The Review of Regional Studies

The Official Journal of the Southern Regional Science Association

Sample Selection in Appalachian Research*

Stratford M. Douglas and Anne W. Walker

Department of Economics, West Virginia University, USA

Abstract: The Appalachian Regional Commission's definition of the Appalachian region is the one used most often by researchers, politicians, and the popular press. The uncritical use of this definition of Appalachia raises issues of both selection bias and excess heterogeneity in regression analysis of Appalachian income and growth. The ARC was created as part of President Johnson's war on poverty, and the geographical extent of its purview has been driven by politics and by the geography of poverty, neither of which is exogenous. It is well known that the use of endogenous variables to choose a sample creates bias and inconsistency in estimation of regression coefficients. To identify the counties that belong to the Appalachian region exogenously we use an algorithm based on three criteria: topography, contiguity, and prevalence of slavery in the 1860 census. We apply our sample to growth regressions using data from 1970 to 2008, addressing the question of the existence of a resource curse from coal extraction. For this model we find strong evidence of excess heterogeneity, but not bias.

Keywords: sample selection, regions, resource curse

JEL Codes: C82, R12, R58, Q32, O18

1. INTRODUCTION

Use of the Appalachian Regional Commission (ARC) definition of Appalachia has become widespread in academia and the popular press. The ARC was created in 1965 as part of President Johnson's war on poverty, and the counties under its purview were deliberately chosen for their low income levels and growth rates. At the same time, the homogeneity of the region attracts social science researchers seeking natural experiments to test hypotheses about the determinants of poverty and income growth rates using regression analysis. Unfortunately, regression coefficients estimated from a sample chosen on criteria that are correlated with a regression's dependent variable will be biased and inconsistent (Heckman, 1979), so the ARC's Appalachian counties constitute a problematic sample for such regressions. To avoid sample selection bias in Appalachian income, poverty, and income growth regressions while maximizing the homogeneity of the sample, we therefore offer an alternative algorithm for identifying the sample that is based on three criteria that are exogenous to income and income growth: topography, contiguity, and prevalence of slavery in the 1860 census.

We demonstrate our method of sample selection by using it to select samples for a set of county-level panel data regressions on per capita income growth that seek evidence of a "resource curse" from coal extraction. Most previous studies of the resource curse have employed international data, which do not allow for sharp identification of the effects of resource

*The authors thank the contributions of the editor, two anonymous referees, and participants in the Research Workshop on Shale Gas and Oil Extraction at Bucknell University, July 31, 2012 for their many useful comments.

Douglas is an Associate Professor and Walker a graduate student in the Department of Economics at West Virginia University, Morgantown, WV 26501. *Corresponding Author:* A. Walker E-mail: anne.walker@mail.wvu.edu

abundance. This difficulty of identification arises from the great diversity among countries, compounded by measurement problems, omitted variables, and multicollinearity.¹ To address these problems in the resource-curse literature, van der Ploeg (2011, p. 381) suggests, “The road forward might be to exploit variation within a country where variables that might confound the relationship between resources and macroeconomic outcomes do not vary and the danger of spurious correlation is minimized.” Thus, if regions are thought of as experimental subjects receiving a random treatment of resource endowments, the experiment will more reliably detect the treatment effect if the subjects are otherwise homogenous.

The counties of the Appalachian region of the United States possess relative uniformity of culture, topography, and history that developed over a period of isolation from the rest of the country. They, therefore, present a nearly ideal population for the type of study suggested by van der Ploeg. The approximately equal number of treatment (i.e., coal-bearing) and control counties within the region also provide a desirable balance of experimental design. The ARC’s definition of the Appalachian region, however, introduces unwanted heterogeneity and raises the possibility of bias and inconsistency due to the endogeneity of its sample selection process. We propose to ameliorate both endogenous sample selection and excess heterogeneity by using exogenous variables to identify counties whose primary source of diversity is the presence or absence of commercially mineable coal. Our objective, quite simply, is to obtain an Appalachian county sample that is as homogeneous as can be obtained using exogenous criteria.

Comparing our regression results from different samples, we find evidence of excess heterogeneity but not bias due to sample selection. Despite the lack of bias in our particular application, the methods and reasoning that we use in selecting our sample have implications for future studies of development in resource-rich regions, including those whose exploitable resource base has recently expanded due to the development of new extraction technologies. Better sample selection techniques will allow more accurate conclusions about the effects of resource abundance on economic growth, which will provide a firmer basis for policy decisions affecting the speed and nature of resource development.

2. SAMPLE SELECTION METHOD

The Appalachian region south of New England possesses a distinct and robust culture that flows from its unique history. The immigrants settling this region in the eighteenth century were primarily of Scotch-Irish extraction, as opposed to the English, German, and African origins of the inhabitants of the Southern lowlands. There is a large literature describing distinctive characteristics of the southern Appalachian highlanders, including their close ties to land and family, distrust of outsiders, pugnaciousness, and independent spirit.² For more than a century after the American Revolution there was little immigration into the region, and a culture of defiant independence grew up among the subsistence farmers living in these isolated uplands. One marker of the region’s distinctive culture and economy was the hostility of its inhabitants to the Confederacy in the Civil War, which led to the wartime secession of West Virginia from

¹ Note that a sample selection criterion that is correlated with an omitted variable may also be endogenous in the sense that it is correlated with the error term in the (misspecified) regression that is being estimated. We seek to avoid this problem by choosing a relatively homogeneous sample. By definition, a homogeneous sample requires fewer variables to control for heterogeneity, which helps to minimize omitted variables bias.

² For example, Fischer (1989) provides an exhaustive treatment of the Appalachian Scotch-Irish culture, with emphasis upon its origins in the violent and impoverished regions on the border of Scotland and England.

Virginia. Despite (or because of) its semi-pariah status among urban Americans, the “hillbilly” culture of the Southern Appalachians has persisted. More than that, it has contributed highly visible elements to mainstream American culture, including the coonskin-capped pioneer icon, the Clampett and Snopes families, roots of bluegrass and country music, and NASCAR auto racing. Immigrants from other cultures that came to work in the Appalachian coal fields and forests in the last century were largely absorbed into this culture, which still predominates in the region. Despite considerable economic development in the past half-century, poverty still predominates in the region as well.

The Appalachian Regional Commission was created in 1965 as part of President Johnson’s war on poverty. Its identification of the counties under its purview was then, and has continued to be, much affected by the politics and geography of poverty. The Kennedy administration established the Presidential Appalachian Regional Commission (PARC) in 1963, which consisted of federal officials and the governors of nine Appalachian states.³ The PARC hired a staff of academics that produced a report identifying 340 Appalachian counties in ten states including Ohio (ARC, 1964). To aid passage of the Appalachian Regional Development Act of 1965, which created the ARC, Congress further expanded the list of Appalachian states to include South Carolina (Bradshaw, 1992). Mississippi and New York were added to the list in 1967, and more counties have been added since to bring the current total to 420 counties in 13 states. Adding counties to the ARC is a purely political process (Bradshaw, 1992), but its programs are designed to help poor counties, and poorer counties are therefore chosen for membership. For example, most Mississippi counties were added in 1967 explicitly because of their poverty (CRS, 2006). Though poor, most of these counties are not topographically, demographically, or culturally part of the Appalachian region. Similarly, many counties in topographically flat but economically depressed northeastern Ohio have been added to the ARC since that region’s economic decline began in the 1970s. Thus, the ARC’s county designations are clearly not exogenous to income growth rates, raising issues of bias and inconsistency in coefficient estimates from growth regressions. Addition of counties that do not share the Appalachian culture and history has added to the heterogeneity of the sample as well, which decreases the precision of estimates.

The Appalachian culture is tightly tied to the rugged topography of the Appalachian countryside. To identify the counties that belong to the Appalachian region we have therefore used an algorithm based on three exogenous criteria that reflect Appalachian landforms and the exogenous constraints they place on economic activity. One criterion is the set of Land Surface Form Topography Codes compiled by the U.S. Geological Survey (USGS, 1970). These codes, which are listed in Appendix 1, range from 1 (Flat Plains) to 21 (High Mountains), and broadly increase in relief as the numbers increase. The progression of codes is not perfectly monotonic, however, which led to some difficult implementation problems. For example, it is unclear whether “plains with low mountains” (code 11) are more or less mountainous than “open hills” (code 14), and it is not obvious whether code 8 counties (“tablelands, very high relief”) should be counted as Appalachian. After some study, it became clear to us that we needed other exogenous criteria to classify some counties.

³ The original nine ARC states were Alabama, Georgia, Tennessee, North Carolina, Kentucky, West Virginia, Virginia, Maryland, and Pennsylvania.

Figure 1: “Map Showing the Distribution of the Slave Population of the Southern States of the United States,” by Edwin Hergesheimer

Source: U.S. Bureau of the Census, 1861.

To supplement our sample selection algorithm we use the percentage of slaves in the county as enumerated in the 1860 Census. Prevalence of slavery within a county indicated its economic system, demography, culture, social relationships, and political loyalties. Because of poor soil and rough terrain, plantation agriculture and slavery never took hold in the Southern Appalachians. The inverse correlation between slavery prevalence and mountainous terrain is readily apparent in Edwin Hergesheimer's famous 1861 "Map Showing the Distribution of the Slave Population of the Southern States of the United States," an official publication of the U.S. Census Bureau that is reproduced in Figure 1.⁴ Our third criterion is geographical contiguity with other Appalachian counties, including information about the continuation and prevalence of ridges and mountainous terrain obtained from the terrain feature of Google Maps maps.google.com.

Our algorithm identifies both a narrow set of 250 "core" Appalachian counties with both mountainous terrain and low slave percentages, and a more broadly defined set of 409 counties with either smoother terrain and fewer slaves or rougher terrain and more slaves. Specifically, the core Appalachian counties include all contiguous counties with USGS topography codes exceeding 14 and 1860 Census enslavement ratios of 10 percent or less. We also include in the

⁴ See http://www.press.uchicago.edu/books/akerman/maps_slavery.html for a brief history of the significance of this map. We downloaded our copy from the Library of Congress, <http://www.loc.gov/item/99447026>, specifically the 15.9 MB jp2 file at <http://lcweb2.loc.gov/gmd/gmd386/g3861/g3861e/cw0013200.jp2>.

Appalachian core counties possessing topography codes of five through nine and 1860 Census slave percentages of 10 percent or less. Counties containing Chattanooga and Knoxville Tennessee, Asheville North Carolina, and Charleston West Virginia were included in the Appalachian core despite their relatively high slave percentages (10 to 17 percent) because of their high topographic relief and their roles as centers of Appalachian culture and commerce.⁵ To designate our more broadly-defined Appalachian counties, we added 159 counties to the 250 in the core. Most of these additional counties have topography codes of five to nine or higher, all have slave percentages above 10 percent, and all are contiguous to other Appalachian counties. The additional counties in the broadly-defined Appalachian region include contiguous counties with topography codes of ten and 1860 slave populations of 20-28 percent, as well as a few counties in South Carolina and Virginia with higher slave percentages but topography codes of 20 (“low mountains”).

Table 1 shows that Appalachian per capita personal income is relatively low, and it drops a little as one goes deeper into Appalachia. The ARC and Narrow definitions of Appalachia are similar in terms of income, poverty rates, and unemployment rates. In addition, the ethnicity and demographic data presented in Table 2 provide some support for the idea that the Appalachian counties are similar to each other and different from non-Appalachian counties. For example, Appalachian populations are about 25 percentage points more ethnically white than the other counties in the ARC states, and their ethnicity is less variable from county to county. Appalachian residents are slightly more likely to claim British or Irish descent than residents of other regions in the ARC states, and are less likely to claim German descent. They are also slightly more likely to be native English speakers, which is not surprising given the relative paucity of Hispanic and other recent immigrants in Appalachian counties. The ARC’s willingness to add fringe counties to its purview on the basis of poverty means that ARC-designated Appalachian counties are less white and more variable in ethnicity than those designated by our algorithm, though they are about equally poor.

Table 1: Descriptive Statistics by Region within the 13 ARC States

	Appalachian Status Definition							
	Not App		Broadly App		Core App		ARC	
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Real Per Capita Income								
Personal Income 1970	11975	3478	10936	2779	10456	2828	10622	2573
Personal Income 1980	13644	4057	12653	2792	12207	2803	12383	2700
Personal Income 1990	16920	5412	14841	3622	13999	3385	14502	3373
Personal Income 2000	19748	6706	17060	4300	16005	3855	16673	4050
White 1960	0.708	0.187	0.964	0.052	0.979	0.032	0.929	0.119
White Non-Hispanic 1990	0.736	0.159	0.965	0.047	0.979	0.026	0.937	0.108
Poverty Rate 1990	0.174	0.087	0.177	0.079	0.194	0.086	0.186	0.078
Number of Counties	660		409		250		420	
Number of Coal Counties	13		174		142		181	

⁵ We excluded Birmingham, Alabama, from the core (but not the broad) Appalachian designation because of Jefferson County, Alabama’s 22 percent slave ratio in 1860.

Figure 2: Appalachian County Designations **Table 2a: Cross-Tabulation of Appalachian Counties and ARC Counties within ARC States**

		Appalachian Broadly Defined					Appalachian Core		
		No	Yes	Total			No	Yes	Total
ARC	No	593	56	649	ARC	No	639	10	649
	Yes	67	353	420	ARC	Yes	180	240	420
	Total	660	409	1069	ARC	Total	819	250	1069

Figure 2 illustrates the differences and commonalities among the ARC, core, and broad definitions of Appalachia, and Table 2a provides cross-tabulations. Table 2b provides some descriptive statistics of the counties that are affected by differences among the various methods of designating the region. Our broad definition of the Appalachian includes 353 (84 percent) of the ARC counties, though we exclude 113 of those common counties from our “core” Appalachian designation. These 113 include some with moderately high slave percentages, concentrated in Alabama, Mississippi, East Tennessee, and central Kentucky, in the New, Holston, and Shenandoah river valleys of Virginia, and in the Carolina Piedmont. They also include some ARC counties with moderate topographical relief in Ohio and New York.

There are significant discrepancies between even our broad definition of Appalachia and the ARC definition of Appalachia. Most of the 67 ARC counties that our

Table 2b: Descriptive Statistics by Non-Overlapping Regions

	Appalachian Status Definition							
	ARC		Broad/NotCore		Core/NotARC		Broad/NotARC	
	/NotBroad							
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Real Per Capita Income	n=67		n=159		n=10		n=56	
Personal Income 1970	10934	2510	11759	2591	12768	1892	13291	2825
Personal Income 1980	12507	2746	13422	2688	13088	1642	14507	2700
Personal Income 1990	15170	3581	16281	3781	16611	2515	17776	4064
Personal Income 2000	17363	4115	18906	4856	17858	2506	20324	4549
White 1960	0.75	0.188	0.939	0.069	0.981	0.016	0.968	0.03
White Non-Hisp. 1990	0.779	0.177	0.939	0.074	0.953	0.04	0.954	0.043
Poverty Rate 1990	0.189	0.72	0.152	0.057	0.12	0.043	0.123	0.053

method excludes entirely counties are overall much less ethnically white than the others, and the standard errors of the from Appalachia (ARC/NotBroad in table 2b) are poor counties located on the fringe of Appalachia in Alabama, Mississippi, and northeastern Ohio. Table 2b shows that these ARC ethnicity variables indicate much more ethnic variability as well. Their poverty levels and incomes are similar to those of the Appalachian region listed in table 1, but their poverty may be more driven by the legacy of slavery or the decline of manufacturing than is the case in Appalachia proper.

We also include 56 non-ARC counties in our definition of Appalachia, of which ten are defined as part of the Appalachian core. The 46 broadly Appalachian, but non-core and non-ARC, counties are concentrated in the middle Tennessee River valley, the coal regions of west central Kentucky, the rugged country around the Ohio River west of Cincinnati, and the eastern ridges of the Alleghenies in Virginia, Maryland, Pennsylvania, and New York. The ten counties included in core Appalachia but excluded from the ARC are concentrated in the Adirondacks and Catskills of New York, but also include Shenandoah County, Virginia and Grayson County, Kentucky, which enslaved only 5.5 percent and 4.4 percent of their inhabitants, respectively, in 1860. Figures in table 2b indicate that these 56 counties are similar in ethnicity to the rest of Appalachia, but their relatively low poverty rates and high incomes probably explain their exclusion from the ARC.

3. APPLICATION: THE COAL RESOURCE CURSE

3.1 Model and Data

In this section we describe a simple panel data regression model designed to search for a coal resource curse effect in Appalachia. As discussed above, we rely upon the relative homogeneity of our sample to minimize the influence of omitted variables, allowing us to specify a parsimonious model with only a few control variables. Our resource abundance measure *Coal0* varies little through time, as it simply indicates whether or not a county contains commercially relevant coal deposits. Since the variation in *Coal0* is almost entirely in the cross-sectional dimension, we employ cross-sectional fixed effects only at the state level.

The search for a resource curse has generated a lively debate in the literature. Although the debate has attracted interest from economists since Adam Smith, a recent upsurge of interest has followed from a series of papers by Sachs and Warner (1995, 1997, 2001). Van der Ploeg (2011) provides the most recent survey of this literature, but useful surveys are also available from Ross (1999), Auty (2001), and Barbier (2005) as well. Most empirical work in this area has used cross-sectional analysis on international data, comparing long-run economic growth rates across large samples of both developed and developing countries. The empirical equations echo those found in the empirical growth model literature, which are derived from the neoclassical Solow growth model.

Very few previous studies have used subnational data to address the resource curse question. Michaels (2011) uses county-level data to investigate the effect of oil resources on a variety of measures, including personal income and employment density, and devotes much attention to the problem of sample selection. In contrast to much of the rest of the empirical resource curse literature, he finds a positive effect of resource abundance on employment density and income growth, which he attributes to agglomeration economies. His paper is also of special interest here because he attempts to limit heterogeneity by drawing his sample only from counties within 200 miles of an oil resource.

James and Aadland (2011) found significant negative growth effects of resource abundance using a national cross section of 3092 counties over the 1980-1995 time period. Papyrakis and Gerlagh (2007) also found a resource curse using state-level U.S. data (specifically, they estimated that a one percent increase in resource abundance decreases growth by about 0.047 percent). There are also a few unpublished papers using sub-national data. They include Caselli and Michaels (2009), which looks at the effect of oil resources in different regions of Brazil, and dissertation chapters by Sen (2010), Corey (2008), and Dunn (2008), which employ U.S. data measured at the state and (in the case of Sen) county level. Like James and Aadland (2011), these latter three studies analyze the United States as a whole and define resource abundance broadly, whereas we focus on Appalachia and coal in order to identify the mechanism sharply.

A few recent papers examine the resource curse specifically in the Appalachian region, and most of the empirical results are broadly in line with our finding of a modest negative effect of coal mining on income growth. Deaton and Niman (2012), and Partridge, Betz, and Lobao (2012) both examine the problem using poverty rates as the dependent variable. Deaton and Niman (2012) examined poverty as a function of mining employment in Appalachia and estimated that a one percent increase in mining employment results in a net increase in poverty of approximately 15 percent. Partridge, Betz, and Lobao also find a mostly positive association between coal mining and poverty, but they also find evidence that the association has weakened recently. Black et al. (2005) examined the effects of coal mining on earnings during the boom and bust cycle of the late 1970s to early 1980s, and found that the down-cycle effects predominated, as earnings grew 5 percent faster during the boom, but 5.5 percent slower during the bust. Harkness (2010) used a sample of Kentucky counties to examine the resource curse and found no conclusive evidence of a resource curse associated with coal.

Economic growth models such as the one we use in this study have been employed in the study of other issues in Appalachia. These include a recent paper by Ziliak (2011), one of many that seeks to measure the efficacy of the ARC in reducing poverty by comparing outcomes in ARC counties to similar counties that are not included in the ARC. Ziliak (2011) compares

Appalachian counties to adjacent and surrounding counties, and to other U.S. rural counties. Isserman and Rephann (1995) used a quadratic distance metric to match Appalachian counties to their twins outside of Appalachia. Glaeser and Gottlieb (2008) compare Appalachian counties to other non-coastal counties within their respective states.

Our specific question involves a comparison of counties within Appalachia. We specify a panel data regression model of annualized growth rates, measured as the difference in log per capita personal income divided by the length of the period, g_{it} for county i in state s at time t :

$$(1) \quad g_{it} = y0_{it}\lambda + \text{Coal}_{it}\alpha + X_{it}\beta + \mu_s + \delta_t + \varepsilon_{it},$$

where $y0_{it}$ is the level of per capita personal income in county i at the beginning of time period t , Coal_{it} is a measure of coal resource abundance, δ_t is a time fixed effect, and X_{it} is a vector of exogenous variables. The standard Solow model calls for a region to exhibit conditional convergence to a balanced growth path, which implies a negative coefficient, λ , on the initial income variable. We include state fixed effects μ_s because institutional characteristics such as income, severance, sales, and inheritance taxes, broad infrastructure and educational policy, judicial system, much law enforcement (including environmental law enforcement), and corruption are all strongly affected by statewide institutions. We do not include county fixed effects because the variation in coal resource abundance is cross-sectional in nature and varies little through time, and county fixed effects would remove all time-invariant, cross-sectional differences.

An ideal resource abundance variable should be exogenous, and it should indicate the presence of the resource in a way that reflects both the resource extraction process and the hypothesized transmission channels for the resource curse. In the current study, we believe that if there is an Appalachian coal resource curse the primary transmission channels for it would be through 1) depressed educational attainment by students who terminate their formal education early to take relatively highly paid coal mining jobs that do not require much formal education, 2) economic disruption and long-term unemployment caused by the boom and bust cycle, and perhaps 3) long-term health effects. The effects of these channels, if they exist at all, would persist over a long time period; therefore our resource abundance variable should not be too sensitive to current production, or it should employ a lag structure of some length. Simplicity of both construction and interpretation are also desirable. We therefore employ an indicator variable $\text{Coal}0$, equal to one if the county has ever produced and sold any coal, as of the initial year of the time period. This dummy variable is likely exogenous because the Appalachian coal fields are mature: nearly all of the counties, including all of the major coal producing counties, initiated coal production before our sample period.⁶ The samples are well balanced with respect to this variable: as of 1990, 46 percent of the ARC counties, 61 percent of our core, and 47 percent of our broadly defined Appalachian counties had produced at least some coal.

Our choice of control variables x_{it} is limited by our research goal of estimating a reduced form equation. Past researchers, seeking to explore possible mechanisms or channels for the resource curse, have included a variety of control variables thought to influence the long-run average growth of an economy, many of which introduced endogeneity concerns. To mitigate endogeneity problems, we limit the number of controls. The drawback of this research strategy

⁶ In this respect, our $\text{Coal}0$ variable resembles the resource abundance variable used by Michaels (2011), which was a measure of the total amount of oil resource that existed prior to exploitation.

Table 3: Descriptions of Variables and Descriptive Statistics for 1069 counties in 13 ARC States

Variable	Description	Mean	Std. Dev.	Min	Max
Dependent Variable	Average annual growth rate of per capita personal income, net of taxes and transfers	0.013	0.013	-0.042	0.093
Initial Income	Log of beginning of period log per capita personal income	9.554	0.331	8.322	11.249
Metro County	= 1 if Beale code ≤ 3 (Metro area)	0.325	0.468	0	1
Rural County	= 1 if Beale code ≥ 7 (Very Rural, or urban population $< 20,000$, not adjacent to an urban county)	0.377	0.485	0	1
(dOil Price) $\times$ (Coal0)	(Percent change in real oil price over period) times Coal0	0.079	0.473	-0.940	1.650
Coal 0	= 1 if county has produced coal commercially, as of initial year of period (220 of 1069 as of year 2000)	0.203	0.402	0	1

is that it does not allow us to determine the channel of operation of the resource curse, if any. On the other hand, the strategy does allow us a cleaner test for the existence of the resource curse, given a homogenous sample.

Appalachian poverty is often attributed to geographic isolation and consequent lack of access to transportation, financial centers, higher education, and other growth-enhancing amenities associated with urban areas. Accordingly, we have included dummy variables derived from the USDA's Rural-Urban Continuum ("Beale") Codes. Beale Codes are available for 1974, 1983, 1993, and 2003, and are based primarily on data from the previous census. So we treat them as exogenous beginning-of-period values for their respective decades. Because they are non-monotonic, we use them to create two dummy variables that indicate rural and metro counties, respectively.

The panel data structure includes four time periods—the decades beginning in 1970, 1980, and 1990, along with the nine-year period 2000-2008, which are used to construct the time fixed effects dummies. The inclusion of time fixed effects is standard in panel growth models, as they essentially transform each variable into a difference from its cross-sectional mean as of time t , which controls for the business cycle and other shocks common to all counties at a given time. To control for the cyclical and secular effects of changes in energy prices on counties that produce coal, we include an interaction variable that is the product of the (log-difference) percentage change in real oil price and *Coal0*.⁷

We obtained the data for our economic growth and initial income variables from the Bureau of Economic Analysis's Regional Economic Accounts. We adjusted the measure by subtracting government transfers, and deflated to constant-year terms using a base year of 2000.

⁷ The average real price of oil in 1970, 1980, 1990, and 2000 was \$19.65, \$102.26, \$39.80, and \$35.76, respectively.

Coal production data were obtained from the U.S. Geological Survey (USGS) and the U.S. Energy Information Administration (EIA), and the Beale codes were obtained from the U.S. Department of Agriculture (USDA). Coal and oil price data were obtained from the EIA.

3.2 Results

We estimated equation (1) using time and state-level fixed effects in panel data regressions, and report the results in Table 4 along with county-cluster-robust standard errors obtained by bootstrapping. There are four columns of results from four different samples. The column labeled “All” includes all counties in the 13 Appalachian states, regardless of whether or not the counties are part of Appalachia. The other columns correspond to the samples chosen as explained in detail above.

The coefficients shown in Table 4 are of expected signs and reasonable magnitudes. The estimated coefficient on beginning of period income (*Log BOP Income*) is negative and significant at the 1 percent level in all specifications, indicating conditional convergence of per capita income. The asymptotic rate of convergence can be inferred (Higgins et al., 2006) from the BOP income coefficient estimate as $c = 1 - (T\beta)I/T$. Coefficient estimates range from -0.013 to -0.023, implying a convergence rate of 1.3 percent to 2.3 percent per year. Appalachian counties are estimated to have faster rates of convergence. As expected, the coefficient estimates for the metro dummy variable is positive and significant. Estimates for the rural dummy variable coefficient are of the expected sign, but not significant. Coefficient estimates for the interaction variable between the change in real oil price and the *Coal0* dummy variable are all positive and significant at the 1 percent level, consistent with energy prices’ well-known short-term positive effect on income in coal-producing Appalachian counties.

The coefficient estimates on *Coal0* suggest a moderate but significant coal resource curse, slowing expected economic growth by an estimated 0.3 percent to 0.4 percentage points per year. Over time, even such small decreases in average annual income growth have a sizeable impact on income levels. For example, a drop in average annual income growth of 0.4 percentage points results in income levels being lowered by about 17.3 percent over forty years, with compounding. Therefore, our estimates indicate that counties with coal mining present were 15 to 20 percent poorer by the end of the forty year period covered by our sample than similarly rural counties within the same state that started with the same income but lacked a coal resource. Following Black, et al. (2005), we also estimated the model with (state)*(time) interaction terms to account for state-specific growth rate variations through time. We report the results in Appendix 2. Generally, this alternative model finds a smaller but still significant negative association between coal mining and economic growth. The estimated effect of the coal resource shrinks enough to become statistically insignificant for the Core Appalachian counties.

Does our sample selection algorithm make any difference in the estimates? Somewhat surprisingly, given the clear endogeneity of the ARC’s sample selection, the coefficient estimates for ARC, Core, Broad, and All county samples are of a similar magnitude. This similarity of results suggests that sample selection bias is not a large problem for this particular data set and model.

The similarity of the standard errors across models, however, indicates possible excess heterogeneity in the ARC sample. Smaller data sets imply larger coefficient standard errors, which should increase according to the square root of the ratio of sample sizes. The Core

Table 4: Growth Regression Results.
Dependent Variable: Annualized Growth Rates of Per Capita Personal Income 1970-2008
(Panel Data, 4 Periods, State and Time Fixed Effects County-level-cluster-robust bootstrap standard errors in italics)

	ALL	APPALACHIAN		
	COUNTIES	BROAD	CORE	ARC
Intercept	0.1309*** (0.0246)	0.1958*** (0.0144)	0.2214*** (0.0300)	0.2033*** (0.0228)
Initial Income	-0.0133*** (0.0026)	-0.0202*** (0.0015)	-0.0227*** (0.0032)	-0.0210*** (0.0024)
Metro	0.0049*** (0.0006)	0.0046*** (0.0009)	0.0040*** (0.0009)	0.0045*** (0.0008)
Rural	-0.0011 (0.0007)	-0.0009 (0.0007)	-0.0018** (0.0007)	-0.0010 (0.0009)
dOilPrice*Coal0	0.0063*** (0.0011)	0.0066*** (0.0012)	0.0070*** (0.0013)	0.0064*** (0.0012)
Coal0	-0.0044*** (0.0009)	-0.0045*** (0.0013)	-0.0035** (0.0015)	-0.0033*** (0.0011)
#Counties	1069	409	250	420
sigma_u	0.0023	0.0026	0.0033	0.0025
sigma_e	0.0108	0.0102	0.0104	0.0100
rho	0.0431	0.0592	0.0892	0.0577
R-Square				
Within	0.3050	0.3522	0.3133	0.3721
Between	0.1021	0.0406	0.0793	0.1306
Overall	0.3027	0.3374	0.2868	0.3504

Appalachian data set contains many fewer counties than the other three, (about 60 percent of the ARC and Broad samples, and 23 percent of the All sample). These sample size effects suggest that standard deviations of the Core estimates should be approximately 1.3 times those of the ARC estimates and twice those of the All sample estimates. Instead, the standard deviations of coefficients for the three sample sizes are of similar magnitude for most coefficients, consistent with our assertion that the Core data sets are more homogeneous.

To further explore this issue of the efficiency gain from using our method of selecting Appalachian counties we performed a bootstrap sampling experiment to see whether our Core sample's estimated standard errors are low by chance, or because the sample was well-chosen.⁸ In each of 10,000 bootstrap repetitions we drew randomly (with replacement) a sample of 250 counties from the set of 420 ARC counties. The sample size of 250 counties was chosen to reflect the number of counties in the Core Appalachian sample, so that the sample represented a Core-size sample of counties drawn at random from the population of ARC counties. In each case all four time-series observations were used from each county, so that the total size of the

⁸ We thank an anonymous referee for suggesting the bootstrap experiment.

data set in each sample was $4n = 1000$. If the Core sample is no more homogeneous than the ARC sample, we would expect Core coefficient standard errors to rank at about the 50th percentile of the standard errors from all the bootstrap samples. For each bootstrap repetition b we applied our empirical model to the bootstrap sample to estimate the k -vector of regression coefficients $\hat{\beta}_b^*$ from that sample, along with its $k \times k$ coefficient covariance matrix $Cov(\hat{\beta}_b^*)$. We then compared the size of $Cov(\hat{\beta}_b^*)$ to the comparable covariance matrix obtained from our sample of 250 Core Appalachian counties $Cov(\hat{\beta}_{core})$. To do so, we calculated $Cov(\hat{\beta}_b^*) - Cov(\hat{\beta}_{core})$ and tested to see if the result was positive definite, which would indicate that our Core sample's covariance matrix was smaller than the bootstrap sample's covariance. We found that our Core sample's covariance matrix was smaller than all but 12 percent of the bootstrap samples' covariance matrices drawn from the ARC counties, which indicates a high degree of uniformity of the Core sample relative to the ARC sample.

4. CONCLUSIONS, POLICY IMPLICATIONS, AND FUTURE DIRECTIONS

The empirical literature on the determinants of economic growth has long been impeded by problems of unobserved heterogeneity, multicollinearity, and endogeneity found in international data sets. The resource curse literature has exactly the same set of problems, and for the same reasons. The answers to many important questions are not identified in the data because we don't have enough exogenous structural variables or good instruments to control for all of the heterogeneity of the world's nations. For many years researchers in this area have sought to solve the problem by using more sophisticated econometric methods and choosing instruments more cleverly.

We have addressed these problems by implementing a novel sample selection approach. Rather than trying to model and control for heterogeneity and endogeneity, we use exogenous criteria to construct a data set that minimizes heterogeneity. We have laid out the reasons why we regard the Appalachian region of the U.S. as a good candidate for this style of inquiry (uniform data collection, cultural homogeneity, balanced treatment design, and a set of questions that has attracted the keen interest of policy makers for half a century), and we have gathered a panel data set that has allowed us to isolate the presence of coal production as a source of exogenous variation. The results for our particular application suggest that any problems created by using the ARC definition of Appalachia for estimation of the resource curse arise more from the introduction of sampling noise rather than from bias due to endogenous sample selection. Partly, this result may be driven by the fact that few of the counties that are included in the ARC Appalachian sample but excluded from the Core Appalachian sample possess coal resources, meaning that heterogeneity in the economic growth variable in the excluded ARC counties arises from sources other than the resource abundance variable.

The lessons of this study are much clearer with respect to the possible efficiency gains from careful attention to sample selection. Our bootstrap experiment indicates that judicious use of exogenous cultural and geographical indicators to refine the sample selection algorithm can substantially reduce sampling variation, even within a sample such as the ARC counties that is already relatively homogeneous. It may be possible to improve upon the simple geographic contiguity and proximity measure used for sample selection by Michaels (2011), for example, if the researcher has sufficient knowledge of the region, creativity, and access to appropriate exogenous variables. Researchers seeking to study the impact of Marcellus Shale development

may be able to apply some form of the sample selection algorithm developed here. Useful sample selection algorithms for studies of the Bakken or resources farther afield must be tailored to the characteristics of those regions, but this study may be of some use as an example of the kinds of variables to seek when developing such algorithms.

Although lagged income, rurality, topography, and location of mineral deposits are not endogenous, they are not policy variables either, which unfortunately limits the policy implications that may be drawn from our reduced-form empirical results. That does not mean that there are no policy implications of our analysis. Our results so far cast some doubt on the notion that encouraging coal extraction is tantamount to encouraging growth. They provide some support for those who suggest that well-considered regulatory impediments to unbridled coal extraction could actually encourage economic growth in the regions where the resources are found. This is a point that is not obvious to many policy makers, focused as they are on the benefits arising from exploitation of new sources of energy such as oil shale, tar sands, and shale gas. On the other hand, as Michaels (2011) demonstrated, those who assume that resource abundance and exploitation is always a curse may also be mistaken. Determining the conditions under which resource extraction may be conducted to the benefit of those who live where the resource is abundant is the task of social science researchers. The first step in getting clear answers from statistical analysis is to choose an appropriate sample that minimizes noise and bias.

REFERENCES

- ARC. (1964) "Appalachia: A Report by the President's Appalachian Regional Commission," USGPO. Downloaded from <http://www.arc.gov/>.
- Auty, Richard M. (2001) *Resource Abundance and Economic Development*. Oxford: Oxford University Press.
- Barbier, Edward B. (2005) *Natural Resources and Economic Development*. Cambridge and New York: Cambridge University Press.
- Black, Dan, Terra McKinnish, and Seth Sanders. (2005) "The Economic Impact of the Coal Boom and Bust," *Economic Journal*, 115, 449–476.
- Bradshaw, Michael. (1992) *The Appalachian Regional Commission: Twenty-Five Years of Government Policy*. University Press of Kentucky: Lexington, KY.
- Caselli, Francesco and Guy Michaels. (2009) "Do Oil Windfalls Improve Living Standards? Evidence from Brazil," *CEP Discussion Papers dp0960*, Centre for Economic Performance, London School of Economics.
- Congressional Research Service (CRS). (2006) *Federal Regional Authorities and Commissions: Their Function and Design*.
- Corey, Joab. (2008) "Three Essays on Resources, Institutions, and Development across U.S. States," unpublished dissertation, West Virginia University.
- Deaton, B. James and Ekaterina Niman. (2012) "An Empirical Examination of the Relationship Between Mining Employment and Poverty in the Appalachian Region," *Applied Economics*, 44, 303–312.

- Dunn, Leslie. (2008) "Three Essays on Development in Specialized Economies: The Resource Curse in U.S. States and Small Island Economies," unpublished Dissertation, West Virginia University.
- Fishcer, David Hackett. (1989) *Albion's Seed: Four British Folkways in America*. Oxford University Press: New York.
- Glaeser, Edward and Joshua Gottlieb. (2008) "The Economics of Place Making Policies," *Brookings Papers on Economic Activity*, 1(Spring), 155–253.
- Harkness, Kristen A. (2010) "Escaping the Resource Curse: Lessons from Kentucky Coal Counties," unpublished paper presented at APSA 2009 Toronto Meeting. Available at SSRN: <http://ssrn.com/abstract=1450181>.
- Heckman, James J. (1979) "Sample Selection Bias as a Specification Error," *Econometrica*, 47, 157–161.
- Higgins, Matthew J., Daniel Levy, and Andrew T. Young. (2006) "Growth and Convergence across the U.S: Evidence from County Level Data," *Review of Economics and Statistics*, 88, 671–681.
- Isserman, Andrew and Terrance Rephann. (1995) "The Economic Effects of the Appalachian Regional Commission," *Journal of the American Planning Association*, 61, 345–365.
- James, Alex and David Aadland. (2011) "The Curse of Natural Resources: An Empirical Investigation of U.S. Counties," *Resource and Energy Economics*, 53, 440–453.
- Michaels, Guy. (2011) "The Long Term Consequences of Resource-based Specialisation," *Economic Journal*, 121, 31–57.
- Papyrakis, Elissaios and Reyer Gerlagh. (2007) "Resource abundance and economic growth in the United States," *European Economic Review*, 51, 1011–1139.
- Partridge, Mark, Mike Betz, and Linda Lobao. (2012) "Natural Resource Curse and Poverty in Appalachia," *Working Paper MPRA Paper No. 38290* (April). Downloaded from http://mpa.ub.uni-muenchen.de/38290/1/MPRA_paper_38290.pdf.
- Ross, Michael L. (1999) "The Political Economy of the Resource Curse," *World Politics*, 51, 297–322.
- Sachs, Jeffrey D. and Andrew M. Warner. (2001) "The Curse of Natural Resources," *European Economic Review*, 45, 827–838.
- _____. (1997) "Sources of Slow Growth in African Economies," *Journal of African Economies*, 6, 335–376.
- _____. (1995) "Natural Resource Abundance and Economic Growth," *NBER Working Paper 5398*. National Bureau of Economic Research, Inc.: Boston, MA.
- Sen, Arzu. (2010) "Essays on Housing Markets and Economic Growth," Unpublished Ph.D. Dissertation, West Virginia University.
- USGS. (1970) *National Atlas of the United States of America*, U.S. Geological Survey, US Dept of Interior: Washington, DC.

- van der Ploeg, Frederick. (2011) "Natural Resources: Curse or Blessing?," *Journal of Economic Literature*, 49, 366–420.
- Ziliak, James P. (2011) "The Appalachian Regional Development Act and Economic Change," in James P. Ziliak (ed.), *Appalachian Legacy: Economic Opportunity after the War on Poverty*. Brookings Institution Press: Washington, DC, pp. 19–44.

Appendix 1: Land Surface Form Topography Codes, USGS (1970)

Plains	Open Hills and Mountains
1 = Flat plains	13 = Open low hills
2 = Smooth plains	14 = Open hills
3 = Irregular plains, slight relief	15 = Open high hills
4 = Irregular plains	16 = Open low mountains
	17 = Open high mountains
Tablelands	
5 = Tablelands, moderate relief	Hills and Mountains
6 = Tablelands, considerable relief	18 = Hills
7 = Tablelands, high relief	19 = High hills
8 = Tablelands, very high relief	20 = Low mountains
	21 = High mountains
Plains with Hills or Mountains	
9 = Plains with hills	
10 = Plains with high hills	
11 = Plains with low mountains	
12 = Plains with high mountains	

Appendix 2: Growth Regression Results with State–Time Fixed Effect.**Dependent Variable: Annualized Growth Rates of Per Capita Personal Income 1970–2008**

(Panel Data, 4 Periods, State and Time Fixed Effects, cluster–robust standard errors in italics)

	ALL	APPALACHIAN		
	COUNTIES	BROAD	CORE	ARC
Intercept	0.1091*** (0.0295)	0.1476*** (0.0191)	0.1838*** (0.0335)	0.1688*** (0.0272)
Initial Income	−0.0110*** (0.0030)	−0.0156*** (0.0020)	−0.0179*** (0.0035)	−0.0169*** (0.0028)
Metro	0.0043*** (0.0007)	0.0036*** (0.0008)	0.0030*** (0.0009)	0.0037*** (0.0009)
Rural	−0.0010 (0.0007)	−0.0006 (0.0008)	−0.0011 (0.0007)	−0.0005 (0.0009)
dOilPrice*Coal0	0.0054*** (0.0011)	0.0056*** (0.0016)	0.0045*** (0.0012)	0.0055*** (0.0015)
Coal0	−0.0038*** (0.0009)	−0.0036** (0.0013)	−0.0022 (0.0014)	−0.0027** (0.0011)
#Counties	1069	409	205	420
sigma_u	0.0048	0.0074	0.0093	0.0065
sigma_e	0.0101	0.0093	0.0096	0.0093
rho	0.1875	0.3828	0.4833	0.3257
R–squared				
Within	0.3983	0.4702	0.4353	0.4691
Between	0.0321	0.0343	0.1786	0.0661
Overall	0.3253	0.2503	0.2340	0.3283