Employment Impacts of the ARC's POWER Initiative Grants in Coal Communities

Emiri Morita

Junior Independent Work

Advised by Orley Ashenfelter

April 21, 2021

This paper represents my own work in accordance with University Regulations.

Abstract

Has the economic stimulus funded through the 2015 Partnerships for Opportunity and Workforce and Economic Revitalization (POWER) Initiative led to increased employment in the previously coal-dependent Appalachian region? In this paper, I examine the employment impacts of this project-based grant initiative aimed to boost local economic development and job creation by looking at county-level unemployment data between 2010 and 2019 from the Bureau of Labor Statistics' Local Area Unemployment Statistics (LAUS) combined with grant roll-out information provided by the Appalachian Regional Commission (ARC). Comparing grant-recipient counties to non-grant-recipient counties in the ARC region with similar socio-demographic characteristics, I perform an event-study analysis in addition to a two-way fixed-effects difference-in-difference analysis on matched unemployment data to quantify the policy impact. The analyses reveal that overall unemployment rates were on a downward trajectory in the ARC region even before the POWER Initiative grants were rolled out, and no clear systematic or isolated impact could be identified for the grant policy. However, the grant policy may have effectively targeted a subset of counties within the ARC region that were experiencing slower employment growth relative to others.

I. Introduction

Ever since the coal bust in the 1980s and the subsequent energy transition, communities in the coal industry-dependent Appalachian Region of the United States have struggled to regain economic prosperity (Black et al., 2005). With U.S. coal production forecasted to decline further through 2050, with the steepest decline to occur by the mid-2020s (U.S. Energy Information Administration, 2020), initiatives to manage negative economic effects on the region have become a top priority for the U.S. government. The Partnerships for Opportunity and Workforce and Economic Revitalization (POWER) Initiative, launched in 2015 under the Obama Administration, has provided grants through the Appalachian Regional Commission (ARC) to projects that aim to support local economic development, job creation, and training programs for displaced workers (Cecire, 2019). While there is existing literature on the overall economic and labor force impacts of the coal bust and boom in the Appalachian Region, the employment impacts of this particular grant initiative has not been analyzed in a quantitative manner.

In this study, I aim to analyze the impacts of the ARC's POWER Initiative on county-level unemployment rates of grant-recipient counties. For the outcome variable of interest, I use annual county-level Local Area Unemployment Statistics between 2010 and 2019 from the Bureau of Labor Statistics. I also use 5-year average county-level socio-demographic variables from the American Community Survey as control variables.

My empirical analysis consists of an event study and a difference-in-difference analysis, both performed after matching counties that received a grant ("treated county") to a county that did not receive a grant but had similar characteristics before the allocation of the grant ("control county"). Using information provided by the ARC about the rollout of the grants and socio-demographic covariates, I create two datasets using propensity score matching. I create the first dataset using many-to-one matching to include all 334 grant-recipient counties in the treatment group and matching them to 63 non-grant recipient counties as their counterpart. I create the second, more balanced dataset, using one-to-one matching limiting to 86 grant-recipient counties that are socio-demographically similar counties in the control pool.

For these two datasets separately, I first perform an event-study using time periods relative to the grant rollout year in each county to visually check for parallel trends in the pre-treatment period, in addition to any visible treatment effects. Then, I perform a two-way fixed-effect difference-in-difference analysis, controlling for year, county and pair fixed effects to identify whether or not the grant had any impact on unemployment rate.

The analysis on the many-to-one matched data reveals the possibility of selection bias in the allocation of the grants to counties that were already on a slower improvement rate in terms of unemployment. However, because of this possible selection bias and lack of parallel trends in the pre-treatment period, no clear policy effect could be identified. In the one-to-one matched data case, there are clear parallel trends in the trajectory of unemployment rates in the pre-treatment period, followed by some deviation upon the rollout of the grant. However, again, no statistically significant treatment effect is identified. One possible interpretation given these observations is that the grants successfully targeted counties experiencing slow employment growth, but that it is simply too early to see results. The study calls for a more robust impact analysis in the future when all grants have been closed and heterogenous effects can be studied amongst grant-recipient counties based on the larger total grant amount funded.

The remainder of this paper is organized as follows. Section II provides institutional background on the POWER Initiative. Section III reviews the literature on related topics. Section IV presents the data used in this study followed by Section V on methodologies. Section VI lays out the analysis results. Section VII concludes and summarizes the findings.

II. Institutional Background of the POWER Initiative

The Partnerships for Opportunity and Workforce and Economic Revitalization (POWER) Initiative, launched in 2015 by the Obama Administration as a multi-agency effort, was aimed to ease the economic effects of energy transition in coal industry-dependent communities in the

United States. While certain aspects of the larger POWER Plus Plan were never enacted or funded, elements of the POWER Initiative continue under the Trump Administration, mainly the Assistance to Coal Communities program within the Economic Development Administration, the POWER Initiative under the Appalachian Regional Commission[1] (the only program to retain the original branding), and a funding program for abandoned mine land reclamation. The ARC's POWER Initiative in particular, which my study aims to analyze, is the largest of the economic development programs, having funded nearly $238 million in projects since its first launch in FY2016. It has reportedly leveraged $1.1 billion in private investment into the Appalachian regional economy and helped create or retain more than 26,000 jobs. It has also reportedly prepared thousands of workers and students for jobs in entrepreneurship, broadband development, tourism, and other industry sectors (*POWER Award Summaries by State*, 2020).

As part of the POWER Initiative, the ARC funds three classes of grants: (1) implementation grants with awards of up to $1.5 million; (2) technical assistance grants with awards of up to $50,000; and (3) broadband deployment projects with awards of up to $2.5 million. State and local agencies and governmental entities, local governing boards, nonprofit organizations, Indian tribes and higher education institutions are eligible to apply for the ARC's POWER grants (Appalachian Regional Commission, 2019).

POWER investments are subject to the ARC's grant match requirements which are linked to the commission's economic distress hierarchy (*Classifying Economic Distress in Appalachian Counties*, 2020). The economic distress designations are as follows:

1. distressed (80% funding allowance, 20% grant match);

[1] The Appalachian Regional Commission is a federal-state partnership established in 1965 to address economic distress in the Appalachian Region, and its jurisdiction spans 420 counties in Alabama, Georgia, Kentucky, Ohio, New York, Maryland, Mississippi, North Carolina, Pennsylvania, South Carolina, Tennessee, Virginia, and West Virginia.

2. at-risk (70%);
3. transitional (50%);
4. competitive (30%); and
5. attainment (0% funding allowance).

POWER investments are also aligned with the ARC's strategic plan which prioritizes five investment goals ("Investing in Appalachia's Future," 2015):

1. entrepreneurial and business development;
2. workforce development;
3. infrastructure development;
4. natural and cultural assets; and leadership and community capacity.

The ARC has designated an annual $50 million for POWER activities, and over $238 million in investments have been made since FY2016 through 293 projects in 353 of its 420 counties. **Figure 1** represents the total number of ARC's POWER Initiative projects by state as of February 2020 (the most recent cohort of projects announced in 2020 are not included).

Figure 1. Distribution of ARC POWER Initiative Projects

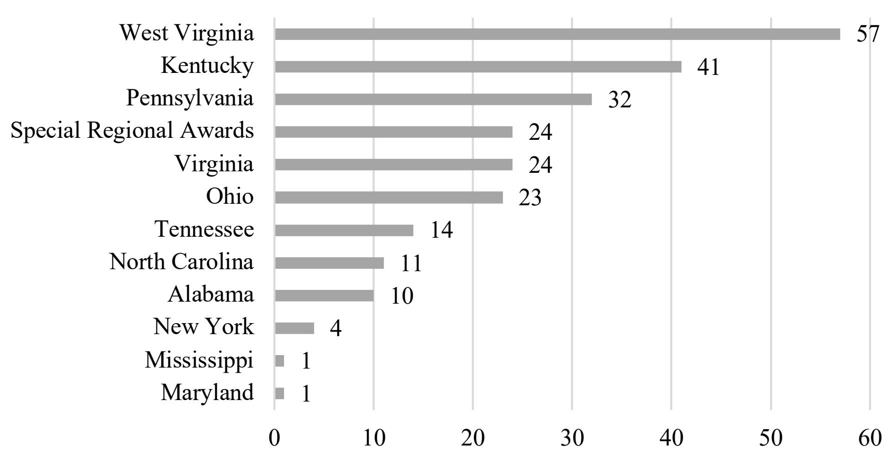

** Special Regional Projects are projects that are undertaken to provide regional benefits and may include collaboration with organizations outside of the ARC's service area.*

III. Literature Review

There have been numerous studies looking at the economic effects of the coal boom and bust in the United States and the Appalachian region. For example, Black et al. (2005) examines the local labor market impacts of the 1970s coal boom and 1980s coal bust in Kentucky, Ohio, Pennsylvania and West Virginia. They find that the spillover effects to non-coal industries were larger for the bust than the boom—where for each 10 new coal sector jobs created during the boom, less than two new jobs were created in the local goods sectors during the boom and 3.5 jobs were lost during the bust (Black et al., 2005). While this study in its conclusion suggests that attracting industrial employment may help the area's economy albeit with modest spillover effects to other sectors, they stop short of empirically examining the real effects of economic stimulus to these regions in the post-bust period. My project thus aims to empirically analyze whether the POWER Initiative, an economic stimulus effort in the Appalachian region in the post-bust period, has had any real impacts on the region's economy.

Betz et al. (2015), on the other hand, assesses the winners and losers of coal development and whether the natural resources curse—the phenomenon that long-run growth rates are lower over the boom-bust cycle in resource intense locations—applied to contemporary American coal communities. While they find no strong evidence of a resources curse, they find that coal mining had a consistent inverse association with measures linked to population growth and entrepreneurship, and thereby future economic growth (Betz et al., 2015). Since the ARC's POWER Initiative has strategically targeted projects that aim to foster business and entrepreneurship in Appalachia's traditionally coal dependent regions, my project will look at whether the grants have been effective at improving these economic indicators despite Betz et al.'s findings. Betz et al. also find that during both the boom and bust periods, ARC counties tended to fare worse economically relative to other U.S. counties, and the benefits engendered from the coal industry to lower and middle-income households in other regions of the U.S. did not appear to hold in the ARC region. Thus, I also aim to examine whether the POWER Initiative, which aims to strengthen the region's economy away from its dependency on coal, has had any positive distributional effects to lower and middle-income households.

As for similar economic development efforts in the Appalachian region, Kline and Moretti (2014) studies the impact of the infrastructural development investments made through the Tennessee Valley Authority (TVA) to the Tennessee Valley region from 1934 to 1958. They find that while gains in agricultural employment were lost after the funding ended, gains in manufacturing employment continued to intensify, suggesting the presence of agglomeration economies in manufacturing (Kline & Moretti, 2014). While these findings may be insightful in forecasting the impact of other economic development initiatives in the region, the nature of the investments made through the TVA were notably different from the ARC's POWER Initiative—TVA investments funded large public infrastructure projects like hydroelectric dams, navigation canals, and extensive road networks, whereas ARC's POWER Initiative targets a mix of projects, some infrastructural like the expansion of broadband access, but also projects that are more community-based like curriculum development for local community colleges, support for career training programs, and entrepreneurial enterprises. Thus, my project will add to the existing literature by analyzing how these more community-based grants will impact employment outcomes in the economically distressed Appalachian region.

As for studies on the ARC's POWER Initiative itself, Chamberlin and Dunn's (2019) early impact analysis studies the effect of investments made up to FY2018. Through a survey of 88 selected grantees using questionnaires and in-person interviews, they identify common longterm visions of the projects, setbacks projects faced within their communities and recommendations for the ARC to improve the program. However, they do not use economic indicators to quantitatively analyze the impacts of the POWER Initiative, mainly due to the fact that the initiative was still in its early stages. Thus, my project aims to fill this gap and analyze the real regional economic impacts of the initiative.

IV. Data

a) Identifying Control and Treatment Groups

To identify the appropriate counties to be used as the control and treatment groups in this study, the initial hope was to use counties that applied but never received the grants as the control group and counties that received the grant as the treatment group. However, after inquiring the ARC's Director of Research & Evaluation with my data request, it became clear that information on previous applicants cannot be disclosed. Thus, to identify counties appropriate for the control group, I use socio-demographic data in a propensity score matching process to identify counties with similar characteristics to the treatment group.

b) Unemployment Data

For my unemployment data, I use the Bureau of Labor Statistics' Local Area Unemployment Statistics (LAUS) from 2010 to 2019, which has county-level unemployment rate estimates at monthly and annual frequencies. While this data is based on predictive modelling, there is data available for every county in the ARC region. **Tables 1** and **2** below show the summary statistics of the average annual county-level unemployment rates for each ARC state (including non-ARC counties) (U.S. Bureau of Labor Statistics, 2020).

Table 1. 2015 County Unemployment Rate Averages % by State (LAUS)

state	N	mean	sd	min	max
AL	67	7.0	.018	.042	.146
KY	120	6.5	.022	.035	.158
MD	24	5.7	.016	.038	.106
MS	82	7.6	.022	.041	.169
NC	100	6.4	.014	.043	.119
NY	62	5.5	.009	.040	.078
OH	88	5.4	.012	.033	.100
PA	67	5.6	.009	.038	.079
TN	95	6.6	.013	.040	.097
VA	133	5.1	.013	.028	.109
WV	55	7.7	.021	.041	.132

Table 2. 2019 County Unemployment Rate Averages % by State (LAUS)

state	N	mean	sd	min	max
AL	67	3.4	.009	.022	.071
KY	120	5.1	.014	.031	.110
MD	24	4.0	.011	.027	.074
MS	82	6.4	.018	.039	.155
NC	100	4.3	.009	.030	.086
NY	62	4.3	.006	.032	.060
OH	88	4.5	.011	.026	.083
PA	67	4.7	.008	.068	.068
TN	95	4.0	.008	.024	.060
VA	133	3.2	.008	.019	.056
WV	55	5.8	.017	.031	.130

Additionally, **Table 3** and **4** below show the average county-level unemployment rates filtered for ARC counties for the pre- and post-treatment periods amongst the control and treatment groups identified by propensity score matching. **Table 3** is the output from a many-to-one (m:1) matching approach with disproportionately more datapoints in the treatment group relative to the control group, while **Table 4** is the output from a balanced one-to-one (1:1) matching approach with an even number of counties in the treatment and control groups. Overall, unemployment rates in the region can be seen to be declining from the pre- to post-treatment periods. **Tables 3** and **4** suggest treated counties had higher average unemployment rates in both periods relative to the control counties

Table 3. County-level Unemployment Rate % for Control vs Treatment Groups (m:1)

Post	Treat	N	mean	sd	min	max
0	0	421	8.50	2.48	2.5	16.4
0	1	2,236	8.57	2.57	2.7	20.5
1	0	209	4.62	1.10	2.4	8.8
1	1	1,104	5.30	1.74	2.2	19.9

Table 4. County-level Unemployment Rate % for Control vs Treatment Groups (1:1)

Post	Treat	N	mean	sd	min	max
0	0	574	8.58	2.54	2.8	16.4
0	1	574	8.81	3.08	2.8	20.5
1	0	286	4.59	1.14	2.4	8.8
1	1	286	4.98	2.04	2.2	19.9

For my demographic control variables, I use the following county-level 5-year averages from the American Community Survey of 2010 to 2019—average household size, educational attainment (percent of population with a bachelor's degree or higher), veteran status (percent of veterans in civilian population 18 years or over), foreign nationals (percent of total population who are foreign born), language spoken (percent of population who speak only English at home), population size, sex (percent of population who are male), median age, and race (percent of population who are white). I use the 5-year averages since the 1-year averages does not provide data for all of the counties of interest (United States Census Bureau, 2020).

c) Controlling for Additional Statewide Grants

To better isolate the effect of the POWER grants on employment, I would have ideally controlled for the effects of other similar grant programs in the region. However, upon inquiring the ARC about other similar grant initiatives in the region, I was notified that there are too many similar economic support programs and funds to be precisely identified and controlled for at the county or state levels in the ARC region. Thus, for future reference, I will point to three other similar grants to the POWER Initiative that would ideally be isolated.

One is the EDA's Assistance to Coal Communities (ACC) Program which has awarded a total of $115 million to EDA-defined economic distressed areas. Second is the Abandoned Mine Land (AML) Reclamation Investments funded a total of $315 million into the Appalachian states with the greatest unfunded needs. The last is the ARC's Investments Support Partnerships in Recovery Ecosystems (INSPIRE) Initiative, which is a relatively recent $10 million initiative to address substance abuse crisis in the region.

V. Methodology

a) Propensity Score Matching

First, I use propensity score matching to identify the appropriate counties within the ARC region that did not receive a grant to use as the control group. By using this approach, I attempt to compare treated counties to counties with similar socio-demographic characteristics and thus similar probabilities of having received the grant. This is an alternative method to comparing grant-recipient counties to counties that applied and did not receive the grant, as this data was not available.

The propensity score matching method used in this study estimates a maximum likelihood model based on a logit function of the conditional probability of treatment given a list of socio-demographic variables, and uses the predicted values from that estimation to collapse those covariates into a single scalar called the propensity score. Because I want to capture county-level characteristics through this method, I focus on largely time-invariant covariates in year 2014 (the year before any POWER grant was rolled out) to estimate the propensity score, excluding unemployment rate. Specifically, I use county-level average household size, educational attainment, veteran status, percent of population who are foreign nationals, language spoken, population size, sex, age, and race as the covariates to estimate the propensity scores.

I produce two matched datasets. First, I use a many-to-one (m:1) matching approach where many grant-recipient counties with similar propensity scores are matched to one non-grant recipient counterpart. Since there is only a small number of counties in the ARC region that have not received a POWER Initiative grant, thus limiting the pool for the control group, this many-to-one matching approach is one way of including all treated counties in the sample for analysis. From the ARC-provided data on the roll-out of the grant, 334 counties within the ARC region are identified to have received the grant. In this many-to-one matching, 334 grant-recipient counties are matched to 63 non-grant recipient counties in the control group.

Separately, I also produce a dataset using a balanced one-to-one (1:1) matching approach with an even number of counties in the treatment and control groups. This produces a dataset that compares across the most similar counties socio-demographically amongst the grant-recipient and non-grant recipient counties. Of the 334 grant-recipient counties, 86 are matched in a one-to-one matching process to a control county with a propensity score in a nearby range. Thus, after the one-to-one propensity score matching process, I end up with 86 counties in the treatment group and 86 counties in the control group.

Characterizing the 86 treated counties that got matched in the one-to-one dataset, the regression output in **Table 5** shows the existence of selection bias in the matching process. Using 2014 data of all grant-recipient counties, I regress each of the socio-demographic variables used in the matching process on a dummy variable "Matched" which equals one if the county was matched and included in the one-to-one dataset.

Table 5. Selection Bias in the m:1 Matching Process

Variables	(1) FamilySize	(2) CollegeDegree	(3) Veteran	(4) Foreign	(5) Language
Matched	0.166***	1.546	-1.114***	1.847***	-1.865***
	(0.0191)	(0.961)	(0.249)	(0.301)	(0.451)
Constant	2.955***	15.76***	10.09***	1.426***	96.89***
	(0.00906)	(0.392)	(0.115)	(0.0781)	(0.223)
Observations	334	334	334	334	334
R-squared	0.200	0.010	0.064	0.180	0.050

	(6) Population	(7) Gender	(8) Age	(9) Race
Matched				
	-7,638	-0.740***	-2.329***	-9.621***
Constant	(8,842)	(0.178)	(0.468)	(1.652)
	61,049***	49.88***	42.77***	95.85***
	(6,750)	(0.138)	(0.203)	(0.306)
Observations				
R-squared	334	334	334	334
Variables	0.001	0.027	0.083	0.191

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Relative to the 248 treated counties excluded from the dataset, the 86 treated counties included have on average a larger family size, lower percentage of veterans, higher percentage of foreign nationals, lower percentage of English-only-speaking households, lower percentage of males, lower percentage of whites and tend to be younger in age. Thus, the counties included in the one-to-one matched dataset are systematically different from the rest of the treated counties.

Additionally, I check and confirm that the one-to-one matching process produces treatment and control groups that are systematically no different socio-demographically from each other preceding the treatment. Using 2014 data for all counties included in the one-to-one matched dataset (86 treated and 86 control), I regress each of the socio-demographic variables used in the matching process on a dummy variable "Treat" which equals one if the county received a grant.

Table 6. Treatment vs Control Group Characteristics (1:1)

Variables	(1) FamilySize	(2) CollegeDegree	(3) Veteran	(4) Foreign	(5) Language
Treat	0.00326	0.122	-0.0360	-0.321	0.823
	(0.0274)	(1.171)	(0.297)	(0.525)	(0.752)
Constant	3.118***	17.19***	9.007***	3.594***	94.20***
	(0.0216)	(0.772)	(0.198)	(0.437)	(0.641)
Observations	172	172	172	172	172
R-squared	0.000	0.000	0.000	0.002	0.007

	(6) Population	(7) Gender	(8) Age	(9) Race
Treat	-11,184	0.0488	0.278	1.355
	(13,135)	(0.159)	(0.615)	(2.361)
Constant	64,596***	49.09***	40.17***	84.87***
	(11,821)	(0.111)	(0.446)	(1.710)
Observations	172	172	172	172
R-squared	0.004	0.001	0.001	0.002

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 6 shows that none of the coefficients on the socio-demographic variables are statistically significant, confirming that there is no systematic difference in the socio-demographic characteristics between the treatment and control groups identified by one-to-one matching.

b) Event-Study Analysis

Proceeding separately for the two different matched datasets explained above, I start the analysis by performing an event-study to visually check for parallel trends in the county-level unemployment over time before the rollout of the grants between the control and treatment groups, in addition to changes to the pattern after the rollout of the grants. Below is the specification for the event-study analysis:

$$(1)\ unemployment_{c,t} = \sum_{\tau=-10,\tau\neq-1}^{4} \beta_\tau * \left(I_{c,t=\tau} * treatment_c\right) + \delta_c + \gamma_t + \rho_p + \varphi * X_{c,t} + u_{c,t}$$

$I_{c,t=\tau}$ is a binary variable for each year within a 10-year window before and 4-year window after a grant, other than the year that precedes the funding of the grant (τ = -1) and $treatment_c$ is a dummy variable indicating treatment (i.e. =1 if the county received the POWER grant). If multiple grants were awarded to the same county across multiple years, I use the first of such years as the treatment year for the county. For the control group, I use the treatment year of its matched counterpart in the treatment group as its corresponding treatment year.

I additionally included δ_c to control for county fixed effects (all time-invariant county-level characteristics), γ_t to control for time fixed effects (such as common shocks to the U.S. economy), ρ_p to control for pair fixed effects (using pair IDs identified in the propensity score matching process) and $X_{c,t}$, a vector of county-level characteristics that vary over time and that may affect employment.

c) Difference-in-Difference Analysis

Then, for my main analysis, I use a matched two-way fixed-effect difference-in-difference model with annual county-level unemployment as the outcome variable, given by the specification below:

$$(2)\ unemployment_{c,t} = \beta_1 * \left(post_{c,t} * treatment_c\right) + \kappa_c + \omega_t + \sigma_p + \phi * X_{c,t} + \epsilon_{c,t}$$

The main independent variable of interest is an interaction term between $post_{c,t}$, an indicator variable that equals 1 if the time period is after the treatment year, and $treatment_c$, a dummy variable indicating treatment. I also included κ_c to control for county fixed effects, $\omega_t$to control for time fixed effects, σ_p to control for pair fixed effects, and $X_{c,t}$, a vector of county-level characteristics that vary over time and that may affect employment. I compare the outputs of this model to that of a standard difference-in-difference model (replacing the fixed effects with a post dummy and a treatment dummy), as well as specifications including only some of the fixed effects.

VI. Analysis

a) Many-to-One (m:1) Matched Dataset

Before getting into the analysis methods outlined above, using the many-to-one matched dataset—which includes all 334 grant-recipient counties but a disproportionately small number of non-grant recipient counties—I visualize the pre- and post-treatment trends amongst the treatment and control groups by simply plotting the average unemployment rates in each year relative to treatment year as shown below in **Figure 2**. First and foremost, average unemployment rates fall drastically in both the treatment and control counties throughout the 15-year time frame presented. Looking in more detail, while visually very similar in trends, there appears to be some selection bias in the allocation of the grants and the parallel trends assumption does not appear to quite hold between the treatment and control groups preceding the treatment. More precisely, counties in the treatment group which received the grant were on a slower improvement trajectory in terms of unemployment rates compared to counties in the control group even before the treatment.

Figure 2. Average Unemployment Rates Using Matched Data (m:1)

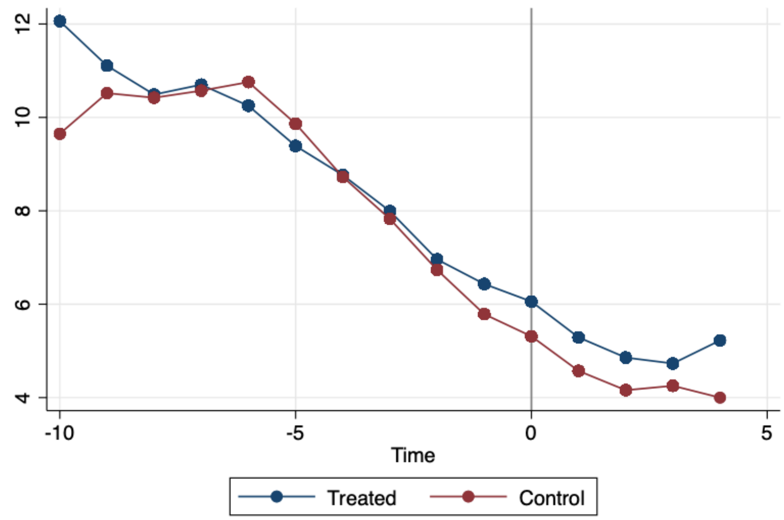

I note here that the trends between the treated and control counties appear to diverge from t=-2 onwards, suggesting possible anticipation about the grant allocation before the treatment period. This observation is addressed in **Figure 6** of the Appendix with regards to the appropriate choice of the omitted year in the following event-study analysis.

Now, running the event-study analysis using the many-to-one matched data (matched on county-level average household size, educational attainment, veteran status, percent of population who are foreign nationals, language spoken, population size, sex, age, and race), I confirm the above notion of selection bias and lack of parallel trends. The event studies plot in **Figure 3** absorbs time, county and pair fixed effects, in addition to controlling for socio-economic characteristics (using the same list of covariates as the matching process). This analysis makes the selection bias for treatment even clearer. The plot now shows that surprisingly, counties that received the grant were faring better in terms of unemployment rate compared to the control group in the pre-treatment period, but those unemployment rates were getting progressively worse relative to the control group as time moved forward.

Figure 3. Event Study (m:1)

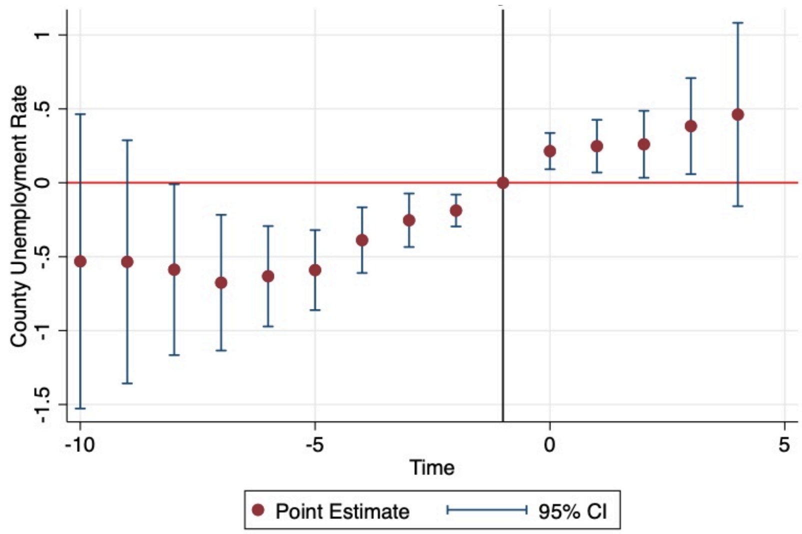

This figure combined with the previous average unemployment visualization in **Figure 2** suggests that while unemployment rates were falling in both the control and treatment counties throughout all time periods with the treatment group starting from a lower unemployment rate at the onset, this gradual fall was slower in the treatment group both before and after the treatment, such that there was a crossover point when unemployment rates in the treated counties became higher than that of the control group. However, carefully studying **Figures 2** and **3**, there appears to be a possibility that this stagnant fall in unemployment rates amongst the treatment counties may have been improved or accelerated slightly after the treatment. Namely, the slope of the datapoints in **Figure 3** flatten out slightly after time period 0.

Lastly, the two-way fixed-effect difference-in-difference analysis on the many-to-one matched data gives the below result. I first run a standard difference-in-difference model replacing the fixed effects for a treat dummy and post dummy as a reference case (labeled here as the Base Case). Then, I run the regression model (2) outlined the methodology section, first without any of the fixed effects (labeled here as the Base Case), then absorbing just the time fixed effects, adding the county fixed effects, and finally adding pair fixed effects from the matching process. Standard errors are clustered at the county level for all specifications (to allow for serial correlation over time within counties).

Table 7. Difference-in-difference Regression Output (m:1)

Dependent variable: Average county-level unemployment rate (%)				
	(1)	(2)	(3)	(4)
Specifications	Base Case	Time FE	Time & County FE	Time, County & Pair FE
Treat	-0.260			
	(0.202)			
Post	-2.605***			
	(0.222)			
TreatPost	0.803***	0.366***	0.385***	0.385***
	(0.191)	(0.133)	(0.106)	(0.107)
FamilySize	0.359	0.213	-0.160	-0.160
	(0.516)	(0.500)	(0.313)	(0.315)
CollegeDegree	-0.0217***	-0.117***	0.0390***	0.0390***
	(0.00393)	(0.0100)	(0.0133)	(0.0134)
Veteran	9.34e-05**	0.000122***	5.26e-05	5.26e-05
	(4.04e-05)	(2.80e-05)	(4.14e-05)	(4.17e-05)
Foreign	-0.272***	-0.100***	-0.00366	-0.00366
	(0.0404)	(0.0344)	(0.0373)	(0.0376)
Language	-7.70e-05**	-0.000107***	-4.38e-05	-4.38e-05
	(3.29e-05)	(2.37e-05)	(3.48e-05)	(3.50e-05)
Population	-1.77e-06**	-1.00e-06	-1.08e-06	-1.08e-06
	(8.05e-07)	(7.19e-07)	(1.18e-05)	(1.19e-05)
Gender	0.0338	0.00334	0.108***	0.108***
	(0.0307)	(0.0271)	(0.0374)	(0.0377)
Age	2.33e-05	0.000144***	0.000116***	0.000116***
	(2.00e-05)	(4.55e-05)	(4.08e-05)	(4.11e-05)
Race	0.0150***	-0.0283***	-0.0301***	-0.0301***
	(0.00173)	(0.00680)	(0.00463)	(0.00467)
Constant	5.694**	11.28***	3.507	3.507
	(2.459)	(2.300)	(2.181)	(2.199)
Observations	3,970	3,970	3,970	3,970
R-squared	0.428	0.685	0.890	0.890

Standard errors clustered at the county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

First, the base case (1) produces a result consistent with **figures 2** and **3**, but seemingly inflated with a large positive and statistically significant coefficient on the interaction term between the treatment and post-period dummies at 0.803. Note here that the R-squared value in the base case is the lowest amongst all models at 0.428.

Next, adding time fixed effects (2) produces a less inflated, but still statistically significant coefficient on the interaction term at 0.366. The R-squared value dramatically improves to 0.685 relative to the base case. Adding county fixed effects (3) increases the coefficient on the interaction term to 0.385 and produces an even higher R-squared value of 0.890. Finally, adding pair fixed effect (4) shows no impact on the size and significance of the treatment effect or the R-squared value. This last point about the pair fixed effect is consistent with the fact that the control and treatment groups are already matched using socio-demographic factors, and thus including pair fixed effects explicitly should not change the coefficients largely.

However, I point out here that since the parallel-trend assumption does not hold in the pre-treatment period, as made clear by the event study analysis, these coefficients on the interaction term cannot be straightforwardly interpreted as a pure treatment effect.

Making some additional observations about the other statistically significant and magnitude-wise relevant coefficients on socio-demographic control variables to check for intuitive consistency in the analysis, the coefficient on “CollegeDegree” is negative for models 1 and 2, while it turns positive for models 3 and 4. Intuitively, I would expect the coefficient on educational attainment here to be negative (i.e., higher educational attainment levels of a county leading to lower levels of unemployment), however, this was not the case for models 3 and 4.

Meanwhile, the coefficients on “Race” (i.e., percent of population who are white) was negative for all models except for the base case. The negative coefficients on race here makes intuitive sense as I would expect unemployment rates to be lower for the majority race of the region relative to that of minorities.

Lastly, the coefficient on sex (percent of population who are male) are positive and statistically significant in models 3 and 4, suggesting higher unemployment levels for counties with more males in the population than other counties. This seems consistent with the context that coal-related jobs lost from the decline of the coal industry disproportionately effects males compared to females.

b) One-to-One (1:1) Matched Dataset

Now, using the one-to-one matched dataset—limiting the analysis to 86 counties in both the control and treatment groups that are most socio-demographically similar—I perform the same set of analysis methods as above. Plotting the average unemployment rates in time period relative to treatment year, I get the plot in **Figure 4**. Again, average unemployment rates fall drastically in both the treatment and control counties throughout the 15-year time frame presented. However, unlike in the many-to-one (m:1) matched data case, it appears reasonable to assume parallel-trends in the unemployment rates in the pre-treatment period, especially between t=-6 and t=-1. This parallel trend diverges upon the treatment year t=0, however, not in any clear systematic way.

Figure 4. Average Unemployment Rates Using Matched Data (1:1)

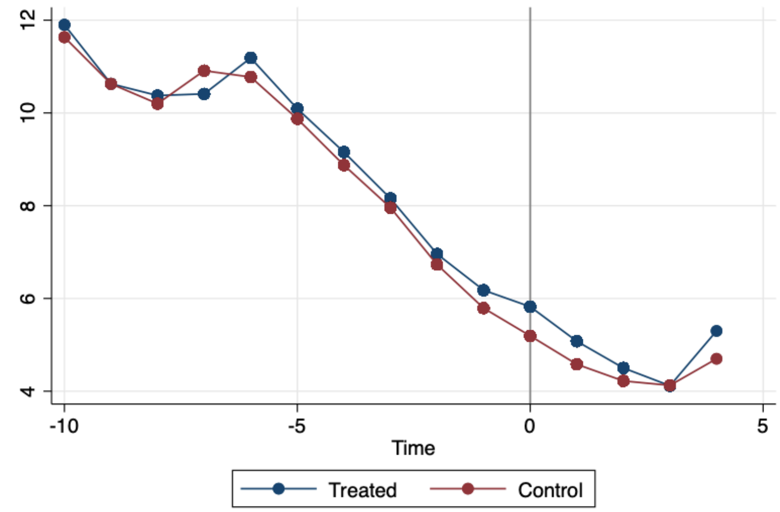

Now, running the event-study analysis using the one-to-one matched, I confirm the above notion of parallel trends in the pre-treatment period followed by some divergence in the post-treatment period.

Figure 5. Event-Study (1:1)

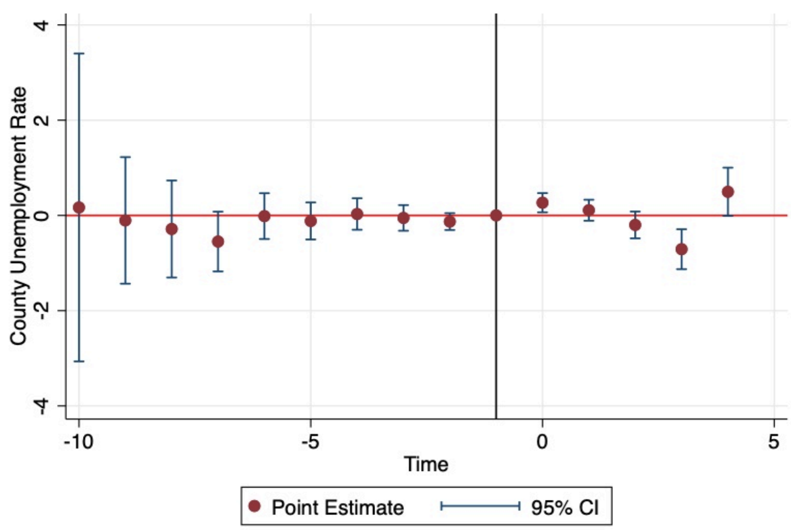

The event-study plot in **Figure 5** absorbs time, county and pair fixed effects, in addition to controlling for socio-economic characteristics. Based on this figure, the slope on the unemployment rates of the treatment group in the pre-treatment period, especially between t=-6 and t=-1, appears no different from that of the control group. This then diverges in small magnitude thereafter the treatment year, both in the positive and negative directions depending on the time period. More precisely, the improvements in unemployment rate appears to slow down slightly in t=0 and t=1 relative to the control group, followed by an acceleration in the drop in unemployment rates in t=2 and t=3. Unemployment rates rise for both groups in t=4, but to a larger extent for the treatment group.

Lastly, the two-way fixed-effect difference-in-difference analysis gives the below result.

Table 8. Difference-in-Difference Regression Output (1:1)

Dependent variable: Average county-level unemployment rate (%)				
	(1)	(2)	(3)	(4)
Specifications	Base Case	Time FE	Time & County FE	Time, County & Pair FE
Treat	0.162			
	(0.248)			
Post	-2.631***			
	(0.263)			
TreatPost	0.166	0.319	0.0563	0.0563
	(0.226)	(0.212)	(0.166)	(0.170)
FamilySize	-0.157	0.281	-0.243	-0.243
	(0.503)	(0.437)	(0.410)	(0.421)
CollegeDegree	-0.0348***	-0.0900***	0.0152	0.0152
	(0.00501)	(0.0130)	(0.0216)	(0.0222)
Veteran	0.000107**	-5.25e-05	7.57e-05*	7.57e-05*
	(4.69e-05)	(3.35e-05)	(4.10e-05)	(4.20e-05)
Foreign	-0.146***	-0.0461	0.0137	0.0137
	(0.0522)	(0.0358)	(0.0343)	(0.0352)
Language	-8.05e-05**	3.87e-05	-6.12e-05*	-6.12e-05*
	(3.62e-05)	(2.66e-05)	(3.32e-05)	(3.40e-05)
Population	-3.11e-08	-2.16e-06	2.67e-05**	2.67e-05**
	(2.36e-06)	(1.64e-06)	(1.11e-05)	(1.14e-05)
Sex	-0.164*	-0.0587	0.0664	0.0664
	(0.0846)	(0.0787)	(0.0687)	(0.0705)
Age	-0.000368***	0.000194***	0.000139**	0.000139**
	(0.000113)	(5.76e-05)	(6.65e-05)	(6.82e-05)
Race	0.0104***	-0.0292***	-0.0208***	-0.0208***

	(0.00333)	(0.00751)	(0.00445)	(0.00457)
Constant	17.43***	13.51***	4.082	4.082
	(4.501)	(4.074)	(3.565)	(3.658)
Observations	1,720	1,720	1,720	1,720
R-squared	0.451	0.753	0.924	0.924

Standard errors clustered at the county level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Overall, while adding more fixed effects increases the fit of the model as can be seen in the R-squared values, none of the models produced a statistically significant coefficient on the "TreatPost" interaction term, suggesting that there was no significant treatment effect on unemployment rate induced by the grant policy in question. This is consistent with the observations made on **Figures 4** and **5**, that there appears to be no clear systematic change in the trend in unemployment rates, neither in magnitude nor direction, for the treatment group in the post-treatment period. Any small positive effect seen in **Figure 5** at the beginning of the post-treatment period is cancelled out by the negative effect later on.

Making some additional observations about the other statistically significant and magnitude-wise relevant coefficients on socio-demographic control variables to check for intuitive consistency in the analysis, the coefficient on "CollegeDegree" is negative for models 1 and 2. This is consistent with the intuition that higher educational attainment levels of a county would lead to lower levels of unemployment. However, controlling for county and pair fixed effects in models 3 and 4 removes statistical significance for the coefficient on "CollegeDegree." This is likely due to the small within-county variation in average educational attainment levels over time.

Meanwhile, the coefficient on "Foreign" (i.e., percent of population who are foreign born) is negative and relatively large in magnitude for the base case. This is a curious result, calling for further investigation into the demographic characteristics (e.g., educational attainment level) of foreign nationals in the relevant counties. However, statistical significance no longer holds once fixed effects are absorbed in the models 2, 3 and 4, perhaps for a similar reason as the educational attainment case.

Lastly, the coefficients on "Race" (i.e., percent of population who are white) is statistically significant and negative for all models except for the base case. The negative coefficients on race here makes intuitive sense as I would expect unemployment rates to be lower for the majority race of the region relative to that of minorities.

c) Comparing the analyses on the m:1 and 1:1 datasets

Comparing the analyses on the many-to-one matched data including all grant-recipient counties and the analyses on the one-to-one matched data limiting to treated counties with socio-demographically similar counterparts in the control pool, I make one key observation. Namely, the unemployment rates of the 248 counties included in the many-to-one dataset but excluded from the one-to-one dataset (i.e., the grant-recipient counties that were not socio-demographically too divergent from the control counties to be matched one-to-one) likely have a different unemployment rate trend trajectory compared to the rest of the counties. More precisely, since the inclusion of these treated counties skewed the unemployment rate trend to be on a slower improvement trajectory relative to the control group, they likely had an even slower rate of improvement in the unemployment rate than the rest of the treated counties.

VII. Conclusion

The analyses reveal that i) overall unemployment rates were on a downward trajectory in the ARC region even before the POWER Initiative grants were rolled out, and that ii) no clear systematic or isolated impact could be identified for the grant policy. However, I identified that

iii) the grant policy may have effectively targeted some counties that were on a slower improvement trajectory relative to others. Additionally, while in small magnitude and varying directions, iv) the trend in unemployment rates deviated from previous trends in the post-treatment period in the grant-recipient counties, suggesting some potential impact of the grants that varied between counties.

One possible limitation of my study is that the analysis may have been performed too early to see real employment impacts of the grant initiative, since roll-out was staggered and it had only been one or two years since some of the treatment counties received their grants. Additionally, the limited number of counties available for the control group relative to the number of grant-recipient counties posed a challenge in carrying out a balanced, yet robust analysis. This was due to the fact that a substantial number of counties in the ARC region had received at least one grant under the POWER Initiative and few were unaffected. While I took the approach of comparing the analyses on two separate matching approaches, in the future, one potential approach to address this issue could be to include counties from outside the ARC region in the pool for potential control counties when performing propensity score matching. Although ideally, comparisons would be made within the formerly coal-dependent region. Lastly, this analysis also did not distinguish between types of grants or grant amounts and treated all grant-recipient counties with equal weight since the distinction of which counties received the grant was already not the clearest cut (e.g., some grants funded projects that broadly covered multiple counties while others were location specific).

Hopefully, this study creates a basis for future study of the POWER Initiative and opens a door for a more robust analysis in the future. Revisiting this policy five to ten years down the line when more counties close out their grants and programs funded under the initiative took into more robust fruition may produce clearer results. A heterogeneity analysis amongst grant recipient counties may also be possible once the total grant amount given out is larger.

VII. Appendix

To check for robustness of the many-to-one event-study plot in **Figure 3** in section VI(a), I produce below the same plot using t=-2 as the baseline year instead of t=-1.

Figure 6. Event-study (m:1) [Using t=-2 as the Baseline Year]

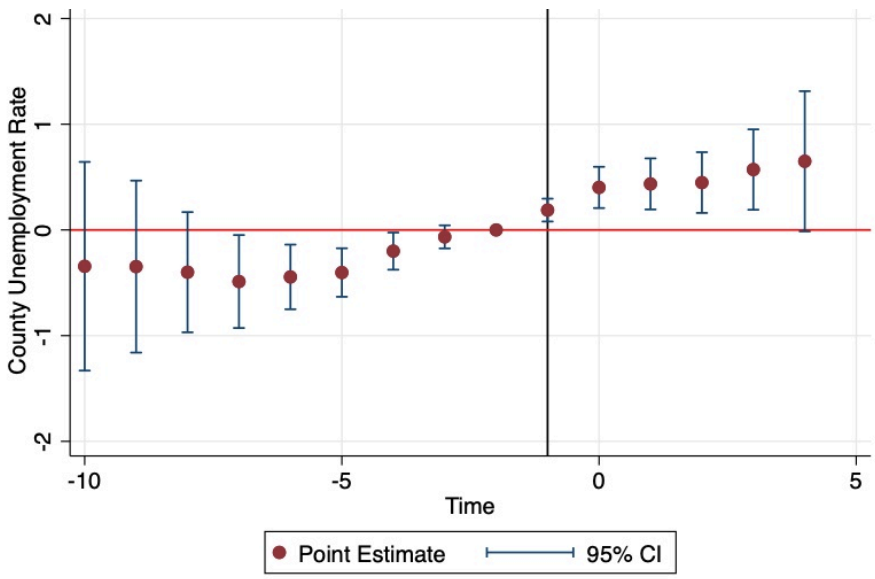

This alternative event-study plot presents similar trends to the original plot, suggesting that even if there was some anticipation about the policy preceding the treatment year, similar observations could be made about the trends in unemployment rates.

VIII. References

Appalachian Regional Commission. (2019). *About ARC Grants*. https://www.arc.gov/about-arc-grants/

Betz, M. R., Partridge, M. D., Farren, M., & Lobao, L. (2015). Coal mining, economic development, and the natural resources curse. *Energy Economics*, *50*(Journal Article), 105–116. https://doi.org/10.1016/j.eneco.2015.04.005

Black, D., McKinnish, T., & Sanders, S. (2005). The Economic Impact of the Coal Boom and Bust. *The Economic Journal (London)*, *115*(503), 449–476. https://doi.org/10.1111/j.1468-0297.2005.00996.x

Cecire, M. (2019). The POWER Initiative: Energy Transition as Economic Development. *Congressional Research Service*, 16.

Classifying Economic Distress in Appalachian Counties. (2020). https://www.arc.gov/classifying-economic-distress-in-appalachian-counties/

Investing in Appalachia's Future. (2015). *Appalachian Regional Commission*, 48.

Kline, P., & Moretti, E. (2014). LOCAL ECONOMIC DEVELOPMENT, AGGLOMERATION ECONOMIES, AND THE BIG PUSH: 100 YEARS OF EVIDENCE FROM THE TENNESSEE VALLEY AUTHORITY. *The Quarterly Journal of Economics*, *129*(1), 275–332.

POWER Award Summaries by State. (2020).

United States Census Bureau. (2020). *American Community Survey (5-Year Estimates)*. https://www.socialexplorer.com/explore-tables

U.S. Bureau of Labor Statistics. (2020). *Local Area Unemployment Statistics*. https://www.bls.gov/lau/#cntyaa

U.S. Energy Information Administration. (2020). *U.S. Energy Information Administration, Annual Energy Outlook 2020: With Projections to 2050. Total Energy Production: Coal*. U.S. Energy Information Administration. https://www.eia.gov/outlooks/aeo/data/browser/#/?id=1-AEO2020&region=0-0&cases=ref2020&start=2018&end=2050&f=A&linechart=~ref2020-d112119a.6-1-AEO2020&ctype=linechart&sid=ref2020-d112119a.6-1-AEO2020&sourcekey=0