THE UNIVERSITY OF CHICAGO

ESSAYS ON THE RESPONSES TO LOCAL LABOR MARKET SHOCKS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE IRVING B. HARRIS
GRADUATE SCHOOL OF PUBLIC POLICY STUDIES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
YIMING LI

CHICAGO, ILLINOIS
JUNE 2018

To my parents, Lingjiao Li and Linjun Dai. For their love and sacrifices.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I thank my dissertation advisors Kerwin Charles, Bruce Meyer, and Damon Jones for guidance and support. This dissertation would not have been possible without their mentorship. I am also indebted to Bob LaLonde and Will Manning for many methodological conversations. Further thanks to numerous seminar participants at the University of Chicago for their helpful feedback. I thank Cynthia Cook-Conley for her support in administrative matters, and Fay Booker for her help in technical matters. And last, I thank my wife, Ruixue Du, for the amazing company along this incredible journey.

ABSTRACT

This dissertation studies the impacts of local labor market changes on the US family structures and disability benefit take-up. The dissertation uses unexpected time series changes in energy prices, together with the pre-existing county exposure to the unexpected time series changes, to identify causal links between the changes in county's economic conditions and changes in marriage outcomes, fertility outcomes, and disability benefit payments.

The first chapter is motivated by the decline in marriage and the rise in out-of-wedlock births, which are among the most important changes in American lives over the last half century. A notable hypothesis to explain both trends, sometimes dubbed as the "unmarriageable men hypothesis," attributes these changes to the disappearance of well-paying industrial jobs for the less-educated men. Exploiting county-level variation in oil-producing and coal-producing areas from the shocks to world oil and coal prices in the 1970s and 1980s, I find that a decrease in men's earnings leads to lower marriage prevalence, lower marital fertility rates, and higher non-marital fertility rates, confirming the hypothesis. When men's earnings increase, however, marriage outcomes and non-marital births respond differently across regions: while higher earnings lead to higher marriage prevalence and lower non-marital births in coal-producing counties, I find little marriage response and more non-marital births in oil-producing counties during the oil boom, which is inconsistent with the implications of the unmarriageable men hypothesis.

The second chapter, coauthored with Kerwin Charles and Melvin Stephens Jr., examines the extent to which US Social Security Disability Benefits (SSDI) and Supplemental Security Income (SSI) payments, the two largest Federal programs that provide assistance to people with disabilities, are affected by local labor market earnings from 1970 to 2011. To circumvent the potential reverse causality problem in the literature, we exploit county-level variation in oil-producing areas from shocks to world oil and gas prices to identify the causal effects. We extend well-known previous work using a similar research design by analyzing a different price shock, a larger, more representative set of labor markets, and a more recent period marked by skyrocketing disability payments. Our estimated elasticity for SSDI payments with respect to earnings of -0.29 is surprisingly similar to earlier findings. Our preferred SSI elasticity estimate of -0.16 is smaller than previous findings, but we show that SSI programmatic changes explain most of the difference.

CHAPTER 1

MEN'S EARNINGS, DECLINING MARRIAGE, AND RISING NON-MARITAL BIRTHS

1.1 Introduction

The steep decline in marriage and rise in out-of-wedlock births are among the most important changes in American lives over the last half century. In 2015, 39% of women aged 18–39 in the U.S. are married, half of the rate in 1960.[1] During the same period, births to unmarried mothers increased dramatically from 5% in 1960 to over 40% in 2015, with higher rate of 62% for women with a high school degree or less. These trends in marriage and non-marital fertility have profound implications for parents and especially children's well-being (See Kearney and Levine (2017) for a recent review). What is driving the decline in marriage and the rise in out-of-wedlock births? Wilson et al. (1986) and Wilson (1987) proposed a notable hypothesis—sometimes dubbed as the "unmarriageable men hypothesis"—suggests that the declining economic prospects of less-educated men had made them "unmarriageable" from women's perspective.

Although conceptually straightforward, it has been and remains to be empirically challenging to identify the causal link between men's economic situation and marriage. Comparing marriage and fertility outcomes of high-income men with those of low-income men cross-sectionally misses the possibility that high-income men might have different preferences for marriage and fertility and are facing different sets of prices and constraints from low-income men. In addition, positive assortative mating (Becker, 1973, 1974; Pencavel, 1998) makes it harder to control for women's earn-

1. Sources: 2015 Current Population Survey ASEC sample and 1960 Census 5% sample.

ings while comparing men's. In addition to absolute levels of income and education, the male-female relative earnings gap also contributes to marriage decisions due to persistent gender norms (Bertrand, Kamenica and Pan, 2015). A host of other factors, such as cohort sizes, housing costs, child care costs, and household labor-saving technologies, may also confound the causal relationship between men's earnings and marriage decisions.

An alternative approach, following Schultz (1985), is to use a large and sustained exogenous shock to men's income in some local areas but not others, and then evaluate the impacts by comparing differentially affected areas. Recent work by Black et al. (2013), Autor, Dorn and Hanson (2017), and Kearney and Wilson (2017) illustrate the promise of this approach in establishing the causal link between men's economic situation and fertility and marriage outcomes. Black et al. (2013) use sustained increase and decrease in coal prices to examine the effects on marital fertility in the Appalachian coal-mining region. They find that an increase in husband's income leads to higher marital fertility rate. Autor, Dorn and Hanson (2017) use differential trade exposure of the U.S. regions to China's export in the 1990s and 2000s to examine the effects on family structures in the U.S. They find that the decline in male-dominated manufacturing jobs leads to a decline in the marriage market values of men, thus a decline in marriage and rise in non-marital births. Kearney and Wilson (2017) focus instead on the marriage and fertility effects of the *improved* economic conditions of men. They document large positive fertility response but no marriage response in the regions that benefited from fracking boom during the 2000s. The authors contrast this finding in fracking boom with the effects of full coal boom and bust cycle during the 1970s and 1980s in the Appalachian area, when both marital birth rate and non-marital birth rate were negatively associated with men's earnings. The authors therefore suggest that the contrast in family formation responses to economic conditions is due to the gradual changes in social norms from 1970s to 2000s.

In this paper, we exploit county-level variation in coal-producing and oil-producing areas from shocks to world energy prices to study the effects of men's economic conditions on marriage and fertility. Our study complements the recent evidence by examining *both* the improvement and the deterioration of men's economic prospects associated with the boom and bust of the energy industries. We are able to examine the geographical differences in the marriage and fertility responses by contrasting coal and oil areas during the *same period* instead of 30 years apart in Kearney and Wilson (2017), therefore shedding additional light on the social norms explanation.

The energy booms and busts in the 1970s and 1980s affected the counties that had large coal or oil industries but did not have much impacts on counties that had no coal or oil. By comparing affected and unaffected counties in the same region, we measure the effects of energy booms and busts on family formation outcomes. We argue that the first order economic impacts of those shocks are on the less-educated men who dominated the mining industries, therefore under certain identifying assumptions we are able to interpret the results as the causal effects of less-educated men's earnings change on marriage and fertility.

We find evidence that is consistent with the unmarriageable men hypothesis. During both coal bust and oil bust in the 1980s, we find significant decline in marriage prevalence and marital birth rates but increase in non-marital birth rates. Our baseline estimates suggest that 10% decline in county earnings per capita reduced the marital birth rate by 6% and 9% in big coal-producing and oil-producing counties respectively, relative to non-coal and non-oil counties in the same region. Ten percent decline in earnings per capita, however, led to 7% and 5% increase in non-marital birth rates in coal and oil counties respectively.

We also find evidence that is at odds with an implication of the unmarriageable men hypothesis. If the disappearance of jobs led to a decline in marriage and more non-marital births, wouldn't reversing the economic trends for the less-educated men slow down the retreat from marriage and reduce non-marital births? We find evidence that confirms this implication in Appalachian coal region during the coal boom, which came after decades of decline in the region. However, we find conflicting results in the Southern and Midwestern oil-producing regions during the 1970s oil boom. Our estimates suggest that 10% increase in oil and gas related earnings is associated with 7% increase in marital fertility rate and 14% *increase* in non-marital fertility rate. This pattern is robust to various specification and robustness checks. The evidence suggests asymmetry in marriage and fertility responses to local economic shocks—reversing a declining economy may not undo all the damages caused. The contrast between coal and oil regions is also consistent with a differential social norms explanation.

This paper is structured as follows: Section (1.2) introduces the energy shocks and documents their differential impacts on local economies. Section (1.3) lays out our empirical strategy. Section (1.4) provides detailed information about the data and measurement. Section (1.5) presents and interprets our main results. Section (1.6) discusses several concerns about the main results, and Section (1.7) concludes with a summary of our findings.

1.2 The Impacts of Coal and Oil Booms and Busts on Local Economies

1.2.1 Energy Booms and Busts in the 1970s and 1980s

Due to regulatory changes and the Organization of Petroleum Exporting Countries (OPEC) oil embargo, there was a large and lasting increase in coal prices during the 1970s, followed by a sustained decline in the price of coal since early 1980s due to reduced demand. The coal boom and bust had dramatic and lasting economic impacts on Appalachian regions of Kentucky, Ohio, Pennsylvania, and West Virginia. Areas with coal reserves saw increasing employment and earnings while nearby areas without coal reserves suffered economic downturn along with other parts of the U.S. economy. The left panels of Figure (1.1) show the real price of coal and the number of national employment in coal industry from 1965 to 1990. A 20% increase in coal price between 1969 and 1970 brought coal price to a post-World War II high, but the price increase due to the 1973 OPEC oil embargo was even larger. The real price of coal increased about 70% between 1973 and 1974, and then remained relatively stable until 1979, when it began a steady decline that continued through the 1980s and 1990s. Indeed, between 1979 and 1993, the real price of coal declined in each year. Closely following the changes in coal price, the national total employment in coal industry experienced huge increase from 1969 to 1977, stayed stable for a few years, and started to decline in 1983. For this paper, we define the coal boom as the period of consecutive increase in national employment (1969–1977), and the coal bust as the period of consecutive decline in national employment (1983–1988).

There were similar oil and natural gas price cycles in some other parts of the U.S. From 1970 to the 1990s, there were at least two large exogenous shocks to the world oil supply: the 1973–74 OPEC oil embargo following the Yom Kippur War, and the period from the end of 1979 to early 1981, following the overthrowing of the Shah of Iran and the start of the Iran-Iraq War. These events affected both the real prices of oil and natural gas, as well as the employment in these industries in the United States, as illustrated in the right panels of Figure (1.1). Real oil price doubled between 1973 and 1974, stayed stable for several years, then tripled over a three-year period. Price fell sharply over the next 5 years to the level in the mid-1980s that were slightly lower than those of the mid-1970s. Real natural gas price followed a very similar pattern to that for oil price: a six-fold increase between 1970 and early 1980s, then a decline of more than half over the next 6 years. The national employment in the oil and natural gas industries closely tracked the movement in prices. We define the oil boom as the period of consecutive increase in national employment starting in 1973 oil crisis (1973–1982), and the oil bust as the period of consecutive decline in national employment (1983–1988).[2]

A key advantage of using exogenous coal and oil shocks to study the link between less-educated men's economic prospects and family formation decisions is that the local coal and oil boom and bust cycles had first order effects on the potential earnings of less-educated men. Table (1.1) presents characteristics of all U.S. workers, workers in the oil and gas industry, and workers in the coal industry in 1970, 1990, and 2010. The numbers for 1970 and 1990 are computed using the Public Use Micro Sample of the Decennial Censuses in those years. The 2010 numbers are computed using 2008–2012 American Community Survey 5-year data files (Ruggles et al., 2015). It is well known that female labor participation rate has been increasing for decades since

2. According to Figure (1.1), the national employment in oil and gas industry started to grow since 1972. Here we use 1973 as the first year of oil boom because of 1973 first oil crisis, a watershed year for oil cycles. Using 1972–1982 as oil boom does not change our results in later sections.

Figure 1.1. Energy Shocks in 1970s and 1980s

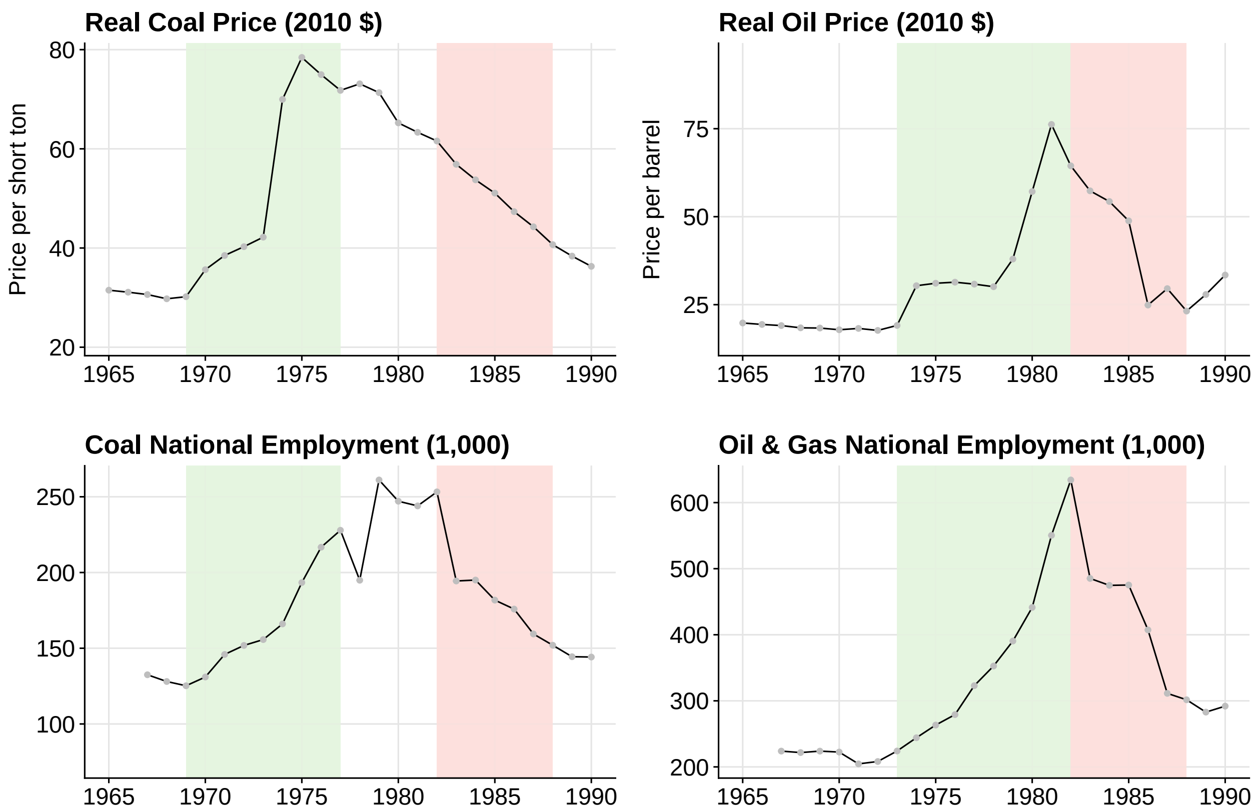

Notes: Prices are from the Energy Information Administration's Annual Energy Review. Coal prices are total price per short ton; oil prices are the US average first purchase price per barrel. National coal, oil, and natural gas industry employment are from County Business Patterns (CBP). We define coal boom as 1969–1977, coal bust as 1982–1988, oil boom as 1973–1982, oil bust as 1983–1988.

Table 1.1: Characteristics of Workers by Industry

	All Workers			Oil & Gas Workers			Coal Workers		
	1970	1990	2010	1970	1990	2010	1970	1990	2010
Age									
25–29	14	17	13	15	15	16	11	8	12
30–39	24	33	24	26	40	28	21	40	23
40–49	26	25	26	29	24	24	30	31	23
50 and over	36	26	36	31	21	33	38	21	41
Education									
Less than high school	41	13	8	38	13	11	72	25	11
High school	34	33	33	30	33	43	23	47	58
Some college	12	29	25	13	26	22	4	20	21
College and above	14	25	34	18	28	24	1	8	10
Sex									
Male	62	54	53	88	81	85	98	94	94
Female	38	46	47	12	19	15	2	6	6
Race									
White	89	83	77	97	90	86	96	97	95
Black	10	10	11	2	5	6	3	2	2
Other	1	7	12	¡1	6	8	¡1	2	2
Occupation									
Managerial and Professional	22	28	33	24	30	27	4	9	10
Technical, Sales, and Administrative	26	31	28	18	21	16	4	7	7
Precision Production, Craft, and Repairers	14	12	10	48	32	41	67	55	57
Operatives and Laborers	22	15	12	9	16	15	24	27	21
Service	12	12	15	1	1	1	1	2	5
Farming, Forestry, and Fishing	4	3	2	¡1	¡1	¡1	¡1	¡1	¡1

Notes: Author's compilations from the 1970, 1990 Public Use Micro Sample of the Decennial Censuses, and 2008–2012 ACS 5-year data files. Census's industry coding for oil and natural gas industry only include the extraction sector. Sample consists of individuals of age 25 and above who have worked positive weeks during the last year. Shares may not add to 100 because of rounding.

the beginning of the 20th century. By 1970, 38% of employed workers were women. In contrast, only 12% of oil workers and 2% of coal workers were women in 1970. Twenty years later in 1990, when the national work force saw almost half-half split between men and women, the oil and coal industries had not changed much in terms of gender composition, only slightly up at 19% and 6% women respectively.

Other features of the mining workforce show that coal and oil workers are significantly less skilled than the average of other industries. In 1970, 1% of the coal workers had a college degree, and by 1990, the share is only 8%, less than a third of national average. Oil workers had a similar education composition to all workers in 1970 but started to lag behind since 2000s. Occupational composition also indicate that the mining workers are less-skilled than average. Mining workers were much less likely to work in managerial, technical, and professional occupations, which typically requires more skills. Also note that coal and oil workers are predominantly white, at over 90% for most years in study period. Therefore, we do not study family formations separately by racial groups.

1.2.2 Uneven Geographical Impacts of Booms and Busts

The economic impacts of the coal and oil shocks on the local economies were not spread evenly. Using coal reserves data form Black, McKinnish and Sanders (2005), Figure (1.2) displays all the counties within the four coal states by coal reserves. In the figure we group counties into 3 buckets: the darkest shaded counties have at least a billion tons of coal reserves; the lightest shaded counties have less than 100 million tons of coal reserves or have no coal reserves at all; other counties have moderate coal reserves from 100 to 999 million tons. The figure shows 3 major coal seams in the region. The first major seam lies in the Appalachian Basin, including eastern Kentucky, most of West Virginia, eastern Ohio, and western Pennsylvania. The second major coal seam is located in eastern Pennsylvania. And the third major seam is at western Kentucky, a part of the Illinois coal basin. The coal boom and bust affected the counties that had large coal reserves and large coal industries but did not have much impact on counties that had no coal to mine. By comparing affected and unaffected counties in the region, we measure the effects of the coal boom and bust on family formation outcomes.

Figure 1.2. Coal Reserves: Kentucky, Ohio, Pennsylvania, and West Virginia

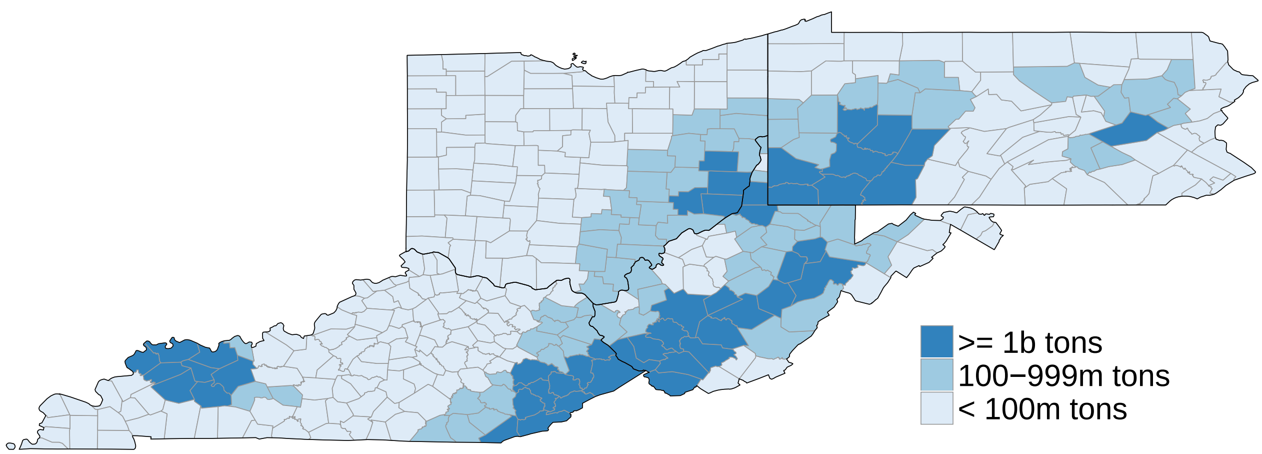

Notes: County-level coal reserves data from Black, Daniel and Sanders (2002). See their data appendix for details of how coal reserves data is constructed.

To focus on the counties that were deeply affected by the coal boom and bust cycle, following Black, McKinnish and Sanders (2005), we define the "treated" counties to be those coal-producing counties that derived at least 10% of their total earnings from the coal industry in 1969. Note that it is important to measure the "exposure" to treatment or propensity of treatment before the treatment actually happened. Of the total 34 treatment counties, 15 are in Kentucky, 13 are in West Virginia, 4 are in Pennsylvania, and 2 are in Ohio. The median fraction of earnings from the coal industry is 24% and the mean is 29%. In the left panel of Table (1.2), we summarize the average annual growth in real mining earnings and mining employment in the treatment counties. We use county-level earnings and employment data from Regional Economic Information System (REIS) data provided by the Bureau of Economic Analysis (BEA, US Department of Commerce). Here we casually interpret the difference in logarithm of earnings and employment as percentage change. During the coal boom (1969–1977), the treatment counties experienced 12% annual growth in real mining earnings and 7% annual growth in mining employment on average, which was extraordinary economic performance by any standard, especially during a period of stagflation. During the coal bust (1983–88), real mining earnings declined by 6% annually and mining employment declined by over 8% annually on average in the treatment counties. During the relatively stable period from 1978 to 1982, treatment counties saw small but statistically insignificant declines in both real mining earnings and mining employment.

In order to assess causal impacts of the coal boom and bust to the treatment counties, we need a proper comparison group. We want to avoid the comparison between big coal-producing counties with moderate or small coal-producing counties, because the categorization based on imprecise measures of earnings or employment from coal industry can be arbitrary. Instead, we limit the comparison counties to be

Table 1.2: Fast Growth During the Boom and Fast Decline During the Bust; Treatment Counties, 1969–88

Average annual growth in:	Treatment (Coal)			Treatment (Oil)		
	N = 646	Coef/SE	P-value	N = 2341	Coef/SE	P-value
Mining earnings						
Boom:	1969-77	0.119		1974-82	0.111	
		(0.013)	[0.000]		(0.008)	[0.000]
Bust:	1983-88	-0.060		1983-88	-0.122	
		(0.015)	[0.000]		(0.010)	[0.000]
Stable:	1978-82	-0.019		1969-73	-0.013	
		(0.017)	[0.257]		(0.012)	[0.286]
Mining employment						
Boom:	1969-77	0.069		1974-82	0.091	
		(0.009)	[0.000]		(0.006)	[0.000]
Bust:	1983-88	-0.084		1983-88	-0.070	
		(0.011)	[0.000]		(0.007)	[0.000]
Stable:	1978-82	-0.006		1969-73	-0.024	
		(0.012)	[0.601]		(0.009)	[0.007]

Notes: Author's calculations based on county-level earnings and employment data from Regional Economic Information System (REIS) data provided by the Bureau of Economic Analysis (BEA). This table shows the average of year-to-year differences in natural log of earnings from mining, and employment in mining. Numbers in parantheses are standard errors; numbers in brackets are p-values.

those without any coal reserves. In addition, since the large coal-producing counties tend to be rural and thinly populated, we also restrict the population range of the comparison counties to be between 8 to 225 thousand, similar to the population range of the treatment counties. This results in a comparison sample of 139 counties: 58 are in Kentucky, 11 are in West Virginia, 17 are in Pennsylvania, and 53 are in Ohio. These comparison counties are affected by similar cultural and institutional trends and restrictions as the treatment counties.

Figure (1.3) shows the geographical distribution of the treatment and comparison counties for the coal states. Compared with the coal reserves map in Figure (1.2), we can confirm that all treatment counties are in the major coal seams. And since we exclude the moderate and small coal-producing counties from the comparison group, we see a buffer area (unshaded counties in Figure (1.3)) separating the treatment and comparison counties. This buffer area would alleviate the concerns about geographic spillover from treatment to comparison counties.

In Table (1.3) we compare the local economic conditions between the treatment and comparison counties during the booms and busts. We estimate the following multiple-period differences-in-differences model:

$$\Delta E_{cst} = \sum_{k=1}^{3} \beta_k \text{Treatment}_c \times \text{Period}_k + \boldsymbol{\theta} \left(\text{State}_s \times \text{Year}_t\right) + \boldsymbol{X}_{cst}\boldsymbol{\beta}_1 + \varepsilon_{cst}, \quad (1.1)$$

where ΔE_{cst} is the year-to-year change in the natural log of county total earnings or total employment, which measures the growth or decline in local economic conditions; Treatment_c is a binary indicator of whether a county is in the treatment group, or otherwise in the comparison group; $\text{Period}_k, k = 1, 2, 3$ are the binary indicators of the boom, bust, and stable periods. $\boldsymbol{X}_{cst}$ is a set of control variables, including level and change of log population, change in the share of black population, change in the

Figure 1.3. Coal Treatment and Comparison Counties

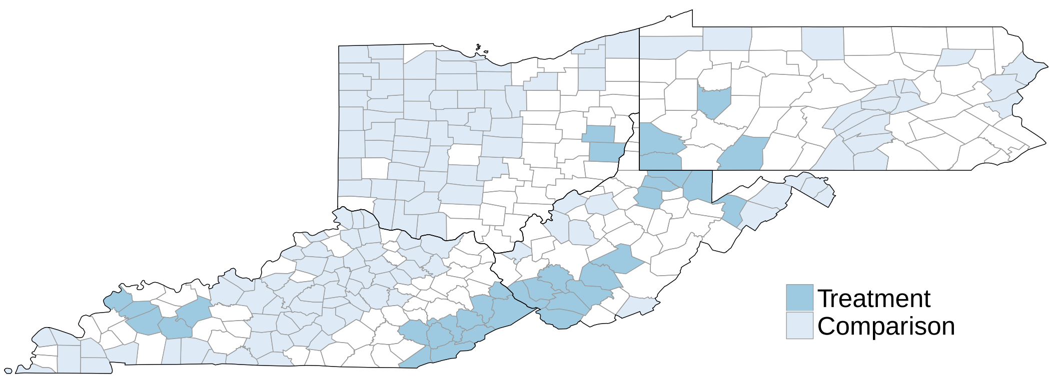

Notes: Treatment counties are defined as counties with 10% or more earnings from coal mining in 1969, the beginning of the study period. Teh comparision counties are counties with no coal reserves. Both treatment and comparison counties' populations are restricted to be between 8,000 to 225,000. There are 34 treatment counties and 139 comparison counties.

Table 1.3: Treatment Counties' Earnings and Employment Grew Faster During the Boom and Decline Faster During the Bust, Relative to Comparison Counties

Average annual growth in:	Treatment - Comparison (Coal)			Treatment - Comparison (Oil)		
	N = 3287	Coef/SE	P-value	N = 10220	Coef/SE	P-value
Total earnings						
Boom:	1969-77	0.046		1974-82	0.022	
		(0.005)	[0.000]		(0.006)	[0.000]
Bust:	1983-88	-0.049		1983-88	-0.047	
		(0.005)	[0.000]		(0.007)	[0.000]
Stable:	1978-82	0.000		1969-73	-0.003	
		(0.006)	[0.999]		(0.009)	[0.718]
Total employment						
Boom:	1969-77	0.019		1974-82	0.012	
		(0.002)	[0.000]		(0.002)	[0.000]
Bust:	1983-88	-0.028		1983-88	-0.022	
		(0.003)	[0.000]		(0.002)	[0.000]
Stable:	1978-82	0.000		1969-73	-0.003	
		(0.003)	[0.994]		(0.003)	[0.255]

Notes: Author's calculations based on county-level earnings and employment data from Regional Economic Information System (REIS) data provided by the Bureau of Economic Analysis (BEA). This table shows the average of year-to-year differences in natural log of total earnings and total employment. Numbers in parantheses are standard errors; numbers in brackets are p-values. Regressions include state by year fixed effects.

share of other race, change in the share of female population. We also include state by year fixed effects dummies to control for any state-specific nonlinear time trends in the local economies, which implies that the direct comparison is made within the state-year cells.

Table (1.3) presents the results of estimating Equation (1.1). In the left panel for coal states, we find that during the coal boom, large coal-producing counties' real total earning grew 4.6% faster than the average comparison county, and total employment grew 2% faster. Conversely, during the bust, real earnings in treatment counties dropped 5% faster than comparison counties, and 3% faster decline for total employment. We also compare two groups of counties during the relative stable period of coal prices to serve as a placebo test, and we find no evidence of differential growth or decline between the treatment and comparison counties.

Now we turn from the analysis of coal to oil. We restrict our attention to 11 states in the South and Midwest, where over 1% of state employment worked in mining industries in the early 1970s and a substantial share of the mining establishments are related to oil and gas.[3] See Table (1.4) for the filtering criterion. Similar to the coal analyses, we define a treatment group and a comparison group. Since county-level oil reserves data is not available, we rely on the share of county total earnings from oil and gas extraction industry in 1969 (REIS data), as well as the share of total employment working in oil and gas extraction industry in 1970 (CBP data). We categorize a county as "treated" if the county had 10% or more earnings from oil and gas industry in 1969 *or* had 10% or more employment in oil and gas industry in 1970. We combine these two exposure measures primarily because there are many missing values for the earnings share measure in 1969 for oil states. We also impose population restrictions on the comparison counties to make appropriate comparison with the treatment counties. See Section (1.4) for more details in constructing these measures. These criteria yield 127 treatment counties and 411 comparison counties, as shown in Figure (1.4). Unlike the case in coal states, we have some comparison counties adjacent to one or more treatment counties. In Section (1.6) supplementary analysis, we exclude those comparison counties and check the robustness of the main results.

3. These 11 oil and gas states are Colorado, Kansas, Louisiana, Mississippi, Montana, New Mexico, North Dakota, Oklahoma, Texas, Utah, and Wyoming. Employment share calculated based on Census Bureau's County Business Patterns (CBP) data.

Figure 1.4. Eleven Oil & Natural Gas-Producing States

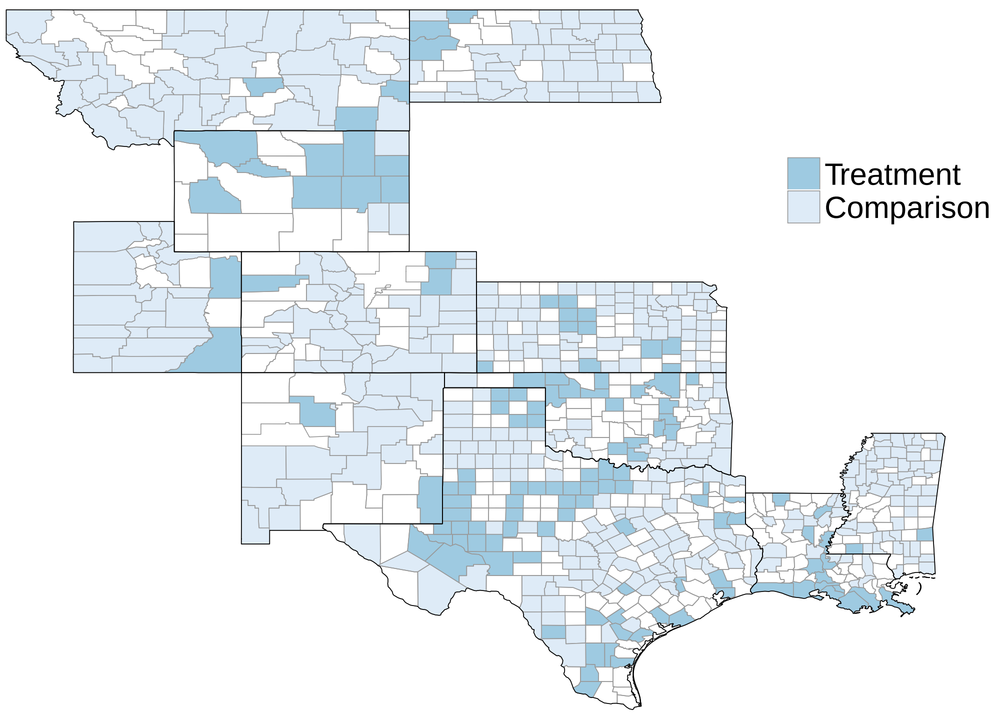

Notes: Oil states are Colorado, Kansas, Louisiana, Mississippi, Montana, New Mexico, North Dakota, Oklahoma, Texas, Utah, and Wyoming. Darker counties are the treatment counties, lighter counties are the comparision counties, and the unshaded counties are not used in our main sample. See the data section for the definitions.

The right panel of Table (1.2) documents the growth and decline of the 127 treatment counties during the oil boom and bust. Similar to the patterns found above about the big coal-producing counties, we find large average annual growth rate for both mining earnings (11%) and mining employment (9%) during the boom, and just as large declines during the bust. The relatively stable period for oil and gas cycle was 1969-73, the end of a decade-long slow decline of oil prices. We find small decline in mining employment during this period but no evidence for the mining earnings. The right panel of Table (1.3) compares the growth of treatment and comparison counties. The table shows similar stories to the coal region: during the boom, we find faster growth for big oil-producing counties than the comparison counties; and faster decline during the oil bust. Again, for the relatively stable period, there is no discernible difference in the growth in earnings and employment between two groups of counties.

1.3 Empirical Approach

Our objective is to assess the effects of men's earnings on marital and non-marital birth rates, we estimate the following equation:

$$\Delta y_{cst} = \beta_0 + \delta \left(\Delta E^*_{cst-1} \right) + \boldsymbol{X}_{cst-1} \boldsymbol{\beta}_1 + \boldsymbol{d}_{st} \boldsymbol{\beta}_2 + \varepsilon_{cst-1}, \tag{1.2}$$

where Δ indicates the first difference within the county; y_{cst} is the fertility outcomes for county c in state s during year t, defined as the number of live births per 1,000 women aged 15–44; $\boldsymbol{d}_{st}$ is a vector of state by year dummy variables for state s in year

t to allow for state-specific differences by year; $\boldsymbol{X}_{cst-1}$ is a vector of lagged control variables; and ε_{cst-1} are idiosyncratic factors of county c in year $t-1$.[4]

The regressor of interest, ΔE^*_{cst-1}, is change in natural log earnings between year $t-1$ and $t-2$. We use one-year lag because conception decisions are roughly one year earlier than birth. The variable is starred to denote the fact that we observe error-ridden versions of the actual variable. Given the log-log specification, the parameter on this regressor, δ, is the elasticity of fertility rate with respect to local economic conditions.

Some concerns arise in estimating using equation (1.2). One concern is that although the first difference form accounts for unobserved time-invariant differences in counties, first-differencing exacerbates the attenuation effects of measurement error (Bound, Brown and Mathiowetz, 2001). Also, as discussed in McKinnish (2007), year-to-year within-county variation often reflects transitory fluctuations that have little effect on long-term behavioral outcomes such as getting married and having babies. Although permanent earnings is the relevant concept for the marriage and fertility decision, the available noisy earnings measure contains both permanent and transitory components.

Another concern is that ΔE^*_{cst-1} is an endogenous regressor since changes in local economic conditions are likely correlated with changes in unobserved local factors that may affect marriage and fertility choices. The introduction section lists several of the potential confounders.

4. We control for the level and change in log population, level and change of the share of black population, level and change of the share of other race, level and change of the share of female. We also control for the level and change in county total receipt of AFDC. County-level AFDC receipts are mixed with many other programs in the REIS dataset. But AFDC is orders of magnitudes larger than other programs, we use it as a proxy for AFDC benefit receipts.

To circumvent the potential endogeneity and measure error problems, we isolate the labor demand component of the local economic changes by instrumenting for the change in county-level earnings. The instrumental variables capture both the cross-sectional variation in exposure to national shocks, as well as the time series variation in global coal and oil prices between 1969 and 1988. As these prices surge (decline), both fuel production as well as efforts to locate and extract new reserves increases (decreases), leading to higher (lower) earnings in the affected counties. After accounting for state-year interactions, our instrumental variables strategy exploits differences across counties within the same state, as counties with higher baseline employment shares in the mining industry experience larger earnings changes in response to fuel price movements, as we've shown in Section (1.2). Our baseline instruments for ΔE^*_{cst-1} are constructed as the interaction between the real fuel price changes and the importance of mining to a county, that is,

$$\Delta R_{t-1} \times \text{Mining Importance}_c, \tag{1.3}$$

where ΔR_{t-1} measures the real fuel price changes between periods $t-1$ and $t-2$, and Oil_c is the pre-existing importance of mining industries in the county as measured by coal reserves or oil and gas employment share at the beginning or prior to the study period. The resulting 2SLS model estimates equation (1.4), with the change in the local labor market measure replaced by the predicted value from the first stage equation

$$\Delta E^*_{cst-1} = \gamma_0 + \alpha_1 \Delta R_{t-1} \times \text{Mining Importance}_c + \boldsymbol{X}_{cst-1}\boldsymbol{\gamma}_1 + \boldsymbol{d}_{st}\boldsymbol{\gamma}_2 + u_{cst-1}. \tag{1.4}$$

An important advantage of using the predicted changes in earnings instead of observed year-to-year changes in earnings is that much of the variation in the observed annual earnings is transitory. The hypothesis we are examining, however, argues that the long-term structural decline of manufacturing and other good-paying industrial jobs for the less-educated men, not short-term fluctuations in wages or employment, is responsible for the decline in marriage and the rise in non-marital births. The predicted earnings and employment series are smoother than the noisy observed ones, due to reduced short-term noises in the changes in labor market conditions. This is largely due to the fact that the energy booms and busts lasted for many years, generating a relatively long-term change in the earnings potential of less-educated men.

1.4 Data Sources and Measurement

1.4.1 Local Earnings and Employment

Annual county-level earnings and employment data are from Regional Economic Information System (REIS) provided by the Bureau of Economic Analysis. Earnings include wage and salary disbursements, other labor income, and proprietors' income. The Bureau of Labor Statistics (BLS) compiles REIS wage and salary disbursements using ES-202 filings 5 collected as part of the state unemployment insurance program. The REIS data on county wages and employment are known to be measured with some error.[5] Earnings is top coded at $100,000 per year. Observations are excluded

5. Some employers with establishments in multiple counties may only report wages and employment ES-202 information at the state-level. These reports are allocated back to counties based on their industry level distribution by county among employers reporting at the county-level, generating some measurement error. In addition, components of other labor income and proprietors' income such as pension plan contributions, health and life insurance contributions, and private worker's

when the earnings per week variable is missing or earnings are less than \$126.875 per week in 2014\$ which is half the federal minimum wage in 2014. All dollar values are deflated by annual average CPI-U.

Population shares of different demographic groups are constructed using REIS data. County-level total welfare receipts is also from REIS.

1.4.2 Birth Outcomes

Data on birth outcomes are from the vital statistics natality files and the U.S. County-Level Natality and Mortality Data, 1915–2007, ICPSR study 36603. As per convention in the literature, fertility rates are calculated per 1,000 women aged 15–44. Women's population data is from the NIH Surveillance, Epidemiology, and End Results (SEER) Program hosted on the NBER website.[6]

We break down total births by county of residence into marital and non-marital births. Marital status of mothers after year 1988 are only available for counties with a population of 100,000 or more, which excludes most of counties that produce coal and oil. We therefore focus our study from 1969 to 1988.

Ohio, Montana, and New Mexico did not report the marital status of mothers until early 1980s. We exclude these 3 states from our main analysis. Texas did not have consistent reporting of marital status during 1977 to 1979, so we drop this time period for Texas.

In the baseline specification, we use the annual changes in the natural log of the birth rates (number of births per 1,000 women aged 15-44). Because the birth rate

---

compensation contributions are only collected at the state-level and also use an allocation rule to determine county-level totals. 6 County-level SSDI payment data was not published in 1981.

6. https://www.nber.org/data/seer_u.s._county_population_data.html, retrieved in October 2016.

measures are based on potentially unreliable county-level population estimates, we also use the natural log of the number of births as a specification check. The results are similar.

1.4.3 National Shocks and Local Exposure

We define the oil and gas states as those contiguous states where the mining industry comprises at least 1% of total employment in the 1974 County Business Patterns (CBP) and a substantial share of the mining establishments are related to oil and gas. We construct a measure of the initial importance of oil and gas industry in a county, oil_c, which is the share of employment in the oil and gas industry computed from the CBP data. These 11 oil and gas states are Colorado, Kansas, Louisiana, Mississippi, Montana, New Mexico, North Dakota, Oklahoma, Texas, Utah, and Wyoming. We exclude West Virginia since less than half of total mining employment is in the oil and gas industry.

Since early 1970s, CBP data have been based on the Census Bureau's Standard Statistical Establishment List. Because of the risk of disclosing firm specific information, exact employment numbers for two-digit industries such as coal and oil and gas are not available at the county level. However, the CBP provides county-level information on both the number of firms in each two-digit industry and the number of firms that fall into a specific firm-size category (e. g., 20 to 49 employees) for these industries. By weighting the number of firms in a firm-size category by the midpoint of the number of employees in that category, we create an estimate of the number of employees in each two-digit industry at the county level. We then create county-level estimated employment shares by industry as the ratio of the estimated industry employment to the estimated total county employment, where the total county employment is also estimated by using the firm-size methodology.

Prior to 1970, we only observe one-digit total mining employment for each county in CBP. As the first large oil price shock occurs in late 1973 and our sample begins in 1970, we would prefer to determine the initial oil and gas importance in some year prior to 1974, if not prior to 1970. The tradeoff, however, is that with an earlier year we can only determine mining employment as opposed to oil and gas employment at the county-level. As shown in Table (1.4), over half of mining establishments in the oil and gas states, with the exception of Utah, are in the oil and gas industry with little change in these shares between 1967 and 1974. Thus, we construct the county-level initial oil and gas importance measure, Oil_c, using the 1970 CBP and we use the 1967 one-digit mining employment as robustness check.

Coal reserves data is from Black, McKinnish and Sanders (2005). We follow them to construct the change in the value of coal reserves instrument by multiplying the change of national coal prices by county-level coal reserves. We obtain data on energy prices from the Energy Information Administration's Annual Energy Review. Coal prices are total price per short ton; oil prices are the US average first purchase price per barrel. National coal, oil, and natural gas industry employment are from CBP. We define coal boom as 1969–1977, coal bust as 1982–1988, oil boom as 1973–1982, oil bust as 1983–1988.

1.4.4 Marriage Outcomes

Public use individual marriage outcomes in the census do not have county identifiers. The marriage analysis therefore use the county aggregates in 1970, 1980, and 1990 Decennial Census data to estimate the impacts of energy shocks on marriage

Table 1.4: Oil and Gas States

	Share of 1974 CBP State Employment in Mining	Oil and Gas Share of All Mining Establishments in: 1974	1967
Wyoming	15.80%	77%	82%
New Mexico	7.80%	82%	76%
Louisiana	5.90%	88%	91%
Montana	4.30%	66%	68%
Oklahoma	3.90%	90%	92%
Utah	3.70%	47%	32%
Texas	3.20%	88%	92%
Colorado	2.20%	58%	50%
Kansas	1.80%	85%	87%
North Dakota	1.20%	65%	61%
Mississippi	1.10%	81%	83%

Notes: CBP stands for County Business Patterns.

outcomes. County aggregate data is downloaded from IPUMS National Historical Geographic Information System (NHGIS). The data include the counts of married, never married (except 1980), and divorced men and women by age groups. In our main marriage analysis in Section (1.5.3), we use the counts of young women from age 20 to age 34 by marital status.

1.4.5 Sample Construction

Our main coal sample consists of Kentucky, West Virginia, and Pennsylvania. Ohio is dropped because of the birth outcomes reporting issues documented above. The treatment counties are defined as those coal-producing counties that derived at least 10% of their total earnings from the coal industry in 1969. Of the total 34 treatment counties in coal states, 15 are in Kentucky, 13 are in West Virginia, 4 are in Pennsylvania, and 2 are in Ohio (excluded from the main analysis). We want to avoid the comparison between big coal-producing counties with moderate or small coal-producing counties. So we limit the comparison counties to be those without any coal reserves. In addition, since the large coal-producing counties tend to be rural and sparsely populated, we also restrict the population range of the comparison counties to be between 8 to 225 thousand, similar to the population range of the treatment counties. This results in a comparison sample of 139 counties: 58 are in Kentucky, 11 are in West Virginia, 17 are in Pennsylvania, and 53 are in Ohio (excluded from the main analysis).

Since county-level oil reserves data is not available, we rely on the share of county total earnings from oil and gas extraction industry in 1969 (REIS data), as well as the share of total employment working in oil and gas extraction industry in 1970 (CBP data). We categorize a county as "treated" if the county had 10% or more earnings from oil and gas industry in 1969, *or* had 10% or more employment in oil and gas industry in 1970. We combine these two exposure measures primarily because there are many missing values for the earnings share measure in 1969 for oil states in the REIS data. We also impose a population range from 1,000 to 100,000 for the comparison counties, since most oil treatment counties' populations fall within this range.

Matching the data from the various sources was facilitated by the use of county FIPS codes. Merging was based on the modified FIPS codes used by the REIS. These county FIPS codes are generally the same as the standard FIPS codes with the exception that many independent cities in Virginia are merged with neighboring counties to create new counties.

To address a small number of counties which are involved in either a merge or a split during our sample period, we only use observations for these counties from or after the merge or split. Counties which split during the sample period are Yuma, AZ (creating La Paz, AZ in 1983) and Valencia, NM (forming Cibola, NM in 1981). Washbaugh, SD and Jackson, SD merged in 1976.

1.5 Main Results

1.5.1 First Stages and Reduced Forms

We have shown in Table (1.3) the differences-in-differences estimates of the impacts of energy shocks on local economies, using the variation from a binary indicator of treatment. Now we turn to a more parametric examination of the impacts of energy cycles on local economies. Table (1.5) displays the impact of energy price shocks on local earnings per capita conditional on local characteristics. Columns (1) and (2) regress the change in log earnings per capita on the instrument, which is the change in national coal prices multiplied by coal reserves. The regression includes state by year fixed effects as well as a full set of demographic and welfare receipts controls. We find strong positive impacts of the coal boom on earnings per capita, compared with counties without coal within state and year. The excluded instrument is reasonably strong according to the F-statistic. The coal bust had an even stronger negative impacts on local earnings per capita. Columns (3) and (4) present the first stage results for the oil and gas states. The instrument for oil regressions is the annual change in national crude oil prices multiplied by the share of county employment in oil and gas extraction industry in 1970. Similar to coal states, the big oil-producing counties' local economy experienced boom and bust along with the national oil prices, although the sizes of the estimates are not directly comparable between coal and oil results. Appendix Tables (A.1) and (A.2) present the same regression results for county employment per capita and county average annual wage. The impacts of energy cycles on local employment and wage are also large.

Building on the evidence that energy shocks have a positive effect on the local economy, especially the economic prospects of less-educated men, we now turn to the fertility responses. We start by comparing the changes in fertility outcomes between fuel-producing regions and regions with little or no fuel production during the boom and bust. Figure (1.5) shows strong reduced form impact of energy shocks on marital birth rates. During the booms (left-hand panels), we see positive association between energy shocks and the growth rate in marital birth rates. Note that the mean growth rate of birth rates is negative for the majority of counties, which reflects the secular trends in the decline of fertility in the U.S. However, the energy shocks lead to slower decline in coal and oil regions. During the bust (right-hand panels), we observe faster decline of marital birth rate in coal and oil regions relative to regions that have no coal and oil production.

Table 1.5: Impact of Coal and Oil Prices Shocks on Local Earnings

States/Shocks:	Coal	Coal	Oil/Gas	Oil/Gas
Years:	Boom	Bust	Boom	Bust
	(1)	(2)	(3)	(4)
Earnings per Capita				
Instrument	.050	.100	.492	.504
	(.017)	(.027)	(.092)	(.071)
F-Statistic for Excluded Instrument	9.3	13.5	28.7	50.6
N	826	708	3,499	2,868

Notes: Each Column represents a regression. The instrument for coal regressions is the change in national coal prices multiplied by coal reserves. The instrument for oil regressions is the change in national crude oil prices multiplied by the share of county employment in oil and gas extraction industry in 1970. Standard errors in parentheses account for arbitrary forms of clustering within counties. Samples are restricted to where fertility rate is not missing. All regressions control for: state by year fixed effects, level and change in log population, level and change in log real total welfare benefits, level and change of the share of black population, level and change of the share of other race, level and change of the share of female. All regressions are weighted by the number of total births in 1969.

Figure 1.5. Larger Coal/Oil Counties' Marital Birth Rates Grew Faster During the Boom and Decline Faster During the Bust

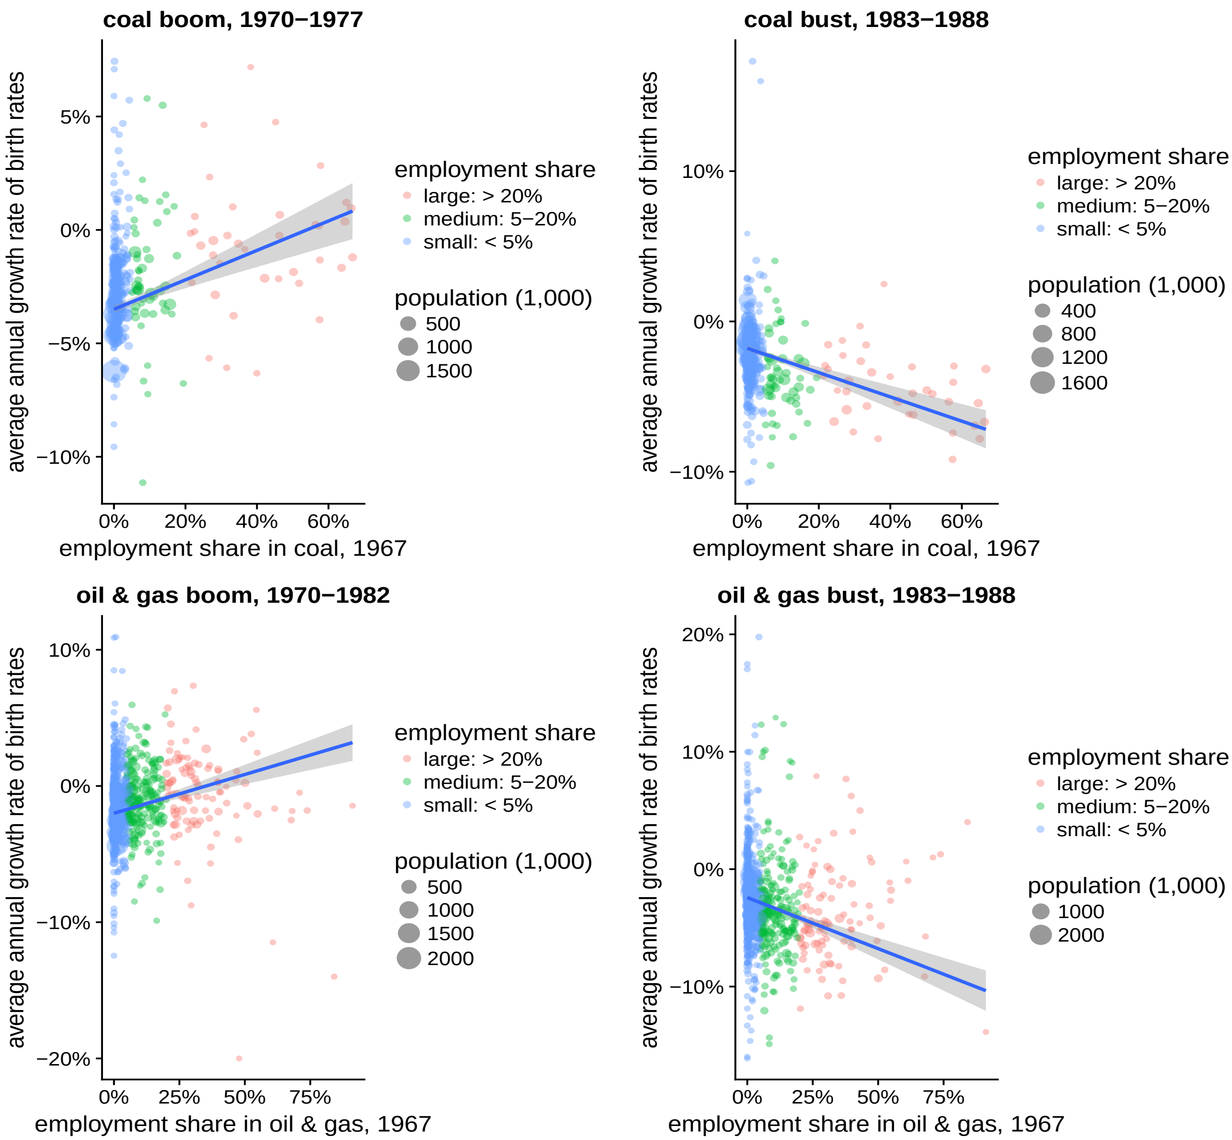

Notes: The horizontal axis represents counties' employment share in coal or oil&gas industries in 1967, calculated using CBP data. The vertical axis represents the mean annual growth rate of marital fertility rate. The straight lines and shaded 95% confidence intervals in each panel are from unweighted bivariate regressions.

Now we turn to Figure (1.6) to examine the reduced form impact on non-marital birth rates. Unlike the consistent patterns across coal and oil regions in terms of martial birth responses, we find evidence for differential responses to exogenous shocks in coal regions versus oil and gas regions. Specifically, we find decreased non-marital birth rates in coal-producing regions during the coal boom (upper left-hand panel). But we find no significant difference in non-marital births responses between oil and non-oil region during oil and gas boom (bottom left-hand panel). This finding is consistent with the findings in Kearney and Wilson (2017).

1.5.2 Fertility Outcomes

We use parametric 2SLS model to further explore the fertility responses. Panel A in Table (1.6) shows the estimated elasticities of total birth rates with respect to real earnings. The estimated elasticities are consistently positive, large, and statistically significant for both coal and oil regions, booms and busts alike. This suggests that increased earnings tend to increase total birth rate: 10% increase in earnings per capita leads to 3.3% increase in total birth rate in coal-producing counties relative to non-coal counties, and 10% decrease in earnings per capita leads to 4.6% decline in total birth rate. The fertility responses were even larger in the oil-producing counties.

If we break down the total births into marital and non-marital births in Panel B and Panel C, we observe different patterns across regions and time for non-marital births, confirming the visual evidence we found above. The marital birth elasticities are largely consistent with the total birth rates but larger in magnitudes. Possibly because the total birth elasticities are attenuated by the negative elasticities of non-marital birth.

Figure 1.6. The Decoupling of Marriage and Fertility in Oil and Gas Regions

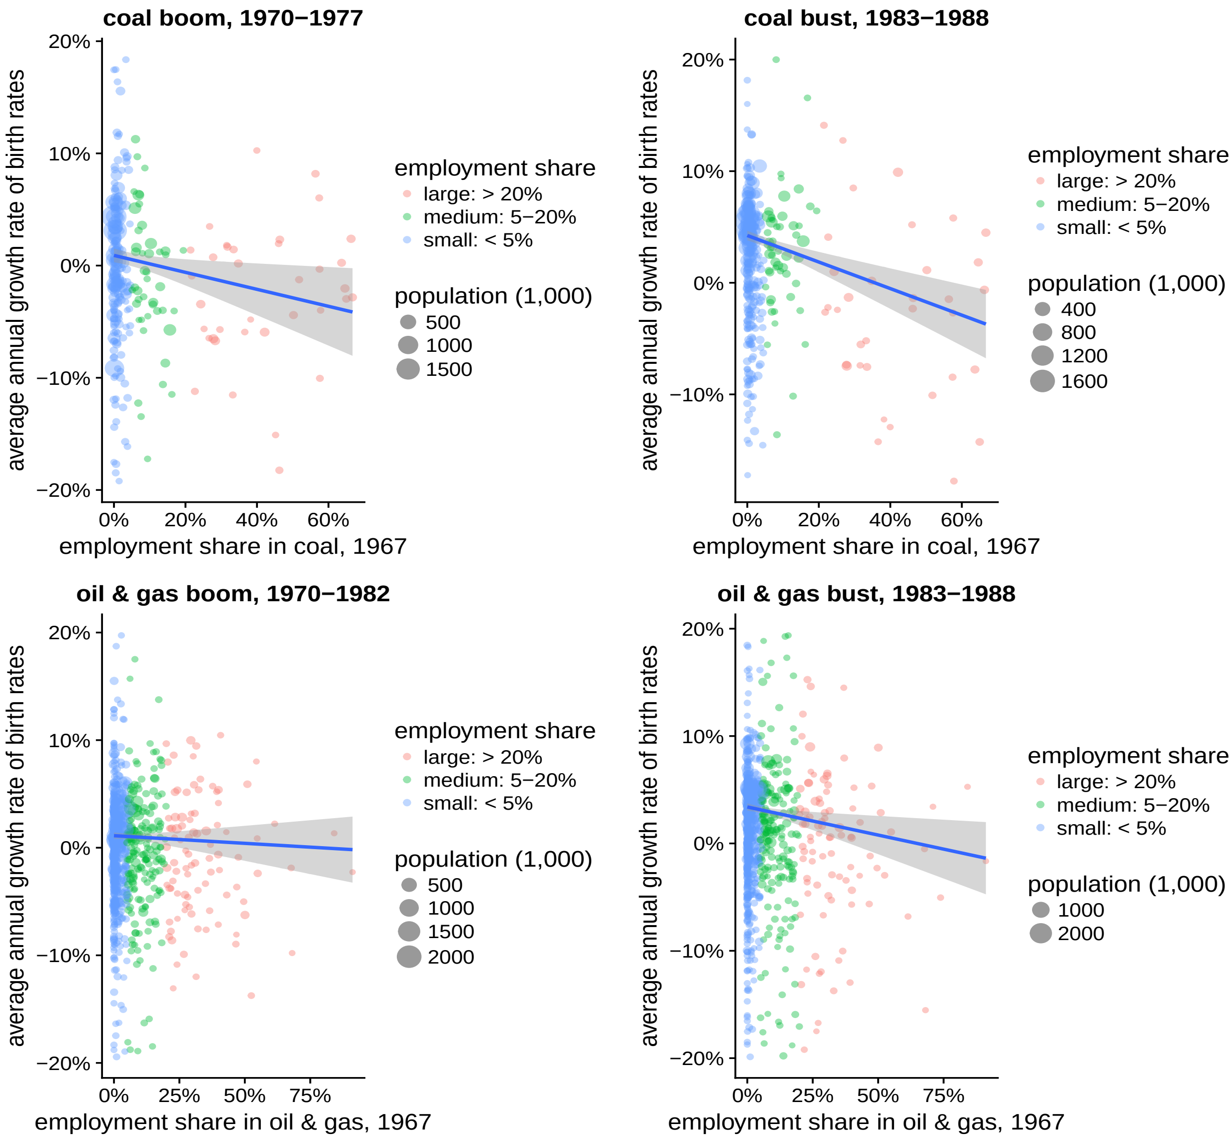

Notes: The horizontal axis represents counties' employment share in coal or oil&gas industries in 1967, calculated using CBP data. The vertical axis represents the mean annual growth rate of non-marital fertility rate. The straight lines and shaded 95% confidence intervals in each panel are from unweighted bivariate regressions.

Table 1.6: 2SLS Estimates of the Impact of Earnings on Fertility

States/Shocks: Years:	Coal Boom (1)	Coal Bust (2)	Oil/Gas Boom (3)	Oil/Gas Bust (4)
	A: Total Birth Rate			
Earnings per capita	.334 (.141)	.460 (.133)	.703 (.203)	.649 (.157)
F-Statistic for Excluded Instruments	9.3	13.5	28.7	50.6
N	826	708	3,499	2,868
	B: Marital Birth Rate			
Earnings per capita	.503 (.186)	.628 (.167)	.699 (.206)	.891 (.176)
F-Statistic for Excluded Instruments	9.3	13.5	28.7	50.6
N	826	708	3,499	2,868
	C: Non-marital Birth Rate			
Earnings per capita	-1.424 (.564)	-.692 (.289)	1.413 (.552)	-.509 (.282)
F-Statistic for Excluded Instruments	9.3	13.5	28.7	50.6
N	826	708	3,499	2,868

Notes: Each cell represents a regression. Standard errors in parentheses account for arbitrary forms of clustering within counties. All regressions control for: state by year fixed effects, log population, change in log population, change in log real total welfare benefits, share of black population, share of other race, share of female, share of population aged 30s, 40s, 50s, 60s, 70 and up. All regressions are weighted by the number of total births in 1969.

For the coal states, the estimated non-marital birth elasticities for both boom and bust are sizable and significantly negative, which implies that when local economy performs well, people tend to have fewer children outside marriage, or slower increase in out-of-wedlock births compared with less prosperous places. During the boom period, the estimated elasticity is -1.4, which translate to over 14% decline in non-marital birth rate for 10% increase in earnings per capita. In Section (1.2) we estimated that on average treatment coal counties grew 4.6% faster in total earnings than non-coal counties over the period of 1969–77, which is about 40% more increase than non-coal counties during the period. During the coal bust, the estimated non-marital elasticity is -0.7, also sizable and significant.

For the oil states, we find similar non-marital fertility response during the bust but opposite responses when economic conditions *improved*. The estimated non-marital elasticity is -0.5, similar in size to the coal counterpart and marginally significant. The estimated non-marital birth elasticity during the boom, however, is 1.4—the opposite sign and large in size. This means the oil-producing counties were having more children both inside and outside marriage during the oil boom, relative to the comparison counties. Using employment instead of earnings to measure the change in local economic conditions, Table (1.7) confirms this finding. It challenges the implication of the "unmarriageable men hypothesis": if the disappearance of jobs led to the decline in marriage and more non-marital births, wouldn't some long-term positive economic shock primarily on less-skilled men slow down the decline in marriage and reduce the non-marital births? This implication is confirmed in the Appalachian coal region but not in the oil-producing regions in the South and Midwest.

Table 1.7: 2SLS Estimates of the Impact of Employment on Fertility

States/Shocks: Years:	Coal Boom (1)	Coal Bust (2)	Oil/Gas Boom (3)	Oil/Gas Bust (4)
	A: Total Birth Rate			
Employment per capita	.657 (.275)	.859 (.275)	1.436 (.414)	1.055 (.275)
F-Statistic for Excluded Instruments	14.2	11.2	25.4	48.4
N	826	708	3,499	2,868
	B: Marital Birth Rate			
Employment per capita	.992 (.355)	1.171 (.355)	1.426 (.428)	1.448 (.306)
F-Statistic for Excluded Instruments	14.2	11.2	25.4	48.4
N	826	708	3,499	2,868
	C: Non-marital Birth Rate			
Employment per capita	-2.807 (1.008)	-1.291 (.571)	2.816 (1.056)	-.824 (.449)
F-Statistic for Excluded Instruments	14.2	11.2	25.4	48.4
N	826	708	3,499	2,868

Notes: Each cell represents a regression. Standard errors in parentheses account for arbitrary forms of clustering within counties. All regressions control for: state by year fixed effects, log population, change in log population, change in log real total welfare benefits, share of black population, share of other race, share of female, share of population aged 30s, 40s, 50s, 60s, 70 and up. All regressions are weighted by the number of total births in 1969.

1.5.3 Marriage outcomes

One explanation of the non-response of out-of-wedlock birth rates during the oil boom is that non-marital birth rate has two components, the women's marriage rate and the average fertility inside and outside marriage. It could be true that marriage rate did increase during the oil boom, but the average fertility among unmarried women grew faster relative to the average fertility among married women, leading to an increase of non-marital birth rate. In this section, we examine the marriage outcomes directly.

Unfortunately, annual county-level data for marriage and other household structures is not available. Since the booms happened primarily in the 1970s and the busts were mostly in the 1980s, we use the county aggregates in 1970, 1980, and 1990 Decennial Census data to estimate the impacts of energy shocks on marriage outcomes. The Decennial Censuses do not reveal individual county-level data in many mining regions because the counties are too small. The Public Use Micro Samples identifies the location of respondents only at the county group level in 1970 and 1980, and at PUMA level in the 1990 Census. In 1970, county groups were collections of "counties" in which at least 250,000 people resided. For the 1980 and 1990 Censuses, county groups or PUMAs were collections of counties in which at least 100,000 people resided. Instead of micro data, we use county-level aggregate counts of women by marital status, obtained from IPUMS National Historical Geographic Information System (NHGIS).

Similar to Autor, Dorn and Hanson (2017), our marriage outcome analysis uses long decadal change between Censuses instead of annual change. Although motivated by the data limitation, the long differences may help in addressing a few concerns. One concern is that year-to-year labor market changes associated with energy shocks does not enter into the long-term decision of marriage and child-rearing. Another concern is that women need time to observe and update their expectations of men's economic situations. The long difference could help us tease out the relatively long-term variations and discard transitory noise.

In Table (1.8), we assess the causal effect of energy shocks on the 10-year change in the fraction of married young women aged 20–34. We fit the following model:

$$\Delta y_{cst} = \beta \Delta x_{cst} + \boldsymbol{X}'_{it}\delta + \varepsilon_{cst}, \tag{1.5}$$

where Δy_{cst} is the county-level decadal change in percent married of women aged 20-34; Δx_{cst} is the change in log earnings per capita; and $\boldsymbol{X}_{it}$ is a vector of control variables, including state fixed effects, change in log population, change in share of black population, change in share of other race, and change in share of female. We estimate both the OLS and 2SLS models for comparison. The difference between the OLS and 2SLS estimates suggest potential endogeneity problems. We also limit the sample to be the treatment and comparison counties. For coal states, we find consistent results with previous findings: marriage prevalence responded positively to earnings changes, that is, when earnings increase (decrease), marriage prevalence increases (decrease). In oil states, however, we again observe asymmetric marriage responses during the oil boom and bust. During the relatively long-term oil decline in the 1980s, percent of married young women decreased—consistent with the unmarriageable men hypothesis. But we find no discernible marriage response for young women during the oil boom in the 1970s.

We also estimate a differences-in-differences model to directly compare the change in young women's marriage prevalence in the treatment counties and the comparison counties. Table (1.9) shows the DID estimates. Relative to the coal comparison counties, the treatment counties in coal states experienced 6.4 percentage point faster

Table 1.8: OLS and 2SLS Estimates of the Impacts of Changes in Log Earnings Per Capita on Changes in Percent Marriage for Women Aged 20–34

538

Notes: Marriage outcomes come from National Historical Geographic Information System (NHGIS). Each cell represents a regression. Standard errors in parentheses account for arbitrary forms of clustering within counties. All regressions control for: state fixed effects, change in log population, change in share of black population, change in share of other race, change in share of female, change in share of population aged 30s, 40s, 50s, 60s, 70 and up. All regressions and dependent means are weighted by county population in 1969.

increase in the percent married of young women. Given that from 1970 to 1980, an average county's fraction of married young women actually decreased by 7 percentage point, the 6.4 percentage point faster increase in the treatment counties is a large effect on marriage. From 1980 to 1990, however, the coal treatment counties' fraction of married young women decreased 1.5 percentage point faster than the comparison counties. Similarly, for the percent divorced outcome in the bottom panel, we see big marriage response during both the boom and bust in coal-producing counties, which is consistent with the unmarriageable men hypothesis. For big oil-producing counties, we do observe higher fraction of married young women relative to non-oil counties, albeit the size of the marriage response is much smaller during the boom than the response in coal counties.

Table 1.9: Changes in Percent Married and Percent Divorced of Women Aged 20–34, Treatment and Comparison Counties, 1970–80 and 1980–90

10-year Difference in:	Treatment - Comparison (Coal)		Treatment - Comparison (Oil)	
	Coef/SE	P-value	Coef/SE	P-value
Percent Married				
1970-80	0.064		0.015	
	(0.008)	[0.000]	(0.005)	[0.004]
Dependent Mean	-0.070		-0.063	
1980-90	-0.015		-0.002	
	(0.004)	[0.000]	(0.004)	[0.604]
Dependent Mean	-0.081		-0.073	
N	173		538	
Percent Divorced				
1970-80	-0.011		-0.003	
	(0.002)	[0.000]	(0.002)	[0.040]
Dependent Mean	0.041		0.031	
1980-90	0.011		0.003	
	(0.002)	[0.000]	(0.002)	[0.064]
Dependent Mean	0.020		0.020	
N	173		538	

Notes: Marriage outcomes come from National Historical Geographic Information System (NHGIS). Each cell represents a regression. Standard errors in parentheses account for arbitrary forms of clustering within counties. All regressions control for: state fixed effects, change in log population, change in share of black population, change in share of other race, change in share of female, change in share of population aged 30s, 40s, 50s, 60s, 70 and up. All regressions and dependent means are weighted by county population in 1969.

1.6 Supplementary Analysis

1.6.1 Selective Migration

One concern about the identification is that people may respond to energy booms and busts by moving in or out of an affected county. If there were large in-migration and out-migration *and* the migrants have systematically different preference for marriage and fertility from existing residents, our identification strategy may not be valid. The literature on regional adjustment to localized labor-market demand shocks suggests that mobility responses are sluggish and incomplete (Blanchard and Katz, 1992; Glaeser and Gyourko, 2005; Artuc, Chaudhuri and McLaren, 2010). In addition, mobility is lower in level and slower in pace among less-skilled workers, lower when information about the local shocks is not readily available.

Black et al. (2013) use the 1980 Census data to show that women aged 18 to 34 residing in coal areas in 1980 were less likely to be new residents than their counterparts in non-coal areas. This evidence is consistent with the claim that there was no large influx of young women into large coal-producing counties during the coal boom, alleviating the concern of selective net migration.

1.6.2 Spillover

In Figure (1.4), we define two groups of counties in oil and gas states: 127 treatment counties and 411 comparison counties in 11 states. By comparing the average outcomes of the treatment counties and the comparison counties, we make causal statement about the effects of oil boom and bust on the family formation outcomes. One might be concerned about the possibility that the shocks experienced by the treatment counties could spillover to the neighboring comparison counties, therefore *understate* the impacts of oil cycles. If we could construct some buffer area between the treatment and comparison counties, we could feel more confident that geographic spillover from the treatment to the comparison was limited.

In Figure (1.7) we compare the maps before and after we exclude 112 adjacent comparison counties. The bottom graph shows the geographically segregated distribution of treatment and comparison counties. Using this new comparison group, we re-run the 2SLS regressions, and found similar results. See Appendix Table (A.3).

Geographic spillover is also possible via state government's redistribution across counties. In many of the coal and oil states, severance taxes imposed on coal mining and oil and gas extraction account for a large fraction of state revenues. Figure (1.8) shows the severance taxes as a share of state total revenues in 1980. Oil states Wyoming, Oklahoma, New Mexico, Texas, Louisiana, and North Dakota had more than 10% of state tax revenues from severance taxes. State governments might then redistribute tax revenues to counties via public transfer programs, according to counties' economic situations. To account for the potential redistribution, we control for the level and change in county total receipt of AFDC, unemployment insurance, and education and training assistance. These programs are the largest public welfare programs that are partly funded by state revenues. The inclusion of these control variables does not change the qualitative results.

Another concern is the spillover into other sectors, therefore spillover to women's wage as well as more-educated men. Such spillover would affect family formation and fertility outcomes, thereby contaminate our identification strategy. By comparing the growth rate of earnings and employment from mining and the total earnings and total employment, we could get a sense of how big the cross-sector spillover effects are. Table (1.10) presents the same regressions as in Table (1.2), except now we use total earnings and employment. The average growth rates are vastly different between

Figure 1.7. Exclude Comparison Counties That Are Adjacent to Treatment Counties

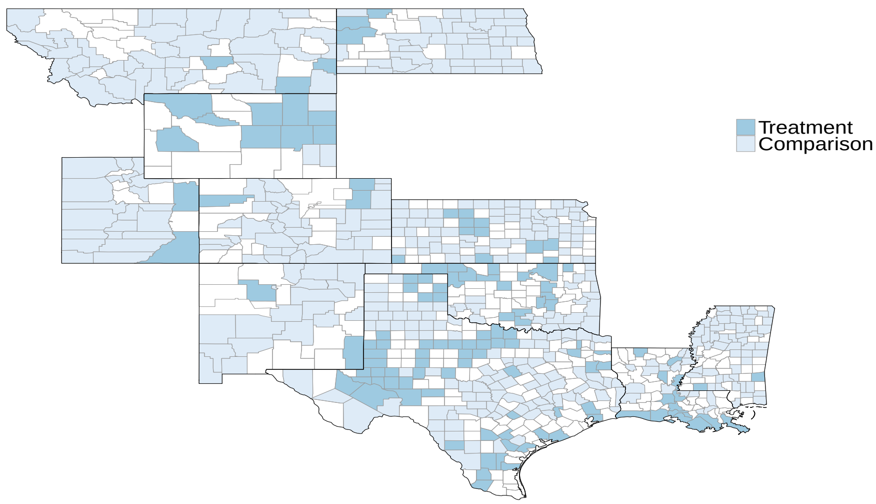

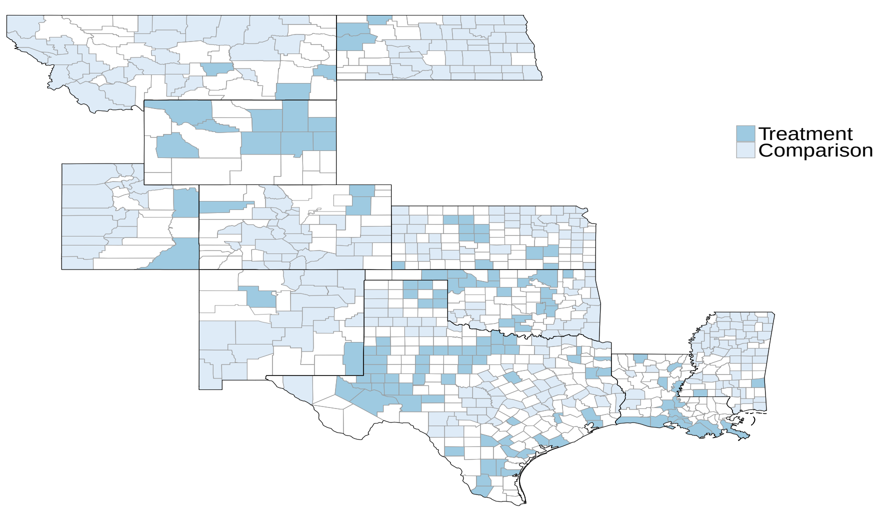

Notes: The top graph is the same as in Figure (1.4). The bottom graph exlucde all the comparison counties that are adjacent to one or more treatment counties.

Figure 1.8. Severance Taxes as a Share of Total Tax Revenues in 1980

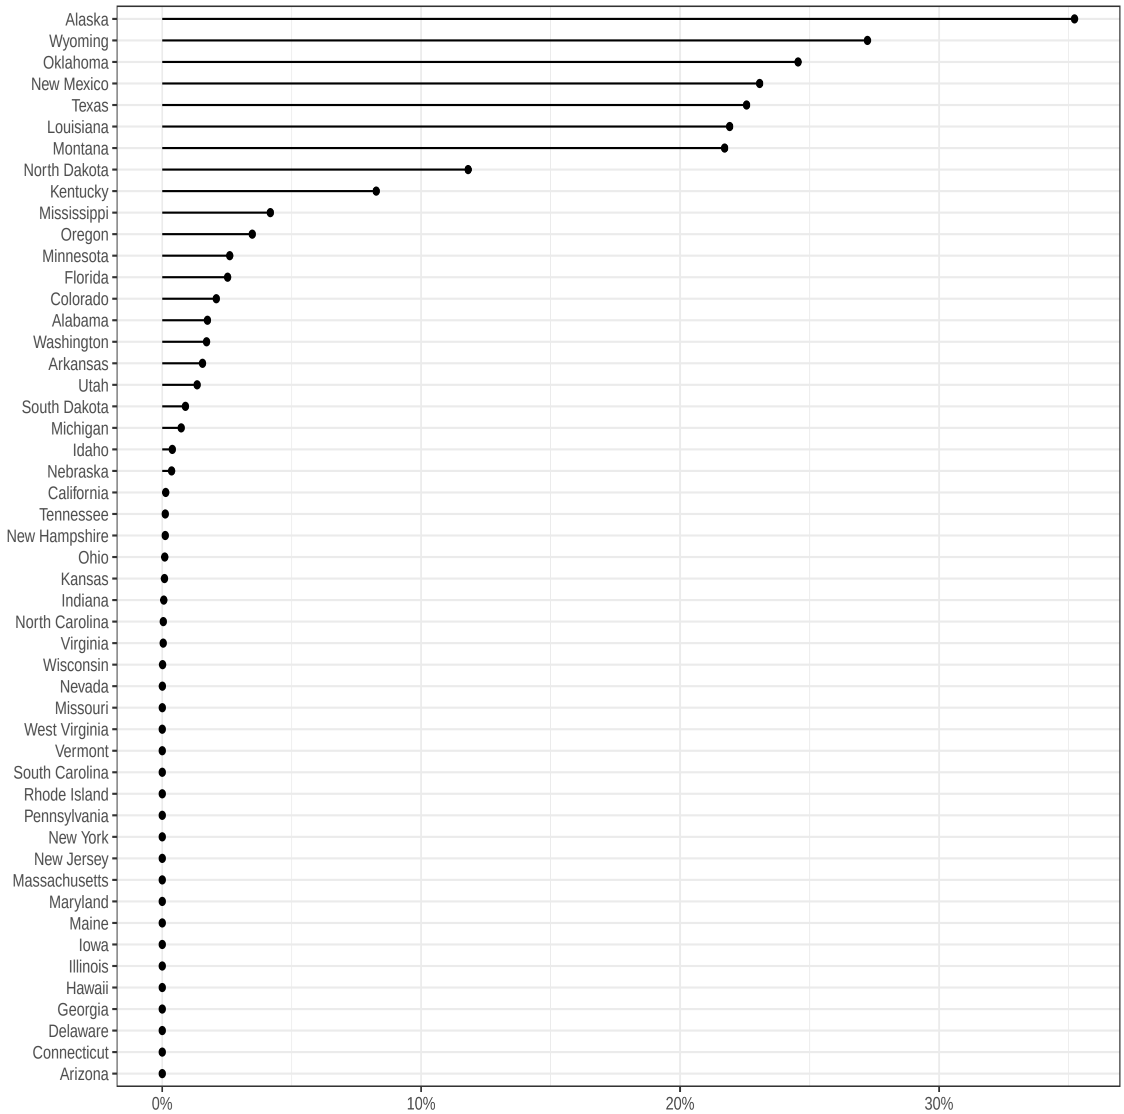

Source: Annual Survey of State Government Tax Collections (STC).

Table 1.10: The Growth of Total Earnings and Total Employment During the Energy Booms and Busts; Treatment Counties, 1969–88

Average annual growth in:	Treatment (Coal)			Treatment (Oil)		
	N = 646	Coef/SE	P-value	N = 2341	Coef/SE	P-value
Total earnings						
Boom:	1969-77	0.069		1974-82	0.016	
		(0.005)	[0.000]		(0.004)	[0.000]
Bust:	1983-88	-0.017		1983-88	-0.035	
		(0.006)	[0.004]		(0.005)	[0.000]
Stable:	1978-82	-0.017		1969-73	0.066	
		(0.007)	[0.010]		(0.006)	[0.000]
Total employment						
Boom:	1969-77	0.033		1974-82	0.032	
		(0.003)	[0.000]		(0.001)	[0.000]
Bust:	1983-88	-0.008		1983-88	-0.014	
		(0.003)	[0.018]		(0.002)	[0.000]
Stable:	1978-82	0.004		1969-73	0.012	
		(0.004)	[0.241]		(0.002)	[0.000]

Notes: Author's calculations based on county-level earnings and employment data from Regional Economic Information System (REIS) data provided by the Bureau of Economic Analysis (BEA). This table shows the average of year-to-year differences in natural log of earnings from mining, and employment in mining. Numbers in parantheses are standard errors; numbers in brackets are p-values.

the mining industries and other industries, suggesting limited spillover into other industries. For example, during the oil boom, large oil-producing counties' earnings from mining grew on average at over 11% per year, while total earnings merely grew at 1.6% per year. During the coal bust, for another example, coal mining employment in big coal counties dropped at more than 8% per year, 10 times the decline rate of total employment.

Black, McKinnish and Sanders (2005) also find limited spillover into other sectors. For instance, among men, who as potential miners should have been affected more than women, mining earnings grew between 1970 and 1980 at a rate of 0.273 log points, but non-mining earnings grew at only 0.058 log points. Black, McKinnish and Sanders (2003) report that female employment in large coal counties actually declined between 1970 and 1980 relative to the growth in non-coal counties.

1.6.3 Housing Cost

Housing is a major cost associated with child rearing, so housing costs will likely affect marriage and fertility outcomes, and the housing market could potentially respond to local labor demand shocks. Ideally, we would control for county-level housing cost in our empirical models, but county-level index of housing cost from Federal Housing Finance Agency dates back to only 1975, at the height of the energy booms. Also, the index is too unbalanced to be useful in our settings.

Omitting housing prices is likely to bias downwards the earnings-fertility estimates, because increases in housing prices during the boom increase the cost of living for renters, therefore reduce renters' propensity to have kids, and more than 75% of 15–44 women are renters (Dettling and Kearney, 2014). In a similar settings of Kearney and Wilson (2017), the authors find that the inclusion of housing prices does not change their estimates, indicating that housing prices are not driving the observed relationship between the local fracking boom and subsequent family formation outcomes.

Moreover, if the changes in the housing prices associated with energy shocks were largely due to changes in earnings, an argument can be made that we may not want to include housing prices, because the impacts of housing prices on family formation outcomes are part of the eanrings impacts we want to estimate.

1.7 Concluding Remarks

We assess the causal effects of less-educated men's earnings on marriage and fertility outcomes. Exploiting exogenous variations from pre-existing geological differences and unexpected world shocks to coal and oil prices, we find strong evidence that a persistent decline in local economic conditions reduces marriage prevalence and marital fertility, but increases non-marital births. We also find evidence that the reverse is not necessarily true: sustained economic prosperity may or may not increase the marriage propensity of less-educated men nor slow down the increase in out-of-wedlock births. Further research is needed to disentangle possible mechanisms of this potential asymmetry of the marriage and fertility responses to changes in local economic conditions.

1.8 Appendix

Table A.1: Impact of Coal and Oil Prices Shocks on Local Employment

States/Shocks: Years:	Coal Boom (1)	Coal Bust (2)	Oil/Gas Boom (3)	Oil/Gas Bust (4)
Employment per capita				
Instrument	.026 (.007)	.054 (.016)	.241 (.048)	.310 (.045)
F-Statistic for Excluded Instrument	14.2	11.2	25.4	48.4
N	826	708	3,499	2,868

Notes: Each Column represents a regression. The instrument for coal regressions is the change in national coal prices multiplied by coal reserves. The instrument for oil regressions is the change in national crude oil prices multiplied by the share of county employment in oil and gas extraction industry in 1970. Standard errors in parentheses account for arbitrary forms of clustering within counties. Samples are restricted to where fertility rate is not missing. All regressions control for: state by year fixed effects, level and change in log population, level and change in log real total welfare benefits, level and change of the share of black population, level and change of the share of other race, level and change of the share of female. All regressions are weighted by the number of total births in 1969.

Table A.2: Impact of Coal and Oil Prices Shocks on Local Average Wage

States/Shocks: Years:	Coal Boom (1)	Coal Bust (2)	Oil/Gas Boom (3)	Oil/Gas Bust (4)
Average Wage				
Instrument	.093	.031	.333	.557
	(.032)	(.018)	(.066)	(.076)
F-Statistic for Excluded Instrument	8.6	3.1	25.9	54.1
N	236	708	2,527	2,857

Notes: Each Column represents a regression. The instrument for coal regressions is the change in national coal prices multiplied by coal reserves. The instrument for oil regressions is the change in national crude oil prices multiplied by the share of county employment in oil and gas extraction industry in 1970. Standard errors in parentheses account for arbitrary forms of clustering within counties. Samples are restricted to where fertility rate is not missing. All regressions control for: state by year fixed effects, level and change in log population, level and change in log real total welfare benefits, level and change of the share of black population, level and change of the share of other race, level and change of the share of female. All regressions are weighted by the number of total births in 1969.

Table A.3: 2SLS Estimates of the Impact of Earnings on Fertility, Removing Comparison Counties That Are Adjacent to One or More Treatment Counties, Oil States

States/Shocks:	Oil/Gas		Oil/Gas	
Years:	Boom		Bust	
	Coef/SE	P-value	Coef/SE	P-value
	A: Total Birth Rate			
Earnings per capita	0.903		0.623	
	(0.266)	[0.001]	(0.128)	[0.000]
F-Statistic for Excluded Instruments	19.5 0		66.2	
N	2,796		2,256	
	B: Marital Birth Rate			
Earnings per capita	0.922		0.792	
	(0.275)	[0.001]	(0.140)	[0.000]
F-Statistic for Excluded Instruments	19.5 0		66.2	
N	2,796		2,256	
	C: Non-marital Birth Rate			
Earnings per capita	1.712		-0.296	
	(0.712)	[0.017]	(0.256)	[0.248]
F-Statistic for Excluded Instruments	18.4		66	
N	2,597		2,197	

Notes: Each cell represents a regression. Standard errors in parentheses account for arbitrary forms of clustering within counties. All regressions control for: state*year fixed effects, log population, change in log population, change in log real total welfare benefits, share of black population, share of other race, share of female, share of population aged 30s, 40s, 50s, 60s, 70 and up. All regressions are weighted by the number of total births in 1969.

CHAPTER 2

DISABILITY BENEFIT TAKE-UP AND LOCAL LABOR MARKET CONDITIONS

Kerwin Kofi Charles

University of Chicago and NBER

Yiming Li

University of Chicago

Melvin Stephens, Jr.

University of Michigan and NBER

2.1 Introduction

Understanding the labor market disincentive effects of government-provided disability insurance benefits has long been of interest to economists. While larger benefit amounts lead workers to apply for benefits and exit the labor force, better labor market opportunities have a countervailing effect on such transitions (Parsons, 1980; Haveman and Wolfe, 1984; Bound, 1989; Bound and Waidmann, 1992). Early work (e.g., Parsons, 1980) finds that increases in the share of earnings replaced by disability benefits lower employment rates although the potential endogeneity of the simple replacement rates in these studies may complicate the interpretation of these early findings (Bound, 1989). One strand of more recent work, examining arguably exogenous local labor demand shocks, finds that worse labor market conditions lead to increased levels of disability payments (Black, Daniel and Sanders, 2002; Autor and Duggan, 2003; Autor, Dorn and Hanson, 2013; Sloane, 2015). Other recent work uses random assignment of decision makers in the disability benefit determination process to estimate the labor supply impact of benefit awards (Maestas, Mullen and Strand, 2013; French and Song, 2014).

This paper examines the extent to which US Social Security Disability Benefits (SSDI) and Supplemental Security Income (SSI) payments, the two largest Federal programs that provide assistance to people with disabilities, are affected by local labor market earnings levels.[1] To circumvent the potential reverse causality problem arising from the possibility that local areas with chronically poorer health likely will have both higher rates of disability claiming and lower earnings, we use variation in local earnings generated by shocks to oil and gas production due to exogenous movements in the world prices of oil and gas. Our research design is akin to the seminal article of Black, Daniel and Sanders (2002), who use county-level movements in coal production in four states in the Appalachian region during the energy price boom and bust in the 1970s and 1980s, although our analysis differs from and extends their work in several key ways.

First, whereas Black, Daniel and Sanders (2002) study the boom and bust cycle in energy prices in the 1970s and 1980s, our analysis covers the period from 1970 to 2011, which spans both the earlier cycles and more recent run up in energy prices during the 2000s. Second, the demographic composition of workers in the oil and gas industry is more representative of U.S. workers overall than are workers in the coal industry studied by Black et al.[2] Third, compared to the four eastern U.S.

1. In December 2013, 8.9 million former workers and 2 million dependent spouses and children received SSDI and 9 million people received federal SSI. Source: OASDI Beneficiaries by State and County, 2013; SSI Annual Statistical Report, 2013. Disability is defined by law as the "inability to engage in any substantial gainful activity by reason of any medically determinable physical or mental impairment". An applicant's education, age, and work experience are considered when determining whether she may still work.

2. Appendix Table (B.1) presents characteristics of all U.S. workers, workers in the oil and gas industry and workers in the coal industry in 1970, 1990, and 2010. The numbers for 1970 and

states from Black et al.'s study, the eleven states with the largest share of oil and gas employment used in our analysis are more broadly representative, as they include places from along the Gulf of Mexico to the central U.S. To the extent that disability benefits offer higher replacement rates to lower-earning workers (Autor and Duggan, 2003), disability take-up for the typical American worker, who more closely resembles the oil and gas workers we study, may exhibit very different responses to earnings shocks that do the coal workers to shocks in their sector. Given these differences, our work provides new and timely evidence on the relationship between disability programs and local labor market conditions.[3]

We find that increases in county earnings lead to significant decreases in both county SSDI and SSI payments. Our baseline estimates of the elasticity of benefit payments to earnings for SSDI and SSI are -0.29 and -0.16, respectively, and are robust to numerous specifications of our instrumental variable. Despite the differences across the studies in the time periods and samples they analyze, our SSDI elasticity is comparable to Black et al.'s estimate. We show that our much smaller estimate for the SSI elasticity compared to Black et al. can be substantially accounted for by the federalization of the SSI program in 1974, which produced uniform program standards across states and led to an immediate 60% increase in total benefit payments.

The remainder of the paper proceeds as follows. Section (2.2) lays out the empirical models we use for estimation while Section (2.3) discusses the construction of

---

1990 are computed using the Public Use Micro Sample of the Decennial Censuses in those years. The 2010 numbers are computed using 2008–2012 American Community Survey 5-year data files (Ruggles et al., 2015).

3. Our analysis differs from the work of Autor, Dorn and Hanson (2013) and Sloane (2015) in several key ways, apart from the fact that we study a very different type of shock. In particular, whereas these papers study more states, our analysis is conducted at the level of the county, exploits annual-level variation, and studies a longer time frame than either of these studies.

our county-level data. Section (2.4) presents our main results and robustness checks and Section (2.5) concludes.

2.2 Emprical Specification

To access the effects of local economic conditions on disability payments, we estimate the following equation:

$$\Delta_c y_{cst} = \beta_0 + \delta\,(\Delta_c E^*_{cst}) + \Delta_c \boldsymbol{x}_{cst}\boldsymbol{\beta}_1 + \boldsymbol{d}_{st}\boldsymbol{\beta}_2 + \varepsilon_{cst}, \tag{2.1}$$

where y_{cst} is the logarithm of real SSDI or SSI payments for county c in state s between year $t-1$ and t; $\boldsymbol{d}_{st}$ is a vector of state-year dummy variables for state s in year t to allow state-specific differences by year; $\boldsymbol{x}_{cst}$ is a vector of control variables; ε_{cst} is idiosyncratic factors of county c in year t; Δ_c indicates the difference between t and $t-1$ within the county.[4]

The regressor of interest, $\Delta_c E^*_{cst}$, is the change in (log) earnings or employment of county c in year t. The variable is starred to denote the fact that we observe error-ridden versions of the actual variable. Given the log-log specification, the parameter on this regressor, δ, is the elasticity of disability payments with respect to local economic conditions.

4. We control for the logarithm of county population, the log growth of county population, whether the county is in a Metropolitan Statistical Area (MSA) for the 1990 Census, and the fraction of 1969 earnings from manufacturing. Population and MSA status are controlled as proxies for better access to health care and amenities such as public transportation, both of which are appealing to individuals with disabilities. If people with disabilities are over-represented in the migrants, then the population growth captures this mechanical impact on the disability payments. We use the fraction of 1969 earnings from manufacturing to control for any impact of industry structure on disability. This last variable is missing for some small counties due to disclosure concerns although our results are quite similar if we drop this variable and include these counties.

Some concerns arise in estimating δ using equation (2.1). One concern is that first difference estimators exacerbate the attenuation effects of measurement error (Bound, Brown and Mathiowetz, 2001). Also, as discussed in McKinnish (2007), year-to-year within-county variation often reflects transitory fluctuations that have little effect on long-term behavioral outcomes such as initiating SSDI benefits. Although permanent earnings is the relevant concept for the SSDI participation decision, the available earnings measure contains both permanent and transitory components. Another concern is that $\Delta_c E^*_{cst}$ is an endogenous regressor since changes in local economic conditions are likely correlated with changes in unobserved local factors that may affect disability participation. For example, if SSA field workers are more sympathetic to unemployed workers when local economy is suffering or if law firms spend more in disability-related advertisements in counties where economic conditions decline, OLS estimates of the relationship between disability payments and earnings will be biased downwards. Alternatively, more generous disability payments (less stringent screening criteria, higher replacement rates, etc.) may decrease labor force participation and bias the estimates in the opposite direction (Parsons, 1980; Autor and Duggan, 2003). The combined impact of these potential biases on the OLS estimate is therefore ambiguous.

We use two stage least squares (TSLS) to address both the measurement error and endogeneity concerns. We use variation in county-level earnings induced by exogenous shocks to international oil prices. As these prices surge (decline), both oil production as well as efforts to locate and extract new oilfields increases (decreases) leading to higher (lower) earnings in the affected counties. After accounting for state-year interactions, our instrumental variables strategy exploits differences across counties within the same state, as areas with higher baseline employment shares in the oil and gas industry experience larger earnings changes in response to oil price movements.

Our baseline instruments are constructed as the interaction between the oil price changes and the importance of oil and natural gas to a county, that is,

$$\Delta R_t \times \mathrm{Oil}_c, \tag{2.2}$$

where ΔR_t measures the real oil price changes between periods t and $t-1$, and Oil_c is the initial importance of oil and gas industry in the county as measured by oil and gas employment share in a year before the sample period. Two lags of the instrument are also included in ΔR_t to account for the setup time of new wells and local businesses adjustment following international shocks. In addition, given the long disability benefit determination periods, disability benefit payments may take time to fully adjust to a given shock. The resulting TSLS model estimates equation (2.2), with the change in the local labor market measure replaced by the predicted value from the first stage equation

$$\Delta_c E^*_{cst} = \gamma_0 + \sum_{\tau=0}^{2} \alpha_\tau \Delta R_{t-\tau} \times \mathrm{Oil}_c + \Delta_c \boldsymbol{x}_{cst} \boldsymbol{\gamma}_2 + \boldsymbol{d}_{st} \boldsymbol{\gamma}_3 + u_{cst}, \tag{2.3}$$

2.3 Data

The Bureau of Economic Analysis' Regional Economic Information System (REIS) provides annual county-level data on earnings and employment beginning in 1969. Earnings include wage and salary disbursements, other labor income, and proprietors' income. The Bureau of Labor Statistics (BLS) compiles REIS wage and salary disbursements using ES-202 filings collected as part of the state unemployment insurance program. The REIS data on county wages and employment are known to be measured with some error.[5] Dan Black kindly provided us with county SSDI payments in Decembers 1970–2001.[6] We collected more recent county SSDI payments from SSA's "OASDI Beneficiaries by State and County" Report (various years). County annual SSI received payments are from REIS. All dollar values are deflated by annual average CPI-U with the base year 2010.

The TSLS part of our analysis focuses on the effect of energy supply shocks in oil and natural gas producing counties in eleven "oil and gas states", spanning 1970 to 2011. We define the oil and gas states as those contiguous states where the mining industry comprises at least 1% of total employment in the 1974 County Business Patterns (CBP) and a substantial share of the mining establishments are related to oil and gas.[7] We construct a measure of the initial importance of oil and gas industry in a county, Oil_c, which is the share of employment in the oil and gas industry computed from CBP data.

Prior to 1974, we only observe one-digit total mining employment for each county in CBP. As the first large oil price shock occurs in 1974 and our sample begins in 1970, we would prefer to determine the initial oil and gas importance in some year prior to 1974, if not prior to 1970. The tradeoff, however, is that with an earlier year we

5. Some employers with establishments in multiple counties may only report wages and employment ES-202 information at the state-level. These reports are allocated back to counties based on their industry level distribution by county among employers reporting at the county-level, generating some measurement error. In addition, components of other labor income and proprietors' income such as pension plan contributions, health and life insurance contributions, and private worker's compensation contributions are only collected at the state-level and also use an allocation rule to determine county-level totals.

6. County-level SSDI payment data was not published in 1981.

7. These eleven states are Colorado, Kansas, Louisiana, Mississippi, Montana, New Mexico, North Dakota, Oklahoma, Texas, Utah, and Wyoming. We exclude West Virginia since less than half of total mining employment is in the oil and gas industry which affects the construction of our 1967 instrument.

can only determine mining employment as opposed to oil and gas employment at the county-level. As shown in Appendix Table (B.2), over half of mining establishments in the oil and gas states, with the exception of Utah, are in the oil and gas industry with little change in these shares between 1967 and 1974. Thus, we construct the county-level initial oil and gas importance measure, Oil_c, using the 1967 CBP and we present results using the 1974 CBP to compute Oil_c as robustness check.[8]

Figure (2.1) depicts the geographic distribution of the oil and gas industry by county within the eleven oil and gas states using CBP data from 1974. This figure illustrates the tremendous variation in the importance of oil and gas across counties in these states.

We obtain data on energy prices from the Energy Information Administration's Annual Energy Review.[9] From 1969 to the 1990s, there were at least three large exogenous shocks to the world oil supply: the 1973–74 OPEC oil embargo following the Yom Kippur War; the period from the end of 1979 to early 1981, following the overthrowing of the Shah of Iran and the start of the Iran-Iraq War; and the 1990–91 First Persian Gulf War. The late 1990s and 2000s saw booming oil demand from newly industrialized countries and stagnant world crude oil production, with occasional exogenous events such as the U.S. attack on Iraq and the turmoil in Nigeria (Hamilton,

8. Because of the risk of disclosing firm specific information, exact employment numbers for two-digit industries are not available at the county level. However, the CBP provides county-level information on both the number of firms in each two-digit industry and the number of firms that fall into a specific firm-size category (e.g., 20 to 49 employees) for these industries. By weighting the number of firms in a firm-size category by the midpoint of the number of employees in that category, we create an estimate of the number of employees in each two-digit industry at the county level. We then create county-level estimated employment shares by industry as the ratio of the estimated industry employment to the estimated total county employment, where the total county employment is also estimated using the firm-size methodology.

9. Oil prices are the US average first purchase price per barrel; natural gas prices are the wellhead price per thousand cubic feet. National oil and gas industry employment are from CBP.

Figure 2.1. Share of County Employment in Oil and Gas Industry in 1974

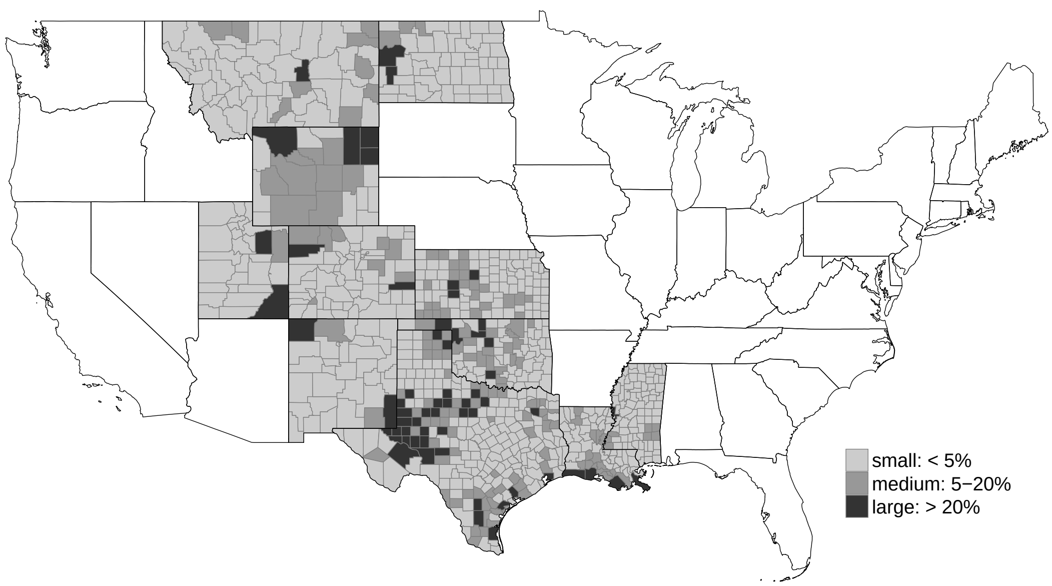

Notes: Share of county employment in oil and gas industry is calculated based on data from CBP. See Data section for details.

2009b,a). During the same period, the advances in hydraulic fracturing and horizontal drilling technologies made shale gas commercially viable. These events and trends affected both the prices of oil and natural gas, as well as the employment in these industries in the United States, as Figure (2.2) illustrates.

Real oil prices doubled between 1973 and 1974, were stable for several years, then tripled over a three-year period. Prices fell sharply over the next five years to levels in the mid-1980s that were slightly lower than those of the mid-1970s. After hovering around the same level for more than a decade, real oil prices started to climb up again in the early 2000s, reaching the all-time high in mid-2008, then collapsed in 2009. Real natural gas prices followed a very similar pattern to that for oil prices: a six-fold increase between 1970 and the early 1980s, then a decline of more than half over the next six years, and later a similar run-up and collapse in the 2000s. Figure (2.2) shows that national employment in the oil and natural gas industry closely tracked the movement in prices, with the main difference that national employment was not kept artificially flat during various periods, as was true for oil prices in two periods in the 1970s because of policy decisions.[10] As part of our robustness checks, we use national oil and gas employment changes as the measure of ΔR_t in place of oil price changes to construct our instruments.

10. Direct federal control of crude oil prices lasted for almost a decade, from August 1971 to January 1981. In addition, since the early 1900s, oil production in the United States has been overseen by various state regulatory boards, such as the Oklahoma Corporations Commission, the Louisiana Conservation Commission and, most importantly, the Texas Railroad Commission. Although the specific language outlining each board's functions and objectives differ from state to state, these agencies set limits on level of extraction and exploration in their particular states so as to stabilize price, and prevent over-exploitation of oil reserves.

Figure 2.2. Effect of International Oil Supply Shocks on Prices and Employment

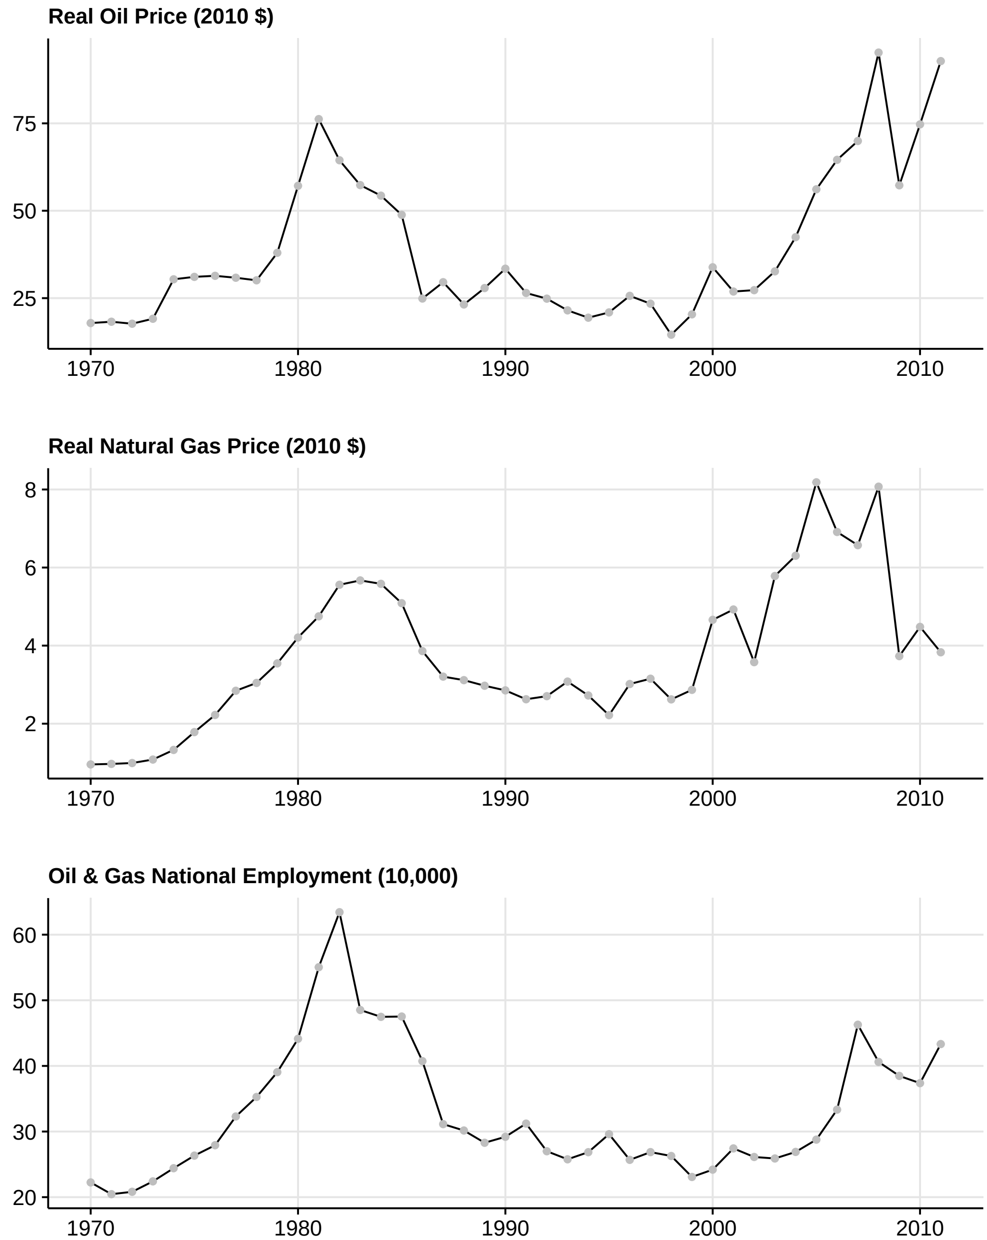

Notes: Prices are from the Energy Information Administration's Annual Energy Review. Oil prices are the US average first purchase price per barrel; natural gas prices are the wellhead price per thousand cubic feet. National oil and gas industry employment are from CBP.

2.4 Results

Table (2.1) presents our results from estimating equation (2.1) with both SSDI (Panel A) and SSI (Panel B) as outcomes using counties in the eleven oil and gas states.[11] Our OLS estimates, presented in column (1), show no significant relationship between annual county-level changes in earnings and the analogous changes in disability payments. However, as we discussed above, there are reasons to believe that these estimates will be inconsistent because the regressor $\Delta_c E^*_{cst}$ may be both being measured with error and endogenous in (2.1).

We next present TSLS estimates of equation (2.1) in which we instrument for the change in log earnings using variation in oil prices and the geographic distribution of oil and gas production as shown in equation (2.3). The F-statistic testing the joint significance of the excluded instruments in the first stage equation, shown below each point estimate, greatly exceeds the conventional threshold of 10.[12] Over the entire sample period, which runs from 1970 through 2011, our TSLS estimates shown in column (2) of Table (2.1) are -0.29 and -0.16 for SSDI and SSI, respectively, and are highly statistically significant. Thus, local disability payments increase significantly when local earnings decline.[13]

---

11. The sample sizes for SSDI payment regressions are larger than for SSI payment regressions because the REIS contains missing payment data for counties in years when these payments are below a nominal \$50,000. Thus, these missing values occur more frequently in earlier sample years and in smaller counties. In addition, we only present results from using a balanced panel for each payment outcome, although the results are quite comparable when we allow for an unbalanced panel. These results are available from authors.

12. The corresponding first stage estimates are found in Appendix Table (B.3).

13. We also estimated the baseline TSLS model separately for the two periods conventionally treated in the literature as being times of distinct oil "boom" (1973–1981) and oil "bust" (1982–1986). For SSDI receipt, the estimated elasticity with respect to earnings of $-.547$ (.132) during the boom and $-.135$ (.124) during the bust are quite different, which suggest that the propensity to participate in the program during bad labor market conditions versus to leave or not apply during good

Table 2.1: The Effect of Local Economic Performance on Disability Payments, 1970–2011

Method:	OLS	TSLS	OLS	TSLS
Economic Condition:	Log Earnings	Log Earnings	Log Employment	Log Employment
	(1)	(2)	(3)	(4)
		A. SSDI		
Economic Condition	0.004	-0.293	0.079	-0.699
	(0.014)	(0.069)	(0.065)	(0.131)
F-Statistic for Excluded Instruments		25		27
N	28,329	28,329	28,329	28,329
		B. SSI		
Economic Condition	-0.019	-0.16	-0.001	-0.36
	(0.012)	(0.079)	(0.030)	(0.168)
F-Statistic for Excluded Instruments		27.1		25.5
N	26,154	26,154	26,154	26,154

Notes: Each point estimate in the table represents results from a different regression. Standard errors in parentheses account for arbitrary forms of clustering within counties. All regressions control for: state*year fixed effects, MSA indicator (1990), log population, change in log population, and fraction of earnings from manufacturing (1969). SSDI balanced sample only includes counties with non-missing SSDI payments for all years except 1981, when all counties are missing. SSI balanced sample only includes counties with non-missing SSI payment for all years. All regressions are weighted by county population in the previous period.

Using earnings as the regressor of interest captures both the impact of employment and wage changes on disability benefit take-up. However, a decline in earnings due to lost employment (e.g., a plant closing) might have a different effect on the decision to apply for disability benefits than lower real wage growth, especially if lack of employment is a requirement for a successful disability benefit award. In the last two columns of Table 1, we estimate models identical to (2.1), except that $\Delta_c E_c st^*$ in these regressions is the change in log employment in order to focus on the extensive work margin. Our TSLS benefit payment elasticity estimates for employment (column (4)) are more than twice as large as the earnings elasticities (column (2)), indicating that a 10% reduction in employment leads to a much larger disability response than a 10% reduction in earnings. Throughout the remainder of the paper we only present results that use the earnings regressor as proportional difference between the estimates from the earnings and employment regressors choices is roughly the same as is found in Table (2.1).[14]

Table (2.2) analyzes the robustness of our results to a number of additional specification variants. The energy price variation that identifies our results comes from two eras: the energy price boom and bust cycle in the 1970s and 1980s as well as more recent sharp price movements that occurred in the 2000s. The first two columns of Table (2.2) split the sample along these time dimensions with the results for 1970–1993 shown in column (1) and the findings for 1994–2011 shown in column (2). For SSDI, the estimate for the boom and bust cycle is very similar to that for the full sample period. The result for the latter period is about 30% smaller than the point estimate

conditions might be asymmetric. For SSI, the estimated effects of − .072 (.149) and −.151 (.05) during the boom and bust are not statistically different, which suggests that for this program the effects might not be asymmetric.

14. These results are available from the authors upon request.

for the full sample although both point estimates fall well within the 95% confidence interval for the other point estimate. For SSI, the boom and bust cycle estimate is slightly larger than the full sample estimate while the latter period estimate is not statistically different than zero.

The third column of Table (2.2) constructs Oil_c based on the 1974 oil and gas employment shares rather than the 1967 mining employment shares. Recall that 1974 is the earliest year in which the CBP data allows us to compute county-level employment shares for the oil and gas industry as opposed to the entire mining industry. Although the measure of Oil_c constructed using the 1974 data is based on information from after the start of the sample period we study, the resulting estimates are comparable to our baseline findings.

The models shown in next column in Table (2.2) use the contemporaneous oil shock as an instrument without also including two lags of the oil shock as in the baseline results. Since local businesses need time to respond to international price shocks, and the disability determination process also takes many months to complete, especially given the time it may take to successfully appeal a claim that is initially denied, the total impact of an oil price shock in a given year will likely take multiple years to be realized. Indeed, in column (4) of Table (2.2), we find smaller estimates when we exclude lagged oil shocks from our instrument set, especially for SSDI payments.

We have noted previously that oil price movements were constrained because of various institutional features prior to the early 1980s. Given the infrequent adjustment of oil prices during this era, movements in oil and gas industry employment likely better represented producers' desires to extract and search for oil than oil price movements. We therefore use national oil and gas industry employment to construct the instrument ΔR_t rather than oil prices. In the final column of Table (2.2), we find

Table 2.2: Robustness of Two Stage Least Squares Estimates

Specification:	1967 Instrument: 1970–1993 (1)	1967 Instrument: 1994–2011 (2)	1974 Instrument 1970–2011 (3)	No Lags of Instrument 1970–2011 (4)	Employment Instrument 1970–2011 (5)
			A. SSDI		
Change in Log Earnings	-0.287	-0.2	-0.272	-0.147	-0.414
	(0.123)	(0.091)	(0.088)	(0.126)	(0.068)
F-Statistic for Excluded Instruments	36.9	12.8	17.6	67.3	23.9
N	15,225	13,104	28,329	28,329	28,329
			B. SSI		
Change in Log Earnings	-0.203	0.042	-0.118	-0.107	-0.374
	(0.087)	(0.167)	(0.068)	(0.080)	(0.082)
F-Statistic for Excluded Instruments	41.6	13	17.7	79	23.4
N	14,940	11,214	26,154	26,154	26,154

Each point estimate in the table represents results from a different regression. Standard errors in parentheses account for arbitrary forms of clustering within counties. All regressions control for: state*year fixed effects, MSA indicator (1990), log population, change in log population, and fraction of earnings from manufacturing (1969). SSDI balanced sample only includes counties with non-missing SSDI payments for all years except 1981, when all counties are missing. SSI balanced sample only includes counties with non-missing SSI payment for all years. All regressions are weighted by county population in the previous period.

somewhat larger responses of SSDI and SSI payments when using this alternative instrument.

Our study, which uses nearly twenty years more recent data, analyses a sample of workers, whose education is higher, who work in a different industry and live in different parts of the country compared to Black et al. How do our estimates compare with those in the earlier study? Our baseline estimates of the elasticity of benefit payments to earnings are −0.29 and −0.16 for SSDI and SSI, respectively. Black et al's corresponding baseline estimate for the SSDI elasticity of −0.345 is quite comparable to ours. However, their baseline estimate of the SSI elasticity of −0.713 is much larger than our finding. On the one hand, across a number of robustness checks, their baseline result for SSI is at the larger end of their range of estimates which spans −0.4 to −0.7.[15] On the other hand, a number of our robustness checks for SSI are closer to −0.2 including our estimates in which we restrict the sample to the 1970–1993 energy price boom and bust period.

The close similarity between our SSDI estimates and those from Black et al. suggests that the differences in the SSI results across the two studies are unlikely to be simply due to the aforementioned observable sample differences.[16] We consider instead the possibility that differences in the SSI estimates between the studies might be due to changes in the SSI program that may have coincided in different ways with the timing of the prices shocks in the two studies.

15. Black et al.'s SSDI estimates have a much tighter range of −0.3 to −0.4.

16. Recent work in the treatment effects literature compares estimates across settings by accounting for differences in observable characteristics with constant sub-group treatment effects (e.g., Hotz, Imbens and Mortimer (2005); Angrist and Fernández-Val (2013)) or by estimating the distribution of treatment effects (e.g., Brinch, Mogstad and Wiswall (2017)). These papers focus on binary treatments as opposed to the continuous endogenous regressor that we use in our analysis and do not easily extend to the current context.

The most obvious candidate programmatic change was the establishment of the federal SSI program in 1974 intended to bring uniform eligibility standards and a national minimum benefit to programs that had previously varied greatly across states along these dimensions. Within one year, total beneficiaries in the federal program increased nearly 25% relative to participation in the state-level SSI antecedents, while total benefit payments increased by almost 60% (Social Security Administration, 1975).[17] However, given heterogeneity in the administration of these programs across different states, the increase in benefit payments varied greatly across states and, quite plausibly, across counties within a given state.

The federalization of SSI in 1974 may have generated outlier observations in terms of county-level SSI benefit growth because of the rapid increase in total SSI payments. We thus re-examine the SSI results in both in ours and the Black et al. samples, after dropping the 1973 to 1974 first difference observations, which were directly affected by the change in the program structure.[18] The first column of Table (2.3) shows our baseline estimate of TSLS estimate of the SSI elasticity with respect to earnings of -0.17 found in Table (2.1).[19] When we drop the 1973–74 first difference observations, our resulting estimate of -0.20, shown in column (2) of Table (2.3), is essentially unchanged from our baseline finding. Column (3) of Table (2.3) shows Black et al.'s original estimate of -0.713 which is found in column (2) of Table 3 in their paper. Our attempt at replicating their result, aided greatly by data provided by Dan Black,

17. The state-level antecedents to the SSI program were Old Age Assistance, Aid to the Blind, and Aid to the Permanently and Totally Disabled.

18. Prior to 1974, the SSI measure used in the analysis contains payments to the state-level SSI antecedents.

19. Following Black et al., we do not weight the analysis in Table (2.3).

Table 2.3: Impact of 1974 SSI Federalization on SSI TSLS Estimate

States/Shocks:	Oil/Gas	Oil/Gas	Coal	Coal	Coal
Years:	1970–1993	1970–1993	1970–1993	1970–1993	1970–1993
Method:	TSLS	TSLS	TSLS	TSLS	TSLS
Instruments:	Employment	Employment	Reserves	Reserves	Reserves
Include 1973–74?:	Yes	No	Yes	Yes	No
	(1)	(2)	(3)	(4)	(5)
Change in Log Earnings	-0.165	-0.195	-0.713	-0.746	-0.382
	(0.107)	(0.108)	(0.134)	(0.140)	(0.098)
F-Statistic for Excluded Instruments	41	36.4	26.5	24.3	22.7
N	14,940	14,318	7,904	7,904	7,577

Notes: Each point estimate in the table represents results from a different regression in which the outcome is the change in log county SSI payments. Standard errors in parentheses account for arbitrary forms of clustering within counties. All regressions control for: state*year fixed effects, MSA indicator (1990), log population, change in log population, and fraction of earnings from manufacturing (1969). SSI balanced sample for oil/gas regressions only includes counties with non-missing SSI payment for all years. Column (3) is taken from column (2) of Table 3 of Black, Daniel, and Sanders (2002). All regressions are unweighted.

yields an estimate of -0.746 as shown in column (4) of Table (2.3).[20] When we omit first difference observations for 1973–74, the point estimate drops nearly in half to -0.382. While this estimate is still larger than our baseline SSI estimate, accounting for the 1974 federalization of SSI explains the majority of the difference between our SSI finding and Black et al's.

20. One likely reason we do not exactly match the results from Black et al. in our replication is that the REIS data is subject to periodic revisions which means that we do not necessarily have exactly the same data on county-level SSI payments, earnings, population, etc. that was used in the prior study.

2.5 Conclusion

We find that worsening labor market conditions in a county cause disability benefit receipt there to increase. Our estimates of the elasticity of benefit payments to earnings during 1970–2011 for SSDI and SSI are -0.29 and -0.16, respectively. Earlier work finds a much larger SSI elasticity, but we show that much of the difference can be accounted for by the federalization of the SSI program in 1974.

Our finding that agents substantially substitute disability benefits for earnings when labor market circumstances worsen shows that this striking result, which was previously documented by Black et al., is not simply a feature of that earlier study's focus on a sample of less-educated men in rural Appalachian coal areas in the time period before 1990. Instead, we find a similar relationship in a more nationally representative set of states, among more-educated persons, and during a much more recent time period. This somewhat surprising result suggests that the response of disability take-up to market conditions is similar across areas and disability programs, despite different targeted population for SSDI and SSI programs.

Our analysis also raises important issues for future research. By focusing on changes in earnings and disability payments measured at the county level, our research design cannot separate disability take-up decisions based on individual earnings fluctuations from those due to movements in market-wide earnings levels. Separately identifying these two potential mechanisms can sharpen our understanding of how earnings changes affect disability take-up. Also, we find that the SSI response for the more recent time period is not significantly different than zero. This result is particularly puzzling as we still find significant responsiveness of SSDI benefits during the latter period. Additional work is needed to explain the divergence of the estimated elasticities between the two disability programs.

2.6 Appendix

Figure B.1. Total Payments to the Disabled by Program, 1970–2012

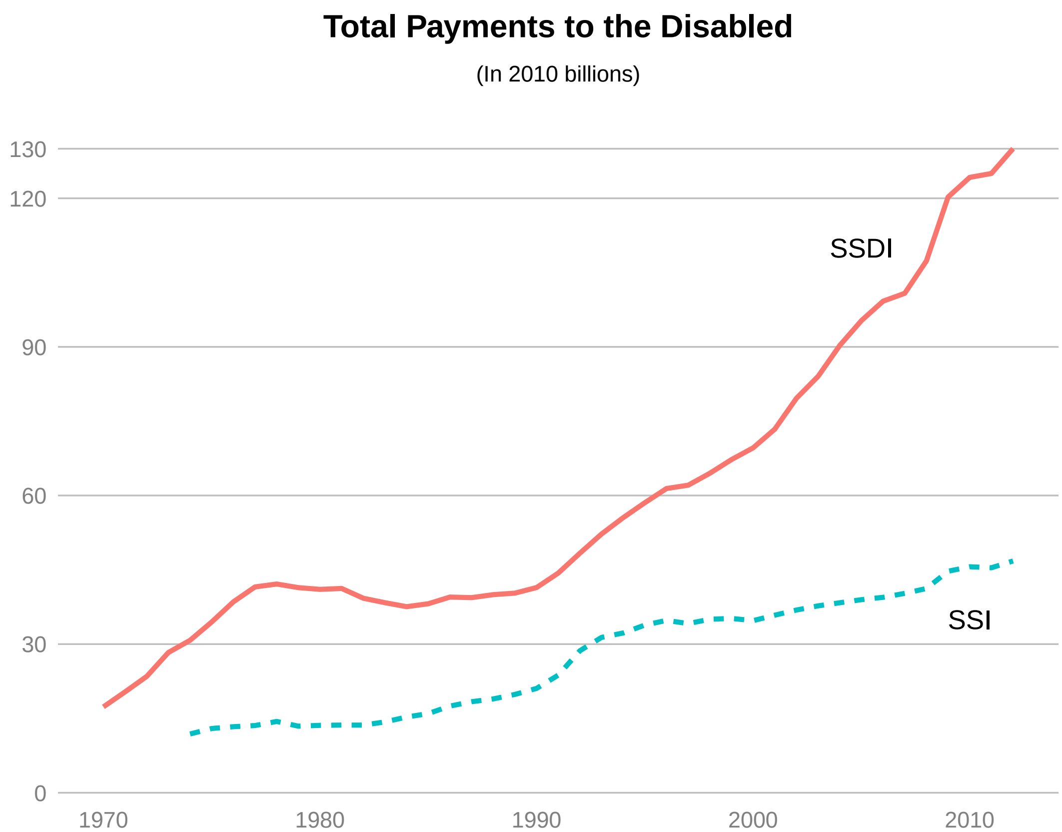

Figure B.2. Number of Disabled Workers by Program, 1970–2012

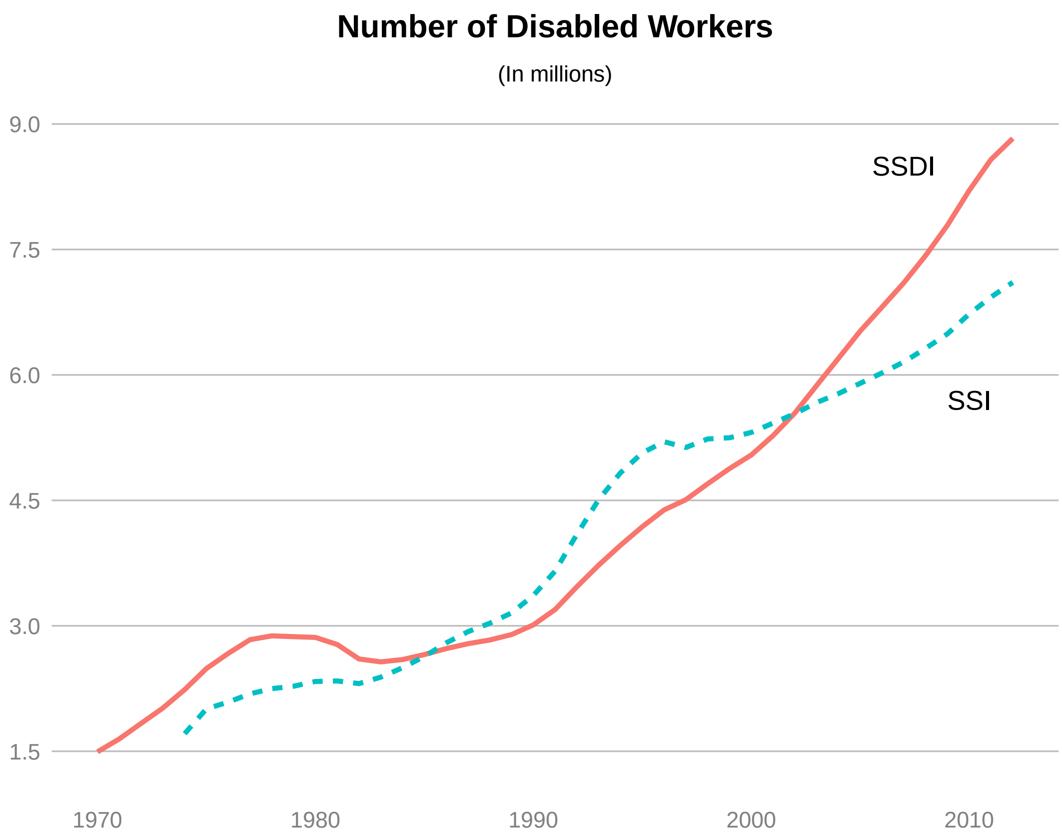

Figure B.3. SSDI Expenditures As Share of Total OASDI

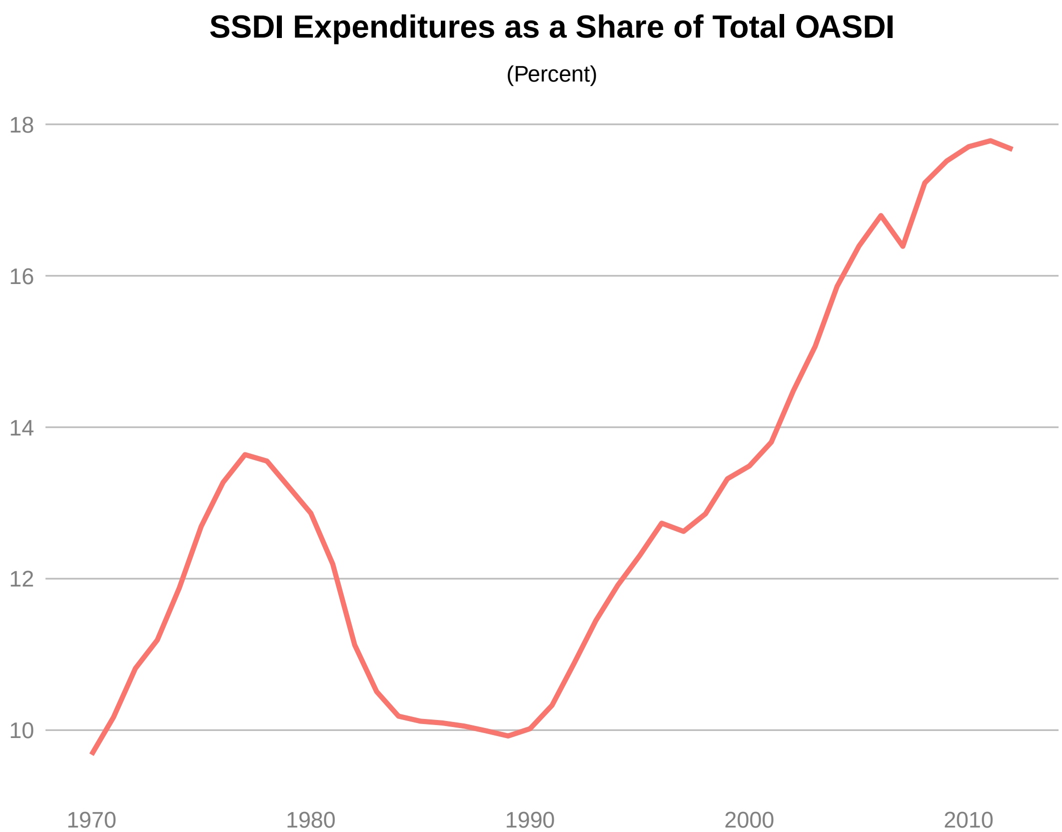

Figure B.4. Share of SSDI New Awards by Diagnosis, 1981–2012

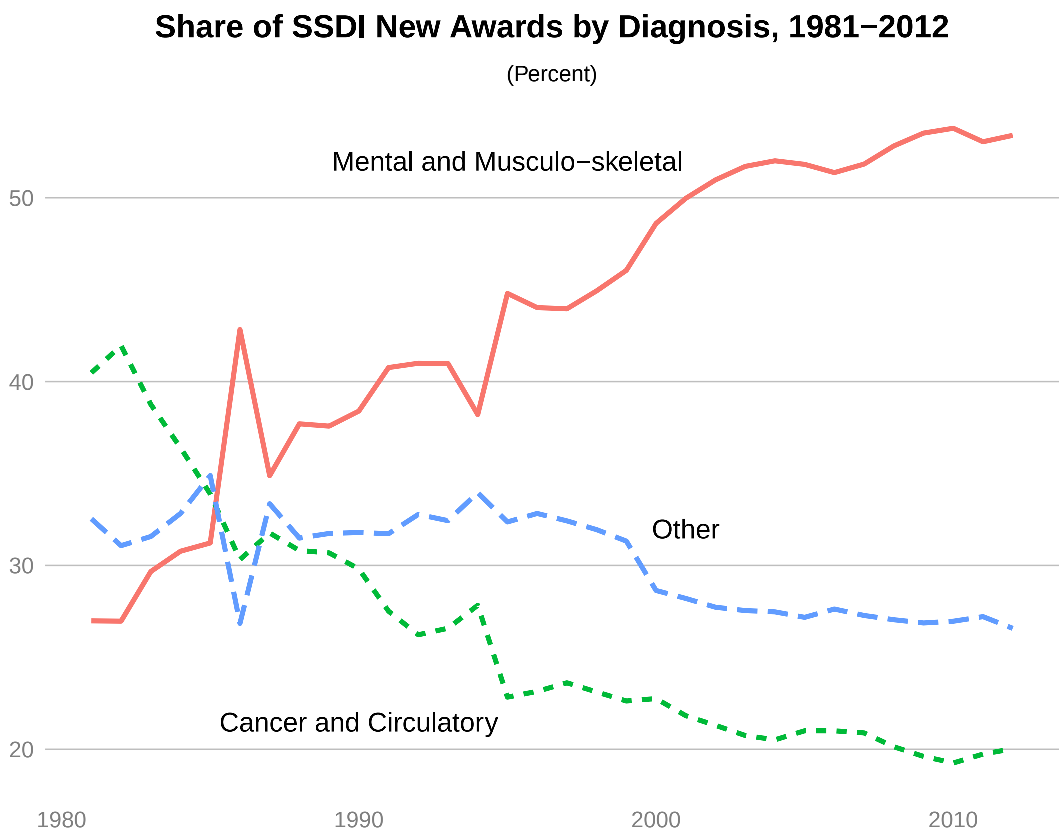

Figure B.5. Co-movement of SSDI Applications and Unemployment Rate, 1981–2012

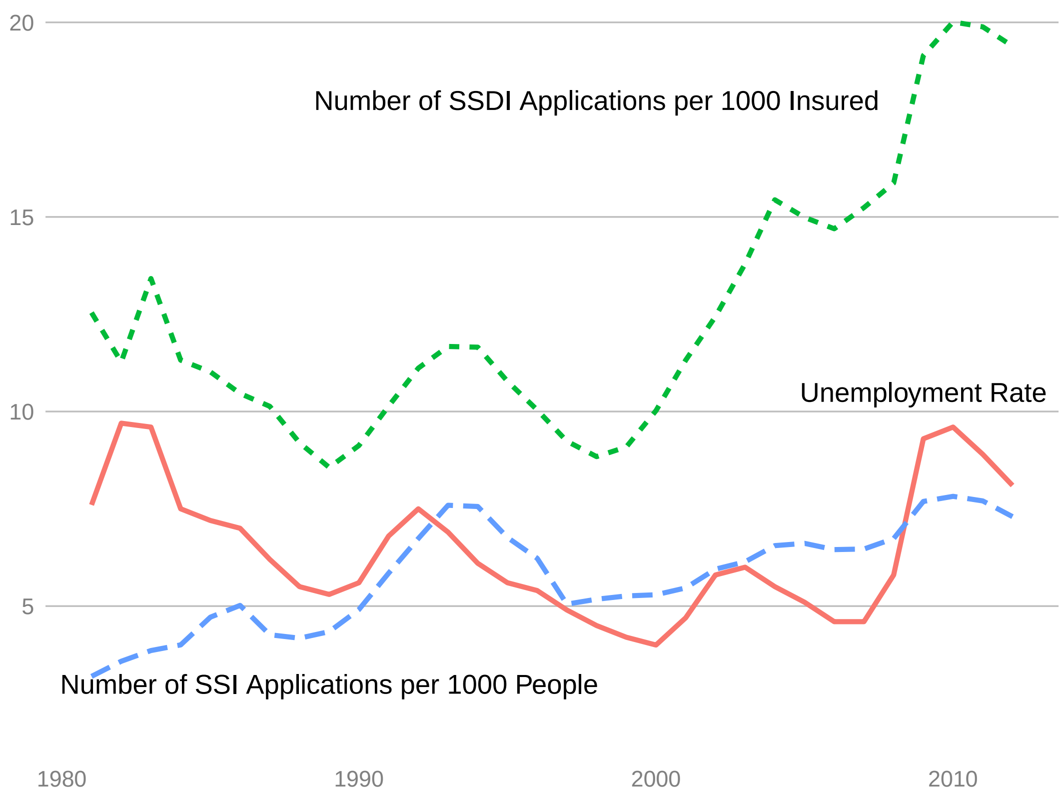

Table B.1: Characteristics of Workers by Industry

	All Workers			Oil & Gas Workers			Coal Workers		
	1970	1990	2010	1970	1990	2010	1970	1990	2010
Age									
25–29	14	17	13	15	15	16	11	8	12
30–39	24	33	24	26	40	28	21	40	23
40–49	26	25	26	29	24	24	30	31	23
50 and over	36	26	36	31	21	33	38	21	41
Education									
Less than high school	41	13	8	38	13	11	72	25	11
High school	34	33	33	30	33	43	23	47	58
Some college	12	29	25	13	26	22	4	20	21
College and above	14	25	34	18	28	24	1	8	10
Sex									
Male	62	54	53	88	81	85	98	94	94
Female	38	46	47	12	19	15	2	6	6
Race									
White	89	83	77	97	90	86	96	97	95
Black	10	10	11	2	5	6	3	2	2
Other	1	7	12	¡1	6	8	¡1	2	2
Occupation									
Managerial and Professional	22	28	33	24	30	27	4	9	10
Technical, Sales, and Administrative	26	31	28	18	21	16	4	7	7
Precision Production, Craft, and Repairers	14	12	10	48	32	41	67	55	57
Operatives and Laborers	22	15	12	9	16	15	24	27	21
Service	12	12	15	1	1	1	1	2	5
Farming, Forestry, and Fishing	4	3	2	¡1	¡1	¡1	¡1	¡1	¡1

Notes: Author's compilations from the 1970, 1990 Public Use Micro Sample of the Decennial Censuses, and 2008–2012 ACS 5-year data files. Census's industry coding for oil and natural gas industry only include the extraction sector. Sample consists of individuals of age 25 and above who have worked positive weeks during the last year. Shares may not add to 100 because of rounding.

Table B.2: Oil and Gas States

	Share of 1974 CBP State Employment in Mining	Oil and Gas Share of All Mining Establishments in:	
		1974	1967
Wyoming	15.80%	77%	82%
New Mexico	7.80%	82%	76%
Louisiana	5.90%	88%	91%
Montana	4.30%	66%	68%
Oklahoma	3.90%	90%	92%
Utah	3.70%	47%	32%
Texas	3.20%	88%	92%
Colorado	2.20%	58%	50%
Kansas	1.80%	85%	87%
North Dakota	1.20%	65%	61%
Mississippi	1.10%	81%	83%

Notes: CBP stands for County Business Patterns.

Table B.3: Selected First Stage Estimates of Impact on Earnings

States/Shocks:	Oil/Gas	Oil/Gas	Oil/Gas
Years:	1970–2011	1970–1993	1994–2011
Weighted?:	Yes	Yes	Yes
	(1)	(2)	(3)
		A. SSDI	
Instrument	0.277	0.315	0.26
	(0.033)	(0.034)	(0.049)
Instrument lagged once	0.18	0.197	0.124
	(0.028)	(0.037)	(0.028)
Instrument lagged twice	0.043	-0.152	0.135
	(0.019)	(0.023)	(0.027)
F-Statistic for Excluded Instruments	25	36.9	12.8
N	28,329	15,225	13,104
		B. SSI	
Instrument	0.302	0.357	0.281
	(0.034)	(0.034)	(0.052)
Instrument lagged once	0.168	0.197	0.128
	(0.027)	(0.034)	(0.030)
Instrument lagged twice	0.027	-0.177	0.15
	(0.020)	(0.026)	(0.030)
F-Statistic for Excluded Instruments	27.1	41.6	13
N	26,154	14,940	11,214

Notes: CBP stands for County Business Patterns.

REFERENCES

Angrist, Joshua D., and Iván Fernández-Val. 2013. "ExtrapoLATE-ing: External Validity and Overidentification in the LATE Framework." In . Vol. 3 of *Econometric Society Monographs*, , ed. Daron Acemoglu, Eddie Dekel and Manuel Arellano, 401–434. Cambridge:Cambridge University Press.

Artuc, Erhan, Shubham Chaudhuri, and John McLaren. 2010. "Trade Shocks and Labor Adjustment: A Structural Empirical Approach." *American Economic Review*, 100(3): 1008–45.

Autor, David, David Dorn, and Gordon Hanson. 2017. "When Work Disappears: Manufacturing Decline and the Falling Marriage-market Value of Men." *National Bureau of Economic Research Working Paper Series*, No. 23173.

Autor, David H., and Mark G Duggan. 2003. "The Rise in the Disability Rolls and the Decline in Unemployment." *The Quarterly Journal of Economics*, 118(1): 157–206.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*, 103(6): 2121–68.

Bailey, Martha, Karen Clay, Price Fishback, Michael Haines, Shawn Kantor, Edson Severnini, and Anna Wentz. 2016. "U.S. County-level Natality and Mortality Data, 1915-2007."

Becker, Gary S. 1973. "A Theory of Marriage: Part I." *Journal of Political Economy*, 81(4): 813–846.

Becker, Gary S. 1974. "A Theory of Marriage: Part II." *Journal of Political Economy*, 82(2, Part 2): S11–S26.

Bertrand, Marianne, Emir Kamenica, and Jessica Pan. 2015. "Gender Identity and Relative Income within Households." *The Quarterly Journal of Economics*, 130(2): 571–614.

Black, Dan A, Kermit Daniel, and Seth G Sanders. 2002. "The Impact of Economic Conditions on Participation in Disability Programs: Evidence from the Coal Boom and Bust." *The American Economic Review*, 92(1): 27–50.

Black, Dan A., Natalia Kolesnikova, Seth G. Sanders, and Lowell J. Taylor. 2013. "Are Children "normal"?" *The Review of Economics and Statis-*

tics, 95(1): 21–33.

Black, Dan A, Terra G McKinnish, and Seth G Sanders. 2003. "Does the Availability of High-wage Jobs for Low-skilled Men Affect Welfare Expenditures? Evidence from Shocks to the Steel and Coal Industries." *Journal of Public Economics*, 87(9-10): 1921–1942.

Black, Dan A, Terra G McKinnish, and Seth G Sanders. 2005. "The Economic Impact of the Coal Boom and Bust." *The Economic Journal*, 115(503): 449–476.

Blanchard, Olivier Jean, and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity*, 1992(1): 1–75.

Bound, John. 1989. "The Health and Earnings of Rejected Disability Insurance Applicants." *The American Economic Review*, 79(3): 482–503.

Bound, John, and Timothy Waidmann. 1992. "Disability Transfers, Self-reported Health, and the Labor Force Attachment of Older Men: Evidence from the Historical Record." *The Quarterly Journal of Economics*, 107(4): 1393–1419.

Bound, John, Charles Brown, and Nancy Mathiowetz. 2001. "Chapter 59 - Measurement Error in Survey Data." In *Handbook of Econometrics.* , ed. James J Heckman Leamer and Edward, 3705–3843. Elsevier.

Brinch, Christian N., Magne Mogstad, and Matthew Wiswall. 2017. "Beyond LATE with a Discrete Instrument." *Journal of Political Economy*, 125(4): 985–1039.

Dettling, Lisa J., and Melissa S. Kearney. 2014. "House Prices and Birth Rates: The Impact of the Real Estate Market on the Decision to Have a Baby." *Journal of Public Economics*, 110: 82–100.

French, Eric, and Jae Song. 2014. "The Effect of Disability Insurance Receipt on Labor Supply." *American Economic Journal: Economic Policy*, 6(2): 291–337.

Glaeser, Edward L., and Joseph Gyourko. 2005. "Urban Decline and Durable Housing." *Journal of Political Economy*, 113(2): 345–375.

Hamilton, James D. 2009*a*. "Causes and Consequences of the Oil Shock of 2007-08." *Brookings Papers on Economic Activity*, 215–278.

Hamilton, James D. 2009*b*. "Understanding Crude Oil Prices." *The Energy Journal*, 30(2): 179–206.

Haveman, Robert H., and Barbara L. Wolfe. 1984. "Disability Transfers and Early Retirement: A Casual Relationship?" *Journal of Public Economics*, 24(1): 47–66.

Hotz, V Joseph, Guido W Imbens, and Julie H Mortimer. 2005. "Predicting the Efficacy of Future Training Programs Using Past Experiences at Other Locations." *Journal of Econometrics*, 125(1-2): 241–270.

Kearney, Melissa S., and Phillip B. Levine. 2017. "The Economics of Non-marital Childbearing and the "marriage Premium for Children"." *National Bureau of Economic Research Working Paper Series*, No. 23230.

Kearney, Melissa S., and Riley Wilson. 2017. "Male Earnings, Marriageable Men, and Nonmarital Fertility: Evidence from the Fracking Boom." *National Bureau of Economic Research Working Paper Series*, No. 23408.

Maestas, Nicole, Kathleen J. Mullen, and Alexander Strand. 2013. "Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of Ssdi Receipt." *American Economic Review*, 103(5): 1797–1829.

Manson, Steven, Jonathan Schroeder, David Van Riper, and Steven Ruggles. 2017. "IPUMS National Historical Geographic Information System: Version 12.0 [Database]. Minneapolis: University of Minnesota."

McKinnish, Terra G. 2007. "Panel Data Models and Transitory Fluctuations in the Explanatory Variable." *Advances in Econometrics*, 21: 335–358.

Parsons, Donald O. 1980. "The Decline in Male Labor Force Participation." *Journal of Political Economy*, 88(1): 117–134.

Pencavel, John. 1998. "Assortative Mating by Schooling and the Work Behavior of Wives and Husbands." *The American Economic Review*, 88(2): 326–329.

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. 2015. "Integrated Public Use Microdata Series: Version 6.0 [dataset]. Minneapolis, Mn: University of Minnesota."

Schultz, T. Paul. 1985. "Changing World Prices, Women's Wages, and the Fertility Transition: Sweden, 1860-1910." *Journal of Political Economy*, 93(6): 1126–1154.

Sloane, Carolyn M. 2015. "Where Are the Workers? Technological Change, Rising Disability and the Employment Puzzle of the 2000s: A Regional Approach." *Working Paper.*

Social Security Administration. 1975. "Growth of the Supplemental Security Income Program in 1974." *Social Security Bulletin*, 38: 32–37.

Wilson, William J. 1987. *The Truly Disadvantaged : The Inner City, the Underclass, and Public Policy.* Chicago:University of Chicago Press.

Wilson, William Julius, Kathryn Neckerman, Sheldon Danziger, and Daniel Weinberg. 1986. "Poverty and Family Structure: The Widening Gap between Evidence and Public Policy Issues." *Fighting Poverty: What Works and What Doesn't.*