{it}) on the interaction of the coal treatment indicator ($Coal_i$) and the post-shock indicator ($Post_t$) with controls for hydraulic fracturing employment, white, Hispanic, foreign-born, college, income per capita (log), poverty, rural, population (log), under 40 years, and female workforce participation (\mathbf{X}). The model interacts the demographic covariates with the post-shock indicator to allow them to have a differential effect pre- and post-shock.

$$y_{it} = Coal_i + Post_t + \delta(Coal_i \times Post_t) + \mathbf{X}_{it}^T \beta + \gamma(Post_t \times \mathbf{X}_{it}^T) + Election_t + State_i + \epsilon_{it}. \quad (1)$$

To keep the units comparable, I estimate this model using the matched counties from the first analysis. Table D5 presents the results from estimating this model. HC1 standard errors are clustered by county.

Table D5: Linear Regression of Two-Party Republican Presidential Vote Share on the Interaction of Coal County Treatment Group Membership and Post-Shale Shock, 1972–2020

	(1)	(2)	(3)	(4)
Treatment and Moderator:				
Coal	−3.48*** (1.24)	−3.73*** (1.45)	−3.19** (1.27)	−3.44** (1.48)
Post-Shale Shock	22.48*** (1.33)	8.78*** (1.09)	26.92*** (1.86)	12.86*** (1.63)
Coal × Post-Shale Shock	3.70*** (1.27)	3.90*** (1.34)	2.83** (1.25)	3.19** (1.30)
Time-Varying Covariates:				
Hydraulic Fracturing Employment	−0.22 (0.35)	−0.34 (0.37)	−0.23 (0.35)	−0.35 (0.37)
Coal Union Share		0.40 (1.83)		0.47 (1.82)
Coal Union Share × Coal Union Share		1.68 (2.09)		1.16 (2.06)
Pretreatment Covariates:				
White	2.14*** (0.80)	2.11** (0.84)	2.01** (0.81)	1.98** (0.85)
White × Post-Shale Shock	2.54** (1.00)	2.56*** (0.90)	2.92*** (1.01)	2.89*** (0.91)
Hispanic	−4.00*** (1.15)	−4.09*** (1.15)	−3.94*** (1.12)	−4.03*** (1.11)
Hispanic × Post-Shale Shock	−1.07 (1.11)	−0.88 (1.00)	−1.24 (1.03)	−1.04 (0.92)
Foreign-born	3.02** (1.26)	3.37*** (1.29)	3.12** (1.25)	3.47*** (1.28)
Foreign-born × Post-Shale Shock	4.36*** (1.36)	3.62*** (1.20)	4.08*** (1.27)	3.37*** (1.13)
College	−0.76 (1.62)	−1.13 (1.62)	−0.87 (1.61)	−1.25 (1.61)
College × Post-Shale Shock	−6.00*** (1.75)	−5.42*** (1.55)	−5.65*** (1.68)	−5.12*** (1.50)
Income per capita (log)	−3.02* (1.72)	−3.06* (1.73)	−2.92* (1.73)	−2.95* (1.74)
Income per capita (log) × Post-Shale Shock	2.13 (1.79)	1.88 (1.65)	1.82 (1.81)	1.64 (1.67)
Poverty	−3.19*** (1.12)	−3.78*** (1.14)	−3.13*** (1.12)	−3.72*** (1.14)
Poverty × Post-Shale Shock	1.70* (1.01)	2.15** (0.93)	1.52 (0.98)	1.99** (0.91)
Rural	−0.88 (1.10)	−0.84 (1.12)	−1.04 (1.11)	−1.02 (1.14)
Rural × Post-Shale Shock	1.65 (1.08)	1.34 (0.98)	2.11* (1.08)	1.77* (0.99)
Population (log)	−0.94 (1.34)	−1.25 (1.36)	−0.92 (1.34)	−1.23 (1.36)
Population (log) × Post-Shale Shock	4.18*** (1.30)	4.31*** (1.18)	4.11*** (1.28)	4.25*** (1.16)
Under 40 years	0.62 (0.78)	0.43 (0.79)	0.61 (0.78)	0.42 (0.80)
Under 40 years × Post-Shale Shock	−2.39*** (0.87)	−2.26*** (0.80)	−2.35*** (0.87)	−2.22*** (0.80)
Female workforce	1.27 (1.20)	1.32 (1.22)	1.12 (1.19)	1.15 (1.22)
Female workforce × Post-Shale Shock	−3.26*** (1.06)	−3.27*** (0.96)	−2.82*** (1.02)	−2.88*** (0.93)
Debt-to-Income			0.87 (0.90)	0.94 (0.90)
Debt-to-Income × Post-Shale Shock			−2.57*** (0.85)	−2.36*** (0.77)
<i>N</i>	2554	2130	2554	2130
Adjusted <i>R</i> ²	0.647	0.671	0.649	0.674
Outcome Mean	56.85	56.58	56.85	56.58
Outcome SD	12.1	11.97	12.1	11.97
State Fixed Effects	Yes	Yes	Yes	Yes
Election Fixed Effects	Yes	Yes	Yes	Yes

Notes: The population of counties in the model is the same as used for matching (Figure D1), those with above 1% of local employment in the coal industry and those within two degrees of adjacency of the coal counties. HCl standard errors clustered by county. **p* < 0.1, ***p* < 0.05, ****p* < 0.01

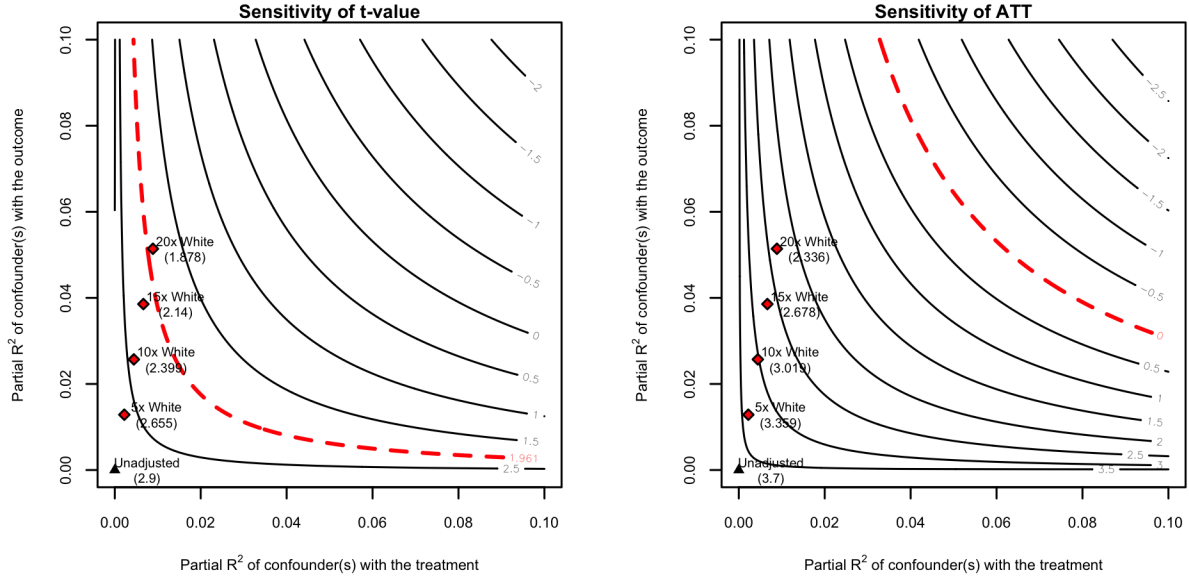
Table D6 presents the results from the sensitivity analysis of the finding in Table D5 with respect to the interaction of post-shock and coal counties. Specifically, the table below shows the maximum strength of unobserved confounders with association with the treatment and the outcome bounded by a multiple of the observed explanatory power of the county share of the white population, our benchmark covariate. I select this covariate given its substantive relevance, and the magnitude of its effect size makes it a formidable test. Interpreting the results in the table, the robustness value indicates that an unobserved confounder (orthogonal to the covariates) that explains more than 5.7% of the residual variance of both the treatment and the outcome would be strong enough to bring the point estimate to 0 (a bias of 100% of the original estimate). In terms of reducing the estimate to a range where it is no longer statistically different from 0 at the 5% significance level, there would have to be an unobserved confounder (orthogonal to the covariate) that explains more than 1.9% of the residual variance of both the treatment and the outcome. The benchmarking exercise presented in Figure D4 indicates that such a confounder is unlikely. It would have to have more than 1,500% the strength of the share of the white population’s correlation with both the outcome and treatment while also being orthogonal to the covariates in the model.

Table D6: Sensitivity Analysis to Unobserved Confounding

Outcome: Two-Party Republican Presidential Vote Share						
Treatment:	Est.	S.E.	t-value	$R^2_{Y \sim D \mathbf{X}}$	$RV_{q=1}$	$RV_{q=1, \alpha=0.05}$
<i>Coal \times Post-Shale Shock</i>	3.699	1.272	2.909	0.3%	5.7%	1.9%
df = 2494	<i>Bound (5x White): $R^2_{Y \sim Z \mathbf{X}, D} = 1.3\%$, $R^2_{D \sim Z \mathbf{X}} = 0.2\%$</i>					

Notes: Sensitivity analysis conducted according to the method proposed by Cinelli and Hazlett (2020). Model subjected to sensitivity analysis from Table D5. Note that this test is a more conservative estimate since the model has state rather than county fixed effects, so there should be greater potential for omitted variable bias in this sensitivity analysis than for an equivalent analysis of the primary results. The reason for using the model with state fixed effects is that it permits the conclusion of benchmark covariates, namely the county share of the white population.

Figure D4: Sensitivity Analysis of Main Results



Notes: Bias contour plots of the t -value (left) and ATT estimate (right). Red diamonds indicate that a confounder up to 1500% as strong as the observed white covariate would not bring the lower bound of the confidence below 0 at the 5% significance level, while a confounder at least 2000% as strong as the observed white covariate would not bring the estimate to 0. Estimates come from a linear regression of two-party Republican presidential vote share on the interaction of the coal treatment indicator and the post-shock indicator with controls for hydraulic fracturing employment, white, Hispanic, foreign-born, college, income per capita (log), poverty, rural, population (log), under 40 years, and female workforce participation. The model interacts the demographic covariates with the post-shock indicator to allow them to have a differential effect pre- and post-shock. Since the aim is to benchmark the ATT against a relevant covariate, the model includes state and year fixed effects (unlike the primary model that includes county fixed effects that would, in this case, be co-linear with the white benchmark covariate). The omission of county fixed effects makes this an even more conservative test since there should be greater potential for omitted variable bias. HC1 standard errors clustered by county. Matched sample employed.

E Time-Varying Shale Shocks

E.1 State-Specific Shale Shocks

Although there was a national structural break in gas prices in 2008, the effects of the shale shock might unfold differentially across space and time depending on constraints in local gas markets, such as the availability of pipelines. Since much of coal is locally traded, with coal-fired power plants often co-locating near mine mouths, it is plausible to assume that changes in state-level gas prices influence demand for coal within a state and associated employment in places where there are mines.⁷

To capture this spatial and temporal heterogeneity—and to help ensure that the results are not spurious with other structural changes that may have occurred in 2008—I estimate the timing of the shale shock in each state. The state is the lowest level of aggregation for which there is reliable gas price data for an extended period. This appendix proceeds by explaining how I estimate the state-specific shale shocks. Then I present results from models that allow for staggered treatment onset.

E.1.1 Estimating State-Specific Shale Shocks

Estimating the precise timing of the shale shock in each state requires high frequency and spatially disaggregated data on gas prices. The state represents the lowest unit of analysis at which there is frequent monthly data on gas prices from the EIA at a long enough interval to estimate structural breaks in prices. The ideal data would be the price that electric utilities pay for gas since it is at these firms where coal-to-gas switching occurs after the shale shock. However, there is a high degree of missingness (around 44%) in the EIA data on the gas price that electric utilities pay. Instead, I use citygate monthly gas price data at the state level, which has better coverage with only 2% of observations missing from 2002-2022. The citygate is where gas is transferred from a pipeline to the local gas utility. Prices can vary at this level of analysis due to different consumer demands, weather, and pipeline coverage.

I conducted a statistical analysis to validate that citygate prices are a valid proxy for the gas price paid by electric utilities. Specifically, I regress the available data on gas prices paid by electric utilities on state-level citygate gas prices from 2002 to 2022. Table E1 reports results from this linear regression with heteroskedasticity robust standard errors clustered by state to account for autocorrelation. Across various specifications, the most stringent being model (4), which includes month and state fixed effects, there is a strong positive correlation between citygate prices and the local price for gas paid by electric utilities. Figure E1 presents visual evidence of this correlation and further evidence of a global price decline after the 2008 shale shock.

I descriptively examine the price data to determine whether there is sufficient variation in state-level citygate prices to warrant the estimation of state-level shale shocks. Figure E2 plots each state’s trajectory of citygate gas prices over time. The grey lines that depict the state-level trajectories show considerable heterogeneity. While there was a common peak in 2008 followed by a decline as markets entered a new price regime, there is still regional heterogeneity.


⁷For example, Haggerty et al. (2018) describes how western states operate a “coal-by-wire” model, where mine-mouth facilities export electricity to urban centers elsewhere.

Table E1: Linear Regression of Monthly State-Level Gas Prices for Electric Utilities on State-Level Citygate Gas Prices, 2002–2022

	Outcome: Gas Prices for Electric Utilities			
	(1)	(2)	(3)	(4)
Intercept	0.566* (0.329)	−0.419 (0.468)	0.555 (0.627)	1.037 (0.640)
Citygate Price	0.788*** (0.053)	0.863*** (0.059)	0.634*** (0.129)	0.507*** (0.140)
N	7210	7210	7210	7210
Adjusted R^2	0.440	0.489	0.531	0.628
State Fixed Effects	No	Yes	Yes	Yes
Year Fixed Effects	No	No	Yes	No
Month Fixed Effects	No	No	No	Yes


Notes: Price data from the EIA. A large share of observations for the gas prices electric utilities pay is missing (44%). This analysis aims to see whether the citygate data, where there is less missingness (2%), have a sufficiently strong correlation with the non-missing electric utility gas price data to be a valid proxy in the subsequent estimation of state-level shale shocks. Heteroskedasticity-robust standard errors clustered at the state level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Figure E1: Average State-Level Gas Price at the Citygate and Electric Utilities, 2002–2022.



Notes: locally estimated scatterplot smoothing (LOESS) line plotted with 95% confidence intervals around the annual state average. The vertical red line denotes 2008, the start of the shale shock. Data from EIA.

Figure E2: State-Specific Trends in Citygate Gas Prices, 2002–2022



Notes: Blue line represents the moving state-level average estimated with LOESS. Each gray line represents the observed state-level trend in citygate gas prices. Gas price units are in dollars per thousand cubic feet. The vertical red line denotes 2008, the start of the shale shock. Data from EIA.

E.1.2 Structural Break Estimation Procedure


Following Andrews (1993) and Holladay and LaRiviere (2017), I estimate structural breaks in each state’s gas market by fitting an intercept and a high dimensional time trend to panel data on citygate gas prices with an indicator variable for all observations after a candidate break point.⁸ Then, I select each state’s change point with the highest Wald statistic. For this analysis, I focus on citygate prices between January 2004 and January 2020, the primary window when the shale shock could occur. After 2020 there is an uptick in gas prices, which would introduce unnecessary noise into the analysis intended to capture a state-level break occurring around 2008.

There is a decisive structural break in citygate gas prices in all but three states. Here, decisive refers to the p -value for the structural break point test statistic being less than 5%. The three exceptions are Alaska ($p < 0.19$), Colorado ($p < 0.1$), and Delaware ($p < 0.17$). The subsequent empirical model uses the implied structural breaks for these states. Since the structural breaks may be less precisely identified, this could introduce noise into our results, possibly introducing bias against our hypothesis. Outside these three states, the median p -value for the test statistic is less than 0.001.

Figure E3 presents the estimated structural break in each state’s gas markets. Baseline prices differ across methane gas markets (e.g., gas is much more expensive in Hawaii), so the plot allows the y -axis to vary with each state to reflect the within-state variation better.

⁸Citygate prices are not available for the District of Columbia.

Figure E3: State-Specific Shale Shocks




Notes: Vertical red line denotes the estimated structural break in city-gate methane gas prices in each state. The solid black line is the moving state-level average estimated with LOESS. The gray line represents the observed state-level trend in citygate gas prices. Gas price units are in dollars per thousand cubic feet. Data from EIA.

E.1.3 Treatment History

Figure E4 plots the treatment onset timing for the counties in the matched sample. Once the state-specific shale shock occurs, counties with coal employment are treated.

Figure E4: State-Specific Shale Shocks Treatment Onset




Notes: Treated units are counties with greater than 1% local employment in the coal industry in the period after the state-specific shale shock.

E.1.4 State-Specific Shale Shocks Results

I estimate the effect of state-specific shale shocks in exposed coal counties compared to their matched neighbors on two-party Republican presidential vote share using the FEct estimator. This method is well-suited when there is staggered treatment onset.

Figure E5 presents the dynamic treatment effects plot alongside an equivalence test for potential violations of the parallel trends assumption. When more precisely measuring the timing of the shale shock in each state, Table E2 shows that the ATT estimate grows to 5.2 percentage points (from around 4.5 percentage points in the primary analysis). The state-specific shale gas shock has a strong, positive effect on Republican presidential vote share in exposed coal counties. With respect to the parallel trends assumption, the p -value indicates equivalence holds with high confidence.

Figure E5: Dynamic Treatment Effect of State-Specific Shale Shocks on Two-Party Republican Vote Share



Notes: The left plot shows the dynamic treatment effects estimates using the FEct estimator. Treated counties are those with coal employment three years before the 2008 shale shock, which are matched to adjacent control counties using socio-economic covariates ($N = 213$). The model includes county and election fixed effects and controls for hydraulic fracturing employment. The bars denote 95% confidence intervals from 2,000 block bootstrap replications clustered at the unit level. The right plot shows the pretreatment average prediction errors and their 90% confidence intervals. The red dashed lines denote the equivalence range set at $[-0.36\hat{\sigma}, 0.36\hat{\sigma}]$ as proposed by Hartman and Hidalgo (2018), whereas the black dashed lines mark the minimum range. The p -value indicates equivalence holds with high confidence.

Table E2: FEct Estimate of the Shale Shock Effect on Two-Party Republican Presidential Vote Share with State-Specific Shale Shocks

	Estimate	S.E.	CI _{2.5%}	CI _{97.5%}	p -value
ATT:					
Observations equally weighted	5.76	1.41	3.00	8.53	0.00
Units equally weighted	5.33	1.44	2.51	8.14	0.00
Covariates:					
Hydraulic Fracturing Employment	0.03	0.14	-0.25	0.30	0.86
Placebo Tests:					
-2 to 0 election interval	0.87	0.73	-0.55	2.30	0.23

Notes: Presidential elections from 1972–2020 in the matched sample of 213 counties ($N = 2,766$). Standard errors and confidence intervals constructed using 2,000 block bootstrap replications clustered at the unit level.

E.2 Coal-to-Gas Price Ratio

I construct a time-varying measure of the shale shock by taking the ratio of annual coal to methane gas prices. The subsequent analysis interacts this ratio with a county’s pre-shock share of coal employment to capture the effect of the shale shock on places most dependent on coal before the shale shock began.

E.2.1 Measurement

Coal data come from a combination of EIA time series on the open market and captive coal sale prices in the United States. Specifically, I use data from the EIA’s Annual Energy Review, Table 7.9, from 1949-2001. Then, I use data from EIA’s Form EIA-7A for 2001-2021. To verify that no bias from measurement is introduced by combining the data from these time series, Figure E6 plots the average annual coal price in both time series alongside each other.⁹ There is a tight visual correlation between the time series in the 10 years when they overlap. Unsurprisingly, the correlation coefficient is 0.99 ($p < 0.001$). I prefer to use the more recent data product for the post-2000 period to maximize the consistency of the data around the time when the shale shock occurs. This helps to make sure that the measurement of the timing of the shale shock reflects actual price changes rather than bias from a change in how the EIA measures coal prices.

To measure gas prices, I use EIA data on the Henry Hub gas spot prices. These data run from 1997-2020. The Henry Hub is a gas pipeline that is the official delivery location for futures contracts on the New York Mercantile Exchange. Since the hub has access to major gas markets in the US, it is the natural benchmark for the North American methane gas market. This price is based on actual supply and demand, unlike in European markets, where gas delivery is more fragmented.

E.2.2 Summary Statistics

I operationalize the ratio by dividing the real price of coal by the Henry Hub gas spot price, which I then standardize for interpretation by subtracting the mean value and dividing by the standard deviation. When the ratio is large, it indicates that coal is more expensive than gas, so there should be market signals for electric utilities to fuel-switch. Figure E7 shows a sizeable positive deviation that began in 2008 at the start of the shale shock.


E.2.3 Time-Varying Shale Shock Results

Using the time-varying measure of the shale shock, I estimate the effect of an increase in the intensity of the shale shock on two-party Republican presidential vote share for counties more exposed to the shale shock due to their higher levels of local coal employment. Specifically, I estimate the following model,

$$y_{it} = Coal_i + Ratio_t + \delta(Coal_i \times Ratio_t) + \mathbf{X}_{it}^T \beta + \alpha_i + \eta_t, \quad (2)$$

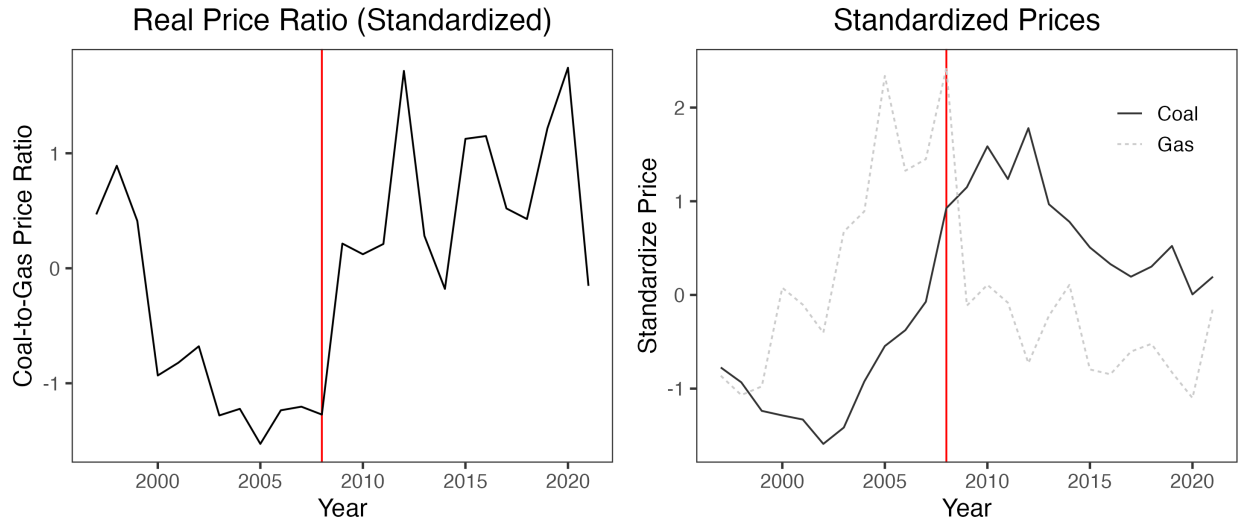
⁹While gas prices begin to increase in the post-2020 period relative to coal prices, this is irrelevant to my analysis which ends in 2020. Further, there is a path dependence with electric power infrastructure. Once a power company decides to build a plant and moves away from coal, those decisions become locked in since they are multi-million dollar assets that cannot be easily replaced.

Figure E6: Temporal Comparison of Average Annual Open Market and Captive Coal Price



Notes: Data represented by the solid black line from 1949-2011 come from the EIA’s Annual Energy Review, Table 7.9. For 1949-2000, prices are for open market and captive coal sales; for 2001-2007, prices are for open market coal sales; for 2008 forward, prices are for open market and captive coal sales. Data represented by the dashed grey line from 2001-2021 come from the EIA’s Form EIA-7A, “Coal Production and Preparation Report.” The line represents the average of average open market and average captive market coal prices.

Figure E7: Ratio of Coal-to-Gas Prices



Notes: The left plot shows the standardized coal-to-gas price ratio. The right plot shows coal and gas prices, independently, each standardized. The vertical red line denotes 2008, the start of the shale shock.

where y_{it} is the two-party Republican presidential vote share in a county-year, $Coal_i$ is an indicator for if a county has more than 1% of local employment in coal before the shale shock, $Ratio_t$ is the standardized coal-to-gas price ratio constructed in this section, δ is our

interaction term of interest, \mathbf{X}_{it}^T is a matrix that contains the time-varying local share of hydraulic fracturing employment, α represents county fixed effects, and η denotes state fixed effects. The careful reader will note that the county fixed effects are colinear with the coal indicator, so the constituent term will drop out during estimation. While I can recover the interaction term of interest, δ , this feature is limiting. Therefore, I also estimate models that use state fixed effects instead while controlling for the same county-level covariates used for matching. As with earlier cross-sectional models, I interact these county-level covariates with a post-shock indicator to allow them to have a differential effect on Republican vote share across elections.

Table E3 presents the results. Across all models, the shale shock in coal counties has a strong, positive effect on two-party Republican presidential vote share. In model (1), which uses the sample from the coal counties and their neighbors with two degrees of adjacency, a standard deviation increase in the intensity of the shale shock causes a 3.5 percentage point increase in Republican vote share. This effect size falls in model (2), which uses the entire nation as the sample but remains positive and precise. Models (3) and (4) replace county fixed effects with state fixed effects and add covariates to see how the magnitude of the shale shock effect size compares with other county-level characteristics. The shale shock effect in coal counties is around 2.6 percentage points, which is larger than the correlation between the share of white county residents in the post-shale shock period, about a 1.5 percentage point increase in Republican vote share.

Table E3: Linear Regression of Two-Party Republican Presidential Vote Share on the Interaction of the Time-Varying Shale Shock Measure and County Vulnerability, 2000–2020

	(1)	(2)	(3)	(4)
Intercept			61.35*** (1.33)	60.97*** (1.29)
Treatment and Moderator:				
Coal			−0.30 (0.76)	
Pre-Shale Shock Coal Employment				−0.04 (0.14)
Coal-to-Gas Price Ratio	3.55*** (0.20)	1.73*** (0.03)	2.61*** (0.06)	2.83*** (0.08)
Coal × Post-Shale Shock	1.47*** (0.28)	3.22*** (0.16)	2.67*** (0.26)	0.29*** (0.07)
Time-Varying Covariates:				
Hydraulic Fracturing Employment	−0.29** (0.13)	0.04 (0.02)	0.21*** (0.07)	0.21*** (0.08)
Pretreatment Covariates:				
White			4.51*** (0.28)	4.57*** (0.31)
White × Post-Shale Shock			1.53*** (0.16)	1.62*** (0.20)
Hispanic			−0.05 (0.32)	−0.42 (0.40)
Hispanic × Post-Shale Shock			−0.59*** (0.17)	−0.68*** (0.21)
Foreign-born			0.91** (0.39)	1.08** (0.49)
Foreign-born × Post-Shale Shock			0.05 (0.20)	−0.07 (0.24)
College			−0.67** (0.33)	−1.23*** (0.41)
College × Post-Shale Shock			−2.58*** (0.22)	−2.62*** (0.28)
Income per capita (log) × Post-Shale Shock			0.47* (0.29)	0.29 (0.35)
Poverty			−2.90*** (0.34)	−3.22*** (0.40)
Poverty × Post-Shale Shock			0.23 (0.21)	0.23 (0.25)
Rural			0.31 (0.26)	0.35 (0.31)
Rural × Post-Shale Shock			0.73*** (0.18)	0.94*** (0.22)
Population (log)			−1.37*** (0.31)	−1.37*** (0.37)
Population (log) × Post-Shale Shock			0.58*** (0.21)	0.93*** (0.25)
Under 40 years			0.82*** (0.22)	1.00*** (0.26)
Under 40 years × Post-Shale Shock			−0.92*** (0.16)	−1.08*** (0.19)
Female workforce			−0.70*** (0.27)	−0.58* (0.32)
Female workforce × Post-Shale Shock			−0.71*** (0.17)	−0.88*** (0.21)
<i>N</i>	1278	18 663	18 658	13 666
Adjusted <i>R</i> ²	0.383	0.053	0.619	0.603
Outcome Mean	58.871	58.706	56.847	56.847
Outcome SD	11.196	9.387	12.099	12.099
Sample	Matched	Nation	Nation	Nation
County Fixed Effects	Yes	Yes	No	No
State Fixed Effects	No	No	Yes	Yes
Election Fixed Effects	Yes	Yes	Yes	Yes

Notes: Model (1) estimates the effect of the time-varying shale shock in the sample of matched counties (Figure D1); model (2) estimates the effect of the time-varying shale shock in all counties; model (3) estimates the effect of the time-varying shale shock in all counties but replaces county fixed effects with state fixed effects and adds county-level controls; and model (4) is the same as the former but instead of a binary indicator for being a coal county before the shale shock, the coal variable is the local share of coal employment in 2000. All covariates are standardized. All models use HC2 standard errors clustered by county. Note that the outcome means and standard deviations differ across the models with county and state fixed effects because they are standardized at the county and state levels, respectively, as recommended by Mummolo and Peterson (2018). **p* < 0.1, ***p* < 0.05, ****p* < 0.01

F Employment Loss as Treatment

F.1 Matching Diagnostics

Table F1: Covariate Balance Before and After Weighting for Coal Employment Loss Counties

	Mean		Standardized Mean	
	Control	Treated	Control	Treated
Original:				
White	0.85	0.88	5.24	5.48
Hispanic	0.06	0.04	0.52	0.33
Foreign-born	0.04	0.02	0.73	0.36
College	0.11	0.09	2.15	1.85
Income per capita (log)	9.75	9.68	46.65	46.32
Poverty	0.14	0.16	2.15	2.56
Rural	0.55	0.57	1.94	2.01
Population (log)	10.23	10.34	7.29	7.36
Under 40 years	0.54	0.54	9.21	9.19
Female workforce	0.24	0.22	7.98	7.46
Balanced:				
White	0.88	0.88	5.48	5.48
Hispanic	0.04	0.04	0.33	0.33
Foreign-born	0.02	0.02	0.37	0.36
College	0.09	0.09	1.85	1.85
Income per capita (log)	9.68	9.68	46.33	46.32
Poverty	0.16	0.16	2.56	2.56
Rural	0.57	0.57	2.01	2.01
Population (log)	10.34	10.34	7.36	7.36
Under 40 years	0.54	0.54	9.19	9.19
Female workforce	0.22	0.22	7.46	7.46


Notes: Treatment is whether a county had coal employment between 2005-2007 and lost coal employment after the shale shock. Covariate balancing propensity score estimated (Imai and Ratkovic 2014). Total of 3110 counties, including treated units and the matching pool of possible control units. Covariate data from the 2000 U.S. Census. Coal employment data from CBP.

Table F2: Covariate Balance After Nearest Neighbor Matching for Coal Employment Loss Counties

	Control (N=129)		Treatment (N=149)			
	Mean	SD	Mean	SD	Diff.	<i>p</i>
White	0.88	0.14	0.88	0.15	0.00	0.92
Hispanic	0.05	0.16	0.04	0.08	-0.01	0.38
Foreign-born	0.02	0.03	0.02	0.02	0.00	0.59
College	0.09	0.04	0.09	0.04	0.00	0.31
Income per capita (log)	9.66	0.22	9.68	0.18	0.02	0.39
Poverty	0.16	0.08	0.16	0.07	0.00	0.91
Rural	0.57	0.24	0.57	0.24	0.00	0.93
Population (log)	10.22	1.07	10.34	0.98	0.12	0.34
Under 40 years	0.54	0.05	0.54	0.05	0.00	0.72
Female workforce	0.22	0.03	0.22	0.03	0.00	0.77

Notes: Treatment is whether a county had coal employment between 2005-2007 and lost coal employment after the shale shock. Covariate balancing propensity score estimated (Imai and Ratkovic 2014). Nearest-neighbor matching using the method in Ho et al. (2007). Covariate data from the 2000 U.S. Census. Coal employment data from CBP.


Figure F1: Two-Party Republican Vote Share Pretrends for Matched Counties with Employment Loss as Treatment



Notes: Treatment counties are defined as those with more than 1% of local employment in the coal industry during 2005-2007 and experienced coal employment losses after the shale shock. Control counties are those matched with coal counties using nearest-neighbor matching. Election data from Leip (2020).

F.2 Treatment History

Figure F2: Post-Shale Shock Coal Employment Loss Treatment Onset



Notes: Treated units are counties that experienced a decline in local coal employment after the shale shock began.

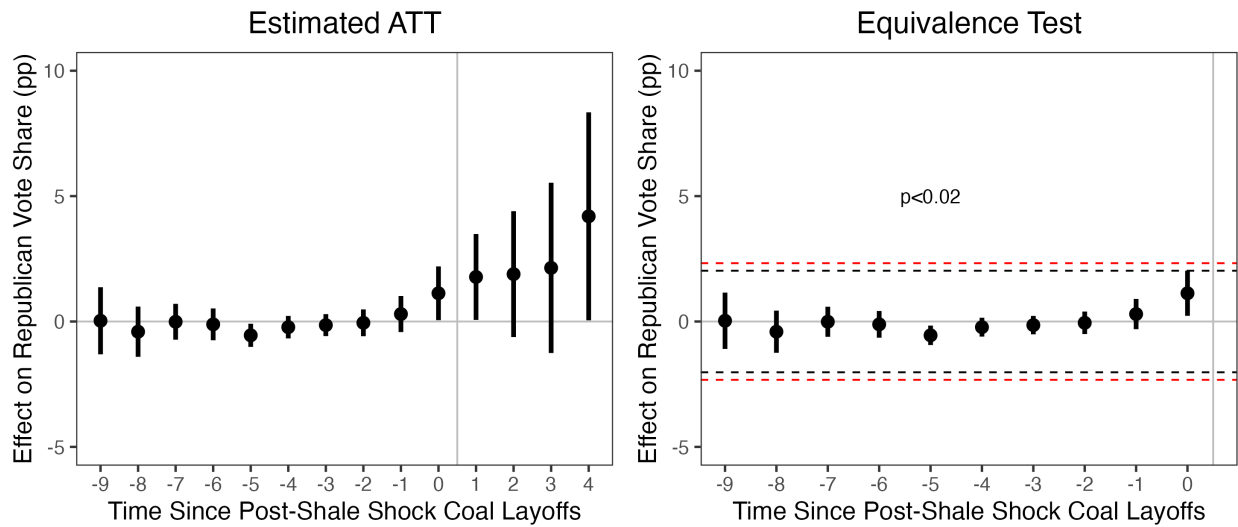
F.3 Employment Loss Results

Table F3: Effect of Coal Employment Loss After the Shale Shock on Two-Party Republican Presidential Vote Share

	Estimate	S.E.	CI _{2.5%}	CI _{97.5%}	<i>p</i> -value
ATT:					
Observations equally weighted	2.19	1.22	-0.19	4.57	0.07
Units equally weighted	2.60	1.13	0.39	4.81	0.02
Covariates:					
Hydraulic Fracturing Employment	0.10	0.11	-0.11	0.32	0.35
Placebo Tests:					
-2 to 0 election interval	1.06	0.84	-0.58	2.70	0.21

Notes: Presidential elections from 1972–2020 in the matched sample of 278 counties ($N = 3,611$). Standard errors and confidence intervals constructed using 2,000 block bootstrap replications clustered at the unit level.

Figure F3: Dynamic Treatment Effect of Coal Layoffs after the Shale Shock on Two-Party Republican Vote Share



Notes: The left plot shows the dynamic treatment effects estimates using the FEct estimator. Treated counties are those with coal employment three years before the 2008 shale shock that experienced layoffs after the shale shock, which are matched to control counties in the rest of the country using socio-economic covariates ($N = 278$). The model includes county and election fixed effects and controls for hydraulic fracturing employment. The bars denote 95% confidence intervals from 2,000 block bootstrap replications clustered at the unit level. The right plot shows the pretreatment average prediction errors and their 90% confidence intervals. The red dashed lines denote the equivalence range set at $[-0.36\hat{\sigma}, 0.36\hat{\sigma}]$ as proposed by Hartman and Hidalgo (2018), whereas the black dashed lines mark the minimum range. The p -value indicates equivalence holds with high confidence.

G Shift-Share Research Design

A shift-share research design is a technique to analyze the sources of change in economic outcomes between different geographic regions. The shift-share specification is increasingly used for many applications (e.g., Bartik 1991; Autor, Dorn, and Hanson 2013). The economic outcome of interest in this study is coal employment. I decompose the change in coal employment into a national shift component that reflects changes in overall economic trends that affect all regions similarly and a county-level component that reflects the extent to which a county is affected by the national trend. To measure a county’s pre-shock exposure, I use the local share of coal employment in 2000. Selecting this year ensures that the local composition of the labor market is not biased by disruptions from coal-to-gas switching in electric fuel stocks.

I estimate the following empirical model using linear regression:

$$y_{it} = S_{i,2000} + S_t + \delta(S_{i,2000} \times S_t) + \mathbf{X}_{it}^\top \beta + Election_t + State_i + \epsilon_{it}.$$

In this equation, y_{it} represents the two-party Republican presidential vote share. $S_{i,2000}$ denotes the share component, the pre-shock county-level share of coal employment. For interpretation, I standardize the share component at the state level to capture the relevant source of variation (Mummolo and Peterson 2018). Standardization at the state level means that for states where there is no coal employment, these observations drop out. For this reason, the sample size in the following table will be smaller than in other analyses. Next, S_t represents the shift component, which captures the overall economic conditions that affect the coal industry similarly in period t . For interpretation, I standardize the national level share of coal employment measure and reverse the direction of this variable, so positive values indicate an increase in unemployment. δ is the interaction term of interest, which indicates the differential effect of changes in national coal employment in counties more exposed to these changes. For controls in \mathbf{X} , I include the same set of pre-treatment county-level socio-demographic covariates as before, in addition to time-varying controls for hydraulic fracturing employment and coal mine unionization. For these time-varying controls, I interact them with a post-shock indicator to allow for hydraulic fracturing and unionization to have heterogeneous effects before and after the shale gas revolution. I employ a heteroskedasticity robust covariance estimator with errors clustered by county.

G.1 Causal Inference Assumptions


There are two approaches to making causal inferences using the shift-share design. One approach is to assume that there are exogenous shares—that is, a county’s employment in an industry is as if randomly assigned. While I use the pre-shock share of coal employment, there is temporal dependency, so it is not guaranteed that there is not a latent factor that explains both earlier coal mining employment, changes in employment trends during the shale shock, and voting behavior. I take two steps to strengthen the plausibility of this assumption. The first strategy is to include covariates in the linear regression predictive of coal mining, such as rurality or poverty, so the treatment is conditionally ignorable. The second strategy is to estimate covariate balancing propensity scores to weight the observations so that covariates are balanced across the distribution of local coal employment shares (Imai and Ratkovic

2014).

The other approach would be to assume the exogeneity of the shift factor—that is, changes in national coal employment are not influenced by any other factor outside of the shale shock that might also affect county-level coal employment. This assumption seems reasonable, albeit harder to falsify.

The final assumption that is required, regardless of whether one assumes that the national shift factor is exogenous or that the pre-shock county shares are exogenous, is common trends. This assumption states that counties with coal employment would have had similar trends in Republican presidential vote share without the national shift in coal employment. To assess the plausibility of this assumption, I estimate the dynamic treatment effect of local coal exposure in elections over time. Figure G1 reports the results, which show that before the shale shock in 2008, counties with equivalent local coal employment shares followed a similar trend.

Figure G1: Dynamic Treatment Effect of Local Coal Exposure on Two-Party Republican Presidential Vote Share



Notes: Estimates from a linear regression model that interacts the pre-shock share of county coal employment with an election year indicator. Point in 1976 represents the baseline election year. The model includes state and election fixed effects, and pre-treatment county-level controls for the white share, Hispanic share, foreign-born share, college-educated share, income per capita (logged), share in poverty, rural share, population (logged), share under 40, share of female workforce participation, and hydraulic fracturing employment. Heteroskedasticity robust standard errors clustered by county.

G.2 Shift-Share Results

Table G1 presents the results from estimating the effect of a one-standard-deviation increase in national coal unemployment in counties with one-standard-deviation greater local coal employment than the rest of their state on two-party Republican presidential vote share. Models (1) through (4) use covariates to satisfy the conditional exogeneity assumption.

The estimates consistently show that the decline of the national coal industry increases Republican presidential vote share in more exposed coal mining counties. Model (5) shows the results without covariates nor balancing weights, whereas model (6) presents the results when using covariate balancing propensity score weights. When using balancing, the results are also strong and consistent. These results increase confidence that the findings in the main text are not dependent on the research design choice.

Table G1: Linear Regression of Two-Party Republican Presidential Vote Share on the Interaction of National Coal Layoffs and the Pre-Shock Share of Coal Employment, 1972–2020

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment and Moderator:						
Pre-Shock Coal Employment	−0.32** (0.15)	−0.35* (0.19)	−0.35** (0.15)	−0.36* (0.19)	−0.50* (0.26)	−1.94** (0.80)
National Coal Layoffs	6.69*** (0.14)	1.83*** (0.15)	6.58*** (0.14)	1.70*** (0.16)	−2.16*** (0.29)	−1.64*** (0.32)
National Coal Layoffs × Post-Shale Shock	0.29** (0.13)	0.50** (0.20)	0.28** (0.13)	0.49** (0.20)	0.47** (0.19)	3.60*** (0.77)
Time-Varying Covariates:						
Hydraulic Fracturing Employment			0.05 (0.10)	0.10 (0.13)	0.28** (0.14)	0.24* (0.14)
Hydraulic Fracturing Employment × Post-Shale Shock			0.43*** (0.08)	0.32*** (0.10)	0.25** (0.10)	0.30** (0.10)
Coal Union Share		−0.62*** (0.13)		−0.62*** (0.13)	−0.67*** (0.14)	−0.75*** (0.13)
Coal Union Share × Post-Shale Shock		1.17*** (0.16)		1.08*** (0.16)	1.10*** (0.19)	0.95*** (0.24)
Pretreatment Covariates:						
White	4.35*** (0.28)	4.74*** (0.30)	4.32*** (0.28)	4.71*** (0.30)		
Hispanic	−0.72** (0.35)	−0.74** (0.37)	−0.73** (0.35)	−0.75** (0.38)		
Foreign-born	1.14** (0.44)	1.13** (0.48)	1.13** (0.45)	1.12** (0.48)		
College	−1.15*** (0.34)	−1.75*** (0.36)	−1.11** (0.34)	−1.71*** (0.36)		
Income per capita (log)	−0.12 (0.45)	0.01 (0.48)	−0.12 (0.45)	0.01 (0.48)		
Poverty	−2.78*** (0.34)	−3.03*** (0.36)	−2.75*** (0.34)	−3.00*** (0.36)		
Rural	0.22 (0.27)	0.46 (0.29)	0.32 (0.27)	0.58** (0.28)		
Population (log)	−0.75** (0.31)	−0.81** (0.33)	−0.64** (0.31)	−0.68** (0.33)		
Under 40 years	0.62** (0.21)	0.69** (0.22)	0.60** (0.21)	0.67** (0.22)		
Female workforce	−0.76** (0.27)	−0.85** (0.28)	−0.72** (0.27)	−0.81** (0.28)		
<i>N</i>	27 330	22 778	27 330	22 778	22 778	22 778
Adjusted <i>R</i> ²	0.460	0.492	0.464	0.496	0.303	0.301
Outcome Mean	58.52	58.52	58.52	58.52	58.52	58.52
Outcome SD	12.846	12.846	12.846	12.846	12.846	12.846
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	No
Election Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Covariate Balancing Propensity Score	No	No	No	No	No	Yes

Notes: HC1 standard errors clustered by county. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

H Heterogeneous Treatment Effects


H.1 Unions

H.1.1 Mine-Level Unionization Rates

I went to great lengths to locate geospatially resolved data on local coal unionization strength. I retrieved relevant data in the EIA-7A form submitted to the U.S. Mine Safety and Health Administration. Coal mine companies are legally obligated to report information about production, company and mine information, operation type, union status, labor hours, and the number of employees. These data date back to 1983. The advantage of these data is that they allow me to create a county-level measure of the share of coal produced by unionized workers.

However, there is one data quality issue that I detected during pre-processing. Figure H1 shows how in 1986 and 1987, there is a suspicious mirror inversion in the number of unionized and non-unionized coal mines. I emailed the EIA about this anomaly, and they replied, “Unfortunately, I do not have any information on this. Anyone who was here in 1986 and 1987 is long gone, and there is no documentation that might explain this.”¹⁰ I further suspect that this is a data anomaly since a comparable decline in unionized coal production is not observed in the CPS data during the same period (Figure H4).

Figure H1: Coal Mine Unionization Data Anomaly in 1986 and 1987




Notes: Data on number of coal mines from the EIA-7A form. The red rectangle indicates the years where there is a suspected data anomaly since the number of unionized coal mines drops, and there is a mirror increase in the number of non-unionized coal mines.

To see if this was a likely data entry problem, I created lagged and leading indicators for whether a coal mine was unionized. If two years ago, a coal mine was unionized, and

¹⁰EIA replied on March 17, 2023. My email asked: “I’ve been using your marvelous detailed data from the EIA-7A and the US MSHA back to 1983. It has the union status of coal mines. I noticed a discrepancy between 1986 and 1987. There seems to be a sudden drop in unions that then reverts in 1988, which makes me think there’s a data issue.”

two years later it was also unionized, but in 1986 and 1987 it was not unionized, it is more likely that for those two years, there was a mistake in how the data were entered because unionization status does not oscillate like that. Indeed, Figure H2 shows that after this correction, the drop in unionization observed in 1986 and 1987 attenuates but does not disappear completely. While not perfect, these data are valuable as the best available source to evaluate the effect that changes in coal unionization over time may have had on vote choice.

Figure H2: Coal Mine Unionization Data after Correction for 1986 and 1987




Notes: Data on number of coal mines from the EIA-7A form. For the years 1986 and 1987, I code the mine as unionized if the two years before and two years after the same mine is otherwise coded as unionized.

H.1.2 Results Controlling for Unionization

I estimate the counterfactual outcoming using FEct with the state-specific shale shocks as the treatment. The model incorporates a control for the time-varying share of unionized local coal production (e.g., takes the value 1 if all local coal produced is at union mines). Figure H3 and Table H1 present the results. The strength of the ATT grows in magnitude without losing precision, which increases confidence in the main results. As expected, greater coal production at unionized mines has a negative association with Republican vote share, but the large p value indicates that change in unionized coal production is not predictive of county-level Republican presidential vote choice.

Figure H3: Dynamic Treatment Effect of State-Specific Shale Shocks on Two-Party Republican Vote Share Controlling for Unionization



Notes: The left plot shows the dynamic treatment effects estimates using the FECT estimator. Treated counties are those with coal employment three years before the 2008 shale shock, which are matched to adjacent control counties using socio-economic covariates ($N = 213$). The model includes county and election fixed effects and controls for hydraulic fracturing employment and the county share of unionized coal production. The bars denote 95% confidence intervals from 2,000 block bootstrap replications clustered at the unit level. The right plot shows the pretreatment average prediction errors and their 90% confidence intervals. The red dashed lines denote the equivalence range set at $[-0.36\hat{\sigma}, 0.36\hat{\sigma}]$ as proposed by Hartman and Hidalgo (2018), whereas the black dashed lines mark the minimum range. The p -value indicates equivalence holds with high confidence.

Table H1: FECT Estimate of the Shale Shock Effect on Two-Party Republican Presidential Vote Share with State-Specific Shale Shocks and Coal Unionization Controls


	Estimate	S.E.	CI _{2.5%}	CI _{97.5%}	p -value
ATT:					
Observations equally weighted	5.86	1.26	3.39	8.33	0.00
Units equally weighted	5.32	1.25	2.86	7.78	0.00
Covariates:					
Hydraulic Fracturing Employment	0.08	0.22	-0.34	0.50	0.71
Unionized Coal Production Share	-0.63	1.34	-3.26	2.00	0.64
Placebo Tests:					
-3 to 0 election interval	0.86	0.63	-0.37	2.09	0.17

Notes: Presidential elections from 1984–2020 in the matched sample of 213 counties. Standard errors and confidence intervals constructed using 2,000 block bootstrap replications clustered at the unit level.

H.1.3 National Trends in Coal Unionization

Data on unionization rates comes from the *Union Membership and Coverage Database*, available at www.unionstats.com (Hirsch and MacPherson 2003). The data are compiled from the Current Population Survey (CPS). I use detailed industry estimates of the percentage of employees in the coal industry who are union members. These data begin in 1983. Figure H4 plots the trend in unionization in the coal industry from 1983 to 2022.

Figure H4: Coal Mining Unionization, 1983-2022



Notes: Union Membership and Coverage Database from the CPS. Data compiled by unionstats.com.

H.1.4 Treatment Heterogeneity by Union Strength

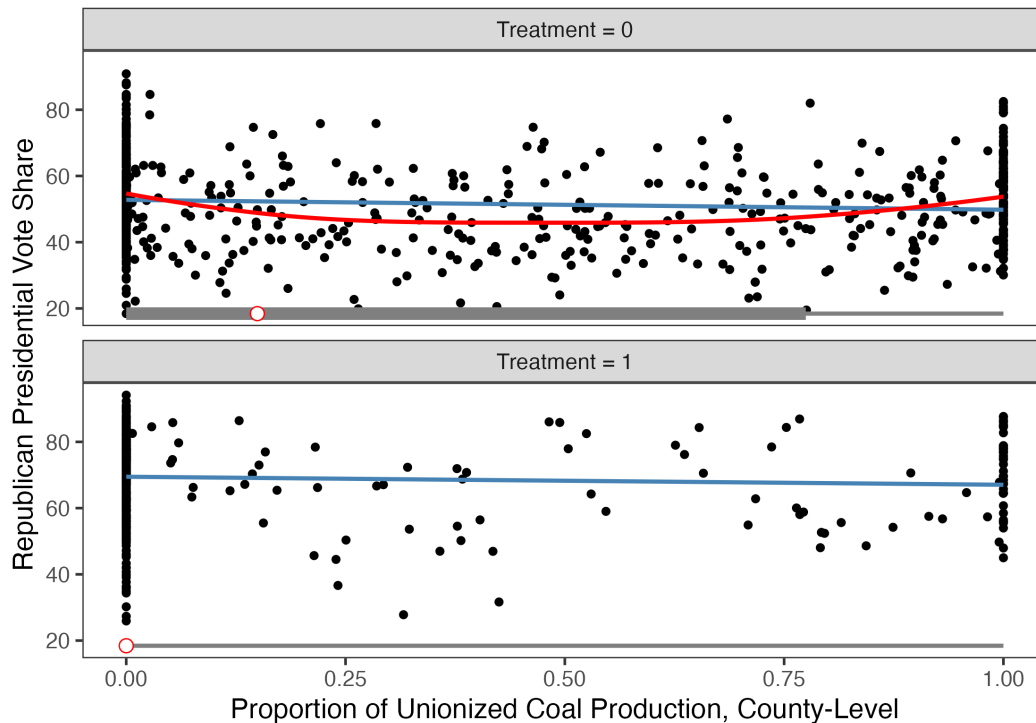
I investigate heterogeneous responses to the shale shock depending on the level of unionized coal production. Since unionization rates are not randomly assigned, it is not possible to draw strong causal claims from this analysis. However, examining variation by unionization levels is nonetheless descriptively informative as to whether the electoral backlash took place in places where coal unions were weaker. This would lend support to an interpretation that the decline of coal unions helped to create the conditions for the partisan reversal triggered by the shale shock.

For this analysis, I examine only the treated units, since there is minimal meaningful variation of the moderator (unionized coal production) in the matched counties that by construction have no or limited coal employment.

To assess whether it is appropriate to estimate a multiplicative linear interaction, I examine the raw distribution of the data in Figure H5 (e.g., Hainmueller, Mummolo, and Xu 2019). This figure reveals that prior to the shale shock, there was a greater share of unionized coal produced, whereas the median proportion of county-level unionized coal production dropped to 0 after the shale shock. This decline in unionization actually is so drastic that

it is not possible to estimate the LOESS fit in the post-shock treatment period. This is a warning sign that there is not sufficient moderator support to use a continuous measure of coal unionization.

Figure H5: Linear Interaction Diagnostic Plots of Shale Shock Effect Moderated by Union Strength



Notes: The plot shows the relationships among the treatment (shale shock), the outcome (two-party Republican presidential vote share), and the union moderator using the raw data. The blue and red lines represent the linear and LOESS fits, respectively.

Indeed, if one were to only look at the traditional linear estimator, one would see that in places with greater coal unionization, there is less of a backlash after the shale shock.¹¹ However, when estimated with the binning and kernel estimators as recommended by Hainmueller, Mummolo, and Xu (2019), problems with the lack of common support emerge. I attempted to convert the moderator into a dichotomous variable to maximize common support, yet the estimation revealed that the model was either rank deficient or indefinite, so the estimates were not reliable enough to be reported. Thus, despite the results at first cut appearing consistent with my theory, I am not confident enough in the validity of these estimates to make strong claims about treatment effect heterogeneity.

¹¹These results are available upon request.

H.2 Transitional Assistance


H.2.1 Measurement

The Partnerships for Opportunity and Workforce and Economic Revitalization (POWER) Initiative is the main initiative by the federal government to assist coal-dependent communities (Lawhorn 2022). The Obama administration launched it in 2015, and certain aspects of it continued under the Trump administration. The initiative provides funding for economic stabilization, social welfare, and environmental remediation. The Appalachian Regional Commission (ARC) administers the POWER Initiative, which exclusively targets the Appalachian region. Western coal communities do not receive these funds.¹²

I collected data on POWER grant distribution through a special request to the ARC, which generously provided data through fiscal year 2022. These data include the counties served by the POWER grants. There are a handful of grants where the recipient counties are not listed, which are trimmed from the data. There are also some grants that aid all ARC counties, which I code as occurring in each of the counties.

Figure H6 plots the count of POWER grants over time. The political response takes place long after the shale shock. There were few projects in 2015 when the program started and more in 2016. The bulk of grant-making occurs during the Trump administration and continues into the Biden administration.

Figure H6: Count of POWER Initiative Grants



Notes: Vertical red line denotes the start of the shale shock in 2008. The POWER Initiative began issuing grants in 2015.

¹²There is also the Assistance for Coal Communities program administered by the Economic Development Administration.

H.2.2 Treatment Heterogeneity by Transitional Assistance

Using data on POWER grants, I examine whether they have a moderating effect on the shale shock in coal counties. One might expect transitional assistance to blunt the effects of the shale shock, so the interaction of receiving POWER funding and the treatment would be negative. However, given that the federal response to the decline of coal has been insufficient to stem the economic losses in communities, it is more likely that there will be no relationship (Roemer and Haggerty 2021). There could even be a positive interaction effect if the Trump administration receives credit for the subsequent aid.

Table H2 presents the results from a linear regression that estimates the average treatment effect of the shale shock in coal counties on Republican presidential vote share, conditional on whether the county received power funding in the election year. There is no causal identification strategy. The aim is an exploratory, descriptive analysis.

The first model examines the effect of a standard deviation increase in POWER funding in an election year; there is no differential effect of the shale shock across different levels of transitional assistance. The second model operationalizes the moderator using the cumulative level of POWER funds a county has received, and here too, there is no differential effect of the shale shock.

The third model uses the count of grants as the moderator, which shows a positive interaction term. This implies that in counties receiving more grants, the shale shock has an even larger, positive effect on Republican vote share. To drill into this finding, the last two models split the sample into two parts: counties that received POWER grants and those that did not. In this split sample, the coefficient for the treatment remains positive, but it is only significant at the 1% level among counties that received POWER funding. This is consistent with the interaction term in model 3.

What explains this positive interaction term? For one, it could be noise since the effect does not appear using the other operationalizations of funding. Another explanation is that the counties most hard hit by coal's decline are the ones that received funds—that were insufficient—so grant-making is simply a proxy for economic vulnerability. Simply receiving funds may also not have been enough, as Bolet, Green, and Gonzalez-Eguino (2023) show how the *process* of just transition matters whereby more success compensation policies involved active negotiation with unions. A third explanation is Trump received credit since many of the grants were made during his time in office.

Since these grants are not assigned at random, I refrain from drawing conclusions about the causal effect of transitional assistance on moderating electoral backlash. Descriptively, this appears to have not been the case, but future research with tailored research designs is needed for credible causal inferences.

Table H2: Moderating effect of federal transitional assistance for coal communities on the shale shock's effect on Republican vote share, 1972–2020

	(1)	(2)	(3)	Funding	
				Yes	No
Treatment	5.11*** (1.58)	5.24*** (1.59)	4.01** (1.60)	5.46*** (1.64)	2.54 (1.97)
POWER Funding	−0.13 (0.21)				
Treatment × POWER Funding	0.44 (0.28)				
Cumulative POWER Funding		0.21 (0.25)			
Treatment × Cumulative POWER Funding		0.32 (0.24)			
POWER Grants			0.30 (0.35)		
Treatment × POWER Grants			0.71** (0.32)		
<i>N</i>	1378	1378	1378	1378	1388
Adjusted R^2	0.865	0.865	0.867	0.865	0.792
County Fixed Effects	Yes	Yes	Yes	Yes	Yes
Election Year Fixed Effects	Yes	Yes	Yes	Yes	Yes

Notes: Heteroskedasticity-robust standard errors clustered by county. Models 1-3 only include counties where there is variation in power funding because the colinearity between the county fixed effects and no variation in funding would otherwise result in a rank-deficient design matrix. Models 5 and 6 split the sample to estimate the effect of the shale shock among counties that did (not) receive POWER grants at any point in time. * $p < 0.1$; ** $p < 0.05$; and *** $p < 0.01$.

I Robustness to Compositional Change

This appendix assesses the extent to which compositional changes in counties explain the results as opposed to conversion and mobilization. The concern about compositional changes would be that the increase in support for the Republican Party as reflected by vote share is actually a consequence of Democratic voters leaving the population through migration or death. While there is undoubtable compositional change over the time period of the study, I argue that much of the change observed before and after the shale shock can be attributed to both formerly Democratic voters changing their minds and non-Democratic voters becoming mobilized by the threat of coal regulations.

I conducted four steps to evaluate the robustness of the results to compositional changes. First, I review a relevant study that provides indirect support for the claim that conversion is partly responsible for the partisan reversal in a significant state in coal country. Second, I discuss how the estimation strategy accounts for common factors causing population changes across the matched counties. Third, I used Census data on population changes to calculate counterfactual vote shares that allow for the estimation of bounds on the possibility of bias from compositional change, finding it would be unlikely to change the conclusion. Fourth, I estimated the ATT of the shale shock on Republican vote share using different baseline periods that progressively minimize possible cohort effects, demonstrating the robustness of the results to compositional changes.

I.1 Study Demonstrating Conversion in Coal States During Study Period

First, a pertinent study from Hill, Hopkins, and Huber (2021) provides indirect evidence of conversion. This study is unique in its ability to assess whether mobilization, conversion, or composition explain vote choice because they have data on 37 million registered voters in six states. One of these states is Pennsylvania, which is relevant because it is a major coal producer. They concluded that “conversion more consistently explains pro-republican Party (GOP) electoral change between” the 2012 and 2016 elections (2). For Pennsylvania in particular, the study estimates that conversion effects were “especially relevant” (2). Composition explained only 8% of electoral change in Pennsylvania. While we cannot make inferences about other coal regions, the fact that conversion and not composition is responsible for the bulk of electoral change in a key state from our analysis provides indirect support for our claim that the partisan reversal happening because of the shale shock is partly the consequence of conversion.

I.2 Reducing Bias through Estimation Strategy

Second, the matched difference-in-differences research design is able to address common compositional change across the treated and control counties. To the extent that there is a common population process over time, the year fixed effects remove this source of bias. There still could be some compositional changes that are not removed if there are time-varying county-specific trends. However, an exploration of population trends in the treatment and control counties indicates that they are quite similar in terms of their population and vote total trajectories; this is by design because of the matching process. For example, 11% of the control counties experience a simultaneous decline in total votes and county population,

which is similar to the 16% of the treatment counties. Likewise, when examining the decline in absolute Democratic votes, 14% of control counties experience a decline in absolute votes, and 17% of treatment counties also experience a decline.

So, there is evidence of a common process of population decline and vote total declines in some of the control and treatment counties, which should be partly handled by the difference-in-differences model. Yet, since there are likely factors specific to the more coal-dependent counties that could cause compositional change, in the next subsection, I calculated bounds on possible compositional bias to assess whether this would alter this study’s conclusions.

I.3 Bounding Compositional Bias Using Counterfactuals

I.3.1 Simple Model of Voting and Population Change

Third, I used population data from decennial censuses to calculate counterfactual vote shares under different assumptions about how much compositional changes are the consequence of Democrats exiting a county. This exercise also clarifies assumptions behind the conditions that would have to hold for compositional change to threaten inference. Some of these assumptions, such as exiting Democrats would not have been converted if they remained, are not self-evident, which instills further confidence in the paper’s interpretation of the electoral change in coal country as the consequence of primarily conversion and mobilization.

Let us consider a basic model of voting and population dynamics. The two-party share (ρ) of Republican voters (R) in an election in period 1 is,

$$\rho_1 = \frac{R_1}{T_1}, \quad (3)$$

where the total number of voters (T) is the sum of Democrats (D) and Republicans,

$$T_1 = D_1 + R_1. \quad (4)$$

The total number of voters in an election is a function of turnout and the population size.

Assumption 1: Constant Turnout We hold turnout constant for tractability. This assumption is obviously false because turnout varies between elections, but it is useful because it will allow for the estimation of bounds under extreme scenarios. Empirically, the difference-in-differences design also accounts for how some elections have higher turnout than others due to candidate characteristics, policy mood, and national economic conditions, so this assumption becomes more tenable given the estimation strategy.

The constant turnout assumption allows us to infer that when the population (P) declines due to migration or death, the composition of the total number of voters can also shift to the extent that the people exiting the population are Democrats or Republicans. We bracket non-partisan and non-voters in the population for simplicity, which introduces bias against detecting an effect of the shale shock because we will be overstating how much of the population decline is composed of Democratic voters—people exiting are likely a mix of Democrats, Republicans, Independents, and non-voters. Thus, the population is defined as

the sum of all Democratic and Republican voters.

$$P = D + R \quad (5)$$

These equalities allow us to compute counterfactual two-party Republican vote shares under different assumptions about the share of population decline due to Democratic voters.

Assumption 2: Constant Partisanship of Exiters The analysis assumes that the Democratic (Republican) voters that leave would have continued to vote for the Democratic (Republican) party had they stayed. This is not automatically apparent since, had Democrats remained, they might have also changed their partisan allegiances due to concerns about the effects of the shale shock on the community. However, to the extent this assumption is not tenable, it demonstrates the validity of the argument about how the shale shock’s economic effects on the community prompted conversion and mobilization.

To formalize how I will assess different levels of potential bias from compositional change, let γ be the share of population decline composed of Democratic voters. If $\gamma = 1$, then all of the population decline is composed of Democrats. If $\gamma = 0$, then none of the population change are Democrats; instead, it is all Republicans. Then, I can define counterfactual quantities under different assumptions about the composition of the changing population. Consider the total number of voters in election 2, T_2 . If the population declined between election 1 and election 2, I can add back the population change to obtain a counterfactual T'_2 for if election 2 had the same number of voters as in the first.

$$T'_2 = T_2 + (T_1 - T_2) \quad (6)$$

To retrieve the counterfactual Republican two-party vote share, ρ' , I can compute the following.¹³

$$\rho'_2 = \frac{R_2 + \gamma(T_1 - T_2)}{T'_2} \quad (7)$$

I.3.2 Model Operationalization

I used the same election data from before in combination with decennial Census data on population to estimate counterfactuals. Since the Census data do not perfectly overlap with the presidential election calendar, I pair the elections with the closest decennial Census year.¹⁴ This means that in the analysis that follows, the post-treatment period is defined as starting in 2010 since I aggregate all data to the Census year level. While this does not perfectly overlap with the 2008 start of the shale gas boom, this is preferable to making the cutoff be 2000 and should introduce bias against the hypothesis since it would include part of the Republican shift in the pre-treatment period.

¹³A careful reader will also note that this abstracts away from the dynamics of population growth, which would add unnecessary complexity with uncertain payoffs.


¹⁴Election years 1972 and 1976 are paired with the 1970 Census, 1980 and 1984 with the 1980 Census, 1988 and 1992 with the 1990 Census, 1996 and 2000 with the 2000 Census, 2004, 2008, and 2012 with the 2010 Census, and 2014, 2016, and 2020 with the 2020 Census.

These population data appear consistent with the more limited data on migration flows (Figure I2). I used population data from the Census because of its broader temporal coverage. The 5-Year ACS began reporting estimates of county-to-county migration flows in the 2005–2009 period. There is also considerable overlap in the 5-year estimates since they share sample data, so it is hard to make comparisons using consecutive years of multiyear estimates.¹⁵ Thus, I relied on population data.

Figure I1 plots the change of the treated coal counties’ population over time. Population increases in 1980 on average, but then falls between 1980 and 1990. Between 1990 and 2010, there is little discernable population change on average, as shown by the black bar in the box plot, which represents the average population change. It is only in the decade from 2010 to 2020 that there is a population decline on average.

This descriptive result, on face, is encouraging because it suggests that for the estimates of the causal effect of the shale shock in at least the 2008 and 2012 elections, compositional change is less likely to be at play. There is not much sign of population decline on average around this time period. The analysis below will systematically evaluate whether the population decline at the end of 2020 completely biases the interpretation of the results.

Figure I1: Population Change in Treated Counties



Notes: Data from the US Census Bureau show no population decline on average from 2000 to 2010 and a slight population decline on average from 2010 to 2020.

I.3.3 Counterfactual Estimates

The results in Table I1 show the ATT estimates of the shale shock on Republican presidential vote share under different assumptions about γ —that is, the percentage of population decline that is composed of Democratic voters. For robustness, I consider two scenarios. In the top panel, I estimated the counterfactual vote shares for counties where there is a simultaneous decline in total votes and population. This is because it does not make sense to calculate

¹⁵ “Understanding and Using American Community Survey Data,” US Census Bureau, <https://bit.ly/3PvfXOL>

the counterfactual for places where there is an increase in population. Across all possible assumptions about γ , the ATT remains positive, though it decreases modestly in magnitude. Based on this analysis, even 100% compositional bias, an extreme scenario, would only change the ATT estimate by 0.8.

The bottom panel repeats the same analysis but computes the counterfactuals for counties where there was a simultaneous decline in Democratic votes and population. This is a more conservative approach than the top panel that used declines in total votes as a criterion because Democrats are only a portion of the electorate, so their decline cannot be offset by increases in turnout by other voters. In this more conservative analysis, the main results hold across reasonable levels of bias. Even if 75% of the population decline is composed of Democrats—which is unlikely since part of population change is among non-voters and people migrate for economic reasons—the ATT would only change by 1.9. The shale shock would still cause a 2.5% increase in Republican vote share due to conversion and mobilization.

Table I1: Analysis of ATT robustness to compositional changes driven by population decline

		Democratic Share of Population Loss (γ)				
	Base	0%	25%	50%	75%	100%
<i>Decline in Total Votes and Population</i>						
ATT	4.41*** (1.51)	3.36** (1.50)	3.43** (1.49)	3.49** (1.49)	3.56** (1.51)	3.63** (1.55)
N	1277	1277	1277	1277	1277	1277
<i>Decline in Democratic Votes and Population</i>						
ATT	4.41*** (1.51)	4.36*** (1.51)	3.73** (1.48)	3.09** (1.48)	2.46* (1.49)	1.83 (1.53)
N	1277	1277	1277	1277	1277	1277

Notes: The outcome is the ATT of the shale shock on Republican vote share. Each model shows the ATT under different assumptions about the share of population decline composed of Democratic voters. Estimates from a two-period difference-in-differences model that pools elections before 2010 as the pre-treatment period and elections after that year as the post-treatment period. The 2010 cutoff is necessary to match the population data from the decennial census and should introduce bias against the hypothesis since it would overstate how Republican the pre-shock period was. * $p < 0.1$; ** $p < 0.05$; and *** $p < 0.01$.

I.4 Reducing Compositional Bias with More Proximate Pre-Treatment Windows

Compositional bias could also arise from cohort effects, where the population changes as people die. There could be a concern that by analyzing a baseline period from 1972 to 2004, some of the people in earlier periods may not be the same individuals who are switching their votes or are being mobilized during the shale gas shock.

However, two factors counterbalance this concern about cohort effects. First, the estimator equally weights each pre-treatment election so no one cohort dominates.

Second, I estimated classical two-period difference-in-difference models with pre-treatment windows that are more proximate to the shale shock. This approach weakens statistical power, but by being closer to the timing of the shale shock, the analysis helps to ensure that the same cohorts are present, increasing the plausibility of conversion and mobilization being behind the electoral change.

Table I2 presents the ATT estimates of the shale shock's effect on Republican vote share in coal counties. Across different baseline periods, the results are robust. As in the main results, the effect of the shale shock is weakest in 2008, which is before its main employment effects are felt. Looking at the difference in the pooled ATT estimates from the full sample and the most restrictive baseline of 2000-2004, there is a difference of 1.17 percentage points. However, the shale shock still increased Republican vote share by 3.52% in affected coal counties, a meaningful effect.


Table I2: Two-period difference-in-difference estimates with different baseline periods

	Post-Treatment Period				
	2008	2012	2016	2020	Pooled
<i>Baseline: 1972-2004</i>					
ATT	2.13 (1.83)	5.24*** (1.91)	6.03*** (1.86)	5.35*** (1.83)	4.69*** (1.15)
N	2127	2127	2127	2127	2766
<i>Baseline: 1988-2004</i>					
ATT	2.62 (1.90)	5.73*** (1.98)	6.52*** (1.93)	5.84*** (1.90)	5.18*** (1.26)
N	1278	1278	1278	1278	1917
<i>Baseline: 1992-2004</i>					
ATT	2.23 (1.96)	5.34*** (2.03)	6.13*** (1.99)	5.45*** (1.96)	4.79*** (1.34)
N	1065	1065	1065	1065	1704
<i>Baseline: 1996-2004</i>					
ATT	1.71 (2.02)	4.82** (2.09)	5.61*** (2.04)	4.93** (2.02)	4.27*** (1.43)
N	852	852	852	852	1491
<i>Baseline: 2000-2004</i>					
ATT	0.97 (2.08)	4.07* (2.15)	4.87** (2.10)	4.18** (2.08)	3.52** (1.51)
N	639	639	639	639	1278

Notes: HC2 standard errors. * $p < 0.1$; ** $p < 0.05$; and *** $p < 0.01$.

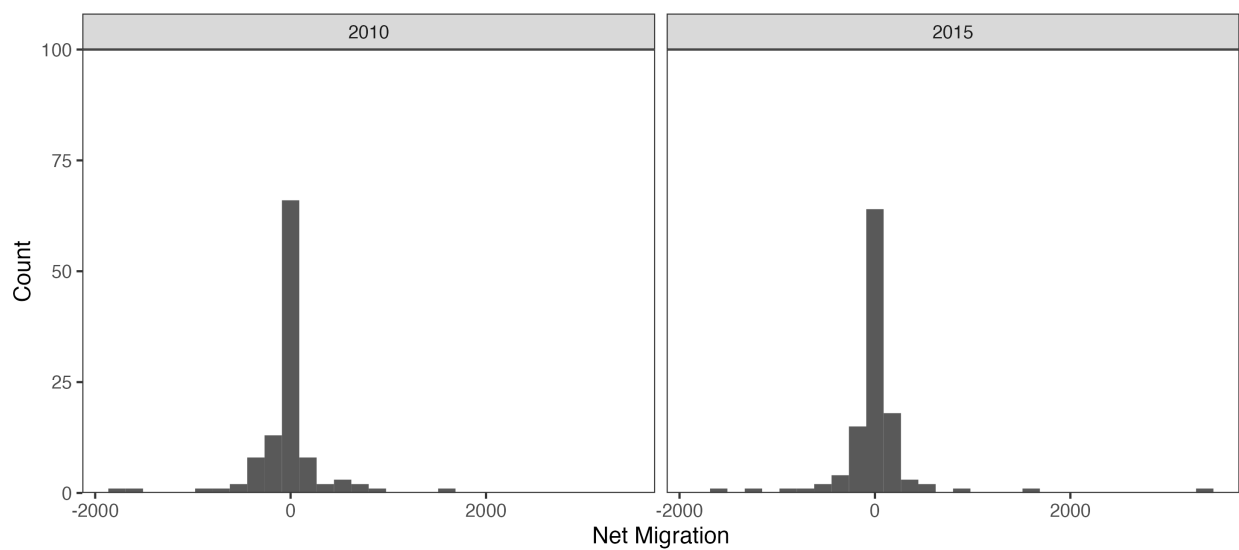
I.5 Migration Data

Figure I2: Net Migration from Treated Coal Counties



Notes: Data from the 5-Year ACS. Median net migration is a loss of 86.5 individuals in 2010 and a loss of 67 individuals in 2015, whereas the average net migration is a gain of 174 people in 2010 and a gain of 175 in 2015.

Figure I3: Net Migration from Treated to Control Counties



Notes: Data from the 5-Year ACS. Median net migration is a gain of 4 individuals in 2010 and a gain of 13 individuals in 2015, whereas the average net migration is a loss of 21 people in 2010 and a gain of 18.2 in 2015.

J Falsification Tests

I conducted several falsification tests, the results of which Figure J1 summarize. The left-most facet of Figure J1 presents the results from different operationalizations of shale shock exposure *within* coal country. The first result is a benchmark showing the ATT from the model in Figure 3. The second falsification test within coal country in Figure J1 compares coal counties that are vulnerable and unexposed to the shale shock. Counties are vulnerable if they mine thermal coal displaced by coal-to-gas switching, whereas counties are unexposed if they mine metallurgical anthracite coal less exposed to electricity markets. The results show that economically vulnerable thermal coal mining counties are likelier to exhibit a partisan reversal (Appendix J.4). The third test within coal country defines the treatment using the timing of observed layoffs in coal counties rather than pre-shock coal employment and the results obtain (Appendix F). The final test within coal counties uses a continuous measure of pre-shock coal employment. The logic is that counties with a larger share of coal employment should have greater economic dependence on the industry and, thus, exhibit a larger electoral response, which is what Figure J1 shows.

Next, the middle facet of Figure J1 reports the results from a series of falsification tests that evaluate the effect of the shale shock in non-coal mining communities. The logic is that if a latent characteristic of the social fabric of mining communities becomes activated post-2008, the shale shock should also cause partisan reversal in these areas. But, to the contrary, the shale shock does not affect these non-coal mining counties.

Lastly, the rightmost facet is a falsification test to see if the results are unique to economic distress in coal country. This test addresses the possibility that it is simply areas with any layoffs in the post-2008 period that cause a shift to the Republican Party for unobserved reasons. The non-coal counties with more than a standard deviation increase in layoffs during an election year represent the treatment group for this falsification test. Figure J1 shows no effect of non-coal layoffs on Republican presidential vote share after the shale shock.

J.1 Non-Coal Mining Communities


I use data from the United States Department of Labor’s Mine Safety and Health Administration’s (MSHA) *Mine Data Retrieval System* (MDRS).¹⁶ The MDRS is a repository of all current and historical databases, which provides mine-level data on all coal and metal/nonmetal mines in the United States. By law, all mine operators must provide data on employment to MSHA, so these data are comprehensive.

For the falsification tests, I construct indicators for whether a county has an iron, silver, or sandstone mine. These are all extractive industries that, while distinctive from coal, should give rise to similar mining communities. These types of mines are also numerous enough to provide for a well-powered statistical test (e.g., there are not enough lithium mines in the United States for that to be a fair comparison to coal).

Table J1 presents balance statistics for mining communities’ main political and socio-demographic covariates. Coal mining communities are similar to iron, silver, and sandstone mining areas, which suggests that these mining communities represent a valid comparison.

¹⁶<https://www.msha.gov/mine-data-retrieval-system>

Figure J1: Falsification Tests of the Shale Shock's Effect on Republican Presidential Vote Share



Notes: The left plot shows the shale shock effect in coal country. Coal vs. matched compares coal counties with matched neighbors (Table D3). Vulnerable vs. unexposed compares vulnerable (thermal coal) with less exposed (anthracite coal) counties (Appendix J.4). Layoffs vs. matched compares counties experiencing coal layoffs with matched counties (Appendix F). +1SD shock exposure is a one-standard-deviation increase in pre-shock coal employment in the matched sample (Table J2). Models control for hydraulic fracturing employment. The central plot presents falsification tests of the shale shock effect in iron, sandstone, and silver mining counties. The right plot shows the effect of a one-standard-deviation increase in non-coal layoffs. falsification tests are estimated with linear regression with county and election fixed effects (Table J2). The reported coefficient is the interaction of the post-shock and falsification group indicators. The bars denote 95% confidence intervals with standard errors clustered by county.

Figure J2 plots the Republican vote share in the non-coal mining and non-coal layoff counties, coal mining counties, and the remaining counties over elections. These plots show how other mining communities generally matched the national trajectory in Republican vote share. In contrast, coal mining counties moved further to the political right following the shale shock in 2008.

J.2 Non-Coal Layoffs

I construct a measure of non-coal layoffs using data from CBP. For each year, I take the average change in overall employment from the previous year. If a county has more than a standard deviation decrease in county-level employment year-over-year, I code that county as experiencing an increase in layoffs. Then, to define the treatment group, treated units are counties without coal that experienced more than a standard deviation decrease in employment year over year.

Table J1: Covariate Balance of Falsification Test Counties

	Coal		Iron		Sandstone		Silver		Layoffs	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Pretreatment:										
White	0.90	0.14	0.86	0.12	0.89	0.12	0.88	0.09	0.81	0.18
Hispanic	0.03	0.07	0.07	0.10	0.06	0.13	0.16	0.21	0.08	0.15
Foreign-born	0.01	0.02	0.04	0.06	0.03	0.04	0.08	0.09	0.04	0.06
College	0.08	0.04	0.11	0.03	0.12	0.06	0.11	0.06	0.10	0.05
Income per capita (log)	9.62	0.17	9.73	0.15	9.76	0.22	9.69	0.21	9.71	0.24
Poverty	0.18	0.07	0.15	0.05	0.13	0.05	0.15	0.06	0.15	0.07
Rural	0.65	0.21	0.56	0.28	0.57	0.26	0.51	0.27	0.59	0.32
Population (log)	10.19	0.79	10.69	1.76	10.64	1.21	9.88	1.56	9.90	1.64
Under 40 years	0.54	0.05	0.53	0.07	0.54	0.05	0.53	0.08	0.54	0.06
Debt-to-income	1.38	0.64	1.68	0.91	1.75	0.83	2.06	0.79	1.63	0.93
Female workforce	0.21	0.03	0.23	0.03	0.24	0.03	0.22	0.03	0.23	0.03
Time-Varying:										
Gas plant distance (1992–2020)	0.21	0.99	0.26	1.00	0.02	1.00	0.07	1.13	0.02	1.00
Hydraulic fracturing employment	0.85	2.39	0.27	0.89	0.49	2.49	0.10	0.51	0.99	3.90
Coal employment	11.16	13.73	1.58	5.56	0.25	1.35	0.02	0.30	0.19	1.77
Power employment	0.76	2.96	0.42	2.12	0.12	0.55	0.06	0.15	0.09	0.55
Two-party Republican vote share	57.81	15.18	56.47	14.84	58.13	14.25	60.53	13.97	58.19	16.08
Observations:										
Counties	116	NA	26	NA	100	NA	19	NA	621	NA
<i>N</i>	1505	NA	338	NA	1297	NA	247	NA	8064	NA

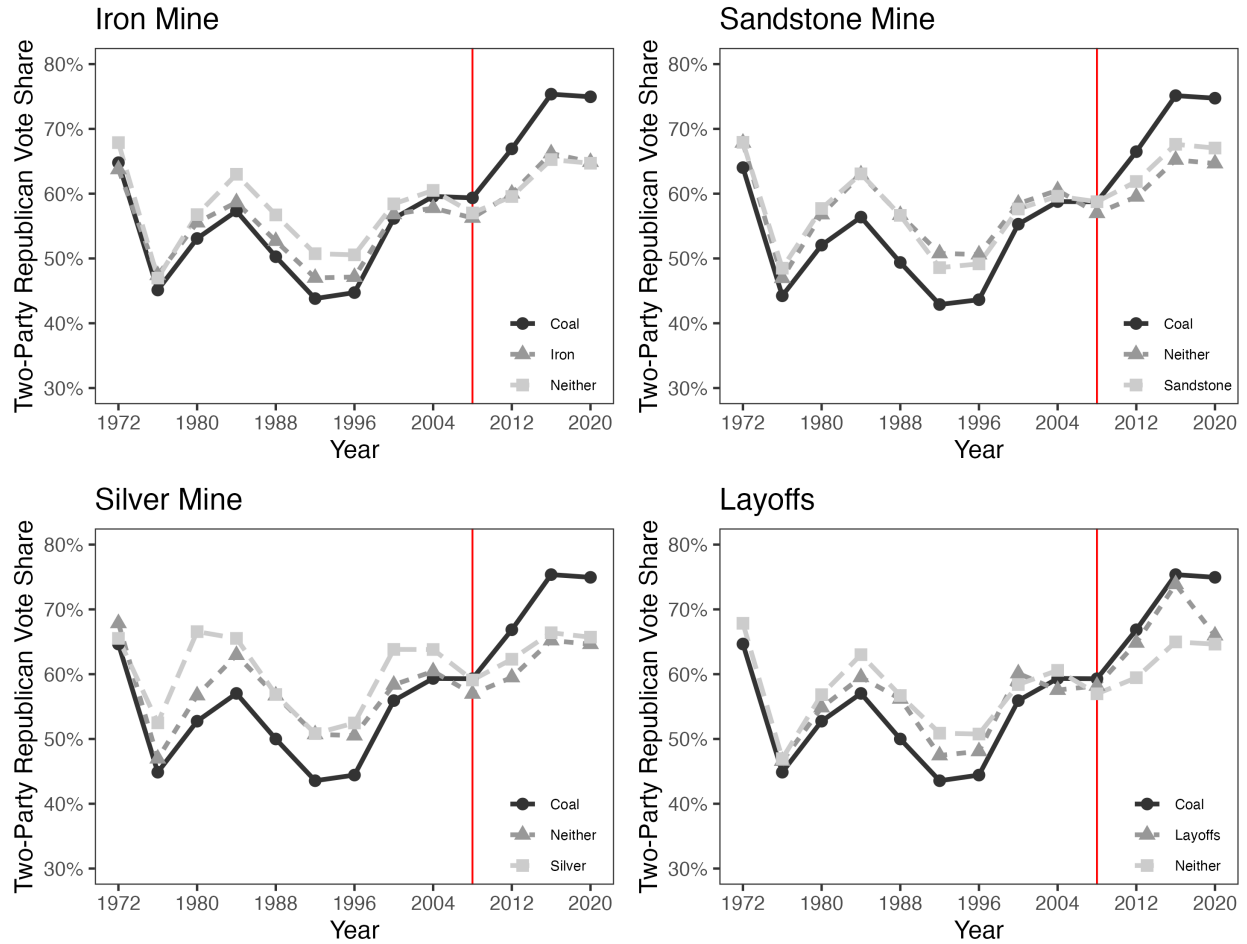
Notes: Coal counties are those with greater than 1% of local employment in coal during 2005–2007. Iron, silver, and sandstone mining counties are those with a mine, as recorded by MSHA. Layoff counties are those without coal and greater than a standard deviation decrease in employment in a year since 2008. Pre-treatment socio-demographic measures come from the 2000 U.S. Census. Data on the debt-to-income ratio are from 2006 and are based on Mian and Sufi (2011). County-level employment data come from the CBP reports that have had missing values imputed according to Eckert et al. (2020). Election data come from Leip (2020).

J.3 Falsification Test Results

Figure J1 in the main text plots the results of the falsification tests in the middle and rightmost facets. In addition, this appendix provides the table for the linear regressions underlying those results for the non-coal mining counties and the non-coal mining layoffs. These empirical models are estimated using a linear regression of two-party Republican presidential vote share on the interaction of an indicator for if the county has non-coal mining or non-coal layoffs and the shale shock indicator. The model also includes county and election year fixed effects, which statistically removes confounding from time-invariant factors.

Table J2 presents the regression results for the falsification tests using non-coal mining counties and non-coal layoffs. As is apparent in Figure J1, the coefficients for the interaction term lack statistical significance at the 5% significance level. While the model includes a shale shock indicator and an indicator for being in the placebo group, these drop out during estimation due to collinearity with the fixed effects. What is most relevant is the interaction term. Unsurprisingly, the adjusted R^2 is low due to the number of fixed effects contributing to the degrees of freedom adjustment and, more importantly, the slight variation explained by the shale shock in these non-coal placebos. In all, these falsification test results indicate that shale shock has a unique effect on coal country rather than representing a latent characteristic of mining communities or places with non-coal layoffs that is activated post-2008.

Figure J2: Falsification Test Pretrends



Notes: Average county-level two-party Republican presidential vote share in counties with iron mining, silver mining, sandstone mining, and substantial layoffs. Data on mining come from the EIA. Data on presidential elections come from Leip (2020). Data on employment changes come from the CBP with missing values imputed by Eckert et al. (2020).

Table J2: Regression Results for Falsification Tests of Shale Shock's Effect on Two-Party Republican Presidential Vote Share in Non-Coal Mining Counties and Non-Coal Layoff Counties

	Iron	Silver	Sandstone	Layoffs
Post-Shale Shock	5.58*** (0.21)	5.59*** (0.21)	5.53*** (0.21)	6.89*** (0.21)
Iron Mine x Post-Shale Shock	-0.62 (2.43)			
Silver Mine x Post-Shale Shock		-1.97 (2.09)		
Sandstone Mine x Post-Shale Shock			1.16 (1.38)	
Non-Coal Layoffs x Post-Shale Shock				0.84 (0.70)
Non-Coal Layoffs				-0.39 (0.36)
N	38919	38919	38919	35921
Adjusted R^2	0.00	0.00	0.00	0.04
County Fixed Effects	Yes	Yes	Yes	Yes
Election Fixed Effects	Yes	Yes	Yes	Yes

Notes: HC2 standard errors clustered by county. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

J.4 Anthracite Coal Falsification Test

The other falsification test compares anthracite coal-producing counties with non-anthracite coal counties. Anthracite is a “rank” of coal, which is a measure of the degree of “coalification.” Anthracite coal has the highest carbon content and fewest impurities, so it is often used in metallurgy due to its “caking” ability—caking is the ease with which coal can be converted to coke, which is used in oxygen furnaces for steel-making, for example. Other ranks of coal include “bituminous” and “sub-bituminous,” which are sometimes referred to as “thermal” coal and are often used in electrical power generation.

I classify counties as producing anthracite if, in the post-2005 period, the county has more than 50 employees working in anthracite. This threshold ensures that there is a meaningful share of local employment in the industry. County-level data come from the CBP. The NAICS code 212113 tracks anthracite mining employment, which enables the precise identification of where there is employment in this part of the coal industry.

The logic for this falsification test stems from insights gained in fieldwork. One conversation with a state bureaucrat in a county with anthracite coal clarifies how this type of metallurgical coal has been shielded from the effects of the shale shock.

Author: ...as I’m sure you know, coal is central to Pennsylvania. Some think I need to move away from coal. I’m curious if you have a sense of what should be done.

Local Official: Well, I mean, you’re always gonna need coal, and most of the coal around here is steel-quality coal. It’s not power plant coal. Power plant coal is sort of a low-grade coal... So, the market for the coal around here, I’d say, is always going to be there... Some of it from here still even goes to Japan for steelmaking.¹⁷

This conversation was one of the major impetuss for this empirical investigation. A state lawmaker representing coal interests also distinguished anthracite from other types of coal.

Author: So there’s this narrative that coal is on the decline... How true is that for the counties that you represent?

State Lawmaker: I think [the] decline is more pronounced today in the bituminous fields. ...but understand that, you know, the difference between anthracite and bituminous. The bituminous fields [are in] much, much, much greater [decline] than anthracite.¹⁸

J.4.1 Empirical Validation of Anthracite Resilience to Shale Shock

As a systematic demonstration of the difference between unexposed (anthracite) and exposed (non-anthracite) coals, Table J3 presents the results from an autoregressive distributive lag model of annual coal production by rank on gas production. I estimate the following

¹⁷Interview A5 with local government bureaucrat in Southwest Pennsylvania.

¹⁸Interview A9 with Pennsylvania state lawmaker.

specification:

$$Coal_t = Gas_{t-1} + \sum_{n=1}^5 Coal_{t-n} + \epsilon_t \quad (8)$$

As before, data come from the EIA and run from 1949 to 2021. Autocorrelation plagues most time series, and this model is no exception. As a correction, the specification includes two lags of coal production, which the p -value from the Durbin-Watson tests suggests is a successful remedy. The coefficient for anthracite coal production regressed on gas production is indistinguishable from zero, whereas, for the non-anthracite coal ranks, increased gas production decreases coal production. The point of this analysis is correlative, but one could advance arguments for why gas production causes changes in coal production in the “Granger” causality sense.

Table J3: Autoregressive Distributed Lag Model of Coal Production by Rank, 1949–2021

	Anthracite		Non-Anthracite	
	(1)	(2)	(3)	(4)
Intercept	−0.072 (0.121)	−0.090*** (0.029)	0.036 (0.387)	0.007 (0.025)
Gas _{<i>t</i>−1}	−0.781** (0.296)	−0.009 (0.009)	0.467 (0.291)	−0.074* (0.042)
Anthracite _{<i>t</i>−1}		0.748*** (0.176)		
Anthracite _{<i>t</i>−2}		0.009 (0.163)		
Anthracite _{<i>t</i>−3}		−0.210 (0.291)		
Anthracite _{<i>t</i>−4}		0.317 (0.246)		
Anthracite _{<i>t</i>−5}		−0.003 (0.111)		
Non-Anthracite _{<i>t</i>−1}				0.843*** (0.117)
Non-Anthracite _{<i>t</i>−2}				0.131 (0.131)
Non-Anthracite _{<i>t</i>−3}				0.076 (0.167)
Non-Anthracite _{<i>t</i>−4}				0.208 (0.198)
Non-Anthracite _{<i>t</i>−5}				−0.258 (0.156)
<i>N</i>	72	68	72	68
Adjusted <i>R</i> ²	0.620	0.987	0.188	0.960
Durbin-Watson test	0.088	1.968	0.072	1.997
<i>p</i> -value	0	0.346	0	0.395

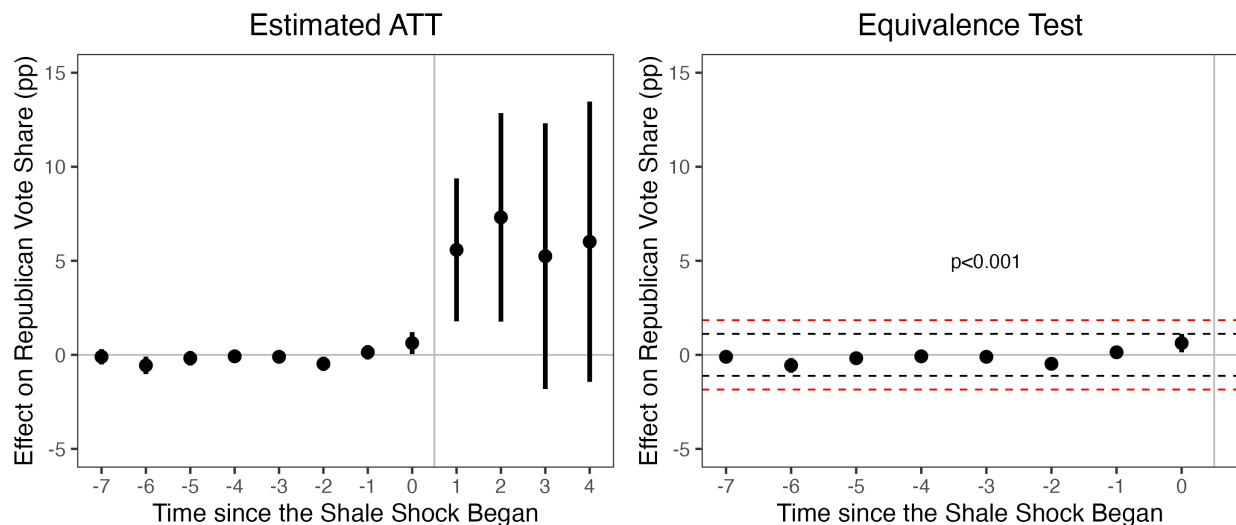
Notes: Anthracite coal has metallurgical applications, so it is less vulnerable to the shale shock, whereas non-anthracite coal ranks primarily have thermal applications, so they are vulnerable to the shale shock. Production values are standardized for interpretation. Models (1) and (3) estimate the correlation between lagged gas production and anthracite and non-anthracite coal production, respectively, without correcting for autocorrelation. Models (2) and (4) add five lags of coal production, which the Durbin-Watson tests indicate ameliorates autocorrelation. Coal production data come from EIA’s 2022 Annual Coal Report (1949-2021). Gas production data come from the EIA’s Natural Gas Gross Withdrawals and Production time series (1936-2022). Standard errors are estimated with the HAC covariance estimator. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

J.4.2 Anthracite Falsification Test Results

I estimate the effect of the shale shock in anthracite counties versus exposed coal counties using the FEct estimator. The treatment indicator takes the value 1 if a county has greater than 1% of local employment in the coal industry before the shale shock. I call these non-anthracite coal mining counties, “exposed,” since they are vulnerable to coal-to-gas switching in electricity markets. By contrast, the control group consists of counties with anthracite mining employment.

Figure J3 plots the estimates using FEct. Following the shale shock, exposed coal mining counties became more likely to vote for the Republican candidate than unexposed anthracite mining counties. The equivalence test in the righthand side panel shows no detectable violation of the parallel trends assumption, indicating that these counties are valid comparison groups for causal inference.

Figure J3: Dynamic Treatment Effect of the Shale Shock in Exposed Coal Mining Counties vs. Resilient Anthracite Mining Counties on Two-Party Republican Vote Share



Notes: The left plot shows the dynamic treatment effects estimates using the FEct estimator. Treated counties are those with non-anthracite coal mining employment greater than 1% in the three years before the 2008 shale shock, whereas control counties are those with anthracite coal mining employment. The model includes county and election fixed effects and controls for hydraulic fracturing employment. The bars denote 95% confidence intervals from 2,000 block bootstrap replications clustered at the unit level. The right plot shows the pretreatment average prediction errors and their 90% confidence intervals. The red dashed lines denote the equivalence range set at $[-0.36\hat{\sigma}, 0.36\hat{\sigma}]$ as proposed by Hartman and Hidalgo (2018), whereas the black dashed lines mark the minimum range. The p -value indicates equivalence holds with high confidence.

K Moderating Effect of Visibility

This appendix describes the measurement of coal-to-gas switching visibility, the procedure to construct covariate balancing weights, the estimation strategy for the visibility regressions, the results from this empirical analysis, and the results from a sensitivity analysis for omitted variable bias.

K.1 Data and Measurement

The data on new power plant construction comes from the EIA-860 Form, a comprehensive survey by the EIA of existing and planned generators at electric power plants with one megawatt or greater of combined nameplate capacity.¹⁹ These data are the inputs to many EIA publications like the Annual Energy Review. Since reporting is mandated by law, the data are comprehensive. For data before 2001, I use the EIA-860A (Utility) Form. In all, the data run from 1990 through 2021.

For each election year, I map the plant-level database, which contains geographic information, to the general-level database, which includes fuel type and activity information. Then, I classify whether the fuel type is gas and whether the plant came online in the last four years. I focus on the previous four years since that is the relevant period before the election. It is consistent with my visibility argument that any construction during that period should have shaped beliefs about the causes of coal’s decline.

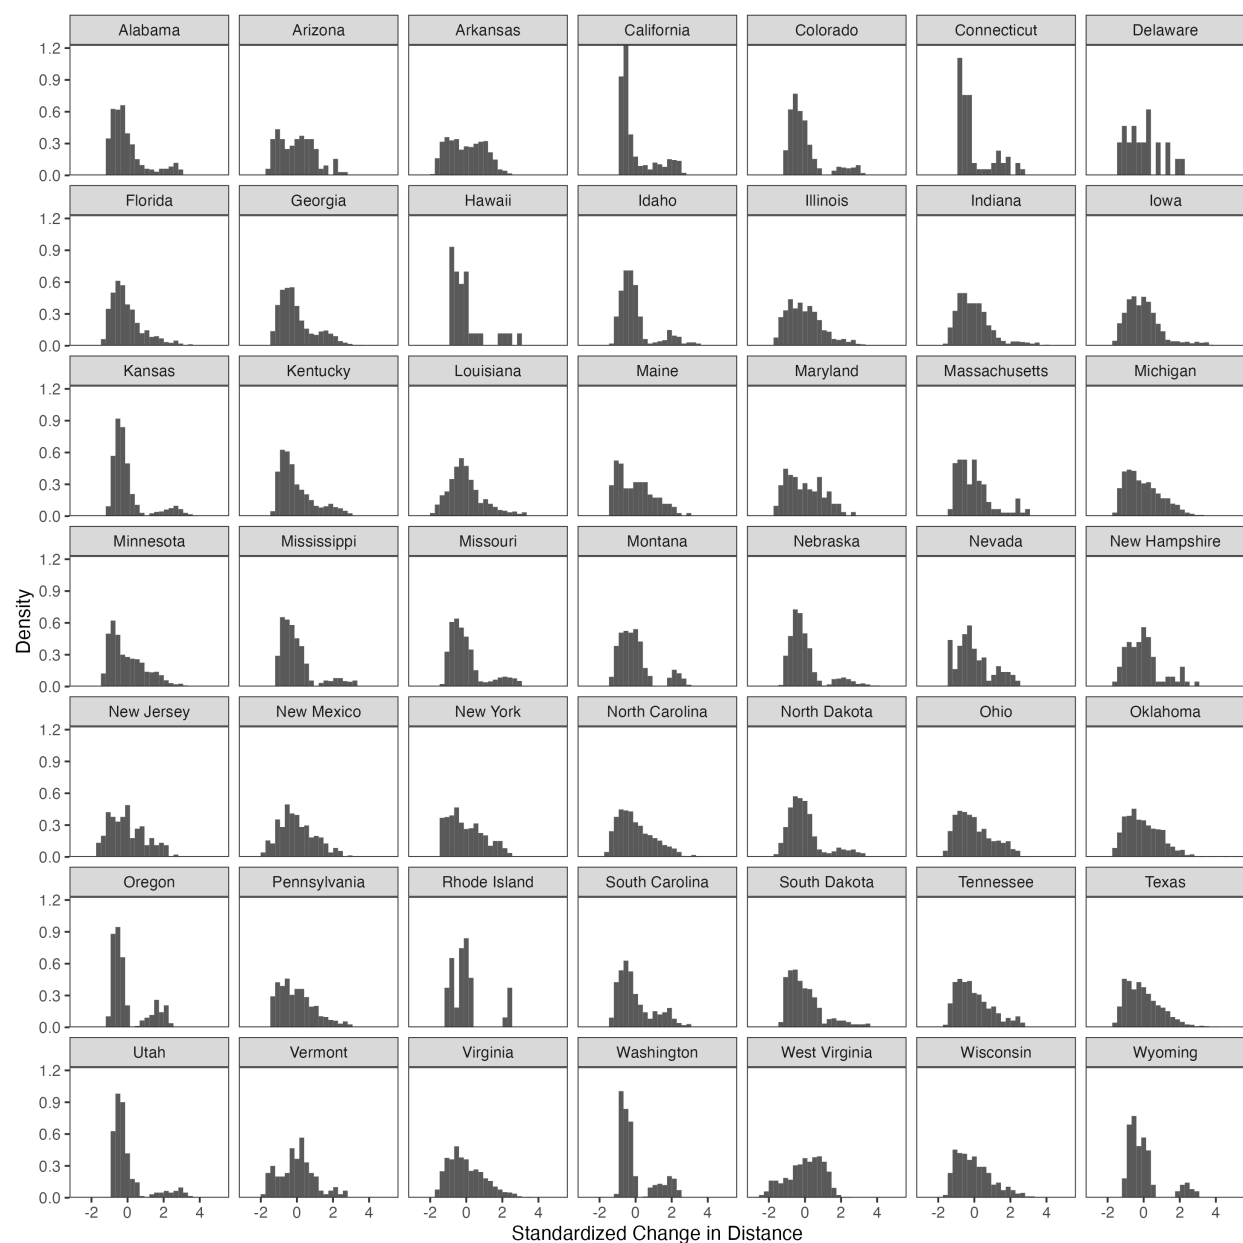
Each power plant is mapped to a complete list of counties throughout the United States. I also employ “fuzzy” matching of county names using the Levenshtein distance with a maximum distance of 3 since county names sometimes have typos when entered by electric power generators.

Then I use the ZIP codes provided for each plant to retrieve its longitude and latitude coordinates. For instances where no match is found between the ZIP code and longitude-latitude coordinates, I manually search the coordinates of the generator’s county.

Then for each county, I calculate the shortest distance between the county’s centroid and the newest gas-fired power plant. I define the shortest distance as being between two points, also called “great circle distance” or “as the crow flies” according to the Vincety ellipsoid method, which is more computationally intensive but very accurate. Figure K1 plots the distribution of within-county change in distance to new gas-fired power plants over the sample period. There is considerable within-unit variation over time, resulting from the tectonic shifts in electricity markets over the study period.

¹⁹<https://www.eia.gov/electricity/data/eia860/>

Figure K1: Within-State Deviation in a County's Distance to New Gas Power Plant between Elections, 1992–2020




Notes: The plot shows the distribution of with-state deviation in a county's distance since that is the relevant source of variation in the empirical model that includes state fixed effects (Mummolo and Peterson 2018). Note that some density values are greater than 1, but their support is limited, so the whole distribution still integrates to 1. Data from the EIA-860 Form with distance calculations done by the author.

K.2 Covariate Balance

Figure K2 reports the treatment-covariate correlations, where the treatment is defined as the interaction of the indicator for if a county has more than 1% of local employment in coal and the continuous measurement of within-state deviations in the distance to new gas-fired power plants. The population of units is those that are within two degrees of adjacency to coal counties, as in the primary analysis. Figure K2 shows mild imbalances by white, Hispanic, foreign-born, college, poverty, rural, population, age, and female workforce participation. However, the covariate balancing propensity scores succeed in reducing the treatment-covariate correlation to near zero.

Figure K2: Covariate Balance by Continuous Treatment Status in the Adjusted and Unadjusted Samples



Notes: Treatment status is defined as the interaction of the indicator for if a county has more than 1% of local employment in coal and the continuous measurement of a standard deviation increase of within-state distance to a new gas-fired power plant. Weights estimated using covariate balancing propensity scores (Imai and Ratkovic 2014).

K.3 Estimating Equation

The primary time series cross-sectional model that I estimate has the following equation:

$$y_{it} = \alpha + \gamma Coal_i + \lambda d_t + Visible_{it} + \delta_1(Coal_i \times d_t) + \delta_2(Coal_i \times Visible_{it}) + \delta_3(d_t \times Visible_{it}) + \tau(Coal_i \times d_t \times Visible_{it}) + \mathbf{X}_{it}^\top \beta + State_i + Election_t + \epsilon_{it}. \quad (9)$$

In this equation, y_{it} is the two-party Republican presidential vote share. $Coal_i$ indicates if a county has more than 1% of local coal employment before the shale shock. d_t is an indicator for the post-shale shock period. $Visible_{it}$ is the continuous measurement of the standardized within-state distance to a new gas-fired power plant. τ is the interaction term of interest, which is the differential effect of the shale shock in coal counties that are closer (or farther) from visible coal-to-gas switching in electricity markets. \mathbf{X} is a matrix of pre-treatment and time-varying covariates. $State_i$ is a state fixed effect, though some model specifications below replace it with county fixed effects. $Election_t$ is an election year fixed effect. Heteroskedastic-robust standard errors are clustered by county.

K.4 Visibility Regression Results

Table K1 presents the results of estimating the effect of the shale gas shock in coal counties, conditional on the visibility of coal-gas switching in downstream electricity markets. Models (1) and (2) include state fixed effects and satisfy the conditional ignorability assumption by the covariates included in the model. The triple interaction term of relevance is positive, as hypothesized, and not sensitive to including a control for power sector employment. Model (3) turns to within-county variation by using county fixed effects.²⁰ The coefficient of interest strengthens in magnitude while becoming more precisely estimated. Model (4) is the same as (3), except that it also includes covariate balancing propensity scores as weights, increasing the estimate’s magnitude. Lastly, Model (5) uses covariate balancing propensity scores as weights in a model with state fixed effects, which returns stronger results than the other models with state fixed effects but without balancing weights. Overall, the findings point in a consistent direction regardless of the approach to satisfying conditional ignorability—with covariates or balancing weights. The effect of the shale shock is greatest when visibility is lower.

²⁰The pretreatment covariates drop out since they would be colinear with the county fixed effects.

Table K1: Moderating Effect of Visibility on Republican Vote Share, 1992–2020

	(1)	(2)	(3)	(4)	(5)
Treatment and Moderator:					
Coal	−5.82*** (1.11)	−5.81*** (1.11)	−9.32*** (0.88)	−9.61*** (1.06)	−6.06*** (1.48)
Post-Shale Shock	19.95*** (0.55)	19.95*** (0.55)	20.19*** (0.58)	19.77*** (0.59)	19.50*** (0.58)
Gas Plant Distance	−1.13*** (0.26)	−1.13*** (0.26)	−0.52** (0.21)	−0.62*** (0.21)	−1.27*** (0.28)
Coal × Post-Shale Shock	7.95*** (1.13)	7.96*** (1.13)	8.50*** (1.20)	9.65*** (1.60)	8.84*** (1.51)
Coal × Gas Plant Distance	−0.27 (0.44)	−0.27 (0.44)	−0.79*** (0.27)	−1.14*** (0.37)	−0.71 (0.58)
Coal × Post-Shale Shock × Gas Plant Distance	0.92* (0.51)	0.92* (0.51)	1.57*** (0.35)	2.04*** (0.47)	1.70** (0.73)
Time-Varying Covariates:					
Hydraulic Fracturing Employment	−0.13 (0.17)	−0.13 (0.17)	0.04 (0.05)	0.03 (0.05)	−0.15 (0.20)
Power Employment		−0.01 (0.08)	0.05 (0.07)	0.13* (0.08)	0.14 (0.13)
Coal Union Share	0.55 (1.61)	0.56 (1.61)	−0.92 (1.32)	−2.06 (1.44)	−0.26 (1.98)
Coal Union Share × Post-Shale Shock	1.37 (1.87)	1.38 (1.87)	−2.00 (2.03)	−0.66 (2.68)	4.37 (2.88)
Pretreatment Covariates:					
White	4.19*** (0.53)	4.19*** (0.53)			
Hispanic	−1.21* (0.62)	−1.21* (0.62)			
Foreign-born	0.83 (0.70)	0.83 (0.70)			
College	−1.72*** (0.63)	−1.72*** (0.63)			
Income per capita (log)	−1.76** (0.89)	−1.76** (0.89)			
Poverty	−3.27*** (0.61)	−3.27*** (0.61)			
Rural	0.96* (0.52)	0.96* (0.52)			
Population (log)	0.69 (0.57)	0.68 (0.57)			
Under 40 years	−0.28 (0.39)	−0.28 (0.39)			
Female workforce	−0.29 (0.57)	−0.29 (0.57)			
<i>N</i>	6192	6192	6192	6192	6192
Adjusted <i>R</i> ²	0.580	0.580	0.829	0.829	0.420
Outcome Mean	61.68	61.68	61.07	61.07	61.07
Outcome SD	12.4	12.4	9.35	9.35	9.35
County Fixed Effects	No	No	Yes	Yes	No
State Fixed Effects	Yes	Yes	No	No	Yes
Election Fixed Effects	Yes	Yes	Yes	Yes	Yes
Covariate Balancing Propensity Score	No	No	No	Yes	Yes


Notes: Model (1) estimates the moderating effect of the within-state change in gas plant distance on two-party Republican presidential vote share; (2) adds a covariate for fossil fuel power plant employment; (3) uses county fixed effects to estimate the within-county change; (4) includes covariate balancing propensity scores as weights; and (5) includes covariate balancing propensity scores as weights but uses state instead of county fixed effects. The population of counties across all models is the same as used for matching (Figure D1), those with above 1% of local employment in the coal industry and those within two degrees of adjacency of the coal counties. HC1 standard errors clustered by county. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

K.5 Interaction Effect Diagnostics

The result that there is an attenuating effect of more visible coal-to-gas switching on Republican vote share in counties exposed to the shale shock rests on a multiplicative interaction effect – that is, the interaction between the post-shale shock indicator, the treatment indicator, and the visibility moderator (Figure 4). There are two primary assumptions that must be satisfied for this approach to be valid: linearity and common support. In this appendix, I conduct the diagnostic tests recommended by Hainmueller, Mummolo, and Xu (2019). Since the interaction term is a triple interaction, I subset the data into the pre- and post-shale shock periods and conduct the diagnostic tests on each sub-sample. Subsetting the data in this way also has the useful property of permitting the covariates to have a conditional effect by period, which helps to minimize possible unmodeled interactions (Blackwell and Olson 2022).

K.5.1 Linearity Assumption Diagnostic Tests


Figure K3: Linear Interaction Diagnostic Plots



Notes: The plot shows the relationships among the treatment (coal counties), the outcome (two-party Republican presidential vote share), and the visibility moderator (within-county standardized distance to new gas power plants) using the raw data. Since this is a triple interaction with the post-shale shock indicator, I subset the data into the pre- and post-shale shock periods. The blue and red lines represent the linear and LOESS fits, respectively. The plot reveals that the conditional expectation function of the outcome given the treatment is well-approximated by a linear model across the values of the moderator (Hainmueller, Mummolo, and Xu 2019).


K.5.2 Flexible Estimation Strategies

Figure K4: Binning Estimator



Notes: The plot shows the conditional marginal effects of the coal treatment at different levels of the visibility moderator from the binning estimator proposed by Hainmueller, Mummolo, and Xu (2019). The red error bars are 95% confidence intervals around the point estimates from the tercile bins of the moderator. The black line and shaded bands are estimates from a conventional linear interaction model. Standard errors clustered by county. This plot corresponds with the specification of model (5) in Table K1, except that it uses the binning estimator.

Figure K5: Kernel Estimator



Notes: The plot shows kernel-smoothed estimates of the marginal effect of the treatment at different levels of the visibility moderator in the pre- and post-shale shock periods as proposed by Hainmueller, Mummolo, and Xu (2019). I use 10-fold cross-validation to select the bandwidth. Shaded bands denote 95% confidence intervals using 2,000 county-clustered bootstrap samples. This plot corresponds with the specification of model (5) in Table K1, except that it uses the kernel estimator.

K.6 Sensitivity Analysis

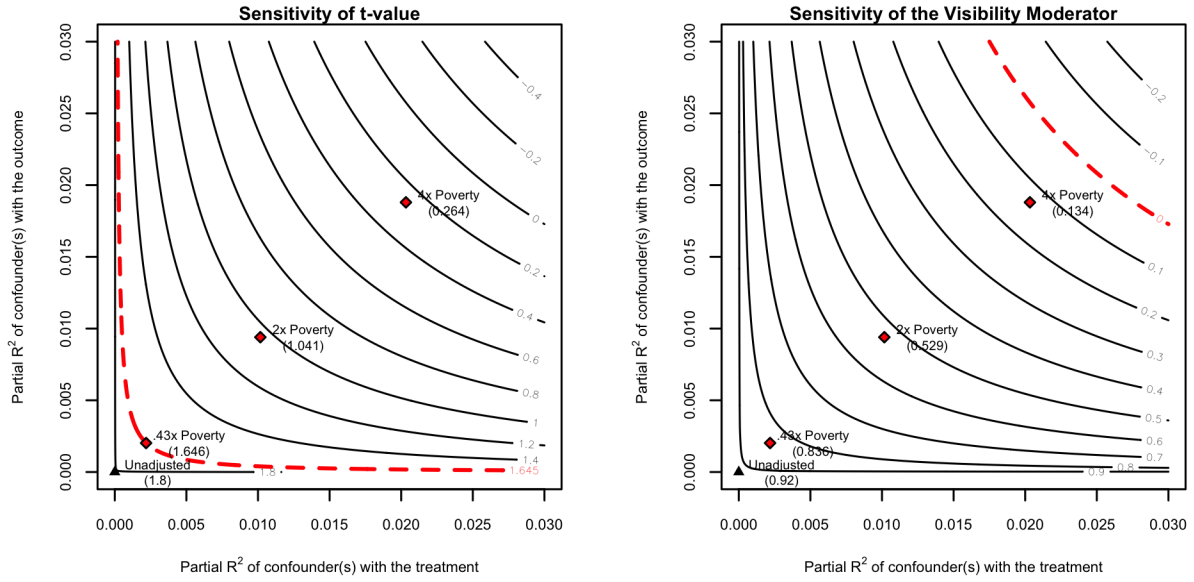
I conduct a sensitivity analysis to see how large a confounder that explains the variation in both the moderator and the outcome, that is orthogonal to the covariates, would have to be to change the conclusions. As a benchmark covariate, I use the county-level pre-shock share of poverty. The reason is that this covariate is one of the strongest predictors of distance to a new gas power plant, even when controlling for other covariates like population and income. A standard deviation increase in county poverty correlates with being 0.08 standard deviations closer to a new gas-fired power plant ($p < 0.001$). Table K2 presents the results from the sensitivity analysis of the finding in Table K1 with respect to the interaction of post-shock, coal indicator, and the gas plant distance moderator. Specifically, the table below shows the maximum strength of unobserved confounders with association with the moderator and the outcome bounded by a multiple of the observed explanatory power of the standardized share of the county in poverty, our benchmark covariate. Interpreting the results in the table, the robustness value indicates that an unobserved confounder (orthogonal to the covariates) that explains more than 2.3% of the residual variance of both the moderator and the outcome would be strong enough to bring the point estimate to 0 (a bias of 100% of the original estimate). In terms of reducing the estimate to a range where it is no longer statistically different from 0 at the 10% significance level, there would have to be an unobserved confounder (orthogonal to the covariate) that explains more than 0.2% of the residual variance of both the moderator and the outcome. The benchmarking exercise presented in Figure D4 indicates that such a confounder is unlikely. It would have to have more than 43% the strength of poverty's correlation with new gas plant construction and the outcome while also being orthogonal to the covariates in the model.

Table K2: Sensitivity Analysis of Visibility Moderator to Unobserved Confounding

Outcome: Two-Party Republican Presidential Vote Share						
Treatment:	Est.	S.E.	t-value	$R^2_{Y \sim D \mathbf{X}}$	$RV_{q=1}$	$RV_{q=1, \alpha=0.1}$
<i>Coal \times Post-Shale Shock \times Gas Plant Distance</i>	0.919	0.508	1.811	0.1%	2.3%	0.2%
df = 6134	<i>Bound (.43x Poverty): $R^2_{Y \sim Z \mathbf{X}, D} = 0.2\%$, $R^2_{D \sim Z \mathbf{X}} = 0.2\%$</i>					

Notes: Sensitivity analysis conducted according to the method proposed by Cinelli and Hazlett (2020). Model subjected to sensitivity analysis from Table K1.

Figure K6: Sensitivity Analysis of Visibility Moderator Results



Notes: Bias contour plots of the t -value (left) and the interaction of the estimate of the interaction between the ATT and the new gas plant distance moderator (right). Red diamonds indicate that a confounder up to 43% as strong as the observed county poverty covariate would not bring the lower bound of the confidence below 0 at the 10% significance level, while a confounder at least 400% as strong as the observed county poverty covariate would not bring the estimate to 0. Estimates come from a linear regression of two-party Republican presidential vote share on the interaction of the coal treatment indicator, the post-shock indicator, and the distance to a new gas plant moderator, while controlling for hydraulic fracturing employment, power plant employment, unionized coal employment pre- and post-shale shock, white, Hispanic, foreign-born, college, income per capita (log), poverty, rural, population (log), under 40 years, and female workforce participation. HC1 standard errors clustered by county.

K.7 New Gas Plants Do Not Offset Tax Revenue Loss

Table K3 shows how there is no correlation between distance to a new gas power plant and changes in local tax revenue. This is likely because any new tax revenue is not sufficient to offset the decline of coal. Indeed, there is an imprecise negative correlation between the shale gas shock and a decline of revenue in coal counties; however, there is no difference in the magnitude of this decline conditional on the distance to new gas plants, which indicates there is no detectable revenue offsetting.

Table K3: Linear regression of county revenue (log) on the standardized distance to new gas power plants, 1972–2012

	Total Revenue (log)			Total Taxes (log)		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	9.850*** (0.023)	9.850*** (0.023)	9.849*** (0.023)	8.925*** (0.018)	8.925*** (0.018)	8.925*** (0.017)
Gas Plant Distance	0.014 (0.011)	0.014 (0.011)	0.016 (0.015)	−0.019* (0.010)	−0.019* (0.010)	−0.015 (0.016)
Coal County		1.158*** (0.000)	1.158*** (0.000)		1.203*** (0.000)	1.203*** (0.000)
Post-Shale Shock		0.000 (0.000)	0.000 (0.000)		0.000 (0.000)	0.000 (0.000)
Gas Plant Distance × Coal County			−0.004 (0.019)			−0.006 (0.021)
Gas Plant Distance × Post-Shale Shock			0.000*** (0.000)			0.000 (0.000)
Coal County × Post-Shale Shock		0.000 (0.000)	0.000 (0.000)		0.000 (0.000)	0.000 (0.000)
Gas Plant Distance × Coal County × Post-Shale Shock			0.000			0.000 (0.000)
<i>N</i>	1834	1834	1834	1834	1834	1834
Adjusted R^2	0.958	0.958	0.958	0.973	0.973	0.973
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
County Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Data from the Government Finance Database (ICPSR 37641), which run through 2015. Revenue data aggregated to the election year level to correspond with the other empirical analyses. Heterogeneity-robust standard errors clustered by county. * $p < 0.1$; ** $p < 0.05$; and *** $p < 0.01$.


K.8 Post-Treatment Bias

This appendix explains why conditioning on employment created by new gas plants helps to identify the visibility mechanism and does not introduce selection bias from controlling for a post-treatment variable. I first constructed directed acyclic graphs (DAG) that illustrated the conditions when conditioning on a post-treatment variable would introduce selection bias, also known as intermediate variable bias. I contrast this with a DAG showing how the theorized process should not be vulnerable to selection bias, providing theoretical and empirical reasons to support this assumption. Lastly, I employ the methods proposed by Acharya, Blackwell, and Sen (2016), including sequential- g estimation, to parse out the causal mechanism, demonstrating how the direct effect of new gas plant constructions operates entirely through a non-economic channel—that is, this is evidence consistent with the claim that the moderator as measured is capturing the visibility of the shale shock.

K.8.1 DAGs Assessing Post-Treatment Bias

Figure K7 demonstrates the conditions when conditioning on a post-treatment variable would introduce selection bias. This occurs when there is an omitted variable correlated with both the post-treatment measure, M , and the outcome, Y .

Figure K7: DAG of How Conditioning on Post-Treatment Variables Induces Selection




Notes: D is the treatment, U is an omitted confounder that affects the mediator, M , and the outcome, Y . In this case, conditioning on M would introduce serious selection bias, also called intermediate variable bias (Acharya, Blackwell, and Sen 2016).

Unlike the situation above, where conditioning on M is problematic, there are theoretical and empirical reasons to expect that the post-treatment variable—employment from gas power plants—is not caused by an omitted intermediate factor that would induce selection bias. Figure K8 depicts a DAG of the theorized process. The shale shock, D , is responsible for the creation of new gas power plants, denoted X . These new gas power plants are both visible to closer communities, denoted V , and have localized economic benefits, denoted E . Since our interest is in the moderating effect of visibility, the empirical model controls for E . This means that the interaction term of the treatment and X has the interpretation of the moderating effect of new gas plants *net* their local economic effects—that is, the effect of visibility.

There would be the potential for selection bias if an omitted intermediate variable, U , correlated with the local economic benefits of new gas power plants, E , and changes in GOP vote share, Y . However, there are theoretical reasons to expect that these jobs are driven by the shale shock, D , as opposed to an omitted confounder. The empirical record supports this reasoning, given the studies cited throughout this paper indicating that cheaper gas is responsible for the rise of new gas power plants that displaced coal.

Further, the research design accounts for possible confounding, U , by balancing counties on factors that predict the creation of a new gas power plant, which would also positively correlate with their local economic effects. Additional model specifications include control variables instead of balancing weights. Lastly, I conducted a sensitivity analysis to assess how large an omitted confounder would have to be to alter the conclusions, which indicates that such an extreme confounder is unlikely given our contextual knowledge of the determinants of power plant siting (Appendix K.6).

Figure K8: DAG of Theorized Shale Shock Effect



Notes: D is the shale shock treatment. Y is Republican vote share. X is the construction of a new gas power plant, which affects the visibility (V) of the shale shock and has localized economic effects (E). This paper argues that V moderates the effect of D on Y . Since there are strong theoretical and empirical reasons to expect that E is the consequence of D and not an omitted confounder, conditioning on E should not introduce the type of selection bias illustrated in Figure K7. Instead, conditioning on E will mean that the estimate of X reflects the moderating effect of new gas plant construction *net* the economic effect of these plants, which I argue is the moderating effect of visibility. The estimation strategy also includes weights and controls to account for possible omitted variables, U , that could affect E and Y through a pathway unrelated to D .

K.8.2 Sequential-g Estimation

In this sub-appendix, I employ the sequential- g estimation method to estimate the controlled direct moderating effect of distance to new gas power plants, net the local economic effects of their construction. Under identification assumptions described in Acharya, Blackwell, and Sen (2016), this procedure *demediates* the moderating effect of changes in gas power plant employment from the outcome, thus retrieving an estimate of the moderating effects of distance to new gas power plants, which I interpret as capturing the visibility of coal-to-gas switching in electricity markets.

Table K4 presents the estimates of the moderating effect of distance to new gas power plants with and without demediation of local employment from gas power plants and hydraulic fracturing.²¹ If economic changes, as opposed to visibility, drive the moderating effect, then the interaction coefficients in model (2) should shrink to 0 because they subtract out the mediating effect of economic changes correlated with distance to new gas power plants. However, there is no change to the point estimates in model (2), which indicates that the moderating effect of distance to new gas power plants is not driven by local economic changes such as increased employment. This is evidence consistent with the claim that visibility is responsible for the moderating effect.

²¹These models use covariates instead of balancing weights because the bootstrapping procedure for constructing the standard errors for the sequential-g estimates introduces complexities with weighting.

Table K4: Sequential g -estimates of the moderating effect of distance to new gas power plants on Republican vote share in coal counties after the shale shock

	(1)	(2)
Coal	−5.813*** (1.110)	−5.813*** (1.103)
Post-Shale Shock	19.954*** (0.553)	19.954*** (0.551)
Gas Plant Distance	−1.130*** (0.256)	−1.130*** (0.258)
Coal × Post-Shale Shock	7.957*** (1.131)	7.957*** (1.106)
Coal × Gas Plant Distance	−0.274 (0.437)	−0.274 (0.431)
Coal × Post-Shale Shock × Gas Plant Distance	0.919* (0.508)	0.919* (0.509)
Hydraulic Fracturing Employment	−0.126 (0.165)	
Power Employment	−0.013 (0.084)	
Coal Union Share	0.563 (1.614)	0.563 (1.628)
Coal Union Share × Post-Shale Shock	1.376 (1.868)	1.376 (1.870)
White	4.190*** (0.526)	4.190*** (0.552)
Hispanic	−1.210* (0.618)	−1.210* (0.677)
Foreign-born	0.835 (0.700)	0.835 (0.708)
College	−1.724*** (0.628)	−1.724*** (0.644)
Income per capita (log)	−1.758** (0.895)	−1.758* (0.905)
Poverty	−3.269*** (0.612)	−3.269*** (0.632)
Rural	0.956* (0.518)	0.956* (0.546)
Population (log)	0.683 (0.568)	0.683 (0.579)
Under 40 years	−0.280 (0.390)	−0.280 (0.404)
Female workforce	−0.290 (0.571)	−0.290 (0.576)
N	6192	6192
Adjusted R^2	0.580	0.582
State Fixed Effects	Yes	Yes
Election Fixed Effects	Yes	Yes

Notes: Cluster-robust standard errors for (1). Block-bootstrapped standard errors from 1,000 samples with replacement for (2). Clusters are at the county level. Models (1) and (2) correspond with the first and second stage estimates, where (2) represents the demediated estimates. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

K.9 Microfoundational Evidence

K.9.1 Survey

In August 2023, I fielded a survey in the Southwest Pennsylvania area for a follow-up project on the effects of visibility on voters’ attributions of responsibility for economic changes.²² A subset of respondents ($N = 37$) come from a county where there was a new gas power plant constructed after the shale gas shock.²³ The data collected in this survey are for a parallel project, but I will present one result to demonstrate how citizens are aware of gas power plants in their county. Of the surveyed citizens, 81% said they see the gas power plant at least every year. Less than 19% said they do not recall seeing it. Over 27% said they see the gas power plant every day. This is direct evidence that local contextual factors—new gas power plants—are visible.

Table K5: Summary Statistics for 2023 County Fair Sample, Visible Subset

	Mean	SD	N
Age	43.08	17.06	37
Female	0.54	0.51	37
Non-White	0.05	0.23	37
College	0.24	0.43	37
Coal Household	0.22	0.42	37
Employed	0.46	0.51	37
Democrat	0.27	0.45	37
Republican	0.51	0.51	37

K.9.2 News

To test the face validity of the claim that people could have connected new gas plant construction with the decline of coal, I examined newspaper articles and press releases around the opening of new gas power plants. For case selection, I examined power plants in Pennsylvania because this is the context with which I have the greatest familiarity given my fieldwork. I examined when the power plant was proposed, when it opened, how open to coal counties it was, and the reasons provided for the construction of a gas power plant. These anecdotes provide evidence that when new gas power plants open, they can be linked to a narrative of an energy transition whereby the shale shock is causing new gas power plants to be opened at a more rapid clip. In some cases, this is directly linked to the decline of coal, as in the case of Snyder County below, where a new gas plant directly replaces an old coal plant.

- **Tenaska Westmoreland Generating Station**

1134 MW combined cycle gas power plant in Westmoreland County, Pennsylvania

Initiated in August 2016

²²See Appendix M.3 for more details about the study site and the sampling approach.

²³Table K5 provides a description of this subset.

The county home to the power plant is adjacent to 2 treatment group counties.

A local newspaper connected the construction of a new gas power plant in Southwest Pennsylvania to the shale shock. This report comes from a director at PJM, the regional power transmission organization, at a ceremony to dedicate the opening of a new gas power plant—a public statement about the role of the shale shock. “The Tenaska Westmoreland Generating Station that sits atop a hill in rural South Huntingdon is just the latest of what federal energy officials predict will be a wave of natural gas-powered plants to be built over the next 30 years...With the natural gas-powered plants in the state, which has ‘the third-largest natural gas basin in the world,’ Pennsylvania is ‘at the epicenter of a global energy transition,’ said Denise Brinley, executive director of the state’s Office of Energy.” (Napsha 2019)

Tenaska, the company opening the gas power plant in Westmoreland above, connected the shale shock to the construction of the new plant in its press release: “Tenaska Westmoreland is a valuable generating asset for the region, providing reliable electricity to the PJM market while *utilizing natural gas from the rich Marcellus formation*” (emphasis added).²⁴

- **Panda Liberty Generation**

829 MW combined cycle gas turbine power plant in Bradford, Pennsylvania

Commissioned in 2016

The power plant is far from treated coal counties in Pennsylvania.

An article covering the construction of a combined-cycle gas power plant that began construction in 2013 and had its official ground-breaking in 2014 said that the Pennsylvania plant was made possible with gas “sourced from Marcellus Shale natural gas formation, which also supplies gas for the Panda Power funded Liberty Generating Station” (**powertechnology**).

- **Panda Hummel Station**

1194 MW combined cycle gas turbine power plant in Snyder, Pennsylvania

Construction began in late 2015

The power plant is within two degrees of adjacency to a treated coal county.

In some counties, there used to be a coal-fired power plant that was directly replaced by a gas plant. For example, the Panda Hummel Station in Snyder County, Pennsylvania, began construction as the coal plant was decommissioned, as reported by a local newspaper (Moore 2018).

Another local news report connected the construction of the gas plant with the shale gas revolution: “The plant is powered by a UGI Energy Services gas line, with gas extracted from the Marcellus Shale about 34 miles north...The projects that get built are the ones that get finances. Building is dependent on the market” (**scicchitano**).

²⁴Tenaska, “Benefits of Tenaska Westmoreland Generating Station Celebrated,” July 31, 2019, <https://bit.ly/3R51FWg>, accessed 9/5/2023.

The press release for the plant’s announcement connected the construction of the gas plant, and the replacement of the more labor-intensive coal plant, as part of an energy transition from coal to gas: “We are proud to continue our successful track record with Panda Power Funds with this modern, highly-efficient power plant. As America continues to turn to cleaner-burning, low-cost natural gas, we are pleased to deliver world-class power generation equipment to provide reliable and efficient energy to major power markets in the U.S.,” said Martin Tartibi, Senior Executive Vice President of Energy Solutions Americas at Siemens Power and Gas Division.”²⁵

An engineering company working on the project framed it as part of “Cleaning up ‘coal country.’... Supplying nearly twice the power of the coal plant it replaces and reducing key emissions by more than 90%, Hummel Station will be one of the cleanest natural gas-powered plants in the US. It is the country’s largest coal-to-gas conversion to date. In an area long-dependent on coal, its combined-cycle technology will generate 180% more power, use 97% less cooling water and reduce sulfur dioxide (SO₂) and nitrogen oxide (NO_x) emissions.”²⁶

²⁵BusinessWire, “Siemens to power natural gas-fired plant in Pennsylvania,” October 28, 2015, <https://bit.ly/45K4OPJ>, accessed 9/6/2023.

²⁶Bechtel, “Hummel Combined Cycle Power Plant,” <https://bit.ly/44Hsmn8>, accessed 9/6/2023.

L Long-Standing National Partisan Environmental Issue Cleavage

This appendix provides evidence that there is a long-standing difference between the Republican and Democratic parties when it comes to environmental policy. Environmental policy refers to both climate policy and traditional air pollution controls. I first discuss other studies that document this finding, which I then confirm with an analysis of party platforms, which are high-profile statements of a party’s programmatic intentions. The end of this appendix discusses when climate change, in particular, emerged on the political agenda.

L.1 Studies on US Environmental Policy Polarization

Previous studies document elite partisan polarization on the environment. Dunlap, Xiao, and McCright (2001) examine League of Conservation scores of Senators and Representatives by political party from 1973-2000 and find that through the entire period, Republicans are less pro-environmental on average. This study notes a growing divergence on the environment since the 1970s, which the authors attribute to the Reagan administration.²⁷ These authors reflect, “there is currently a huge partisan cleavage over environmental policies” (30).

L.2 Analysis of Party Platforms

I demonstrate that the parties have conflicting preferences over environmental regulation by coding the party platforms as either pro or anti-environmental. Platform texts presented at presidential conventions come from the American Presidency Project and cover 1984 to 2020. For this paper, it is only necessary to demonstrate that the national parties held incongruent policy preferences in 2008, the time of the shale shock. However, the more extended period of study reveals how the pro- and anti-regulatory division runs deep.

The author employed the following coding procedure:

1. Search the party platform text for “coal,” “climate change,” and “global warming.”
2. Read the surrounding paragraphs for context, then record the platform as supporting or opposing environmental regulations.

There are not many ambiguous cases, which makes coding straightforward. Examples of opposing regulation include the 2016 Republican platform: “Its Clean Power Plan - the centerpiece of the President’s war on coal - has been stayed by the Supreme Court. I will do away with it altogether.” Examples of supporting environmental protections include the 2016 Democratic platform: “our climate change policy will cut carbon emission.” Of course, both parties argue they are pro-environment. I interpret statements that say that environmental quality is best achieved through the free market and incentives as anti-regulatory positions.

²⁷ “...the Reagan administration’s antipathy toward environmental protection represented a marked contrast to the early bipartisan image of environmental policy-making begun in the 1960s and continued through the Nixon, Ford and Carter administrations. More recently the Republican Congress of 1994 and the current Bush administration have reinforced the image of conservative Republicans being opposed to strong environmental protection efforts” (Dunlap, Xiao, and McCright 2001, 23).

All Republican presidential candidates since 1984 have advocated for free market approaches that deprioritize binding regulations on coal extraction and combustion. In contrast, all Democratic presidential candidates since that time have supported more stringent regulatory solutions, especially over the last decade.

Below I provide key excerpts from early party platforms to demonstrate that the Democratic Party has long-articulated support for regulatory efforts to combat climate change and reduce air pollution. In contrast, the Republican Party has preferred a more “free-market” approach.

- In 1988, the Democratic Party platform called for international environmental cooperation to address the “greenhouse effect,” whereas the Republican platform said “[t]he United States enjoys a rich national endowment of enormous supplies of coal which can provide a secure source of energy for hundreds of years.”
- In 1992, the Democratic Party platform said, “[t]he United States must become a leader, not an impediment, in the fight against global warming. I should join our European allies in agreeing to limit carbon dioxide emissions to 1990 levels by the year 2000.” In contrast, the Republican Party platform objected to efforts to address climate change: “a Republican Senate will not ratify any treaty that moves environmental decisions beyond our democratic process or transfers beyond our shores authority over U.S. property. The Democrats’ national candidates, on the other hand, insist the U.S. must do what our foreign competitors refuse to do: abolish 300,000 to 1,000,000 jobs to get a modest reduction in ‘greenhouse gases.’”
- In 1996, the Democratic Party platform echoed its earlier support for international cooperation to address climate change: “We will seek a strong international agreement to further reduce greenhouse gas emissions worldwide and protect our global climate.” The Republican Party platform, by contrast, cast doubt on the scientific consensus: “Despite scientific uncertainty about the role of human activity in climate change, the Clinton Administration has leapfrogged over reasoned scientific inquiry and now favors misdirected measures, such as binding targets and timetables, imposed only on the United States and certain other developed countries, to further reduce greenhouse gas emissions.”
- In 2000, the Democratic Party platform articulated concern about the harm from global warming and advocated further action by the United States to reduce emissions: “We must act now to protect our Earth while preserving and creating jobs for our people. In 1997, we negotiated the historic Kyoto Protocols, an international treaty that will establish a strong, realistic, and effective framework to reduce greenhouse emissions in an environmentally strong and economically sound way. We are working to develop a broad international effort to take action to meet this threat. Al Gore and the Democratic Party believe we must now ratify those Protocols.” In contrast, the Republican Party platform rejected international cooperation to address climate change: “The United States should aggressively pursue its national interest. Unlike the current administration, Republicans do not believe multilateral agreements and international institutions are ends in themselves. The Kyoto treaty to address momentous energy

and environmental issues was a case in point. Whatever the theories on global warming, a treaty that does not include China and exempts ‘developing’ countries from necessary standards while penalizing American industry is not in the national interest.”

L.3 Climate Policy Came on National Agenda in 1980s/90s

While there have been long-standing environmental policy differences between the Democratic and Republican parties, climate policy, in particular, emerged as an issue in the 1980s and 1990s. Mildenberger (2020, 104) describes how

US political attention to climate change intensified in the late 1980s as the science of climate change became more certain. In part, climate change’s increased profile stemmed from high- salience weather events, including a record 1988 drought in the American Midwest and West (Weart 2008). At the same time, June 1988 congressional testimony profiling the climate threat by scientist James Hansen to the Senate Energy and Environment Committee received widespread media coverage. Hansen’s testimony and the year’s extreme weather were linked by the media, resulting in a spike in public attention (Kamieniecki 2006).

Similarly, McCright and Dunlap (2011, 158) say that ““By the early 1990s, the U.S. environmental community—the environmental movement, mainstream climate scientists, and environmental policymakers—had successfully defined global warming as a legitimate problem deserving the attention of policymakers.”

In Mildenberger’s (2020) account of US climate policymaking, he identifies 1988 as the period beginning (failed) climate action. Notably, George H.W. Bush who softened his stance on global warming on the campaign trail ended up reversing course, stalled international climate efforts, and attacked Clinton’s stance on climate change. Mildenberger (2020, 100) describes how

numerous climate policy proposals made it onto the political agenda in the 1990s and early 2000s. The Clinton administration proposed an energy tax as part of its 1993 deficit reduction package that was explicitly framed as increasing carbon pollution costs...²⁸ During Kyoto Protocol debates, Clinton administration officials debated an emissions trading scheme architecture.

²⁸While there is some evidence that the BTU was designed to try to shield coal miners (Mildenberger 2020, 107–109), ultimately congressional representatives from coal-dependent districts raised serious concerns about the proposal and carbon pricing more generally.

L.4 Politicization of Environmental Regulations in Presidential Debates

This appendix provides evidence that Republican presidential candidates sought to politicize the Democratic presidential candidate's support for environmental regulations during nationally televised debates. Debates are relevant to study since they regularly attract national viewership of over 65 million people, making them an essential means for parties to communicate their issue positions.²⁹ Notably, former President Trump drew a stark distinction between himself and Hillary Clinton, the Democratic candidate, on the issue of environmental regulations:

...energy is under siege by the Obama administration, under absolute siege. The EPA — Environmental Protection Agency — is killing these energy companies...

We have to guard our energy companies. We have to make it possible...the EPA is so restrictive that they are putting our energy companies out of business. And all you have to do is go to a great place like West Virginia or places like Ohio, which is phenomenal, or places like Pennsylvania, and you see what they're doing to the people, the miners and others, in the energy business. It's a disgrace.³⁰

This is not just a Trump phenomenon; presidential hopeful Mitt Romney also contrasted himself with Obama on coal in two out of three debates:

And, by the way, I like coal. I'm going to make sure we can continue to burn clean coal. People in the coal industry feel like it's getting crushed by [Obama's] policies.³¹

Talk to the people that are working in those industries. I was in coal country. People grabbed my arms and said, "Please save my job." The head of the EPA said, "You can't build a coal plant. You'll virtually — it's virtually impossible given our regulations." When the president ran for office, he said if you build a coal plant, you can go ahead, but you'll go bankrupt. That's not the right course for America.³²

These quotes demonstrate how Republican Party elites attempted to blame the Democratic Party's support for environmental regulations for the coal industry's decline. These utterances are notable since debate time is precious, so there must be the expectation that these messages would be compelling to some voters.

²⁹ "Which Presidential Debates Drew The Biggest TV Audiences?" <https://bit.ly/3ntTfrn> (Accessed November 22, 2021).

³⁰ Presidential Debate, October 9, 2016.

³¹ Presidential Debate, October 3, 2012.

³² Presidential Debate, October 16, 2012

M Fieldwork

M.1 Study Site Representativeness

This appendix provides information about the communities in which interviews were conducted, with a focus on the characteristics of these communities along the dimensions of my argument. Table M1 summarizes relevant economic, visibility, and socio-demographic data grouped by whether a county is in the “field” where the interviews took place, a treated coal county in the “state” where the interviews, or a treated coal county nationwide.

The first dimension to consider is the economic effects of the shale shock, which depends on a county’s reliance on coal. The second row of Table M1 shows how the study site had greater pre-shock coal production than other counties in Pennsylvania and across the sample. While the county underwent a large decline in coal product, this is comparatively less than other coal counties in the state and nationwide. The reason is likely because production closed first among the smaller, less efficient operators. Nonetheless, there was a substantial decline in county coal employment, going from 32% of local employment to 17% after the shale shock. This is more than the statewide average and the entire sample.

Unlike other coal counties, the study site was home to relatively more employment in hydraulic fracturing. However, the local share of employment in hydraulic fracturing only averaged 4.77%, which is a consequence of the industry being more capital than labor-intensive.

In terms of the measure of visibility, which is whether there is a new gas-fired power plant that makes clear how gas is being used in electricity—displacing coal—the county is further away from new gas plants than the national average but a bit closer than the state average in Pennsylvania. That said, this average is not the most informative since the county itself has no new gas power plants. The visibility of coal-to-gas switching is low in the county.

Turning to the outcome, there is a sizeable shift to the Republican Party across all coal counties. The shift is even larger for the study site, likely because there was a lower baseline of Democratic support.³³

Lastly, the socio-demographic characteristics of the study site counties are similar to other coal counties in Pennsylvania and across the United States. The study site is slightly whiter but otherwise similar in terms of college education, income, poverty, rurality, age, and female workforce participation.

³³Note that the summary statistics average together the entire pre-shock period, which includes some general shift to the Republican Party, which otherwise gets differenced out in the analysis, so the pre-shock average for GOP vote share is about evenly split, whereas Figure 2 clearly shows the Democratic Party’s advantage in coal country.

Table M1: Study Site Compared to State and Entire Sample of Coal Counties

	Field	State	All
Economic Effects of Shale Shock			
Coal Production (log) $\Delta_{Post-Pre}$	-0.10	-0.42	-0.86
Pre-Shock Coal Production (log)	17.03	15.46	14.95
Post-Shock Coal Production (log)	16.93	15.04	14.09
Coal Employment $\Delta_{Post-Pre}$	-192.44	-796.95	-328.95
Pre-Shock Coal Employment (%)	32.39	10.20	13.08
Post-Shock Coal Employment (%)	16.80	3.97	6.75
Hydraulic Fracturing Employment $\Delta_{Post-Pre}$	614.91	466.33	111.49
Pre-Shock Hydraulic Fracturing Employment (%)	0.07	0.19	0.58
Post-Shock Hydraulic Fracturing Employment (%)	4.77	2.04	1.36
Visibility			
Distance to Gas Power Plant (log) $\Delta_{Post-Pre}$	0.19	0.09	0.29
Pre-Shock Distance to Gas Power Plant (log)	12.02	12.04	12.11
Post-Shock Distance to Gas Power Plant (log)	12.21	12.14	12.40
Outcome			
GOP Vote Share $\Delta_{Post-Pre}$	21.90	15.04	16.60
Pre-Shock GOP Vote Share (%)	41.07	52.18	52.57
Post-Shock GOP Vote Share (%)	62.97	67.22	69.17
Socio-Demographics			
White	0.95	0.97	0.90
Hispanic	0.01	0.01	0.03
Foreign-born	0.01	0.01	0.01
College	0.08	0.09	0.08
Income per capita (log)	9.61	9.69	9.62
Poverty	0.15	0.12	0.18
Rural	0.67	0.59	0.65
Population (log)	10.61	11.25	10.21
Under 40 years	0.52	0.51	0.54
Female workforce	0.19	0.22	0.21

Notes: Coal production data begin in 2000. Socio-demographic data is from the 2000 Census. The state includes all coal counties defined as treated in Pennsylvania. All includes all coal counties in the treatment group.

M.2 Interviews

Interviews took place in Pennsylvania, the fourth largest coal producer in the United States. To select counties for interviews, the author took a balanced cluster random sample from a complete list of Pennsylvania’s anthracite and bituminous coal mine operators. The sampling strategy achieved covariate balance along county-level socio-demographic and political measures, such as gender, educational attainment, and 2016 Republican vote share. The author contacted relevant trade unions, local bureaucrats, politicians, coal and oil firms, and civil society groups within selected counties.

After initial interviews in the randomly sampled counties, the researcher picked two counties reliant on the coal industry. The author then spent two weeks physically in these locations during July and August 2021, visiting subjects in pre-arranged interviews and attending county fairs. The author returned again in August 2022 and August 2023 for an additional two weeks of time in the field. Snowball sampling in advance of entering the field and connections made at county fairs organically expanded the interviewee pool. In all, 60 people were interviewed, including 15 fossil fuel workers, five representatives of nongovernmental organizations for labor, business, and the environment, six bureaucrats, 11 politicians ranging from county commissioners to state senators, and 24 longtime residents.³⁴ With multiple visits to the same field site, the author was able to reinterview some of the same respondents, making sure to be respectful of their time, which helped to glean deeper insights through established trust.

Each respondent gave oral consent before participating. A semi-structured interview guide helped facilitate conversation, with modifications depending on the subject’s occupation. Many questions pertained to a separate project. Most interviewees consented to be recorded, which aided recollection. When it was not possible to record, the author took detailed notes.

- A1, Local politician, Random, Video call; In person, 5/7/2021, Reinterviewed
- A2, Local politician, Random, Phone, 6/16/2021
- A3, Local politician, Random, Video call, 6/17/2021
- A4, Local politician, Random, Video call, 6/14/2021
- A5, Local government employee, Random, Phone, 5/27/2021
- A7, Local politician, Random, Video call, 5/29/2021
- A8, Non-governmental organization, Deliberate, Video call, 6/10/2021, Reinterviewed
- A9, Local politician, Random, Video call, 6/4/2021
- A10, Community business owner, Snowball, Video call; In-person, 6/8/2021, Reinterviewed

³⁴I conducted 15 additional interviews after the article’s submission, which have been incorporated into revisions.

- A11, Community business owner, Snowball, Video call; In person, 6/11/2021, Reinterviewed
- A12, Union, Random, Phone, 6/10/2021
- A14, Local politician, Random, Video call, 6/16/2021
- A17, Non-governmental organization, Random, Video call; In person, 6/16/2021, Reinterviewed
- A18, Local politician, Random, Video call; In person, 8/18/2021, Reinterviewed
- B1, Local politician, Random, Phone, 7/20/2021
- B2, Local government employee, Random, In-person, 7/21/2021
- B3, Local government employee, Random, In-person, 7/21/2021
- B5, Resident, Deliberate, In-person, 7/21/2021
- B6, Resident, Deliberate, In-person, 7/21/2021
- B7, Local politician, Deliberate, In-person, 7/21/2021
- B8, Non-governmental organization, Deliberate, In-person, 7/22/2021, Reinterviewed
- B9, Resident, Deliberate, In-person, 7/22/2021
- B10, Resident, Deliberate, In-person, 7/22/2021
- B11, Resident, Deliberate, In-person, 7/22/2021
- B12, Gas, Deliberate, In-person, 7/22/2021
- B13, Gas, Deliberate, In-person, 7/22/2021
- B14, Coal, Deliberate, In-person, 7/23/2021
- B15, Gas, Deliberate, In-person, 7/23/2021
- B15, Resident, Deliberate, In-person, 7/22/2021
- B17, Gas, Deliberate, In-person, 7/23/2021
- B18, Coal, Deliberate, In-person, 7/23/2021
- B19, Resident, Deliberate, In-person, 7/23/2021, Reinterviewed
- B20, Local government employee, Random, Video call; In person, 7/28/2021, Reinterviewed
- B21, Gas, Snowball, In-person, 8/9/2021


- B22, Resident, Snowball, In-person, 8/9/2021
- B23, Resident, Deliberate, In-person, 8/5/2021
- B24, Local government employee, Random, Video call, 7/23/2021
- B25, Local government employee, Random, Video call, 7/23/2021
- B26, Coal, Snowball, In-person, 8/9/2021
- B27, Gas, Snowball, In-person, 8/7/2021, Reinterviewed
- B28, Resident, Snowball, In-person, 8/8/2021
- B29, Non-governmental organization, Snowball, In-person, 8/8/2021
- B30, Resident, Snowball, In-person, 8/8/2021
- B31, Coal, Deliberate, In-person, 8/10/2021
- B32, Coal, Deliberate, In-person, 8/10/2021
- C1, Technical School, In-person, Deliberate, 8/11/2022
- C6, Educator, In-person, Deliberate, 8/11/2022
- C11, Coal, In-person, Deliberate, 8/7/2022
- C3, Educator, In-person, Deliberate, 8/11/2022
- C14, Educator, In-person, Deliberate, 8/11/2022
- C15, Educator, In-person, Deliberate, 8/11/2022
- C16, Educator, In-person, Deliberate, 8/11/2022
- C17, Educator, In-person, Deliberate, 8/11/2022
- D1, Coal/Healthcare, In-Person, Snowball, 8/12/2023, Reinterviewed
- D2, Local Politician, Deliberate, In person, 8/10/2023
- D3, Resident, In-Person, Deliberate, 8/7/2023
- D4, Coal, In-Person, Deliberate, 8/8/2023
- D5, Resident, In-Person, Deliberate, 8/9/2023
- D6, Resident, In-Person, Deliberate, 8/9/2023
- D7, Coal, In-Person, Deliberate, 8/11/2023

M.3 Surveys

I fielded surveys ($N = 606$) in Southwest Pennsylvania to evaluate this claim descriptively. I recruited participants from this hard-to-reach population by taking convenience samples at three county fairs during the summers of 2021 and 2022. County fairs are iconic cultural institutions in the United States and are well-attended by residents. Though a convenience sample, the individuals surveyed match the local population in terms of most meaningful characteristics such as gender, race, and party identification.

To measure the extent that one’s close social network depends on the fossil fuel industry, the survey asked: “Among the 10 people you know best in your community, roughly how many work in coal, oil, or gas?”³⁵ Figure M1 presents a histogram of the responses. The modal answer is that *all* of the ten people an individual knows best in her community work in coal, oil, or gas. Since this is a convenience sample in one region, it cannot be generalized to the rest of coal country (e.g., Wyoming). Nevertheless, the distribution is striking. Most respondents know at least one person who works in fossil fuels, an individual who, from their perspective, could be conceivably impacted by climate policy. Although these results are surprising given the small but important share of local workers in these industries, they indicate a perception that fossil fuels are an essential part of the community.

Figure M1: Perceived Fossil Fuel Worker Proximity in Local Social Networks




Notes: Data from Southwest Pennsylvania county fair surveys fielded in July/August 2021 and August 2022 (total $N = 606$). The red bars show the unweighted distribution of responses. The black bars show the responses weighted to match the population with iterative proportional fitting. Population data come from the 5-Year ACS, specifically the joint distribution of age-sex-education, the joint distribution of age-sex-race, and the distribution of income.

Although convenience samples, the demographics of the individuals surveyed match population values along several meaningful dimensions. Table M2 compares the demographic characteristics of the summer 2021 sample with the population analogs for the counties in the study area, while Table M3 does the same for the summer 2022 sample.

To enhance the sample’s representativeness, I construct sample weights using population data from the 2018 5-year American Community Survey for the 2021 sample and the 2021

³⁵The question asks about fossil fuels broadly since this paper is part of a larger research agenda.

Figure M2: Field Survey Study Area



Notes: Study area shaded in blue covers Southwest Pennsylvania and surrounding counties in West Virginia and Ohio.

5-year American Community Survey for the 2022 sample. Specifically, the sample weights employ the joint distribution of age, education, and sex; the joint distribution of age, race, and sex; and the distribution of income to estimate weights with raking. The choice to use weights makes the results more conservative.

Table M2: Representiveness of 2021 County Fair Sample

	Sample	Population	Weighted
Sex/Age/Education			
Female \times 18-34 years \times College	0.04	0.07	0.06
Female \times 18-34 years \times No college	0.12	0.05	0.06
Female \times 35-64 years \times College	0.10	0.12	0.12
Female \times 35-64 years \times No college	0.19	0.12	0.14
Female \times >65 years \times College	0.03	0.03	0.03
Female \times >65 years \times No college	0.04	0.09	0.07
Male \times 18-34 years \times College	0.03	0.06	0.05
Male \times 18-34 years \times No college	0.18	0.09	0.10
Male \times 35-64 years \times College	0.07	0.10	0.10
Male \times 35-64 years \times No college	0.14	0.17	0.18
Male \times >65 years \times College	0.03	0.04	0.03
Male \times >65 years \times No college	0.03	0.07	0.06
Income			
<\$20,000	0.17	0.18	0.18
\$20,000-39,999	0.13	0.21	0.19
\$40,000-59,999	0.15	0.16	0.15
\$60,000-99,999	0.30	0.24	0.26
>\$100,000	0.25	0.21	0.22
Race			
White	0.93	0.94	0.94

Notes: The table collapses the 18-24 and 25-34, and 35-44 and 45-64 age bins together, respectively, for exposition. Also not shown is the joint distribution of race/age/sex used to construct weights. Population data from the 2018 5-Year ACS and cover the primary study site county.

Table M3: Representiveness of 2022 County Fair Sample

	Sample	Population	Weighted
Sex/Age/Education			
Female \times 18-34 years \times College	0.03	0.07	0.02
Female \times 18-34 years \times No college	0.14	0.05	0.10
Female \times 35-64 years \times College	0.08	0.12	0.06
Female \times 35-64 years \times No college	0.20	0.12	0.18
Female \times >65 years \times College	0.05	0.03	0.03
Female \times >65 years \times No college	0.05	0.09	0.09
Male \times 18-34 years \times College	0.03	0.06	0.02
Male \times 18-34 years \times No college	0.15	0.09	0.14
Male \times 35-64 years \times College	0.05	0.10	0.04
Male \times 35-64 years \times No college	0.16	0.17	0.23
Male \times >65 years \times College	0.02	0.04	0.02
Male \times >65 years \times No college	0.04	0.07	0.08
Income			
<\$20,000	0.13	0.18	0.15
\$20,000-39,999	0.15	0.21	0.19
\$40,000-59,999	0.17	0.16	0.16
\$60,000-99,999	0.33	0.24	0.25
>\$100,000	0.22	0.21	0.25
Race			
White	0.97	0.94	0.97

Notes: The table collapses the 18-24 and 25-34, and 35-44 and 45-64 age bins together, respectively, for exposition. Also not shown is the joint distribution of race/age/sex used to construct weights. Population data from the 2021 5-Year ACS and cover the primary study site county.

N News Media

N.1 Local News Media Coverage

I coded local online newspaper coverage of coal mine closures and the coal industry in Southwest Pennsylvania to understand better how residents might attribute blame for coal's decline. The focus on the Southwest Pennsylvania area allows for the news stories' content to be contextualized with the interviews in this region. The two primary local newspapers are the *Greene County Messenger* (owned by the Herald-Standard) and the *Observer-Reporter*.

A research assistant created a collection of relevant articles by searching for "coal mine closure" from 1 January 2008 to 31 December 2020. This returned 84 articles. I focused on the post-shock period because there was better online coverage. For example, the first relevant news article appeared in 2012. Of the articles returned from the search term, the research assistant coded the reason provided in the article for the coal mine's closure. Some of the articles did not refer to a specific mine closure but instead were commentary about trends in the coal industry. I then reviewed the articles and verified the coding.

There are five categories of reasons: regulations, low gas prices, other market forces, other reasons, and no reason given. Articles would sometimes provide multiple reasons.

There is a total of 72 local news articles. A handful of the articles are national news stories that are reprinted by the local paper, which I include since these stories represent information that would be placed in front of residents with the local newspaper branding.

While these data provide insight into how the local community attributed blame for the closure of a coal mine or power plant, there are limitations. First, in terms of generalizing the findings within the county, not all residents may learn about their county through these newspapers. Readership might reflect the views of more educated and wealthy individuals, which could create editorial pressures to emphasize issues favorable to the Democratic Party. Then in terms of generalizing these findings elsewhere, the Southwest Pennsylvania area has considerable hydraulic fracturing activities, so these articles might overstate the relevance of gas, which may not be the case in other communities. This would bias these results against the hypothesis that coal communities misattributed the primary reason for their industry's decline to regulations rather than the shale gas revolution.

Figure N1 describes the content of local news coverage over time in the top panel and for each coal mine/power plant closure in the bottom panel. First, examining coverage over time, it is apparent that during the Obama Administration from 2008-2016, the local news is more likely to mention the adverse effects of federal regulations when discussing the coal industry than other topics. Lower natural gas prices are sometimes mentioned, including in 2015. However, the relative frequency of this topic is much lower. Moreover, when low natural gas prices are noted, they are not mentioned as the sole cause of coal's decline. Instead, they are often bundled with other reasons for coal's decline, including regulations. Of the 16 times cheap gas prices are referenced, 10 of the times, the article also identifies federal government regulations as partly to blame.

Then, when looking at overall local news coverage, regulations are the most frequently given reason for coal's decline. Breaking down the topics by three coal mine and one coal power plant closures, regulations emerge as an often-mentioned topic. The Emerald Mine and Hatfield Power Station closure, which happened during the Obama administration, have

news articles that blame environmental regulation. For example, one article on 3 August 2013 quoted a local Republican Congressman who blamed Obama’s environmental regulations:

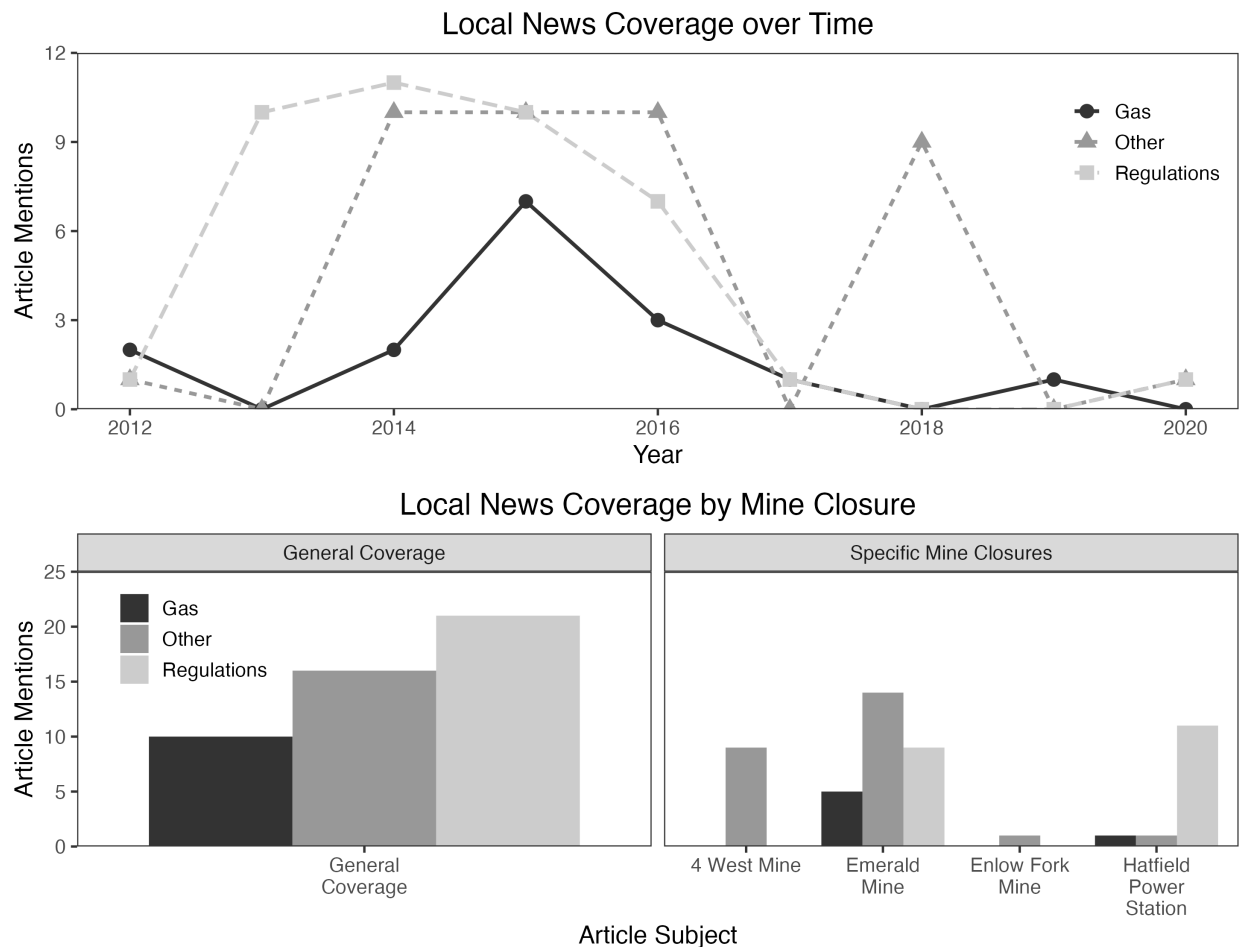
“President Obama recently announced that he intends to bypass Congress and create new limitations on new and existing coal-fired plants using the EPA,” Shuster said. “In recent years, the EPA has overstepped its authority and has been legislating via regulation. Just a few weeks ago, it was announced that two coal-fired plants in southwestern Pennsylvania would close down due to the overwhelming cost of complying with the EPA’s ever-expanding regulations. These are good, clean plants, and their closures are a major blow to the region.”

Notably, Obama’s proposal had not yet passed. The primary regulations the plant had to comply with were standards put in place before Obama’s time, some of which were enacted by the George H.W. Bush administration that signed into law the 1990 Amendments to the Clean Air Act, which dealt with acid rain. Nonetheless, the non-existent Obama administration regulations received blame from this local Congressman.

In contrast, the two coal mine closures during the Trump administration have articles that blame market forces. This may be the consequence of the Republican Party not having the Obama administration’s regulations as a target to politicize after the Trump administration said it had rolled back the supposedly harmful regulations.

Outside of these articles about particular mine closures, there were also op-ed pieces warning about the dangers of EPA regulations. For example, on 14 October 2014, an individual at a “free-market think tank” called the Commonwealth Foundation penned a letter to the editor lambasting the EPA: “Oelbracht and 35 workers at the plant are threatened by extreme carbon-emission regulations proposed by the U.S. Environmental Protection Agency (EPA). The new regulations could force the closure of Westwood and 13 other plants like it in Pennsylvania, as well as five more in other states...McNelly says the plants have been caught up in EPA’s ‘tsunami of anti-coal’ policy that is directed mainly at much larger coal-fired plants that supply 40 percent of Pennsylvania’s electricity.” This practice whereby think tanks often funded by industry publish articles in local news is one channel through which interest groups seek to manipulate public opinion (Oreskes and Conway 2011; Stokes 2020).

Figure N1: Local News Coverage of Coal Mine and Power Plant Closures in Southwest Pennsylvania



Notes: The top panel shows the number of articles that identify “regulations,” “gas,” or “other” as a reason for the coal industry’s decline over time. The bottom panel shows the overall distribution of the reasons given in general coverage and for specific coverage of coal mine and power plant closures. The author coded news articles from the *Greene County Messenger* and the *Observer-Reporter*.

N.2 National News

To see how the national news covers the coal industry’s decline, I leverage data collected from the Stanford Cable TV News Analyzer. This database has compiled closed-captioned transcripts from newscasts on the major cable TV channels. The temporal coverage starts in 2010, so I can only analyze the post-shale shock period. Still, the data help characterize how the national media discusses the coal industry’s decline.

I explore how CNN and Fox News describe coal in two contexts: environmental regulations and fracking. CNN and Fox News are two of the most-watched cable news shows in the United States. It is well-documented how Fox News has a conservative slant (e.g., DellaVigna and Kaplan 2007), which may contribute to how balanced the network is in its coverage of the Democratic Party’s platform on environmental regulations. Fox News’ coverage might reflect the views of Republican Party elites. Unfortunately, I do not have household cable

news viewership data, so I cannot verify how many individuals in coal country tune into which channels. Nonetheless, conservative residents of Southwest Pennsylvania tended to remark that they have greater trust in Fox News.³⁶

I interpret mentions of coal next to either “regulation” or “EPA” as likely indications that the news media is attributing environmental regulations as the cause of coal’s decline. However, a limitation is that the data simply show the frequency that two words co-occur in the same news segment, which is uninformative of the particular message content. So it would not be definitive to say that the newscast is discussing regulations or fracking as the reason for coal’s decline. To provide more interpretive leverage, I select examples of newscasts for each search term combination for an additional qualitative analysis described in this section.

Figure N2 presents a descriptive analysis of the cable news data. The top plot shows the discussion of these keywords over time. The most frequently discussed topic is coal and regulations or the EPA. In the 2010-2016 period, coal is more frequently discussed in reference to regulations or the EPA. After 2016, mentions of coal and fracking also increase. However, an examination of these news segments indicates that they are about the Democratic presidential candidates’ position on environmental regulations to constrain both coal and fracking. Notably, this is generally not a connection between fracking and the decline of coal.

There are noticeable differences between cable news channels. Fox News was much more likely to discuss coal and regulations than CNN. There is some coverage of fracking on Fox News, especially in the post-2016 period, but as we qualitatively unpack below, most of this discussion deals with whether presidential candidate Clinton or Biden would ban fracking and coal.

The bottom plot of Figure N2 shows the overall coverage of the three coal contexts for three presidential administrations. Fox News spent 36.8 minutes talking about regulations/EPA and coal, compared to CNN, which spent 16.2 minutes on the topic. Fox News Spent 16.7 minutes discussing coal and fracking, whereas CNN spent 4.32. While this level of coverage may seem low, it may be considerable given that much of cable news takes place in small soundbites. More importantly, for my purposes, what matters is the relative discussion of regulations and fracking. The answer is clear: regulations or the EPA are talked about more frequently regarding coal than fracking is. This time also coincides with the shale shock and coal’s decline.

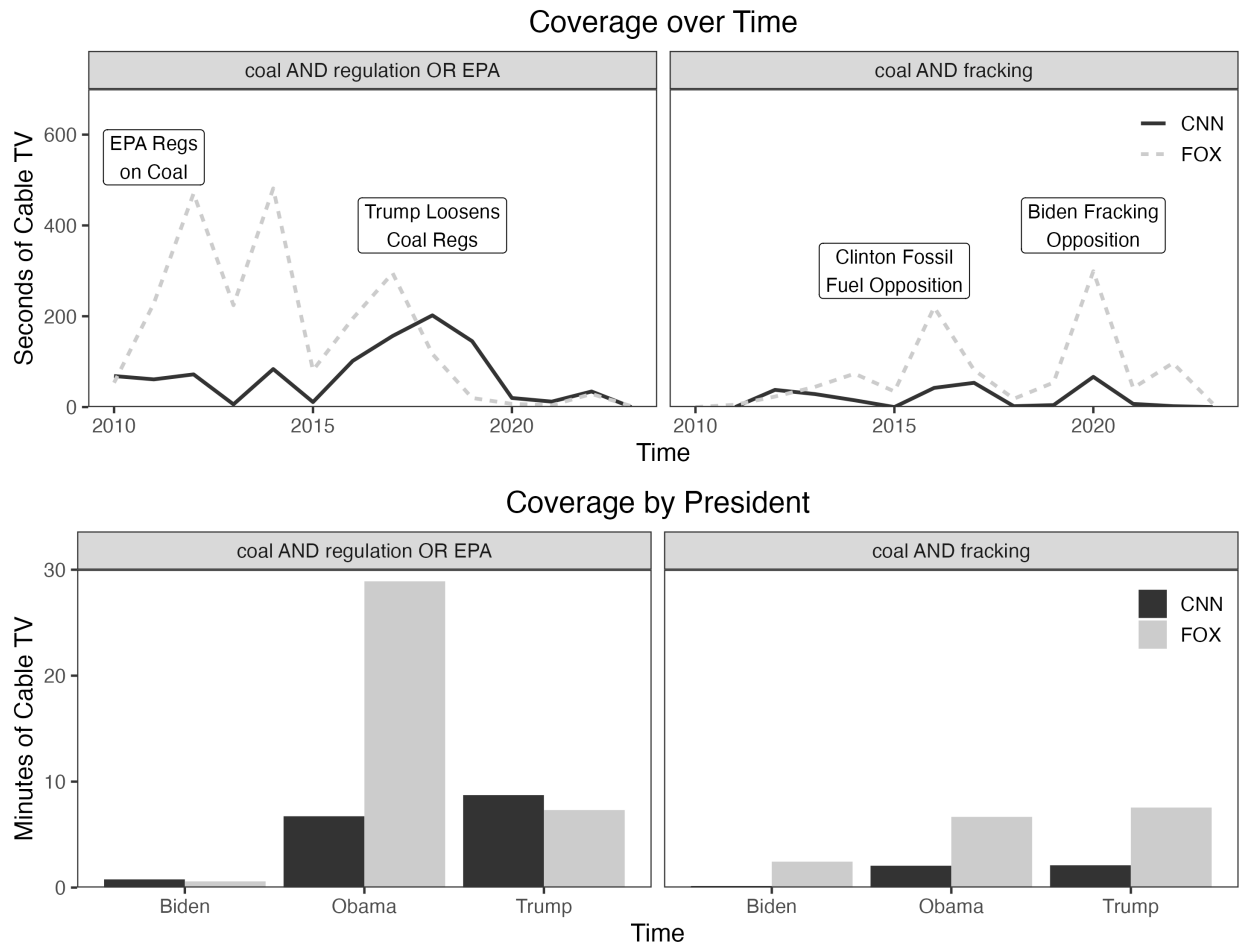
What does the substance of these newscasts look like? I focus on Fox News since this content may influence the partisan reversal in coal country. I select instances where there are peaks in coverage.

- Coal and regulations in January 2014: The newscast in 2014 places blame for the coal industry’s decline at the feet of the Obama administration’s environmental regulations.

Fox News Reporter: ...What action are you taking to stop the EPA from what is nothing less than a heavy-handed approach to destroy the coal industry in America?

³⁶Interview A10

Figure N2: National Cable News Coverage of Coal Industry Decline



Notes: The top plot shows the number of seconds the cable news channels CNN and Fox News spent covering coal in the context of regulations or fracking. The facet labels (e.g., “coal AND regulation OR EPA”) show the search terms used. The labels contextualize the coverage of the topics discussed, which shows when fracking and coal are mentioned simultaneously, it is in reference to the Democratic presidential candidates’ support for environmental regulations that would curtail both coal production and hydraulic fracturing. The bottom plot shows the total number of minutes the channels spent covering coal in the respective contexts during the presidential administrations of Obama, Trump, and Biden. Data from the Stanford Cable TV News Analyzer.

Mitch McConnell: Yeah. Tragic. We have got a depression. Not a recession. A depression in Central Appalachia. We have lost 5,000 out of 18,000 coal mine, coal mining jobs, and in my state, for every coal mining job you lose, you lose three more. A depression created by this administration and the EPA. They’re issuing a regulation that, if it goes through, will guarantee that there will never be another coal-fired plant built in America. That’s 40 percent of our electricity. You know, coal keeps the lights on...

- Coal and fracking in October 2020: The newscast makes clear that the context of the fracking and coal discussion is whether the Biden administration would support these

fuels.

Guest: Joe Biden will not ban fracking. That is a fact, Brian. Be careful with the climate, fracking natural gas, know why are CO2 is down output because of fracking and innovation. In case you think Mike Pence was grabbing things out of the sky and pinning them to Harris and Biden, take a look back to what they said in the past.

Clip: ...states stopping fracking and stopping the pipeline. Yes, new pipeline, exactly. No more, no new fracking. Would there be any place for fossil fuels including coal and fracking in a Biden administration? No. We would work it out. We would make sure it's eliminated and no more subsidies for either one of those.

- Coal and fracking in May 2016: The newscast characterizes presidential candidate Clinton as wanting to put coal and fracking out of business.

News personality: So Hillary wants to put coal miners and coal mining out of business. She doesn't support fracking. She doesn't support drilling. Just if the difference on that one big issue, how many jobs would be created and how good would it be for national security if we move toward energy independence, which Donald Trump has told me repeatedly he wants to do?

O Community-Oriented Preferences

The community-wide consequences of coal’s decline are essential for understanding why people not directly employed in the industry felt they had a stake. First, the industry’s decline undermined revenue for local public goods, which impacted residents without direct employment (DOE 2021; Newell and Raimi 2018).³⁷ The community-wide effects of coal’s decline are well-documented (Carley, Evans, and Konisky 2018; Gazmararian and Tingley 2023; Graff, Carley, and Konisky 2018; Josh Blonz, Brigitte Roth Tran, and Erin E. Troland 2023; Haggerty et al. 2018; Roemer and Haggerty 2021). In related contexts, Broz, Frieden, and Weymouth (2021) show how trade, offshoring, and automation have effects that extend beyond individuals but also hollow out communities. These adverse spillover effects can directly affect an individual’s vote choice through her self-interest.

Second, residents in coal country often have community-oriented preferences that may lead them to see job losses through a collective lens (Gazmararian and Tingley 2023). Sociological studies document how genuine attachment to place and industry rooted in history can lead residents to prioritize community welfare (Bell and York 2010), evidenced by surveys revealing a preference for community-wide redistribution (Gaikwad, Genovese, and Tingley 2022). Likewise, rich qualitative research shows how rural residents make sense of distributive issues through the lens of place-based identity (Cramer 2016).

As evidence of the importance of social ties in fossil fuel communities, I conducted surveys in the field ($N = 606$) at county fairs in Southwest Pennsylvania to evaluate the extent to which coal country residents are embedded in social networks where their friends and family are employed in fossil fuels. The modal respondent believes all ten people she knows best work in fossil fuels (Appendix M.3). The result is consistent with sample weights to enhance representativeness. While this is a possible exaggeration, the answer reflects the perceived social importance of the industry.

³⁷Interviews A1, A11, B24, and B25 raised concerns about the decline of coal and school funding.

P Research Ethics

The study conforms to the APSA Principles and Guidance for Human Subjects Research.

Power I conducted human subjects research exclusively with public officials, business owners, fossil fuel workers, civil society groups, and residents in Pennsylvania. I did not engage with vulnerable populations (e.g., children, prisoners). The questions were not sensitive; the semi-structured interview focused on the subject's views on climate policy and financial assistance from the government.

Consent I obtained voluntary informed consent from all subjects via email or verbally. I transparently communicated my name and affiliation, the general purpose of the research, an explanation of what participation entailed, the potential risks and benefits to participants, how identities and data would be protected, and any other information relevant to the study.

Deception No deception was used.

Harm and trauma No harm or trauma was anticipated or identified.

Confidentiality I clearly communicated assurances of confidentiality during the consent process. I ensured confidentiality by de-identifying responses. Each subject was assigned a number and a letter corresponding with the survey wave. That alpha-numeric combination was used to identify subjects when analyzing data. The file linking the number to individually identifiable subject information was kept in a separate folder on a password-protected computer and a password-protected folder.

Impact No impact on political processes was anticipated or identified.

Laws, regulations, and prospective review The study complied with all relevant laws and regulations. The researchers obtained prospective review by IRB at [[redacted institution]].

References

- Abrajano, M., and Z. L. Hajnal. 2015. *White Backlash*. Princeton University Press.
- Acharya, A., M. Blackwell, and M. Sen. 2016. “Explaining Causal Findings Without Bias: Detecting and Assessing Direct Effects.” *American Political Science Review* 110 (3): 512–529.
- Andrews, D. W. K. 1993. “Tests for Parameter Instability and Structural Change With Unknown Change Point.” *Econometrica* 61 (4): 821.
- Athey, S., et al. 2021. “Matrix Completion Methods for Causal Panel Data Models.” *Journal of the American Statistical Association* 116, no. 536 (2021): 1716–1730.
- Autor, D., D. Dorn, and G. Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103 (6): 2121–2168.
- Bartik, T. 1991. *Who Benefits from State and Local Economic Development Policies?* Upjohn Institute for Employment Research.
- Bell, S. E., and R. York. 2010. “Community Economic Identity: The Coal Industry and Ideology Construction in West Virginia.” *Rural Sociology* 75 (1): 111–143.
- Black, K. J., et al. 2021. “Economic, Environmental, and Health Impacts of the Fracking Boom.” *Annual Review of Resource Economics* 13, no. 1 (2021): 311–334.
- Blackwell, M., and M. P. Olson. 2022. “Reducing Model Misspecification and Bias in the Estimation of Interactions.” *Political Analysis* 30 (4): 495–514.
- Bolet, D., F. Green, and M. Gonzalez-Eguino. 2023. “How to Get Coal Country to Vote for Climate Policy: The Effect of a ‘Just Transition Agreement’ on Spanish Election Results.” Preprint, accessed April 17, 2023. <https://doi.org/10.2139/ssrn.4394195>.
- Brehm, P. 2019. “Natural Gas Prices, Electric Generation Investment, and Greenhouse Gas Emissions.” *Resource and Energy Economics* 58 (2019): 101106.
- Broz, J. L., J. Frieden, and S. Weymouth. 2021. “Populism in Place: The Economic Geography of the Globalization Backlash.” *International Organization*, 1–31.
- Carley, S., T. Evans, and D. Konisky. 2018. “Adaptation, Culture, and the Energy Transition in American Coal Country.” *Energy Research & Social Science* 37:133–139.
- Cinelli, C., and C. Hazlett. 2020. “Making Sense of Sensitivity: Extending Omitted Variable Bias.” *Journal of the Royal Statistical Society Series B* 82 (1): 39–67.
- Clark, R., R. Khoban, and N. Zucker. 2022. “Breadwinner Backlash: The Gendered Political Consequences of Industrial Decline.” Unpublished Manuscript.
- Coglianesi, J., T. Gerarden, and J. Stock. 2020. “The Effects of Fuel Prices, Environmental Regulations, and Other Factors on US Coal Production, 2008-2016.” *Energy Journal*.
- Cramer, K. J. 2016. *The Politics of Resentment*. University of Chicago Press.

- Davis, R. J., J. S. Holladay, and C. Sims. 2021. "Coal-Fired Power Plant Retirements in the U.S." Preprint, Working Paper. Accessed April 2, 2023. <https://doi.org/10.3386/w28949>. National Bureau of Economic Research: 28949.
- DellaVigna, S., and E. Kaplan. 2007. "The Fox News Effect: Media Bias and Voting*." *The Quarterly Journal of Economics* 122, no. 3 (2007): 1187–1234.
- DOE. 2017. *Staff Report to the Secretary on Electricity Markets and Reliability*.
- . 2021. *Initial Report to the President on Empowering Workers Through Revitalizing Energy Communities*.
- Dunlap, R. E., C. Xiao, and A. M. McCright. 2001. "Politics and Environment in America: Partisan and Ideological Cleavages in Public Support for Enviro." *Environmental Politics* 10 (4): 23–48.
- Eckert, F., et al. 2020. *Imputing Missing Values in the US Census Bureau's County Business Patterns*. w26632. National Bureau of Economic Research.
- Frank, T. 2007. *What's the Matter with Kansas?* Picador.
- Gaikwad, N., F. Genovese, and D. Tingley. 2022. "Creating Climate Coalitions: Mass Preferences for Compensating Vulnerability in the World's Two Largest Democracies." *American Political Science Review* 116 (4): 1165–1183.
- Gazmararian, A. F., and D. Tingley. 2023. *Uncertain Futures: How to Unlock the Climate Impasse*. Cambridge University Press.
- Graff, M., S. Carley, and D. M. Konisky. 2018. "Stakeholder Perceptions of the United States Energy Transition: Local-level Dynamics and Community Responses to National Politics and Policy." *Energy Research & Social Science*, Sustainable Energy Transformations in an Age of Populism, Post-Truth Politics, and Local Resistance, 43 (2018): 144–157.
- Haggerty, J. H., et al. 2018. "Planning for the Local Impacts of Coal Facility Closure: Emerging Strategies in the U.S. West." *Resources Policy* 57:69–80.
- Hainmueller, J., J. Mummolo, and Y. Xu. 2019. "How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice." *Political Analysis* 27 (2): 163–192.
- Hartman, E., and F. D. Hidalgo. 2018. "An Equivalence Approach to Balance and Placebo Tests." *American Journal of Political Science* 62 (4): 1000–1013.
- Hill, S., D. Hopkins, and G. Huber. 2021. "Not by Turnout Alone: Measuring the Sources of Electoral Change, 2012 to 2016." *Science Advances* 7 (17): eabe3272.
- Hirsch, B. T., and D. A. MacPherson. 2003. "Union Membership and Coverage Database from the Current Population Survey: Note." *ILR Review* 56 (2): 349–354.
- Ho, D. E., et al. 2007. "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15 (3): 199–236.

- Holladay, S., and J. LaRiviere. 2017. “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.
- Hopkins, D. J. 2010. “Politicized Places: Explaining Where and When Immigrants Provoke Local Opposition.” *American Political Science Review* 104 (1): 40–60.
- Imai, K., and M. Ratkovic. 2014. “Covariate Balancing Propensity Score.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76 (1): 243–263.
- Jardina, A. 2019. *White Identity Politics*. Cambridge University Press.
- Josh Blonz, Brigitte Roth Tran, and Erin E. Troland. 2023. *The Canary in the Coal Decline: Appalachian Household Finance and the Transition from Fossil Fuels*. NBER.
- Joskow, P. L. 2013. “Natural Gas: From Shortages to Abundance in the United States.” *American Economic Review* 103, no. 3 (2013): 338–343.
- Knittel, C., K. Metaxoglou, and A. Trindade. 2016. “Are We Fracked? The Impact of Falling Gas Prices and the Implications for Coal-to-Gas Switching and Carbon Emissions.” *Oxford Review of Economic Policy* 32 (2): 241–259.
- Lawhorn, J. 2022. *The POWER Initiative*. Congressional Research Service, 2022.
- Leip, D. 2020. *Dave Leip’s Atlas of U.S. Presidential Elections*.
- Linn, J., and K. McCormack. 2019. “The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions.” *The RAND Journal of Economics* 50 (4): 733–767.
- Linn, J., and L. Muehlenbachs. 2018. “The Heterogeneous Impacts of Low Natural Gas Prices on Consumers and the Environment.” *Journal of Environmental Economics and Management* 89 (2018): 1–28.
- Liu, L., Y. Wang, and Y. Xu. 2022. “A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data.” *American Journal of Political Science* Forthcoming.
- McCright, A., and R. Dunlap. 2011. “The Politicization of Climate Change and Polarization in the American Public’s Views of Global Warming, 2001–2010.” *The Sociological Quarterly* 52 (2): 155–194.
- Mian, A., and A. Sufi. 2011. “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis.” *American Economic Review* 101, no. 5 (2011): 2132–2156.
- Mildenberger, M. 2020. *Carbon Captured: How Business and Labor Control Climate Politics*. MIT Press.
- Mohlin, K., et al. 2019. “Turning the Corner on US Power Sector CO2 Emissions—a 1990–2015 State Level Analysis.” *Environmental Research Letters* 14 (8): 084049.

- Moore, M. 2018. "Sunbury Generation Plant Will Be Torn down over the next Two Years." The Daily Item, 2018. Accessed September 6, 2023. https://www.dailyitem.com/news/snyder_county/sunbury-generation-plant-will-be-torn-down-over-the-next-two-years/article_3ccc8b32-b38c-5897-8476-fc45c2276583.html.
- Morgan, S., and J. Lee. 2017. "The White Working Class and Voter Turnout in U.S. Presidential Elections, 2004 to 2016." *Sociological Science* 4:656–685.
- Morgan, S. L. 2018. "Status Threat, Material Interests, and the 2016 Presidential Vote." *Socius: Sociological Research for a Dynamic World* 4 (2018): 237802311878821.
- Mummolo, J., and E. Peterson. 2018. "Improving the Interpretation of Fixed Effects Regression Results." *Political Science Research and Methods* 6 (4): 829–835.
- Mutz, D. C. 2018. "Status Threat, Not Economic Hardship, Explains the 2016 Presidential Vote." *Proceedings of the National Academy of Sciences* 115 (19): E4330–E4339.
- Napsha, J. 2019. "Tenaska Plant near Smithton Generating Power, Money." TribLIVE.com, 2019, 6:45 p.m. (-04:00). Accessed September 6, 2023. <https://triblive.com/local/westmoreland/tenaska-plant-near-smithton-generating-power-money/>.
- Newell, R. G., and D. Raimi. 2018. "The Fiscal Impacts of Increased U.S. Oil and Gas Development on Local Governments." *Energy Policy* 117:14–24.
- Oreskes, N., and E. M. Conway. 2011. *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. Bloomsbury Press.
- Roberts, M. E., et al. 2014. "Structural Topic Models for Open-Ended Survey Responses." *American Journal of Political Science* 58 (4): 1064–1082.
- Roemer, K. F., and J. H. Haggerty. 2021. "Coal Communities and the U.S. Energy Transition: A Policy Corridors Assessment." *Energy Policy* 151:112112.
- Sides, J., M. Tesler, and L. Vavreck. 2018. *Identity Crisis*. Princeton University Press.
- Stokes, L. 2020. *Short Circuiting Policy: Interest Groups and the Battle Over Clean Energy and Climate Policy in the American States*. Oxford University Press.
- Watson, B., I. Lange, and J. Linn. 2023. "Coal Demand, Market Forces, and U.S. Coal Mine Closures." *Economic Inquiry* 61 (1): 35–57.