

Discussion Paper Series

IZA DP No. 18466

March 2026

Do Prior Residents Benefit from Energy Booms?

Luyi Han

Pennsylvania State University

John V. Winters

Iowa State University and IZA@LISER

Michael R. Betz

The Ohio State University

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.



Do Prior Residents Benefit from Energy Booms?*

Abstract

The 21st century fracking boom transformed American energy production, but new jobs were often filled by temporary in-migrants and long-distance commuters, possibly reducing economic benefits for prior residents. We use novel restricted-access data from the U.S. Census Bureau to assess fracking impacts on prior residents. We examine impacts on earnings and employment for persons born in non-metropolitan fracking counties. We utilize an event study design to estimate annual impacts during the fracking boom, drilling downturn, and subsequent periods. We find sizable impacts on average log earnings that peaked during the boom and partially persisted during and after the downturn. The fracking boom also increased the probability of being employed but the effect largely disappeared after fracking activity peaked. We also compare our main result for non-metropolitan natives to persons born in metropolitan counties and conduct several other extensions.

JEL classification

Q4, R2, J3

Keywords

fracking, local labor markets, resource boom, rural development

Corresponding author

John V. Winters

winters1@iastate.edu

* This work was supported in part by the United States Department of Agriculture, National Institute of Food and Agriculture (NIFA) under project #2023-51150-41113 and by the Pennsylvania State University and NIFA Multistate/Regional Research Appropriations under project #NE2249. This work is/was also supported by the USDA National Institute of Food and Agriculture and Hatch Appropriations under Project #PEN05090 and Accession #7010745. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 3044. (CBDRB-FY25-P3044-R11547 & CBDRB-FY26-P3044-R12882)

1. Introduction

Fracking has brought significant economic benefits and costs, along with substantial immigration (Hausman and Kellogg 2015; Bartik et al. 2019; Brown et al. 2019; Abboud and Betz 2021; Wilson 2022; Rajbhandari et al. 2024).¹ The current paper investigates whether benefits accrue to long-time residents. We assess fracking impacts on labor market earnings of prior local residents using novel restricted-access data from the U.S. Census Bureau. Specifically, we use the restricted-access decennial census and American Community Survey (ACS) microdata linked with county-of-birth information from the Social Security Administration NUMIDENT file. We define fracking counties based on non-vertically drilled wells and compare treatment group individuals born in fracking counties to control group individuals born in counties with no fracking. We focus our analysis on non-metropolitan counties but also briefly consider metropolitan counties. Non-metropolitan counties have less population and employment than metro counties, and are therefore, likely more heavily influenced by a given magnitude resource shock. We estimate event analysis regressions that allow fracking impacts to vary by year to capture effects during the fracking boom, drilling downturn, and subsequent periods. Our main focus is on earnings, but we also briefly consider impacts on employment.

We find statistically significant positive impacts on average log earnings. Earnings effects are relatively modest in magnitude during the early years of the fracking boom but increase in magnitude as the boom intensified. Earnings effects decrease in magnitude during the drilling downturn but remain positive. We also find positive impacts of fracking on individual probability of employment during the boom period, but these effects shrink in magnitude and are not statistically significant after 2015. We also estimate separate results for

¹ Costs include negative externalities related to water pollution, air pollution, noise, congestion, earthquakes, and aesthetic decline (Muehlenbachs et al. 2015; Black et al. 2021; Blundell and Kokoza 2022; Agerton et al. 2023).

persons born in metropolitan counties. Metropolitan natives experience positive earnings impacts too, but the impacts are larger for persons born in non-metropolitan counties. We also provide evidence that the positive earnings effects are not just for mining industry workers, consistent with spillovers to other industries.

Our main analysis assigns the fracking treatment based on birth county and estimates pooled effects for stayers and leavers. Out-migrants likely receive less actual treatment from the fracking boom in their birth county, so our main estimates are expected to provide lower bounds of the causal impacts of residing in a fracking county for those born in one. Migration decisions are endogenous, and we cannot provide unbiased causal estimates for stayers and leavers separately, but our descriptive analysis in the appendix suggests that earnings impacts are larger for stayers than leavers as expected.

Our research fills important gaps in the literature by using restricted-access data to focus on a specific treatment group of prior residents who were born in fracking counties. We are the first study to use county of birth to assess fracking impacts. Previous research generally reports estimated impacts that do not distinguish among prior residents, in-migrants, and commuters (Munasib and Rickman 2015; Feyrer et al. 2017; Allcott and Keniston 2018; Upton and Yu 2021). While these studies typically report positive estimates on earned income in fracking areas, they may provide little insight into the impacts on prior residents if the new high-paying jobs go primarily to outsiders and they may be distorted by the changing composition of unobservable worker characteristics driven by in-migrants and long-distance commuters. Our study focuses on a specific group of individuals whose composition is not influenced by migration or long-distance commuting in response to the fracking boom and downturn. Again, we find positive effects on earnings that are especially large during the boom but partially persist

during the downturn and subsequent years. Our research provides important insights on whether and how local residents benefit from energy booms in their home area.

Original residents of fracking communities are an interesting population to study for several reasons. First, original residents are a politically important constituency for local policymakers. Many will spend their entire lives in the community, even long after the oil and gas industry has moved on. Fracking has been one of the most impactful developments in rural areas over the past two decades, and local policymakers are interested in how fracking is affecting longtime residents. Second, understanding how fracking has impacted original residents is important from a regional economic development perspective. Many economists favor place-neutral policies, which in many circumstances produce efficient outcomes. However, there is a strong case that a place-based perspective is more appropriate in many rural areas (Austin, Glaeser, and Summers 2018; Rodríguez-Pose 2018; Bartik 2020). The fortunes of longtime residents and the greater community are intrinsically connected and knowing how fracking impacts original residents is important to understanding regional economic growth. Our study's focus on original residents contributes to the place-based versus people-based policy debate by investigating one of the main contentions of supporters of people-based policies that economic development programs with the intention of spurring growth in lagging regions may simply attract migrant workers that absorb all of the employment and wage benefits, leaving original residents no better off than without the program. Our study is able to assess the fortunes of original residents in response to a sizeable economic shock and provides an empirical test of this assertion.

While our empirical approach is unique in focusing on fracking impacts for persons born in fracking counties, a handful of related studies have examined impacts of the fracking boom on

some other group of prior residents. A few studies examine the distribution of employment gains during local fracking booms and find that employment gains primarily go to nonresidents with smaller gains to residents (Wrenn et al. 2015; Gittings and Roach 2020; Guettabi and James 2020). Furthermore, persons who migrate for a new job will often establish a new local address and may be counted among residents, which could cause the estimated job gains for observed residents to be positive even if long-term (e.g. prior to the boom) residents are not affected.

Winters et al. (2021) is one of the few previous studies to address resource boom impacts on earnings of prior residents. However, they use the publicly available ACS, do not know individual birth county, and do not identify direct impacts of fracking on local populations. Winters et al. (2021) measure treatment based on birth state via the share of workers in an individual's birth state employed in the oil and gas industry, and their treatment variable does not differentiate between fracking vs. conventional oil and gas employment. States are often broad geographic areas with considerable variation in fracking activity and labor market impacts across local areas within them. Furthermore, the current study focuses on impacts for persons born in non-metropolitan fracking counties, which experienced the bulk of fracking activity in the US, while the labor market impacts could be largely captured by long-distance commuting urban residents. We find that earnings impacts for non-metropolitan natives are larger than for metropolitan natives, which cannot be assessed via treatment based on birth state. We also differ from Winters et al. (2021) by measuring treatment based on actual fracking activity in an area as opposed to oil and gas employment shares. While Winters et al. (2021) show that birth-state oil and gas booms raised average earnings, the current paper provides the first county-level, fracking-specific, dynamic evidence that prior residents of non-metropolitan fracking counties experienced sizable earnings gains that peak during the boom and partially persist during the

downturn. The focus on the fracking-specific treatment is especially important given the sizable localized negative externalities associated with fracking. Our focus on non-metropolitan area county-level treatment is important for understanding whether those most exposed to fracking and its negative externalities actually experience local labor market benefits.

Jacobsen et al. (2023) is another important study examining longer run impacts of resource booms on a fixed group of residents. They examine earnings impacts of the 1970s and 1980s U.S. oil boom and bust on persons residing in oil producing counties in 1980, which was near the height of the boom.² They use the Panel Study of Income Dynamics and find positive effects on earnings during the boom but negative effects during the bust. However, the post-2000 fracking boom was driven by technological developments, while the 1970s and 1980s oil boom was driven by artificially high prices due to OPEC supply restrictions that eventually collapsed, which likely caused the fracking downturn and subsequent period impacts to differ from impacts during the 1980s oil bust.³

2. Fracking Data and Trends

Hydraulic fracturing (fracking) is a drilling technique that combines non-vertical drilling and fracking fluid injections to extract oil and natural gas deposits trapped in shale and tight rock formations. Fracking was generally rare until the early 21st century because of high costs and low productivity. However, fracking innovations lowered costs, increased profits, and

² They also use an indicator for residing in an oil producing county in 1973 as an instrumental variable.

³ Street (2023) examines impacts of the North Dakota fracking boom on crime and finds reduced criminal behavior for prior residents. Katovich et al. (2026) examine impacts of the Brazilian oil and gas boom driven by energy prices and new oil and gas discoveries and find evidence that the benefits disproportionately accrued to early labor market entrants. Other studies have also examined impacts of resource booms and busts for coal, which likely differ from impacts from oil and gas fracking (Black et al. 2005; Betz et al. 2015; Haywood et al. 2024; Rud et al. 2024).

transformed U.S. energy production in the 21st century (Brown 2014; Munasib and Rickman 2015; Feyrer et al. 2017).

Figure 1 illustrates ups and downs of fracking activity as measured by non-vertical oil and gas wells completed by year since 1990. These data are from Enverus (formerly DrillingInfo). Fracking levels in the U.S. were relatively low and flat during the 1990s and early 2000s before turning upward in the mid-2000s. Fracking boomed in the early 2010s and peaked in 2014. Non-vertical wells drilled decreased considerably in 2015 and 2016 due to falling energy prices. Fracking activity then rebounded in 2017-2019 and then fell again in 2020 with the onset of the COVID-19 pandemic and significantly decreased demand for travel and oil.⁴ 2021 and 2022 saw increased fracking as the world economy began recovering from the pandemic. Russia's invasion of Ukraine in February 2022 further increased energy prices and fracking (Sun et al. 2024). Notably, the number of non-vertical wells never reverted back to pre-fracking boom numbers. For this reason, we refer to the post-boom period of decreased fracking activity as a downturn and not a bust. Our use of the term downturn refers to the downturn in the number of new non-vertical wells drilled and not the volume of oil and gas produced, which also depends on the number of prior wells drilled.

Oil and gas resources and fracking activity are also very unequally dispersed across the U.S. Figure 2 divides U.S. counties into three groups based on fracking activity during the 2006-2014 boom period.⁵ Non-fracking counties are those with zero non-vertical oil and gas wells completed during the entire 2006-2014 period. Most counties in the U.S. have zero fracking. For counties with non-vertical wells, we compute the non-vertical well density as the total

⁴ Oil prices even went negative briefly in April 2020 due to the massive reduction in demand during early days of the pandemic (Reed and Krauss 2020).

⁵ As discussed later, 2006 is chosen as the initial year of the fracking boom for our analysis.

number of non-vertical wells in the county divided by the total geographic area of the county. We next compute the median non-vertical well density across counties and then define the high fracking group as the set of counties with non-vertical well density above the median and low fracking group as the set of counties with positive non-vertical wells but below the median non-vertical well density. The high fracking counties are geographically clustered but also span several states.⁶ The low fracking counties also tend to be clustered near each other and near the high fracking counties. These notions of high fracking, low fracking, and non-fracking counties will help guide our research design discussed in Section 4.

3. Census and ACS Data

We combine the fracking data with individual records from restricted-access versions of the 2000 decennial census and 2005-2022 American Community Survey (ACS). The 2000 census long form survey was administered to roughly one-sixth of the U.S. population and each year of the ACS includes roughly one percent of the U.S. population. These are independently drawn repeated cross-sections and not panels following the same individuals over time. These files contain information on individual earnings, employment, age, sex, race, ethnicity, county of residence⁷, and other characteristics. We merge these files with county-of-birth and year-of-birth information from the Social Security Administration NUMIDENT file for persons born in the U.S., with a match rate greater than 90 percent using the matching algorithm described in Taylor et al. (2016) that builds on work by Isen et al. (2017) and Black et al. (2015). More details on

⁶ These include Alabama, Arkansas, California, Colorado, Kansas, Kentucky, Louisiana, Mississippi, Montana, New Mexico, New York, North Dakota, Ohio, Oklahoma, Pennsylvania, South Dakota, Texas, Utah, Virginia, West Virginia, and Wyoming.

⁷ The publicly available versions of these files only partially identify county of residence. In order to protect individual confidentiality, no geographic area with less than 100,000 people can be identified in the public data; less populous counties are combined via Public Use Microdata Areas (PUMAs). The restricted access data allows much greater geographic identification.

using the NUMIDENT file can be found in Bailey et al. (2021) and Bailey et al. (2024). Persons who are foreign-born or who cannot be matched to a birth county are excluded from the analysis.

We exclude the 2001-2004 ACS from our main analysis because these years contained much smaller samples and number of counties that complicates comparisons, but we do include these years in supplemental analysis discussed below. Our analytic sample also excludes persons born in Alaska and Connecticut because of significant changes to county boundaries during our sample period. However, persons born in other states but residing in Alaska and Connecticut during the ACS are included in the analysis. Alaska and Connecticut contain no counties that would otherwise be in the high fracking treatment group. We limit our main analysis to persons ages 25-54 to focus on the prime-age workforce, but we also conduct supplemental analysis for ages 18-24 and 55-64 reported in the appendix.

Our main outcome of interest is the natural log of annual earned income during the previous year.⁸ Persons with non-positive annual income are excluded from the main earnings analysis.⁹ We also examine individual employment outcomes including a dummy for being employed at all during the reference week and a dummy for being employed in the mining industry.¹⁰

⁸ We do not account for spatial cost of living differences due to limited data and conceptual ambiguity concerning housing costs. In particular, fracking could affect local housing values and housing wealth for prior homeowners in affected areas. Rising housing costs would harm renters but benefit previous owners.

⁹ The appendix includes supplemental results that include persons with non-positive earnings in the analysis.

¹⁰ Mining industry employment includes oil and gas, coal, metals, non-metals, unspecified, and various support activities. We considered examining oil and gas employment separately as an outcome, but there are complications, including disclosure restrictions on small samples, that made it more practical to examine overall mining industry employment instead, which will be heavily driven by oil and gas employment in fracking counties.

4. Research Design

This section describes our empirical approach. We first present our event study specification, then detail our treatment and control definitions, and finally discuss our use of birth county for the treatment measure.

We estimate event analysis regressions of the following form:

$$Y_{ict} = \sum_{t \neq 2000} \theta_t TG_c \times YearDum_t + \beta X_{ict} + \alpha_c + \delta_{r(c)t} + \varepsilon_{ict}.$$

Y_{ict} is an outcome for individual i born in county c and observed in survey year t . The main outcome is log annual earnings. TG_c is a treatment group indicator for persons born in counties treated by high fracking levels. $YearDum_t$ is a dummy variable for survey year t , and $TG_c \times YearDum_t$ is a variable for the interaction of the treatment group indicator and year t dummy. θ_t is a time-varying coefficient on $TG_c \times YearDum_t$ for the average treatment effect on the treated (ATT). We include a full set of $TG_c \times YearDum_t$ interactions in the regression except for year 2000, which we define as the pre-event base year for the analysis. X_{ict} includes detailed dummy variable controls for individual characteristics including age, sex, race, ethnicity, and birth-year, α_c is a set of birth-county fixed effects, $\delta_{r(c)t}$ is a set of birth-region-by-survey-year fixed effects, and ε_{ict} is an error term. We intentionally do not include explanatory variables such as current location, education, and marital status that are potentially affected by local economic conditions during the fracking era because they would be “bad controls” in the language of Angrist and Pischke (2009). Our inclusion of birth-region-by-survey-year fixed effects means that treatment-control comparisons are made within the same

birth region; the four census regions are the Northeast, Midwest, South, and West.¹¹ We cluster standard errors by birth county to account for possible serially correlated errors within counties.

Our event analysis specification is a more general version of a basic difference-in-differences (DID) regression model where the average treatment effect on the treated is the coefficient (θ) on the interaction of a treatment group indicator (TG_c) and a post-treatment time period indicator ($Post_t$):

$$Y_{ict} = \theta TG_c \times Post_t + \beta X_{ict} + \alpha_c + \delta_{r(c)t} + \varepsilon_{ict}.$$

The event analysis regression is more general and a better fit for our setting because it does not force the ATT to be the same over time. Fracking intensity varied over time, so we expect that the ATT would vary over time. A major strength of our event analysis is that it is not heavily reliant on specific choices about when the boom began or ended. Instead, we estimate separate ATT coefficients for each year during 2005-2022. We expect that earnings coefficients will be largest in years with high levels of fracking and decrease during the downturn. We select year 2000 as our pre-event base year for two main reasons. First, it clearly predates the onset of widespread fracking activity in the United States, providing a clean baseline measure of local economic conditions before the fracking boom. Second, the 2001-2004 ACS samples are relatively small and unlikely to provide fully representative samples of persons born in non-metropolitan areas. The restricted sample sizes and incomplete population coverage make the 2001-2004 ACS data less suitable for establishing our baseline period. In contrast, the 2000 Census long-form sample provides comprehensive population coverage and larger sample sizes, making it ideal for our pre-boom reference point. As further consideration, the fracking boom

¹¹ We also considered including birth-state-by-survey-year fixed effects, but long-distance spillovers within states would cause this specification to yield coefficient estimates attenuated toward zero, even more so than with birth-region-by-survey-year fixed effects. Birth-state-by-survey-year fixed effects would also exclude some states with no control group and create a small implicit sample and challenges with disclosure restrictions.

was not widely dispersed but had partially begun in Texas in 2005, making the first full-coverage year of the ACS also less ideal for the base year and leaving year 2000 as the best choice.

Estimating causal impacts of fracking has some complications including two important competing issues. First, counties differed in their actual timing and intensity of fracking activity which could induce heterogeneous impacts. Second, estimating unbiased ATTs in settings with staggered interventions and continuous treatments is difficult and often impractical, especially when effects may depend on the entire path of treatment including lags and leads and involve heterogeneous impacts. Our response to this complexity is to aggregate the range of treatments considered (Callaway et al. 2024). Our event analysis framework uses a binary treatment group indicator and intentionally differs from other studies using continuous treatment measures. We define a set of high fracking treatment group counties and assign them to have the same treatment indicator variable (TG_c) but allow the ATT to vary over time (θ_t) via our event analysis framework. This does not rule out the possibility of heterogeneous treatment effects across treatment group counties in a given year, but we are unable to estimate such heterogeneous effects, in part because Census Bureau confidentiality protections limit what we can disclose for small samples and limit the overall number of estimates we can disclose. Small samples would also mean less statistical power and less reliable estimates overall. Thus, our event analysis framework aggregates high fracking counties into a single treatment group with a binary treatment variable and estimates average treatment effects among treated counties.

Our event analysis research design requires defining treatment and control groups. We define the treatment group as persons born in counties with high levels of fracking during the 2006-2014 fracking boom period. We define 2006-2014 as the fracking boom period based on national trends in non-vertical well completion data as illustrated in Figure 2 and following

established literature (Rickman et al. 2017; Winters et al. 2021; Cascio and Narayan 2022). The boom period extends through 2014, after which drilling activity decreased substantially due to falling global energy prices. Specifically, the high-fracking treatment group includes persons born in counties with fracking density during 2006-2014 above the median fracking density among counties with any fracking. However, we exclude persons born in four high-fracking counties in New York state from the analysis entirely because of the fracking ban implemented in New York state beginning in 2010.¹² We limit the control group to persons born in counties with zero non-vertical wells during the 2000-2022 period.¹³ Persons born in low fracking counties are not appropriate candidates for the control group because they did receive some treatment but less treatment than the high-fracking counties.¹⁴ We also anticipate that non-fracking counties in close proximity to high-fracking counties could experience positive spillover effects. Feyrer et al. (2017) estimate that the bulk of oil and gas labor market spillovers occur within 100 miles. Thus, we also exclude from the control group persons born in any county that is within 100 miles of a high fracking county.

Our main analysis limits the sample to persons born in non-metropolitan counties because non-metro counties have less initial population and employment and are expected to be more impacted by a given magnitude natural resource shock. We use June 2003 metropolitan area definitions. We argue that limiting both the treatment and control group to non-metropolitan

¹² As a practical matter, including or excluding these high-fracking New York counties should be very unlikely to meaningfully affect our ATT estimates because they would be a very small portion of the high-fracking county sample, but we are unable to address this quantitatively because Census Bureau confidentiality protections prevent us from disclosing alternative estimates for similar but slightly different samples that would create a small “implicit” sample from the difference of the two samples because of elevated disclosure risk for the small implicit sample.

¹³ This time period restricting the control group is intentionally broader than the 2006-2014 boom period used to construct the treatment group to ensure that the control group counties have zero fracking during our full period of analysis.

¹⁴ We will later conduct brief analysis comparing low fracking counties to non-fracking counties.

counties will create more accurate comparisons between the two. We do also briefly consider separate analysis for persons born in metropolitan counties.

Table 1 presents selected summary statistics for 2000, our pre-boom base year, separately for our treatment and control groups. Table 1 shows that the treatment group had lower mean log annual earnings than the control group with the difference of 0.05 corresponding to roughly five percent lower earnings in the treatment group prior to the boom. The treatment group also had 5.1 percentage point lower employment rates. Sample means for mining employment are suppressed because of potential disclosure concerns related to small cell sizes.¹⁵ The treatment counties in year 2000 are slightly older, slightly more female, and have less formal education.

4.1 Measuring Treatment via County of Birth

We emphasize that our main analysis constructs treatment and control groups based on county of birth. Most previous studies examining natural resource impacts do not separate the impacts on prior residents from those experienced by in-migrants and non-residents. In-migrant and non-resident workers are likely endogenously induced to work in resource boom areas during the boom and leave during the downturn, which alters the composition of the workforce in ways that hinder interpretation of estimated impacts. Estimated impacts on combined samples of prior residents, in-migrants, and commuters also tell us little about the impacts on long-term residents who are especially likely to bear environmental and health costs from fracking activity. Instead, we focus on a specific group of prior residents whose initial location decisions were determined by their parents long before fracking innovations transformed U.S. energy

¹⁵ We consulted publicly available data and estimate that among all individuals in the 2000 5% Census PUMS ages 25-54 and born in the United States, 0.12 percent were employed in mining. This is higher in some areas than others, but the overall low rate hinders disclosure for our restricted-access data.

production. I.e., we examine the impact of the fracking boom and downturn on persons born in fracking counties.

Admittedly, there are also issues with measuring treatment based on birth county. Specifically, many people move away from their birth county for reasons such as education, employment, family considerations, and quality of life. Among persons born in a fracking county, some will stay, some will move to another fracking county, and some will move to a non-fracking county. Similarly, some people born in a non-fracking county may move to a fracking county, including some who move prior to or independent of the fracking boom. Thus, our main analysis estimates effects of being born in a fracking county, but this is likely not the same as the effects of residing in a fracking county. However, being born in a high fracking county is strongly related to living in a high fracking county later in life, which makes it a useful proxy for the latter. Table 1 also shows that 55.8 percent of the treatment group sample resides in a county in 2000 that becomes a high fracking county during the boom period; only 4.3 percent of the control group reside in a future high fracking county.¹⁶

Because being born in a fracking county is an imperfect proxy for living in a fracking county during a fracking boom, effects for the former likely underestimate effects for the latter somewhat similar to how intent-to-treat (ITT) estimates generally underestimate the ATT. The effects of residing in a fracking county are expected to be larger than the effects of being born in a fracking county because 1) persons who migrate from fracking counties will have less exposure to fracking and likely less positive impacts on wages than persons who are born in a fracking county and continue to live in a fracking county and 2) some persons who are born in non-

¹⁶ These percentages also increase slightly over time as the fracking boom makes fracking counties more attractive labor markets. In results not shown, we also examined the extent to which the treatment vs. control group gap in fracking county residence increases over time; the increased gap never exceeded two percentage points.

fracking counties move to fracking counties and experience earnings increases as a result. The appendix provides a more formal presentation of conditions for which our main estimates provide a lower bound for fracking impacts on long-term residents who are both born in a fracking county and reside in their birth county as adults along with some supporting descriptive analysis comparing stayers and leavers that offers suggestive evidence that impacts are larger for stayers than leavers as expected. However, we emphasize that fracking impacts on persons born in fracking counties are interesting and important in themselves, even independent of our ability to interpret them as a lower bound for persons residing in fracking counties. Thus, our analysis provides direct evidence on fracking impacts for one important group and indirect suggestive evidence on fracking impacts for another group.

4.2 Parallel Trends

Another challenge for our study is that we do not have a long pre-event period to rigorously assess parallel trends between the treatment and control group. The event analysis estimator assumes that treatment and control group counties were on parallel trends prior to the fracking shock to the treatment group and would have continued on parallel trends without the fracking shock. Our primary outcome data are from the 2000 census and 2005-2022 ACS.

We can offer some suggestive evidence related to parallel trends. Our event analysis estimates below can yield insights on parallel trends and the validity of estimates by comparing the event coefficients over time. Figure 3 builds on Figure 1 and shows mean NV wells by year for our high fracking treatment group; recall that Figure 1 reported the national total including wells in low fracking counties. Figure 3 indicates a positive mean for NV wells that was largely flat for 1990-2002. The fracking increase began thereafter but was moderate during the mid- and

late 2000s, so we expect modest effects on earnings early on that should grow over time with the intensifying fracking boom. The strongest boom in fracking activity was during the 2011-2014 period, so we expect the strongest impacts on earnings around these years. We expect smaller impacts on earnings during the downturn and subsequent years that could become not statistically significant or possibly even negative. When we interpret the empirical results, we will consider if the coefficients by year follow the expected pattern. We also conduct supplemental analysis in the appendix that includes the 2001-2004 ACS to offer potential insights on parallel trends.

5. Empirical Results

Our event analysis results are illustrated via figures. For all figures, 2000 is the omitted reference year. Coefficient estimates for each year are from the interaction of the treatment group indicator and given year indicator and indicate how the outcome means changed since 2000 for the treatment group relative to the control group. Annual earnings in the ACS for each individual are for the 12 months prior to the survey, but the ACS is conducted across households throughout the year, so the coefficients by year are ultimately an average of impacts during the current year and immediately preceding year. The employment indicator outcomes are based on the survey week, so results more fully line up with the survey year.

5.1 Baseline Event Analysis Results

Figure 4 presents log annual earnings event analysis regression results for our main sample. The treatment group by year coefficients are statistically significantly different from zero at the five percent level in every year except 2005 and 2021. The small and not statistically

significant coefficient of 0.016 for 2005 is consistent with relatively minimal fracking activity early on. The relatively small coefficient in 2021 is largely attributable to reduced energy prices and production due to the COVID-19 pandemic. The earnings coefficients are largest in magnitude for 2010-2015, which roughly corresponds to the peak of the fracking boom. Earnings coefficients declined as fracking activity declined during the downturn and subsequent periods. The earnings coefficients in Figure 4 peak at 0.093 in survey year 2015, a year in which fracking activity already began declining. The large coefficient for 2015 is partially driven by earnings in calendar year 2014 and could also reflect factors such as nominal wage rigidity and severances packages for laid off workers. Coefficient estimates for 2010-2014 are also consistently above 0.08. These results indicate that the peak fracking boom increased mean annual earnings of persons born in fracking counties by more than eight percent for multiple years. Overall, the pattern of earnings coefficients follows the time pattern for fracking activity in Figure 1. Notably, the downturn and subsequent period coefficient estimates are all positive.

Figure 5 presents event analysis results for the probability of being employed during the survey week. We find significant coefficients for every year 2006-2014 but not for other years. The coefficient peaks at 0.023 for 2009 and is 0.015 in 2014. The relatively large coefficient for 2009 may be due to increased economic resilience during the Great Recession for high fracking counties. The significant positive effects on employment probability disappear during the fracking downturn and do not return, but there are no significant negative coefficients for any period. One possible factor is the substantial increase in labor productivity in the oil and gas sector that began to increase fairly rapidly in 2012 and accelerated during the downturn. As fracking technology matured, the average output per worker increased nearly fourfold for oil (Appendix Figure A33) and 2.5 times for gas (Appendix Figure A34), leading to a large decline

in oil and gas employment (Appendix Figure A35) while overall production continued to increase throughout the period.

Figure 6 takes a specific look at employment in the mining industry as an outcome. Mining employment coefficients are significantly positive in all years 2005-2019. Coefficient estimates for 2020-2022 are small and not statistically significant. Mining is a relatively high-paying industry, especially oil and gas production. Together with the previous figures, these results suggest that access to oil and gas production jobs is likely a partial mechanism for overall increases in earnings and employment. We will consider this further in the next sub-section.

5.2 Sub-Sample Analysis

We next further consider the role of oil and gas employment on earnings by separating the sample into those employed in the mining industry and those employed in other industries. One possibility is that the effects may be dominated by increased earnings for persons in the mining industry with little or no impact on workers in other industries. However, previous literature has often found evidence of spillovers from fracking to other industries, which could increase labor demand and wages for workers outside the mining industry. Figure 7 presents results for mining industry workers and Figure 8 presents results for non-mining workers. Of course, there is likely non-random selection of workers to industries, so splitting the sample in this way is highly imperfect. Still, we view it as suggestive evidence. Also, the sample sizes for mining workers are relatively small, so results are noisier than for the full sample. Figure 7 indicates positive and statistically significant coefficient estimates for mining industry workers in most years but not all. Some of the coefficient estimates for years that are not significant are not small in magnitude, so we are not generally confident in interpreting these as non-effects. The

results in Figure 8 are more precise. Notably, coefficient estimates are positive and statistically significant for all years but 2005 and 2021. Overall, the coefficients in 8 follow the same pattern as the full sample estimates in Figure 4, which is largely to be expected since non-mining workers are the overwhelming majority of the sample. The coefficient estimates for non-mining workers are moderately smaller in magnitude than the full sample in Figure 4, but they are still sizable. E.g., the earning coefficients peak in 2015 at 0.093 for the full sample and at 0.078 for the non-mining sample. Earnings coefficients are never negative and statistically significant.

Two explanations likely underlie why an oil and gas boom might raise non-mining earnings. As the oil and gas sector expands and rapidly attracts a large share of the low-skill workers in the area, the overall local labor market tightens, especially in sectors that also employ low-skill workers. E.g., many businesses that rely on low skill labor (e.g. Walmart or fast food establishments) have substantial staffing shortages as much of the local labor force is absorbed into the mining sector.¹⁷ Migrants can help meet the increase in labor demand, but migrants come primarily for lucrative oil and gas jobs and do not fully alleviate labor market tightening. Consequently, earnings rise across all other sectors. Additionally, the injection of income into a local community during a fracking boom increases demand for local goods and services and can increase labor demand and earnings further. Overall, this suggests that the positive effect of fracking on earnings is not just due to some workers getting high paying jobs in the oil and gas industry. Instead, the benefits spillover via increased earnings in other industries.

5.3 Metropolitan and Weak Treatment Sample Analysis

¹⁷ Additional results in the appendix report earnings impacts for workers employed in the restaurant and retail sector and other sub-samples including by education level, sex, and age group.

Our main analysis is limited to persons born in non-metropolitan counties. We next examine earnings impacts for the sample of persons born in metropolitan counties with results presented in Figure 9. Coefficient estimates are positive and statistically significant for years 2008-2019 but are not significant for the other years. Coefficient estimates for the metro native sample are generally smaller than those for the non-metro native sample in Figure 4. Thus, metro natives benefit from being born in a high fracking county but not as much as workers born in a non-metropolitan high-fracking county.

Figure 10 returns to examining persons born in non-metropolitan counties but alters the treatment group to persons born in low fracking counties, i.e., counties with some fracking but below the median fracking intensity among counties with positive fracking. The control group for Figure 10 is the same as the main analysis in Figure 4. The Figure 10 treatment group is exposed to fracking but receives a weaker dose of the treatment compared to the high-fracking treatment group in the main analysis. Consequently, we expect that impacts in Figure 10 to be directionally similar but smaller in magnitude than those in Figure 4. That is what we observe. Coefficients for Figure 10 are positive and significant for several years but are overall smaller in magnitude and less consistently significant than for the high fracking treatment group. Coefficients are never negative and significant. Thus, having moderate fracking also yields positive earnings impacts, but not as much as more intensive fracking.

5.4 Including the 2001-2004 and Examining Alternative Base Years

While there are important limitations with the 2001-2004 ACS noted above, we also conduct supplementary analysis that includes these years. Results reported in the appendix indicate that our main findings are qualitatively robust. The 2001-2004 coefficients help address

potential concerns about parallel trends. The appendix results indicate that the 2001-2004 coefficients are generally small and not statistically significant, consistent with similar trends between the treatment and control group during 2000-2004 prior to the fracking boom.

We also report results in the appendix examining alternative base years. Year 2000 is our preferred base year for reasons noted above. We explore separately changing the base year to 2002 and 2005. Figure 3 suggests that fracking may have crossed an inflection point around 2003, making 2002 a potentially interesting benchmark. Other scholars have suggested the boom took off around 2006 (Rickman et al. 2017; Cascio and Narayan 2022), making 2005 an interesting benchmark. Results using these alternative benchmark years are qualitatively similar to our main results.

6. Conclusion

Fracking transformed energy production in the United States and altered local labor markets with frackable oil and gas resources. The current study is the first to document the impacts of the fracking boom on persons born in fracking counties, using novel restricted-access data to make this possible. We focus on non-metropolitan areas and find that the fracking boom increased earnings for this important group of prior residents by more than eight percent during 2010-2015. The earnings impacts decreased during the fracking downturn and subsequent periods but remained positive and statistically significant. Earnings growth for 2000-2022 was five percent higher for persons born in high-fracking counties compared to those born in control counties. We also find positive effects on the probability of employment during the boom period that fade out during the downturn and subsequent years. We also find suggestive evidence that the fracking boom impacts earnings via three channels: higher likelihood of mining employment,

higher earnings in mining employment, and higher earnings in non-mining employment. The latter channel suggests an important role for spillover effects from oil and gas to other industries.

While our study documents positive labor market impacts, it does not provide an answer as to whether prior residents are better off overall. Fracking generates numerous externalities, some of which individuals strongly try to avoid and some of which they cannot avoid.

Furthermore, the benefits and costs likely vary across local residents. Some local residents may reap little to no labor market benefits yet bear substantial costs to their physical and mental health from fracking in their area. These are important to consider when crafting policy responses to fracking.

References

- Abboud, A., and M. R. Betz. 2021. The local economic impacts of the oil and gas industry: Boom, bust and resilience to shocks. *Energy Economics* 99:105285.
- Agerton, M., B. Gilbert, and G. B. Upton Jr. 2023. The economics of natural gas flaring and methane emissions in US shale: An agenda for research and policy. *Review of Environmental Economics and Policy* 17 (2):251–273.
- Allcott, H., and D. Keniston. 2018. Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. *Review of Economic Studies* 85 (2):695–731.
- Angrist, J. D., and J. S. Pischke. 2009. *Mostly harmless econometrics: An empiricist's companion*. Princeton, NJ: Princeton University Press.
- Austin, B. A., E. L. Glaeser, and L. H. Summers. 2018. Jobs for the heartland: Place-based policies in 21st century America. NBER Working Paper 24548, National Bureau of Economic Research, Cambridge, MA.
- Bailey, M. J., H. Hoynes, M. Rossin-Slater, and R. Walker. 2024. Is the social safety net a long-term investment? Large-scale evidence from the food stamps program. *Review of Economic Studies* 91 (3):1291–1330.
- Bailey, M. J., S. Sun, and B. Timpe. 2021. Prep school for poor kids: The long-run impacts of Head Start on human capital and economic self-sufficiency. *American Economic Review* 111 (12):3963–4001.
- Bartik, T. J. 2020. Place-based policy: An essay in two parts. W.E. Upjohn Institute for Employment Research Policy Paper 2020-021. Kalamazoo, MI.

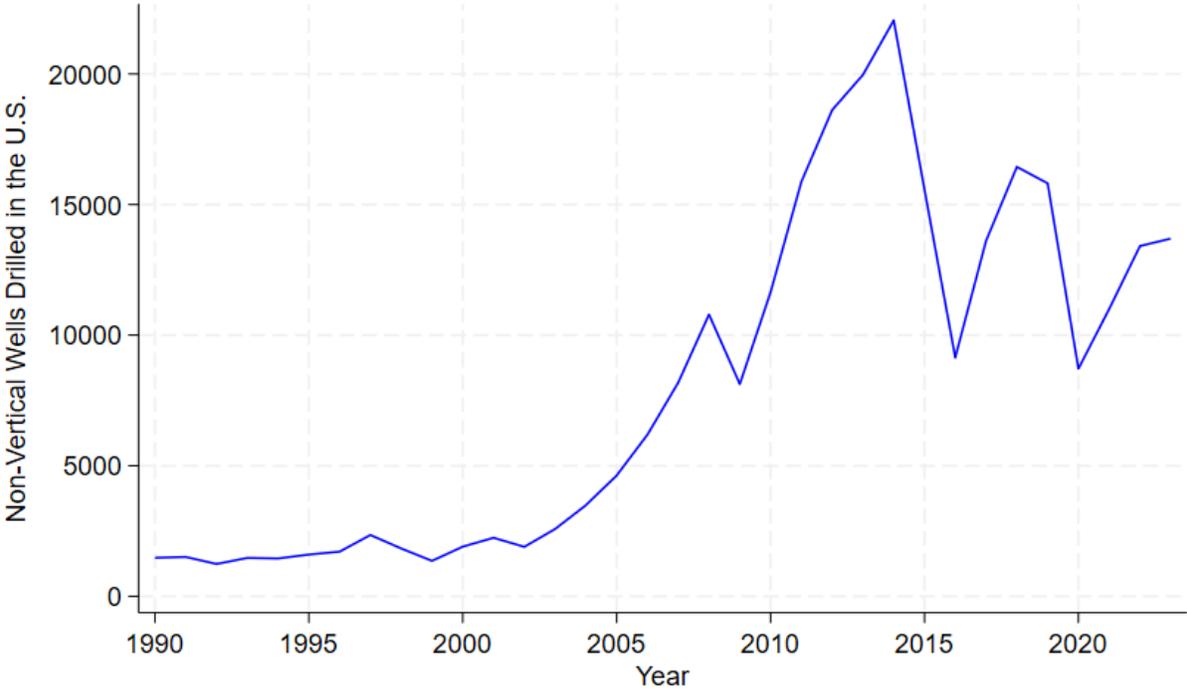
- Bartik, A. W., J. Currie, M. Greenstone, and C. R. Knittel. 2019. The local economic and welfare consequences of hydraulic fracturing. *American Economic Journal: Applied Economics* 11 (4):105–155.
- Betz, M. R., M. D. Partridge, M. Farren, and L. Lobao. 2015. Coal mining, economic development, and the natural resources curse. *Energy Economics* 50:105–116.
- Black, D., T. McKinnish, and S. Sanders. 2005. The economic impact of the coal boom and bust. *Economic Journal* 115 (503):449–476.
- Black, D. A., S. G. Sanders, E. J. Taylor, and L. J. Taylor. 2015. The impact of the Great Migration on mortality of African Americans: Evidence from the Deep South. *American Economic Review* 105 (2):477–503.
- Black, K. J., A. J. Boslett, E. L. Hill, L. Ma, and S. J. McCoy. 2021. Economic, environmental, and health impacts of the fracking boom. *Annual Review of Resource Economics* 13:311–334.
- Blundell, W., and A. Kokoza. 2022. Natural gas flaring, respiratory health, and distributional effects. *Journal of Public Economics* 208:104601.
- Brown, J. P. 2014. Production of natural gas from shale in local economies: A resource blessing or curse? *Economic Review* 99 (1):119–147.
- Brown, J. P., T. Fitzgerald, and J. G. Weber. 2019. Does resource ownership matter? Oil and gas royalties and the income effect of extraction. *Journal of the Association of Environmental and Resource Economists* 6 (6):1039–1064.
- Callaway, B., A. Goodman-Bacon, and P. H. C. Sant’Anna. 2024. Event studies with a continuous treatment. *AEA Papers and Proceedings* 114:601–605.

- Cascio, E. U., and A. Narayan. 2022. Who needs a fracking education? The educational response to low-skill-biased technological change. *ILR Review* 75 (1):56–89.
- Feyrer, J., E. T. Mansur, and B. Sacerdote. 2017. Geographic dispersion of economic shocks: Evidence from the fracking revolution. *American Economic Review* 107 (4):1313–1334.
- Gittings, R. K., and T. Roach. 2020. Who benefits from a resource boom? Evidence from the Marcellus and Utica shale plays. *Energy Economics* 87:104489.
- Guettabi, M., and A. James. 2020. Who benefits from an oil boom? Evidence from a unique Alaskan data set. *Resource and Energy Economics* 62:101200.
- Hausman, C., and R. Kellogg. 2015. Welfare and distributional implications of shale gas. *Brookings Papers on Economic Activity*:71–125.
- Haywood, L., M. Janser, and N. Koch. 2024. The welfare costs of job loss and decarbonization: Evidence from Germany’s coal phaseout. *Journal of the Association of Environmental and Resource Economists* 11 (3):577–611.
- Isen, A., M. Rossin-Slater, and W. R. Walker. 2017. Every breath you take—every dollar you’ll make: The long-term consequences of the Clean Air Act of 1970. *Journal of Political Economy* 125 (3):848–902.
- Jacobsen, G. D., D. P. Parker, and J. B. Winikoff. 2023. Are resource booms a blessing or a curse? Evidence from people (not places). *Journal of Human Resources* 58 (2):393–420.
- Katovich, E., D. Parker, and S. Poelhekke. 2026. Timing is everything: Labor market winners and losers during boom-bust cycles. *Journal of the Association of Environmental and Resource Economists* 13 (1):229–260.
- Muehlenbachs, L., E. Spiller, and C. Timmins. 2015. The housing market impacts of shale gas development. *American Economic Review* 105 (12):3633–3659.

- Munasib, A., and D. S. Rickman. 2015. Regional economic impacts of the shale gas and tight oil boom: A synthetic control analysis. *Regional Science and Urban Economics* 50:1–17.
- Rajbhandari, I., A. Faggian, and M. Partridge. 2024. Migrants and boomtowns: Evidence from the US shale boom. *Regional Studies* 58 (11):2087–2098.
- Reed, S., and C. Krauss. 2020. Too much oil: How a barrel came to be worth less than nothing. *New York Times*, April 20.
- Rickman, D. S., H. Wang, and J. V. Winters. 2017. Is shale development drilling holes in the human capital pipeline? *Energy Economics* 62:283–290.
- Rodríguez-Pose, A. 2018. The revenge of the places that don't matter (and what to do about it). *Cambridge Journal of Regions, Economy and Society* 11 (1):189–209.
- Rud, J. P., M. Simmons, G. Toews, and F. Aragón. 2024. Job displacement costs of phasing out coal. *Journal of Public Economics* 236:105167.
- Street, B. 2023. The impact of economic opportunity on criminal behavior: Evidence from the fracking boom. SSRN Working Paper.
- Sun, M., X. Cao, X. Liu, T. Cao, and Q. Zhu. 2024. The Russia-Ukraine conflict, soaring international energy prices, and implications for global economic policies. *Heliyon* 10 (16):e34712.
- Taylor, E. J., B. A. Stuart, and M. J. Bailey. 2016. Summary of procedure to match NUMIDENT place of birth county to GNIS places. CES Technical Note 16-01, Center for Economic Studies, U.S. Census Bureau.
- Upton Jr, G. B., and H. Yu. 2021. Labor demand shocks and earnings and employment differentials: Evidence from the US shale oil and gas boom. *Energy Economics* 102:105462.

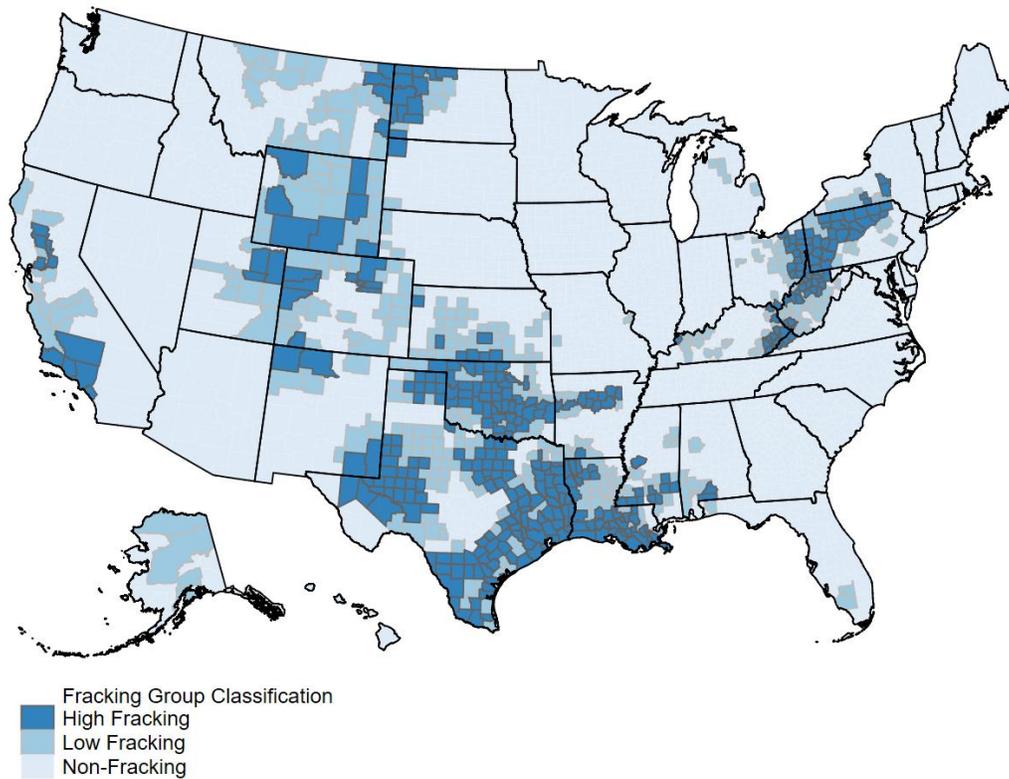
- Wilson, R. 2022. Moving to economic opportunity: The migration response to the fracking boom. *Journal of Human Resources* 57 (3):918–955.
- Winters, J. V., Z. Cai, K. Maguire, and S. Sengupta. 2021. Causal effects of the fracking boom on long-term resident workers. *Journal of Regional Science* 61 (2):387–406.
- Wrenn, D. H., T. W. Kelsey, and E. C. Jaenicke. 2015. Resident versus nonresident employment associated with Marcellus shale development. *Agricultural and Resource Economics Review* 44 (1):1–19.

Figure 1: Oil and Gas Fracking Activity in the U.S. Since 1990



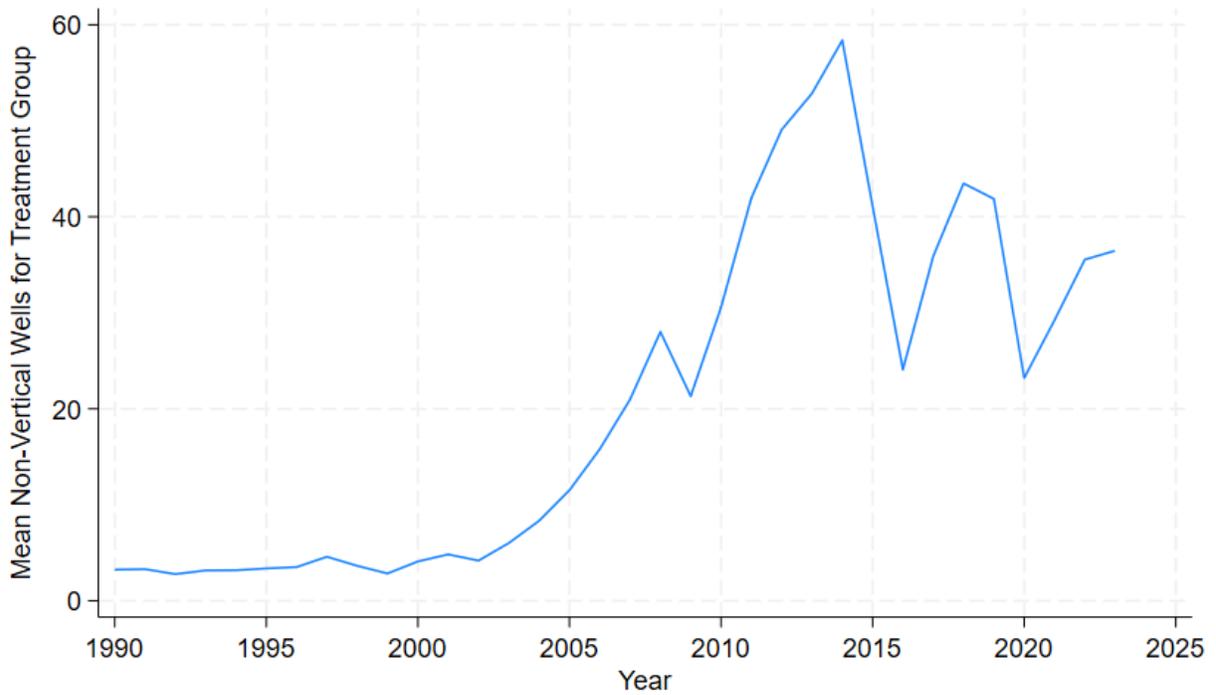
Notes: based on authors' calculations from Enverus data on oil and gas non-vertical wells completed by year.

Figure 2: Fracking Activity across U.S. Counties During 2006-2014



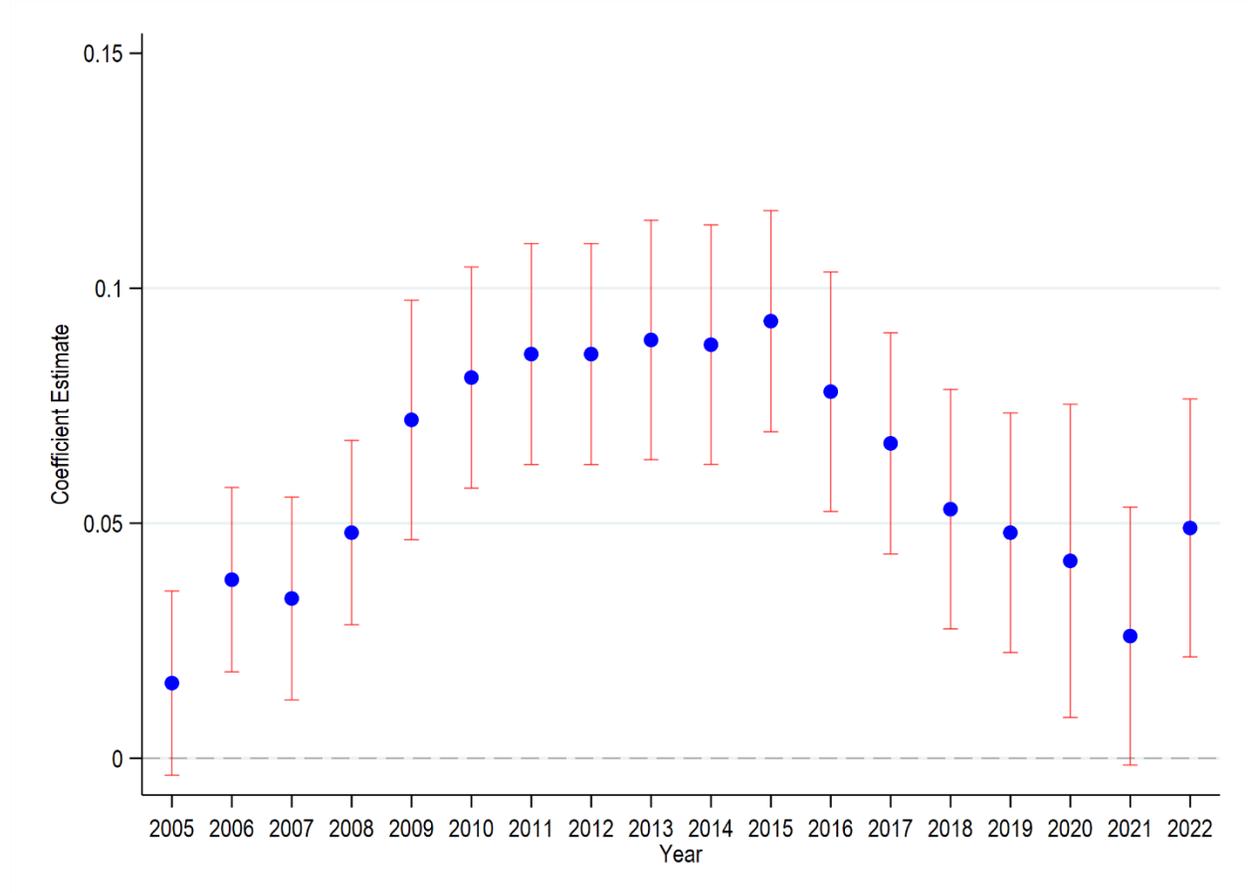
Notes: based on authors' calculations from Enverus data on oil and gas non-vertical wells completed during the 2006-2014 fracking boom period. For counties with non-vertical wells, we compute the non-vertical well density as the total number of non-vertical wells in the county divided by the total geographic area of the county. High fracking counties are those in the top half of non-vertical (NV) well density among those with non-zero NV wells. Low fracking counties have positive fracking but below median NV well density. Non-fracking counties have zero fracking during 2006-2014.

Figure 3: Mean Non-vertical Wells for High Fracking Counties



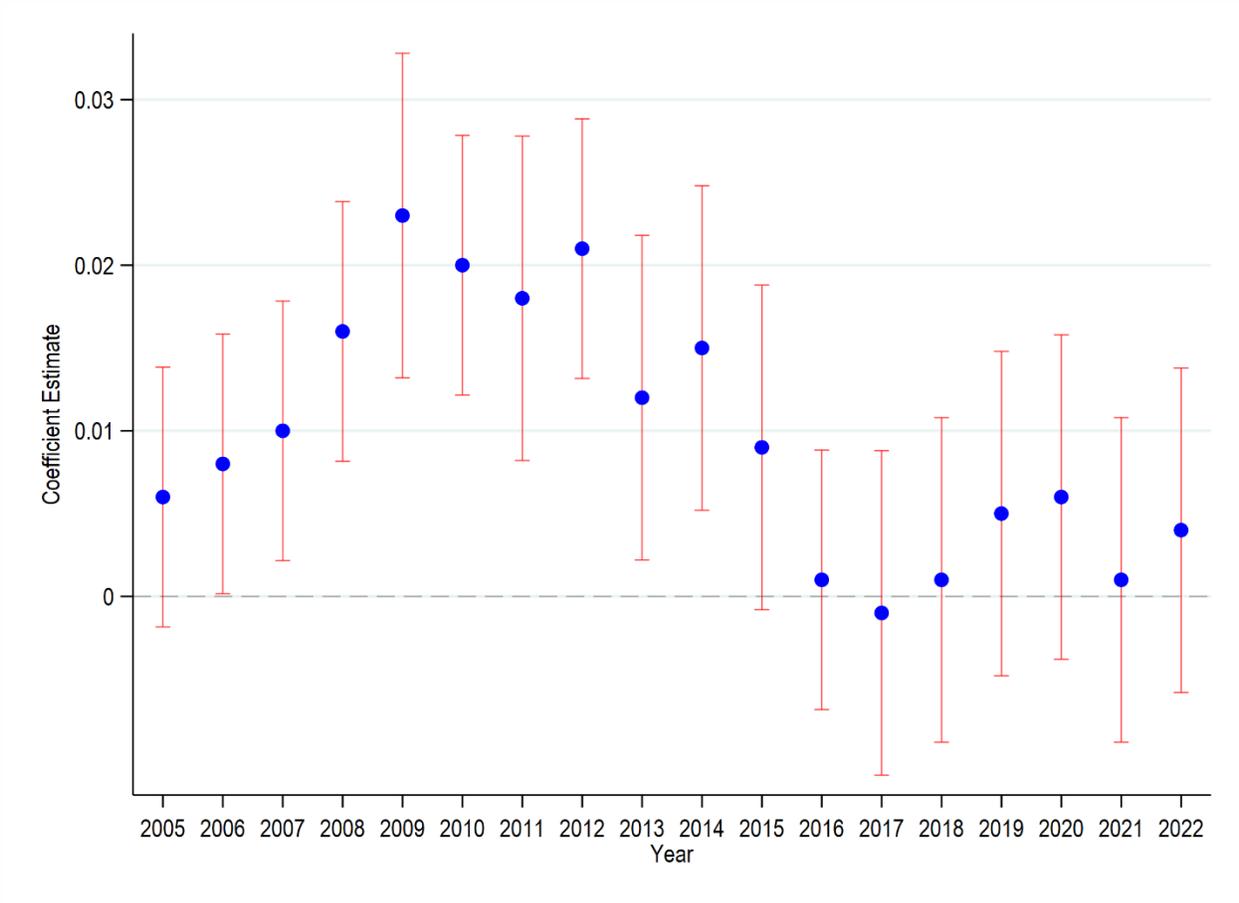
Notes: based on authors' calculations from Enverus data on oil and gas non-vertical wells completed. High fracking counties are those in the top half of non-vertical (NV) well density among those with non-zero NV wells during the 2006-2014 fracking boom period.

Figure 4: Fracking Impacts on Log Annual Earnings for Ages 25-54



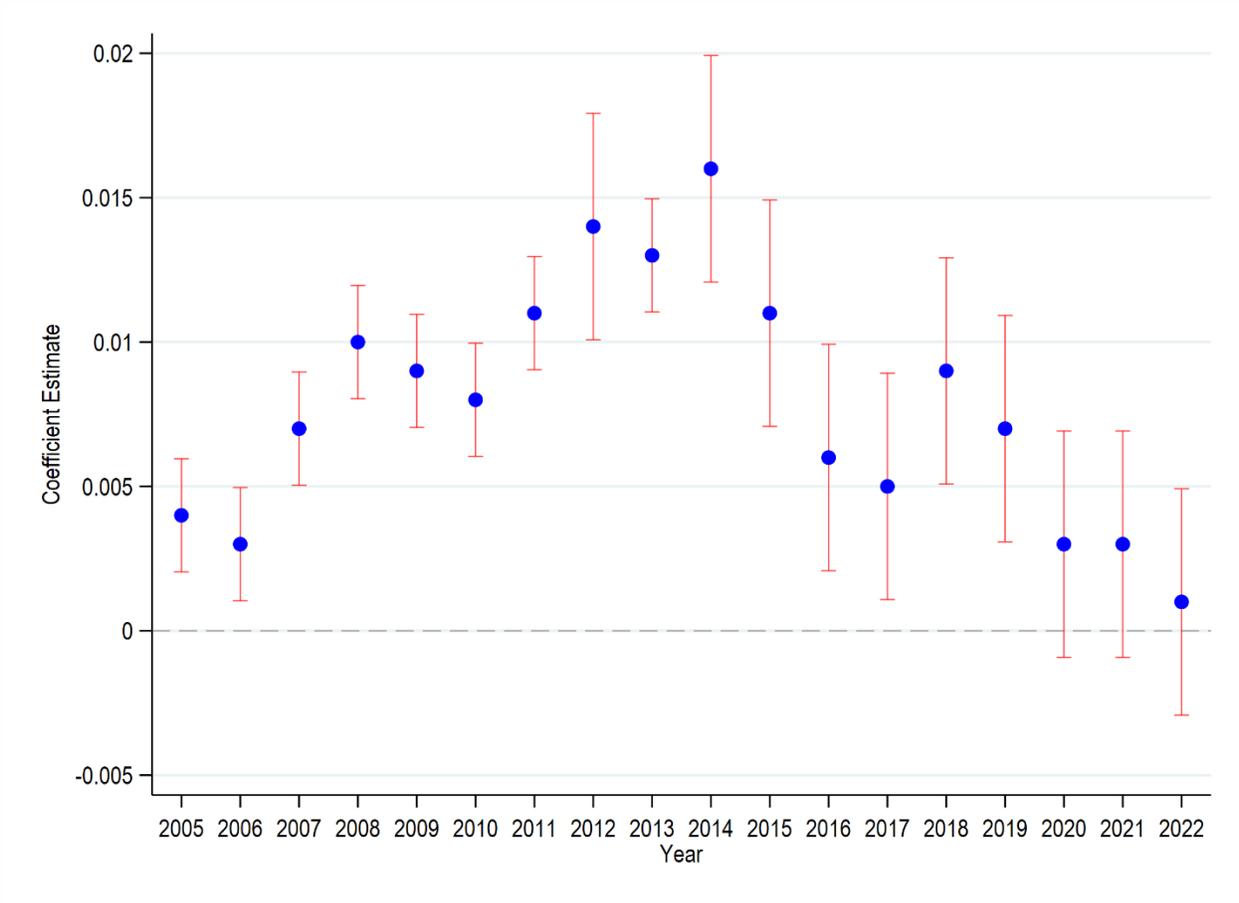
Notes: The figure reports event-study estimates for a sample limited to individuals born in non-metropolitan U.S. counties, excluding Alaska and Connecticut. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. We define the high-fracking treatment group as individuals born in counties in the top half of non-vertical (NV) well density among counties with non-zero NV wells; we exclude four New York counties that would otherwise qualify for treatment due to the state’s fracking ban adopted in 2010. The control group is restricted to counties with (a) zero NV wells during 2000–2022 and (b) locations more than 100 miles from any treatment-county. Regressions include birth-county fixed effects, birth-region \times survey-year fixed effects, birth-year fixed effects, and detailed demographic indicators for age, sex, race, and ethnicity; standard errors are clustered by birth county. There are 5,208,000 observations (rounded) in the sample.

Figure 5: Fracking Impacts on Employment Probability for Ages 25-54



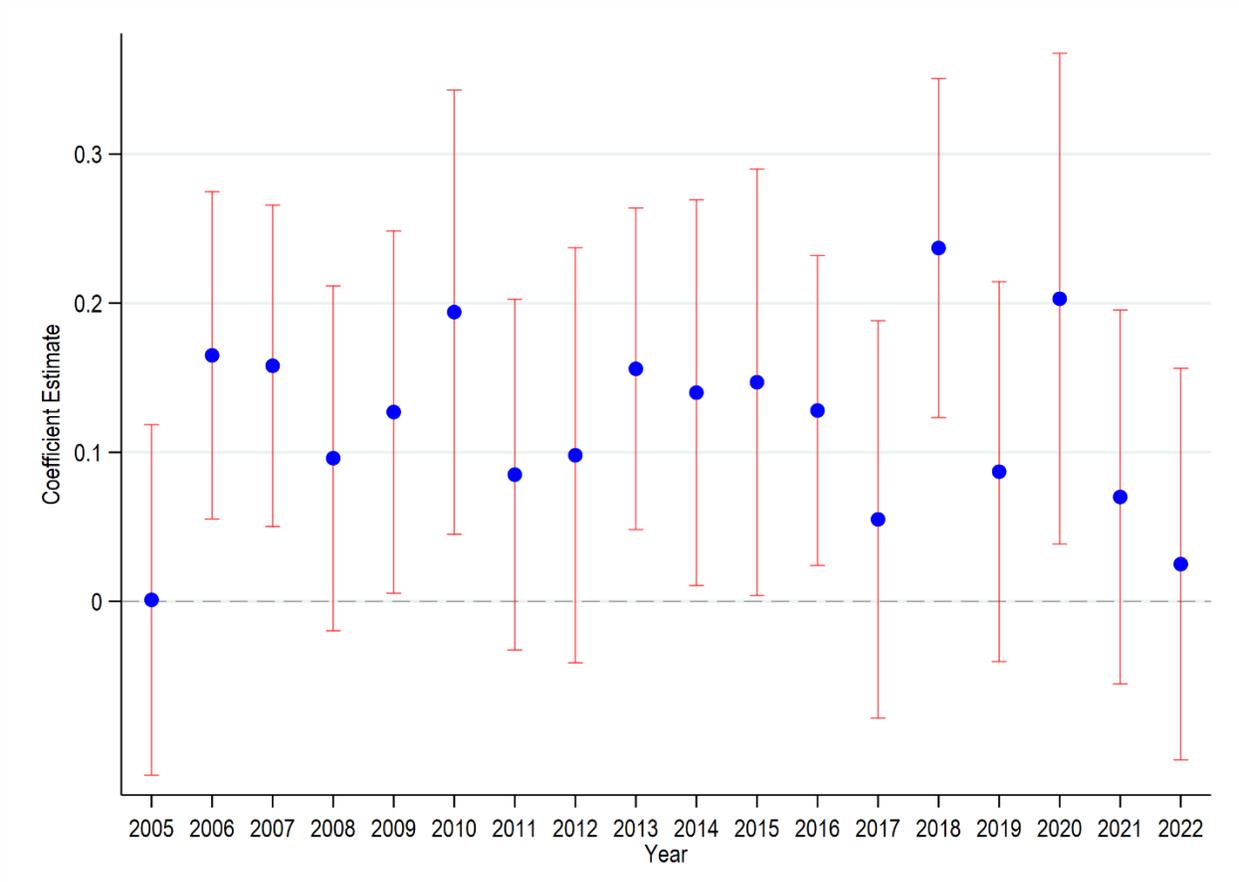
Notes: Sample and measurement are the same as indicated in Figure 4 notes. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure 6: Fracking Impacts on Mining Employment Probability for Ages 25-54



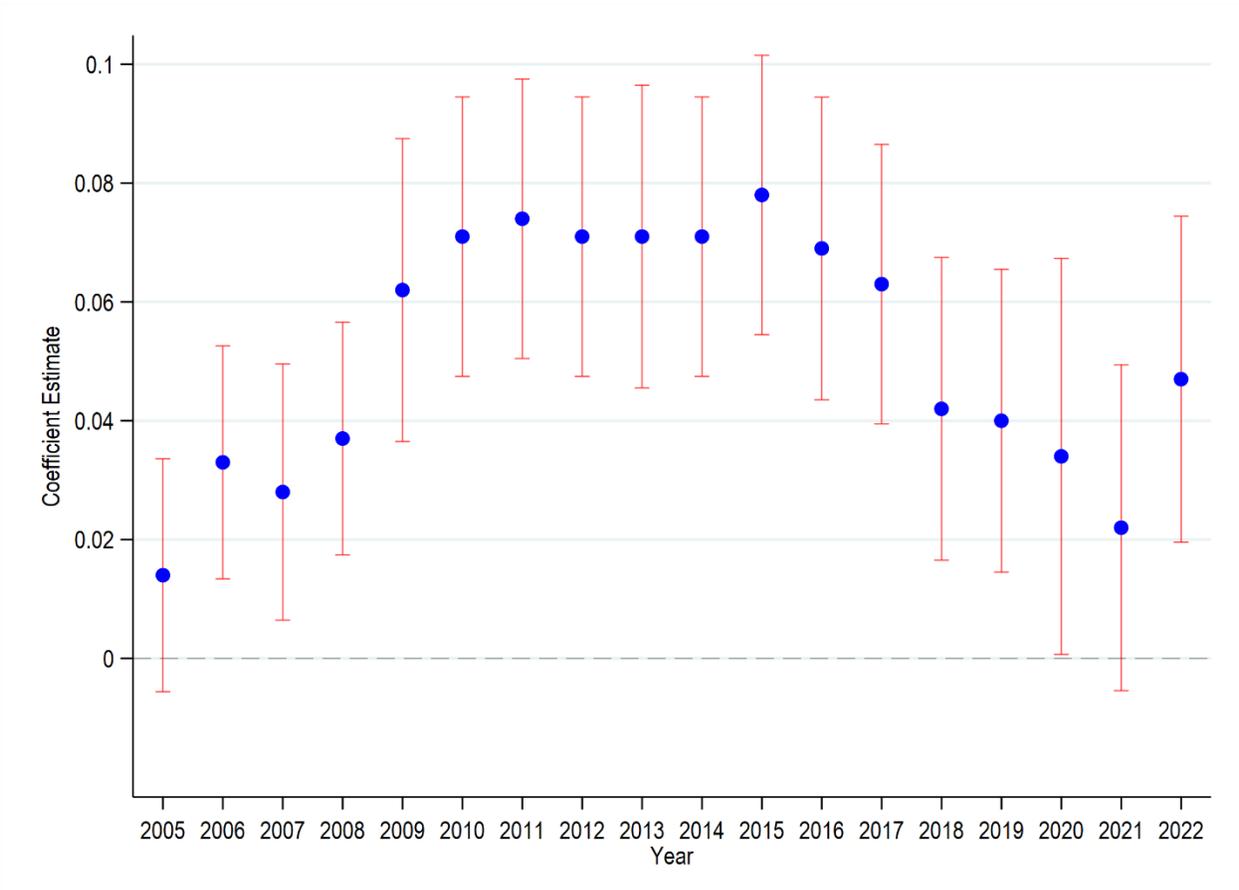
Notes: Sample and measurement are the same as indicated in Figure 4 notes. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure 7: Fracking Impacts on Log Annual Earnings for Mining Industry Workers



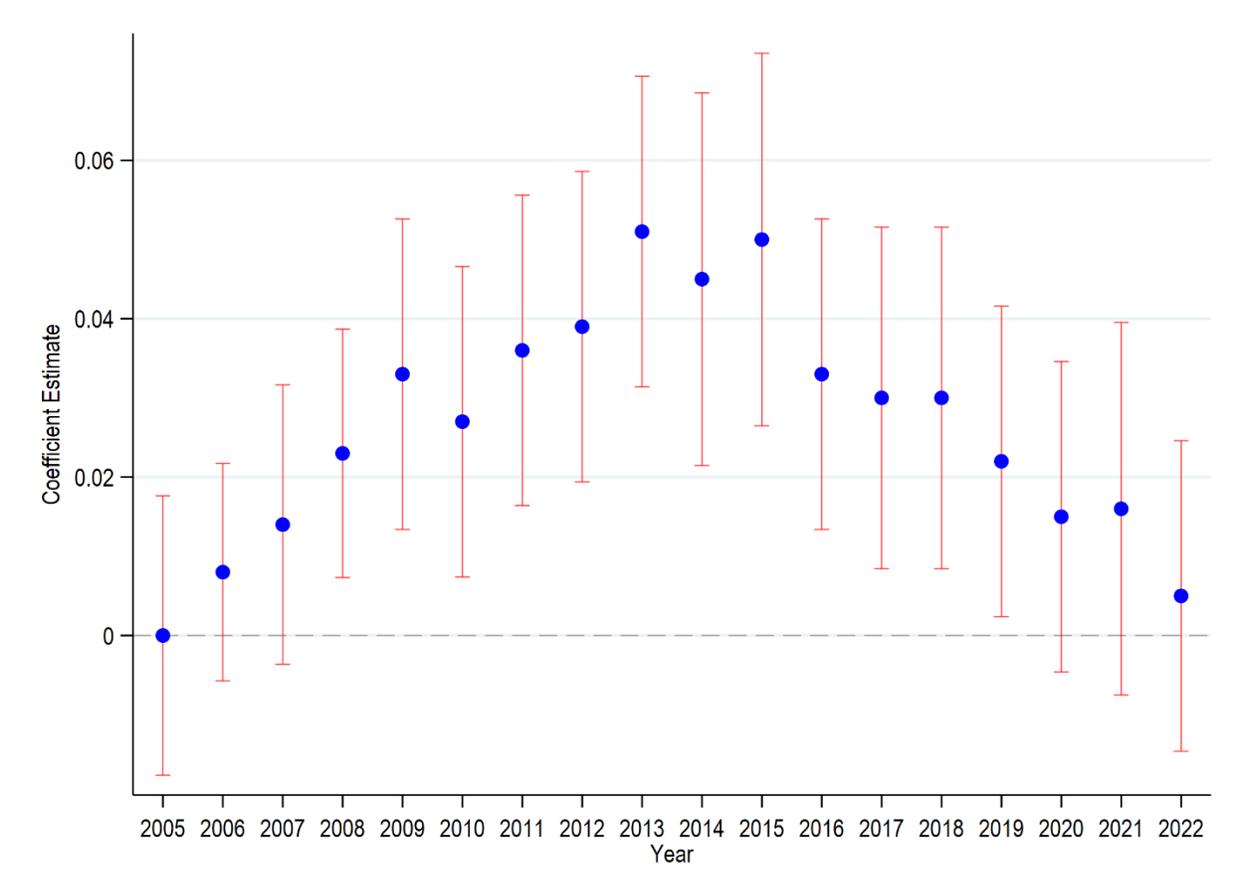
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to workers employed in the mining industry. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 50,000 observations (rounded) in the sample.

Figure 8: Fracking Impacts on Log Annual Earnings for Non-Mining Industry Workers



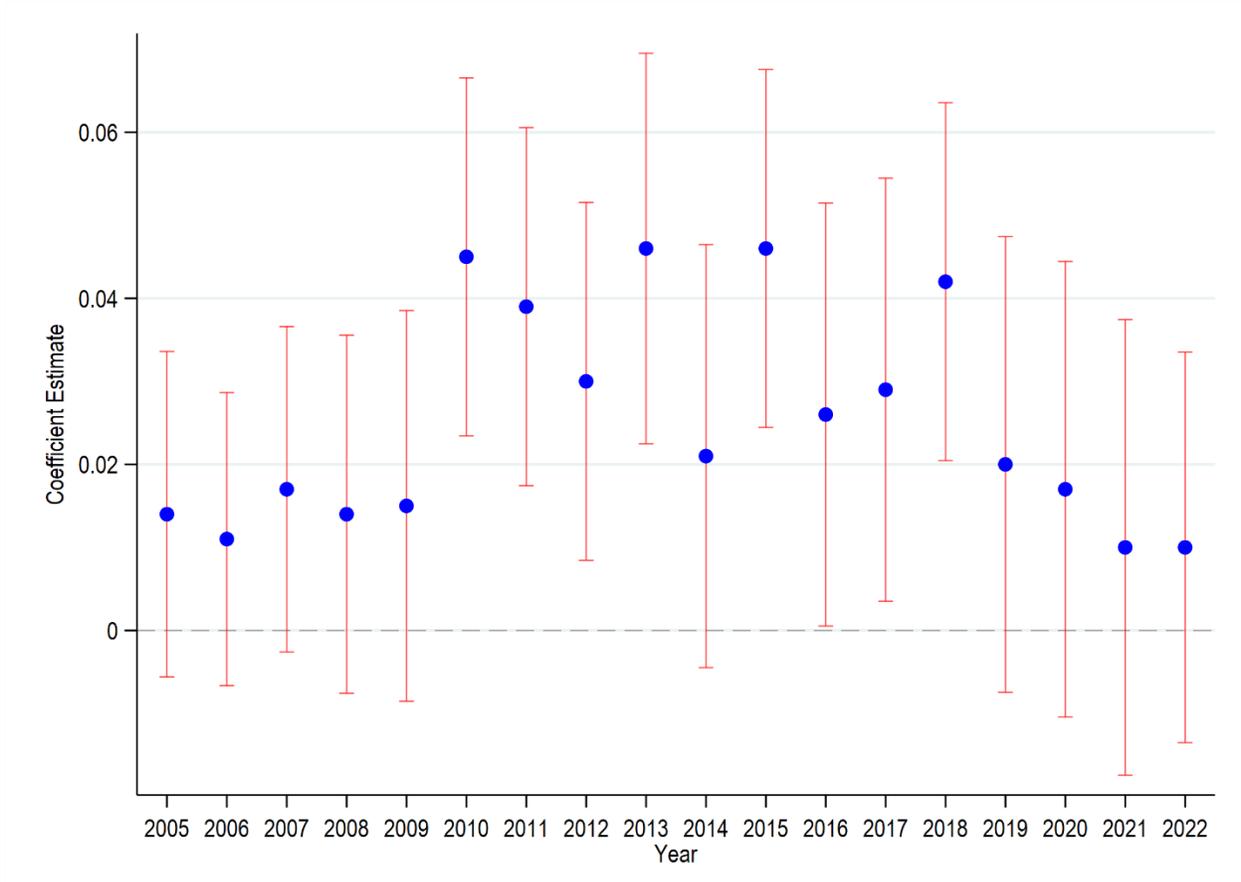
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to workers not employed in the mining industry. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 5,158,000 observations (rounded) in the sample.

Figure 9: Fracking Impacts on Log Annual Earnings for Persons Born in Metropolitan Counties



Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to persons born in metropolitan counties. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 19,110,000 observations (rounded) in the sample.

Figure 10: Impacts on Log Earnings for Persons Born in Low Fracking Non-Metro Counties



Notes: Sample and measurement are the same as indicated in Figure 4 notes except the treatment group includes persons born in low fracking counties instead of persons born in high fracking counties. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 5,289,000 observations (rounded) in the sample.

Table 1: Year 2000 Selected Summary Statistics for Treatment and Control Group

	Control Group	Treatment Group	Difference of Means	P-value for Diff. of Means
Log Annual Earnings	10.08	10.03	-0.05	<0.001
Observations	1,720,000	397,000		
Employed Indicator	0.801	0.750	-0.051	<0.001
Age in Years	40.52	40.94	0.42	<0.001
Female	0.503	0.505	0.002	0.013
White non-Hispanic	0.841	0.844	0.003	<0.001
High School Diploma or Higher	0.888	0.853	-0.035	<0.001
Some College or Higher	0.567	0.522	-0.045	<0.001
Bachelor's Degree or Higher	0.240	0.217	-0.023	<0.001
Residing in High Fracking County	0.043	0.558	0.515	<0.001
Observations	1,954,000	479,000		

Notes: The sample is limited to persons born in non-metropolitan counties of the U.S. excluding Alaska and Connecticut. We define the (high-fracking) treatment group as individuals born in counties in the top half of non-vertical (NV) well density among those with non-zero NV wells. We exclude four counties from New York state that would otherwise qualify for the treatment group from the analysis due to the state's fracking ban adopted in 2010. The control group is restricted to counties meeting two criteria: a) zero NV wells during the 2000-2022 period and b) not within 100 miles of a treatment group county. Year 2000 is the pre-event period. Sample means are reported for selected variables. Means for mining employment are suppressed due to disclosure concerns given the small number of mining workers. The number of observations is rounded to reduce disclosure concerns.

Appendix

A.1 Birth-County Effects as Lower Bounds for Residence-County Effects

Our paper suggests that fracking impact estimates based on birth-county likely provide lower-bound estimates for the effect of fracking for long-term current residents. More formally, consider the ATT we estimate for year t for the sample of persons born in a fracking county (group F):

$$ATT_t^F = E(Y_t^F - Y_{2000}^F) - E(Y_t^N - Y_{2000}^N)$$

, where Y_t^F denotes earnings in year t for persons born in a fracking county and Y_t^N denotes earnings in year t for persons born in a non-fracking control group county. This will generally differ from the ATT for the sample of persons who both were born in a fracking county and reside in a fracking county in year t (group FF):

$$ATT_t^{FF} = E(Y_t^{FF} - Y_{2000}^{FF}) - E(Y_t^{NN} - Y_{2000}^{NN})$$

, where Y_t^{FF} denotes earnings for persons born in a fracking county and residing in a fracking county in year t , and Y_t^{NN} denotes earnings for persons born in a non-fracking county and residing in a non-fracking county in year t .

ATT_t^F is a lower bound of ATT_t^{FF} if two assumptions are satisfied.

Assumption 1: $E(Y_t^F - Y_{2000}^F) \leq E(Y_t^{FF} - Y_{2000}^{FF})$.

Assumption 2: $E(Y_t^N - Y_{2000}^N) \geq E(Y_t^{NN} - Y_{2000}^{NN})$.¹⁸

Intuitively, we expect Assumption 1 to hold because group F includes group FF and a group of individuals born in a fracking county but not living in a fracking county (group FN). We expect

¹⁸ We can rewrite the ATTs as $ATT_t^F = A - B$ and $ATT_t^{FF} = C - D$. The lower-bound condition, $ATT_t^F \leq ATT_t^{FF}$, implies $A - B \leq C - D$, which can be rearranged as $A - C \leq B - D$. Rewrite Assumptions 1 and 2 as $A - C \leq 0$ and $B - D \geq 0$ and combine to get $A - C \leq 0 \leq B - D$.

these fracking county out-migrants in group FN to receive little or no dose of the fracking treatment and limited positive impacts on earnings compared to group FF , so that their inclusion causes average earnings growth for F to be less than earnings growth for FF .¹⁹ We also expect Assumption 2 to hold. Persons born in non-fracking counties (group N) includes both persons living in a non-fracking county in year t (Group NN) and persons living in a fracking county in year t (group NF). Moving to a fracking county is likely to have a positive earnings impact for NF , giving them larger earnings growth than NN , and cause the earnings growth for N to exceed that for NN .

Unfortunately, migration decisions are not random and we are unable to offer unbiased causal evidence on the impacts of living in a fracking county or the validity of Assumptions 1 and 2. However, we report descriptive analysis below for individuals living in a fracking county during the survey period and compare these to our main results for impacts on individuals born in a fracking county to offer suggestive evidence.

A.2 Descriptive Estimates for Non-Migrants and In-Migrants

Our main results in Figure 4 are for persons born in fracking counties, which includes both persons residing in their birth county and persons who do not. As an extension, we next explored coefficients by migration status. Specifically, the specification for Figure A1 limits the sample to individuals residing in their birth county during the survey and is otherwise the same as in Figure 4. This sample includes people who have lived their whole lives in their birth county as well as people who moved away and then came back. Notably, adult location

¹⁹ Fracking county out-migrants may receive positive spillover effects if they move to nearby non-fracking counties. Even so, we expect direct effects for fracking county residents to be larger than spillover effects to nearby non-fracking counties. Furthermore, some fracking county out-migrants are likely to move sufficiently far away that they receive little or no positive spillovers.

decisions can reflect sorting based on unobservable characteristics and complementarity with the oil and gas industry, so we do not interpret these as unbiased causal effects of residing in a fracking county.

Figure A2 examines in-migrants and has three key differences from the main analysis in Figure 4: 1) all spatial explanatory variables are based on county of residence, 2) the sample is limited to high fracking and non-fracking treatment and control groups based on county of residence, and 3) persons residing in their birth county are excluded. Thus, Figure A2 is an analysis of fracking residence for fracking county in-migrants. Again, adult location decisions are endogenous, so these are not unbiased estimates. However, Figures A1 and A2 provide illustrative comparisons.

The coefficients in Figures A1 and A2 follow a similar pattern as Figure 4, though coefficients are somewhat larger in Figures A1 and A2 than in Figure 4. Coefficients in Figure A1 and A2 are often similar to each other with some moderate differences. Again, these are unlikely to be unbiased estimates but they offer two suggestive interpretations: 1) effects of residing in a fracking county appear larger than effects of being born in a fracking county consistent with expectations, and 2) fracking county natives who reside in their birth county during a fracking boom appear to benefit roughly as much as in-migrants. While not definitive, these findings provide suggestive evidence in favor of our main estimates in Figure 4 for the effects of being born in a fracking county as lower bounds for the effects of residing in a fracking county. Given the endogeneity of migration decisions, Figures A1 and A2 are not our preferred results. Our preferred results are those in Figure 4 that define the sample and explanatory variables based on county of birth.

A.3 Descriptive Estimates for Non-Migrants and Out-Migrants

We conducted additional analysis that compares birth-county stayers (non-migrants) to two groups of leavers (out-migrants) that we call “short-distance” and “long-distance”. Short-distance movers are those who live outside their birth county and live in a county meeting any of the following criteria: 1) within 100 miles of their birth county, 2) any fracking county, or 3) within 100 miles of a fracking county. Long-distance movers are those who live outside their birth county and live in a county meeting all of the following criteria: 1) more than 100 miles away from their birth county, 2) a non-fracking county, and 3) more than 100 miles from any fracking county. Our definitions for movers require some choices and groupings. One might prefer even more narrowly defined groups, but census disclosure protections limit our ability to report estimates for small samples, especially when they involve potentially implicit samples of even smaller groupings.

We expect that stayers will receive the largest impact of the fracking boom in their birth county. Short-distance movers are also expected to be impacted but not as much as stayers. Long-distance movers are expected to be the most minimally exposed to and impacted by the fracking boom in their birth county. However, there could be some positive effect for long-distance movers depending on when they left and the skills and wage increases they previously obtained. E.g., increased employment in higher wage jobs when they resided in the boom area could translate to higher wages and employment even after they leave the boom area.

In conducting and disclosing this analysis for stayers and leavers, we encountered an issue with census disclosure protections related to small implicit samples that hindered us from using the exact same sample criteria as before. Therefore, we changed the age range for this analysis from 25-54 to 30-59 so that these new samples do not perfectly overlap with our other

samples and satisfy disclosure review criteria. A priori, there is no reason why this modest change in age range should meaningfully affect our qualitative conclusions. We also compare the stayers results for these two overlapping age groups and they are very similar (see Figures A1 and A3).

Our results for stayers vs. leavers are in Figures A3-A5. Consistent with expectations, we find large effects for stayers. Short distance movers also have positive effects during the boom period that are generally significant but smaller than those for stayers. Coefficient estimates for long-distance movers are generally the smallest and are only statistically significant in a few years. Again, migration decisions are endogenous and we cannot provide unbiased causal estimates separately for stayers and leavers, but this yields suggestive evidence that our main estimates are lower bounds for the effects of residing in a fracking county for individuals born in one.

A.4 Descriptive Estimates Based on Year 2000 Residential Locations

We next present descriptive analysis that leverages county of residence in 2000 rather than birth county for sample criteria and defining treatment and control groups. We expect year 2000 location to be more strongly correlated with contemporaneous exposure to a fracking boom since there is less time for migration between the location observation and exposure to the fracking boom. The earnings results are presented in Figure A6 and suggest positive impacts of fracking exposure that are generally as large as or larger than our preferred estimates based on birth-county treatment in Figure 4. However, year 2000 location may not be exogenous because people may have previously migrated or not in response to idiosyncratic shocks. E.g., the people living in a fracking county in year 2000 before the fracking boom may disproportionately consist

of people who benefited from doing so. These people may have received positive location-individual-specific shocks in 2000. Notably, our results are qualitatively robust to this alternative analysis.

A.5 Earnings Spillovers across Geographic Areas

Prior literature has indicated geographic spillovers from natural resource shocks (e.g. Feyrer et al. 2017). We report illustrative spillover estimates on log earnings in Figure A7 using our data and methods. For this analysis, persons born in high fracking counties are excluded and we include a “spillover treatment group” along with our original control group and conduct our event analysis. The spillover treatment group includes individuals born in a county with no fracking but within 100 miles of a high fracking county. We find evidence of positive earnings spillovers that are non-zero but considerably smaller than the estimated impacts for the actual treatment group of individuals born in high fracking counties. This is consistent with labor market impacts decreasing with distance due to transportation costs, migration costs, and other spatial frictions. Our main analysis excludes the spillover treatment group from both the treatment and control groups.

A.6 Additional Non-Fracking Industry Analysis

Our analysis suggests that fracking increases earnings in both mining and non-mining sectors through increased local labor demand. Income spillovers to non-mining sectors may especially impact industries that produce services consumed locally, employ less-educated workers, and historically pay relatively low wages. We explore this by using our data and methods to document impacts of fracking on earnings for workers employed in restaurants and

retail. More oil and gas workers with higher incomes increases demand for local restaurants and retail outlets and leads to increased worker earnings in these industries. Furthermore, labor market competition in one sector can bid up wages in other sectors. Our supplemental analysis combines restaurant and retail workers into a single sector because sample sizes in each are relatively small and a combined sample provides more statistical power and a clearer summary picture of how the fracking boom affects these traditionally low-wage sectors. Results in Figure A8 show that the fracking boom does increase earnings of restaurant and retail workers with coefficient estimates peaking at 0.116 in 2011. Overall, the earnings generally follow the pattern for the full sample. Increased local labor demand can also cause some long-term resident workers to move into higher-paying industries and out of low-paying jobs in restaurants and retail. Figure A9 estimates the impact of fracking on the probability of employment in restaurants or retail. Coefficient estimates are frequently negative and significant consistent with the fracking boom creating higher-paying local job opportunities that pull prior residents out of low-paying restaurant and retail jobs.

A local fracking boom might also adversely impact some industries via rising labor costs that reduce their competitiveness in national and global markets. Reduced competitiveness and production in other traded industries could lead to a resource curse after the natural resource boom subsides. We explore this by using our event analysis framework to examine the impacts of being born in a fracking county on the probability of employment in wholesale trade in Figure A10. Our wholesale trade employment coefficients are generally small and not statistically significant, suggesting that the short- and medium-run impacts on wholesale trade employment among long-term residents may be minimal. However, there is likely not yet enough evidence or time passed to say confidently that there isn't a long-run resource curse happening.

A.7 Including 2001-2004 ACS in the Analysis and Altering the Base Year

Our main analysis excludes the 2001-2004 ACS because many counties are not included in these years. Furthermore, individuals are sampled based on their county of residence and not county of birth, but our analysis is based on county of birth. Some individuals in the sample are only in the sample because they moved away from their birth county and would not be included if they had remained in their birth county. This means that some birth counties are largely excluded from the sample, but the individuals born there that are included in the sample are out-migrants and thus not well representative of the people who were born and continue to live in the county. The extent of non-random selection that this creates is difficult for us to quantify given the data availability and confidentiality protections. But given the concerns, our main analysis excludes the 2001-2004 ACS.

However, we also conduct supplementary analysis that includes the 2001-2004 ACS while recognizing their limitations due to non-representativeness. Results are reported in Figures A11-A13. The 2001-2004 coefficients help address potential concerns about parallel trends. The 2001-2004 coefficients are generally small and not statistically significant, consistent with similar trends between the treatment and control group during 2000-2004 prior to the fracking boom. We also reiterate that some areas, especially in Texas, saw fracking begin to proliferate in the early 2000s, so seeing a few small positive coefficients during these years is not completely unexpected.

We also examine using alternative base years. Year 2000 is our preferred base year for the reasons noted above. We explore separately changing the base year for event analysis estimates to 2002 and 2005. Figure 3 suggests that fracking may have crossed an inflection point

around 2003, making 2002 an interesting benchmark. Others have suggested the fracking boom took off around 2006 (Rickman et al. 2017; Cascio and Narayan 2022), making 2005 an interesting benchmark. Results in Figures A14-A19 using these alternative benchmark years are qualitatively similar to our main results.

A.8 Log Earnings Analysis Not Conditional on Working the Previous Year

We also examine impacts on the probability of having positive earnings the previous year and the impacts on log earnings not conditional on having positive earnings the previous year. The first outcome is a binary dependent variable similar to the previous employment analysis in Figure 5, but the binary dependent variable now equals one for positive annual earnings and zero otherwise. Results in Figure A20 confirm that fracking affects this margin as well consistent with expectations. We next define additional variables similar to log earnings but including persons without positive earnings. Unfortunately, there is no clear best way to do this, and results potentially could be sensitive to how the non-positive earners are coded, so we use four alternatives: 1) defining non-positive earners to have \$0.01 earnings before taking logs, 2) defining non-positive earners to have \$1 earnings before taking logs, 3) defining non-positive earners to have \$1000 earnings before taking logs, and 4) using the inverse hyperbolic sine (IHS) transformation of annual earnings. Results are in Appendix Figures A21-A24 and confirm that our main results are qualitatively robust to all of these alternatives. The coefficients magnitudes are typically larger with these alternatives than in our main specification in Figure 4, consistent with positive effects on the extensive margin being very influential.

As expected, the magnitudes are somewhat sensitive to the four alternative earnings variables. Assigning non-positive earners \$0.01 yields especially large magnitudes while

assigning them \$1000 yields more moderate magnitudes. This results because of how they are treated by the log transformation and the clear impact of fracking on the extensive margin. For example, consider a year in which fracking induced one percent of the population to move from non-positive earnings to low positive earnings of say \$10,000. Note that $\ln(0.01) = -4.61$, $\ln(1) = 0$, $\ln(1000) = 6.91$, and $\ln(10,000) = 9.21$. A one percentage point extensive margin increase will shift log earnings by roughly 0.092 if non-positive earners are assigned \$1 but only by 0.023 if non-earners are assigned \$1000. E.g., divide the values above by 100 (because of the one percentage point extensive margin increase) and compute differences between non-earners and those earning \$10,000. Setting non-earners to \$0.01 yields an increase of 0.138. Of course, different assumptions for earnings of those pushed into the labor market and those with non-positive earnings would yield somewhat different values. The broader point is that the extensive margin matters quite a bit for any reasonable assumptions, but the quantitative magnitude depends on the assumptions. Again, these alternatives all yield larger impacts than our main specification in Figure 4.

A.9 Impacts on Poverty

In addition to our main analysis, we conduct supplemental analysis examining an indicator variable equal to one if an individual's household income is above the poverty threshold for their household size and composition. This is the opposite of being below the poverty threshold; we code it as being above the poverty threshold so that a positive coefficient represents a favorable outcome similar to higher earnings and employment. This is a useful complement to the annual earnings measure because it is targeted to particularly vulnerable individuals, i.e., those near the poverty line. Figure A25 presents the results. We find positive

and significant coefficients in most, but not all, fracking boom years. Effects are not significant for most downturn and subsequent years, with 2017 and 2020 as exceptions. The coefficient estimates are never negative and significant.

A.10 Analysis by Gender, Age, and Education

We also estimate earnings impacts for alternative samples by education level, sex, and age group in Figures A26-A31. Oil and gas employment offers high pay without requiring high levels of education, and the earnings gains may especially accrue to workers with education less than a bachelor's degree. The oil and gas industry workforce is also disproportionately male, and the same is true for some complementary industries such as construction and transportation. Thus, we expect that labor market impacts of fracking will be stronger for males than females. Impacts may also differ by age. For example, young workers may be more mobile across industries and occupations and benefit more from a boom in the oil and gas sector, while older workers may benefit less.

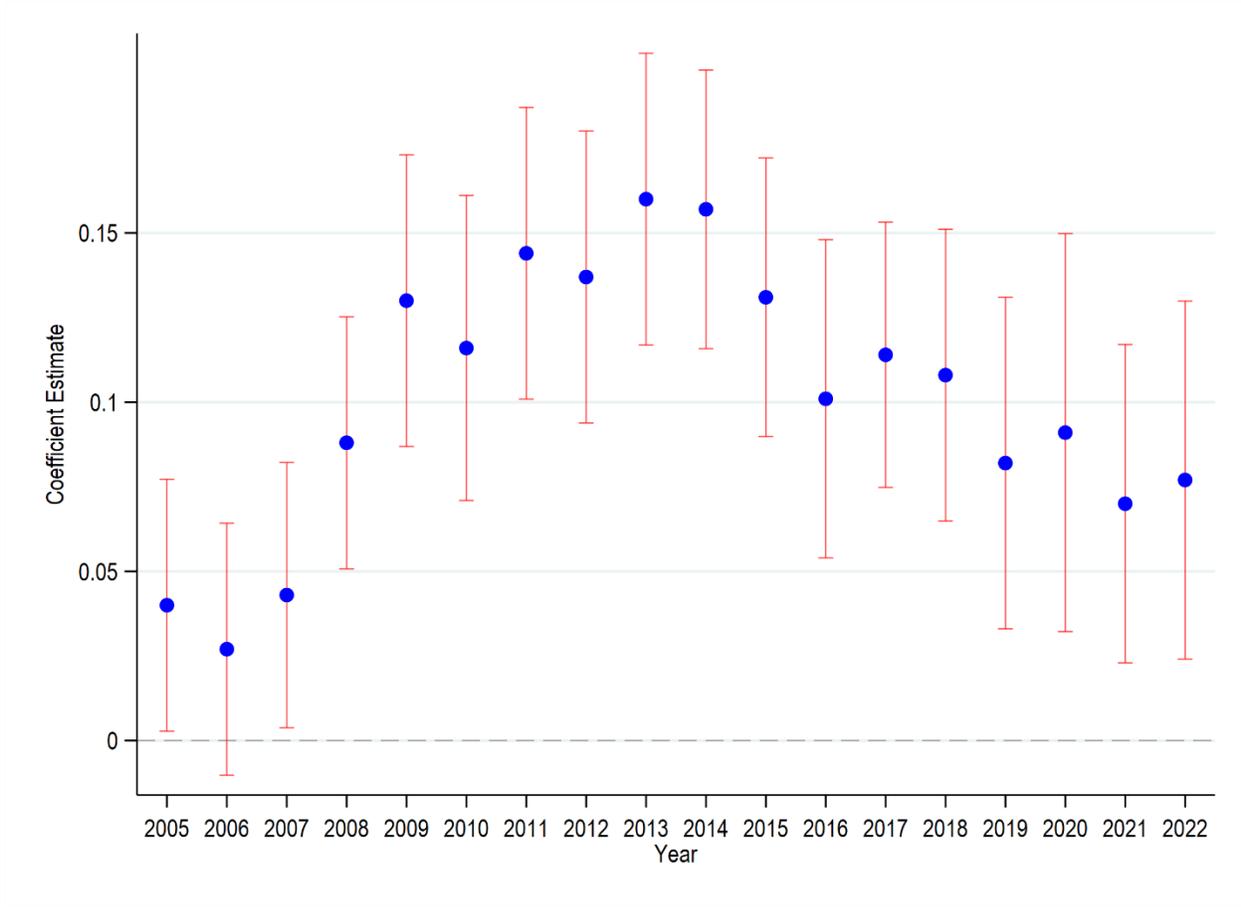
Coefficient estimates generally follow the same pattern as the full sample in Figure 4. Coefficient estimates are generally positive and significant, never negative and significant, and generally larger during the peak boom period. However, the magnitudes do differ some across groups. Coefficients are generally larger for non-college graduates than for college graduates. Coefficients are generally larger for men than women, which is not surprising given that the vast majority of mining workers are male. Coefficients for ages 18-24 are often larger than those for the main sample of ages 25-54 and for the sample of ages 55-64. Overall, these figures suggest that the fracking boom alters local labor market conditions in ways that especially benefit

workers who are young, male, and have education less than a bachelor's degree, but other workers benefit as well.

A.11 Analysis for Birth-Cohorts Born in 1966-1975

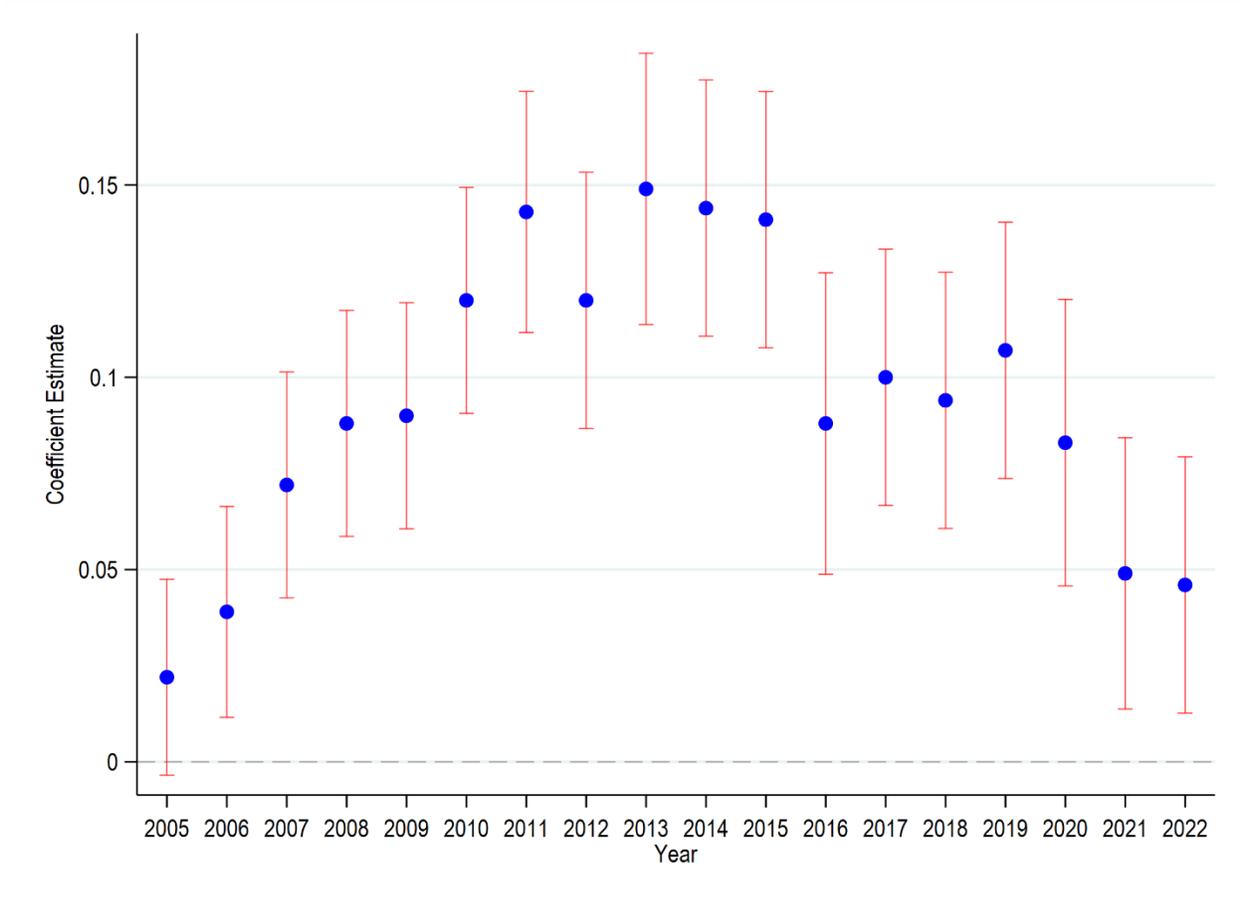
Our main analysis includes persons ages 25-54 in the 2000 Census and 2005-2022 ACS. However, the birth-cohort sample composition changes over times as individuals age into and out of the sample. To confirm that our results are not driven by birth-cohort sample composition changes, we next conduct log earnings analysis following persons over time who were born in 1966-1975. These cohorts were roughly ages 25-34 in year 2000 and ages 47-56 by 2022. The results in Figure A32 are qualitatively similar to our main results for the full sample.

Figure A1: Log Earnings Results for Persons Residing in Their Birth County



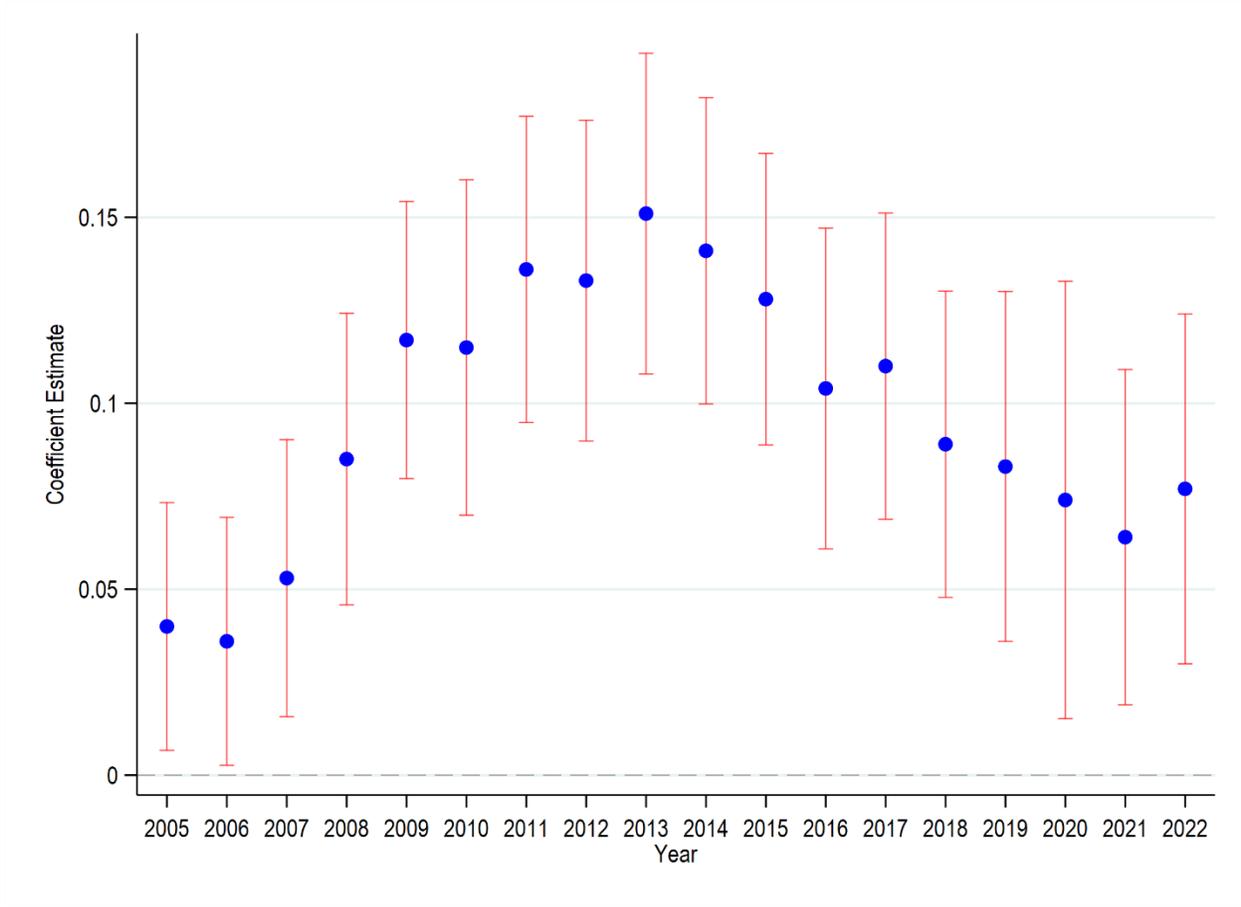
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to persons residing in their birth county at the time of the survey. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 1,623,000 observations (rounded) in the sample.

Figure A2: Log Earnings Results for Domestic In-Migrants



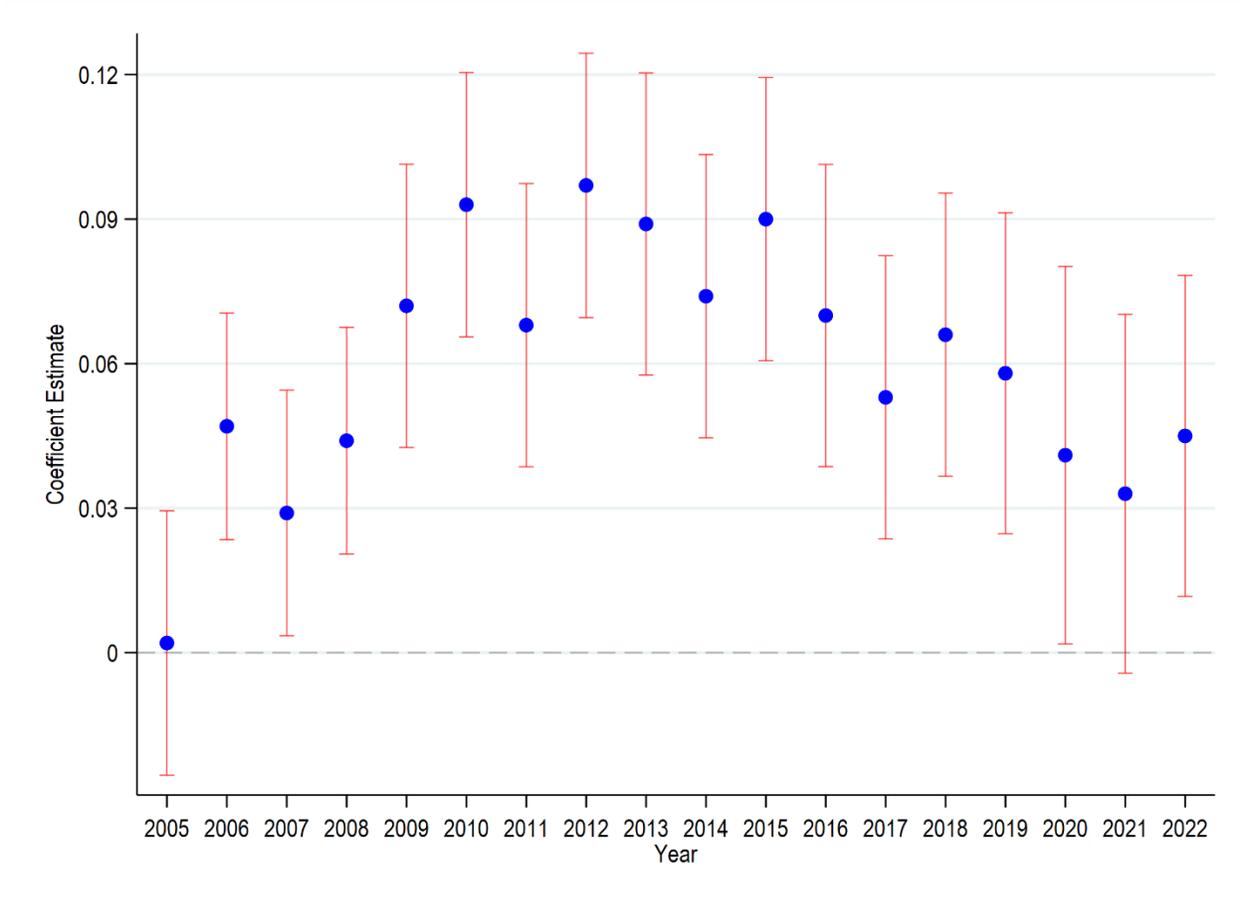
Notes: Sample and measurement are similar to Figure 4 except 1) all spatial variables are based on county of residence, 2) the sample is limited to high fracking and non-fracking treatment and control groups based on county of residence, and 3) persons residing in their birth county are excluded. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 3,854,000 observations (rounded) in the sample.

Figure A3: Log Earnings Results for Ages 30-59 Residing in Their Birth County



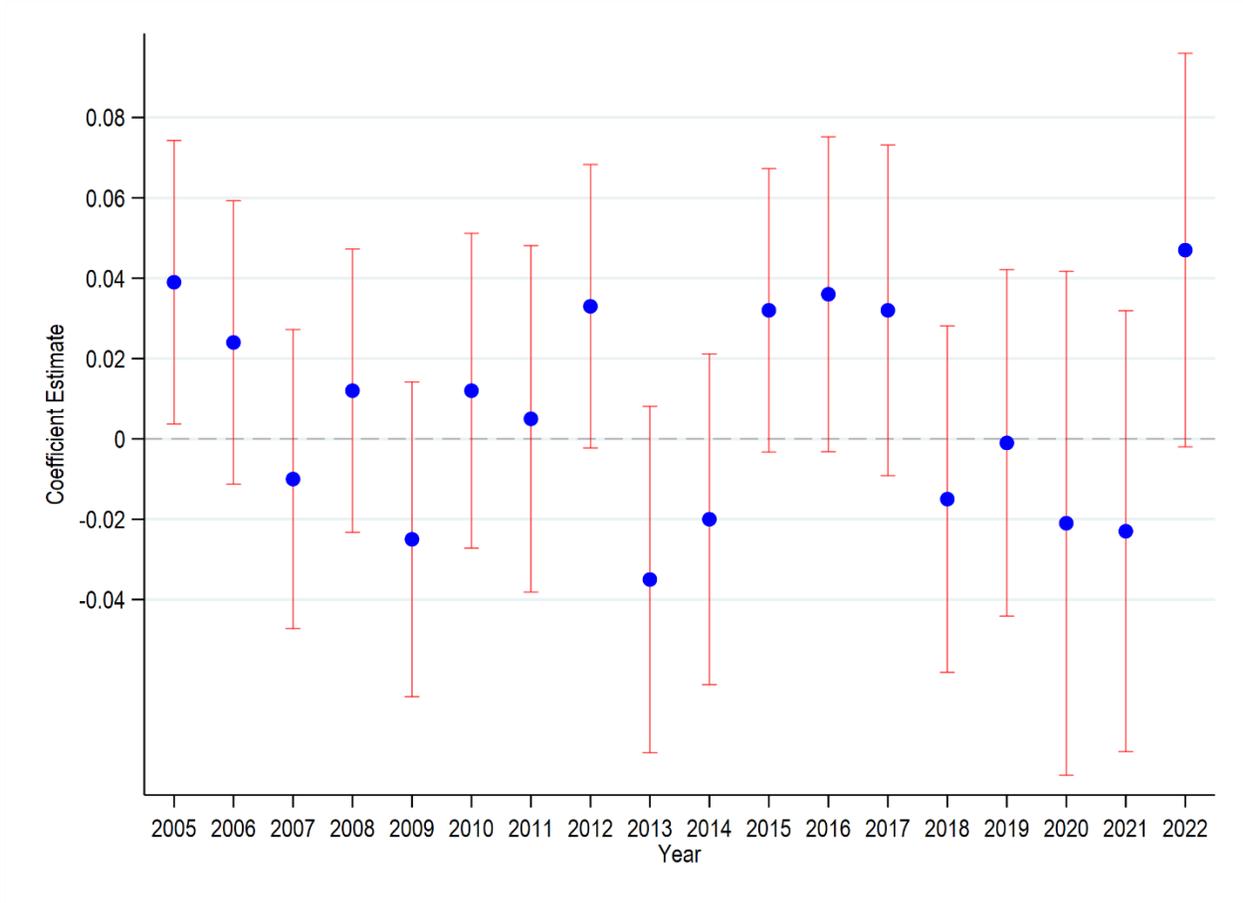
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to persons residing in their birth county at the time of the survey and includes ages 30-59. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 1,689,000 observations (rounded) in the sample.

Figure A4: Log Earnings Results for Ages 30-59 Short-Distance Movers



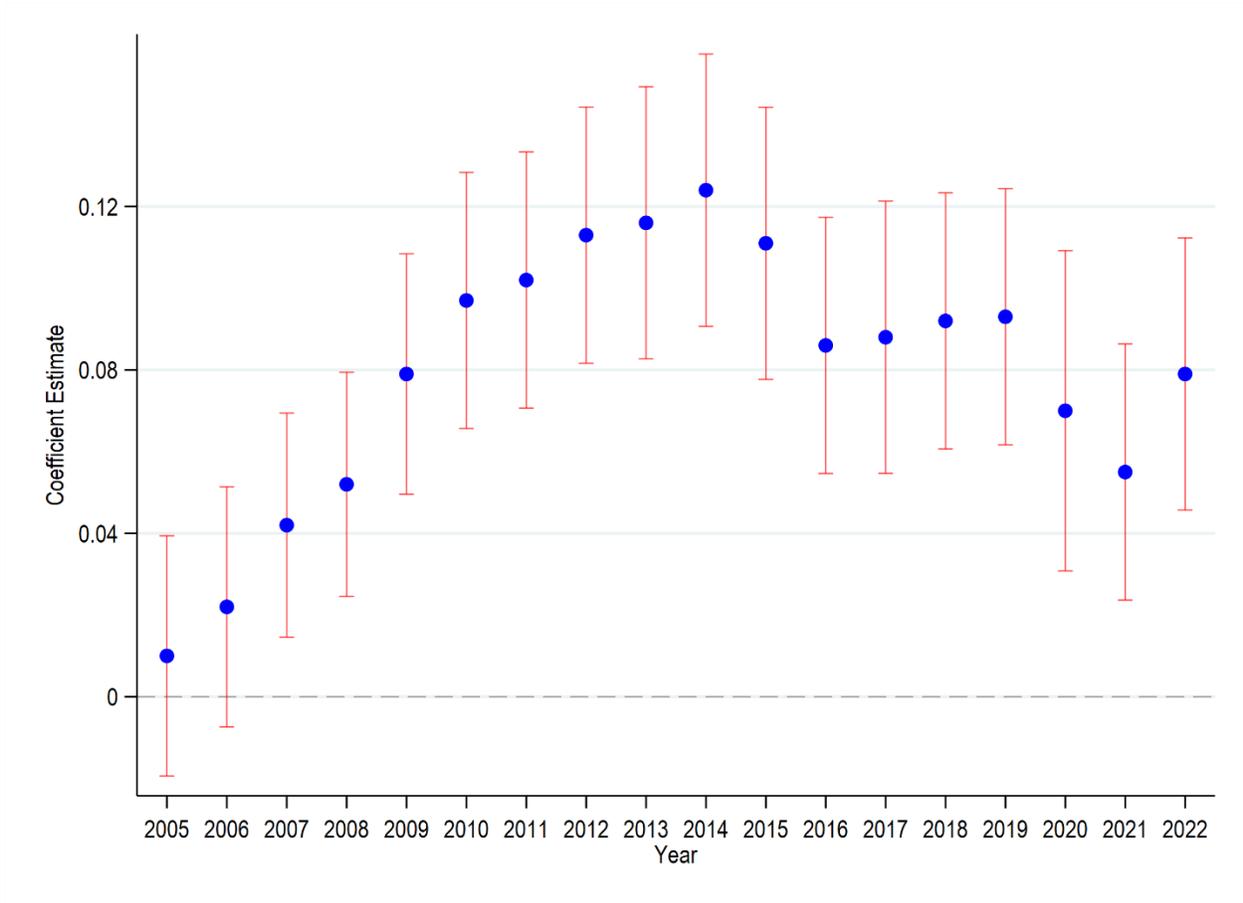
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to individuals ages 30-59 who reside outside their birth-county and live in a county meeting any of the following criteria: 1) within 100 miles of their birth county, 2) any fracking county, or 3) within 100 miles of a fracking county. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 2,486,000 observations (rounded) in the sample.

Figure A5: Log Earnings Results for Ages 30-59 Long-Distance Movers



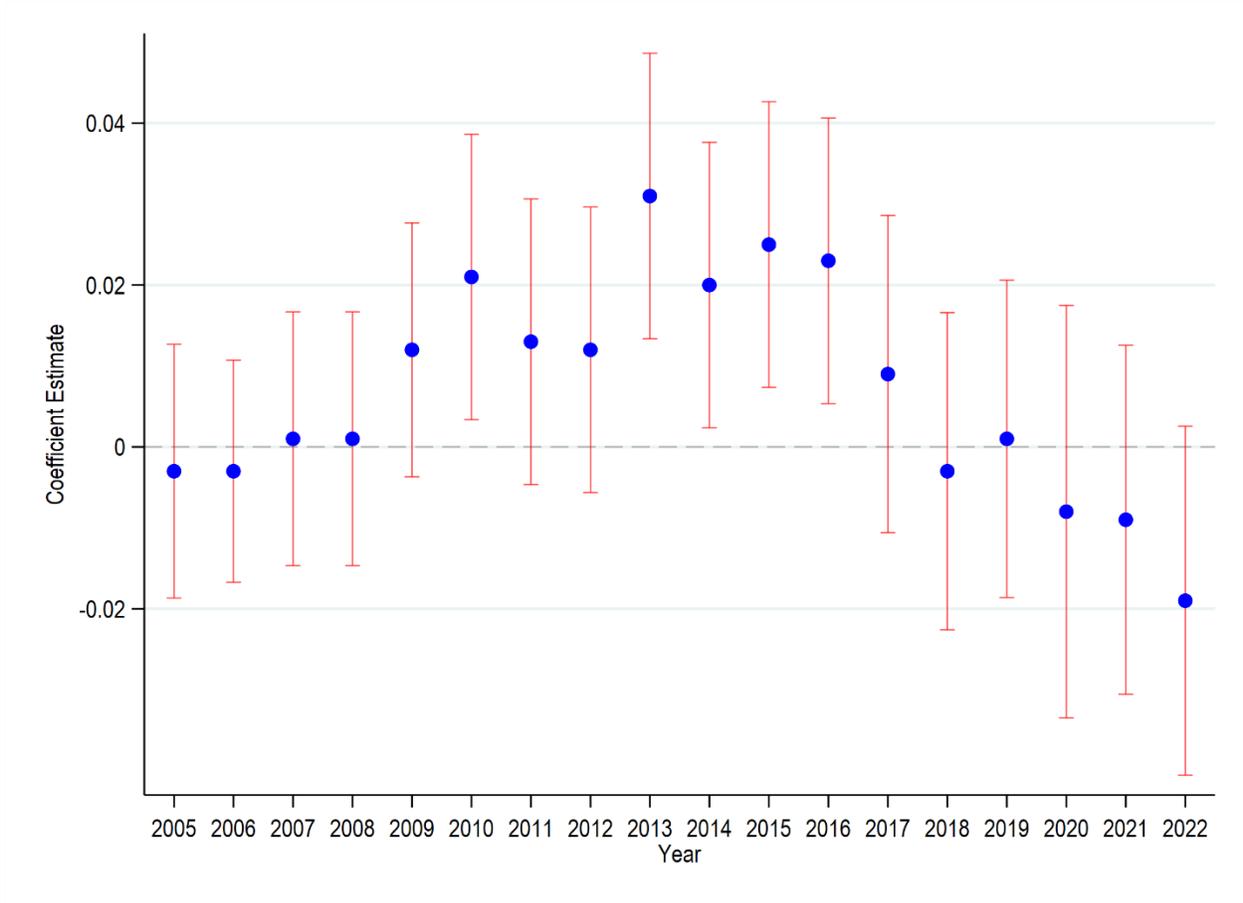
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to individuals ages 30-59 who reside outside their birth-county and live in a county meeting all of the following criteria: 1) more than 100 miles away from their birth county, 2) a non-fracking county, and 3) more than 100 miles from any fracking county. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 1,297,000 observations (rounded) in the sample.

Figure A6: Log Earnings Results Using Year 2000 Residential Locations



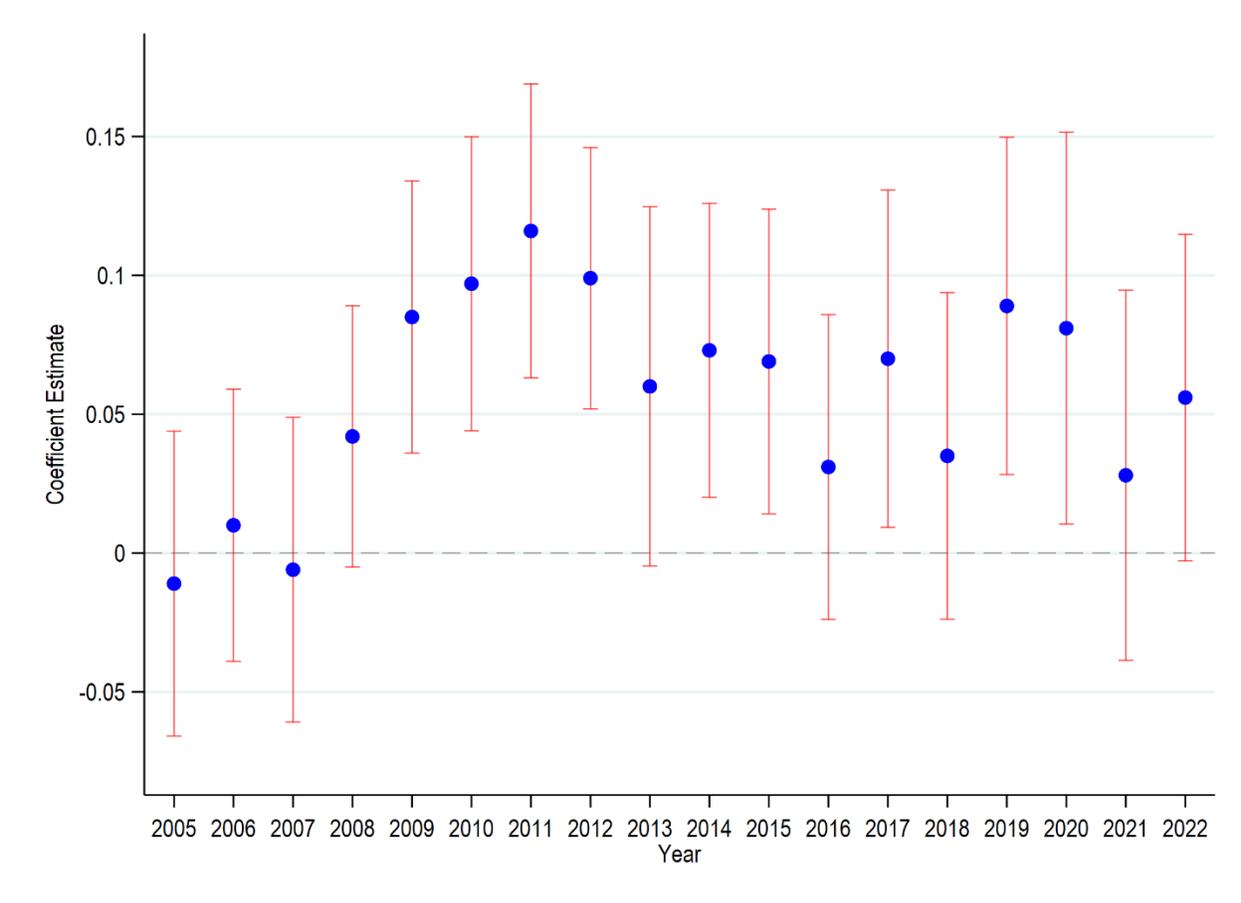
Notes: The sample and variables are constructed similarly to Figure 4 except all location variables are defined based on residential location in the 2000 census (instead of county of birth). Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 5,154,000 observations (rounded) in the sample.

Figure A7: Log Earnings Spillover Estimates for Non-Fracking Neighboring Counties



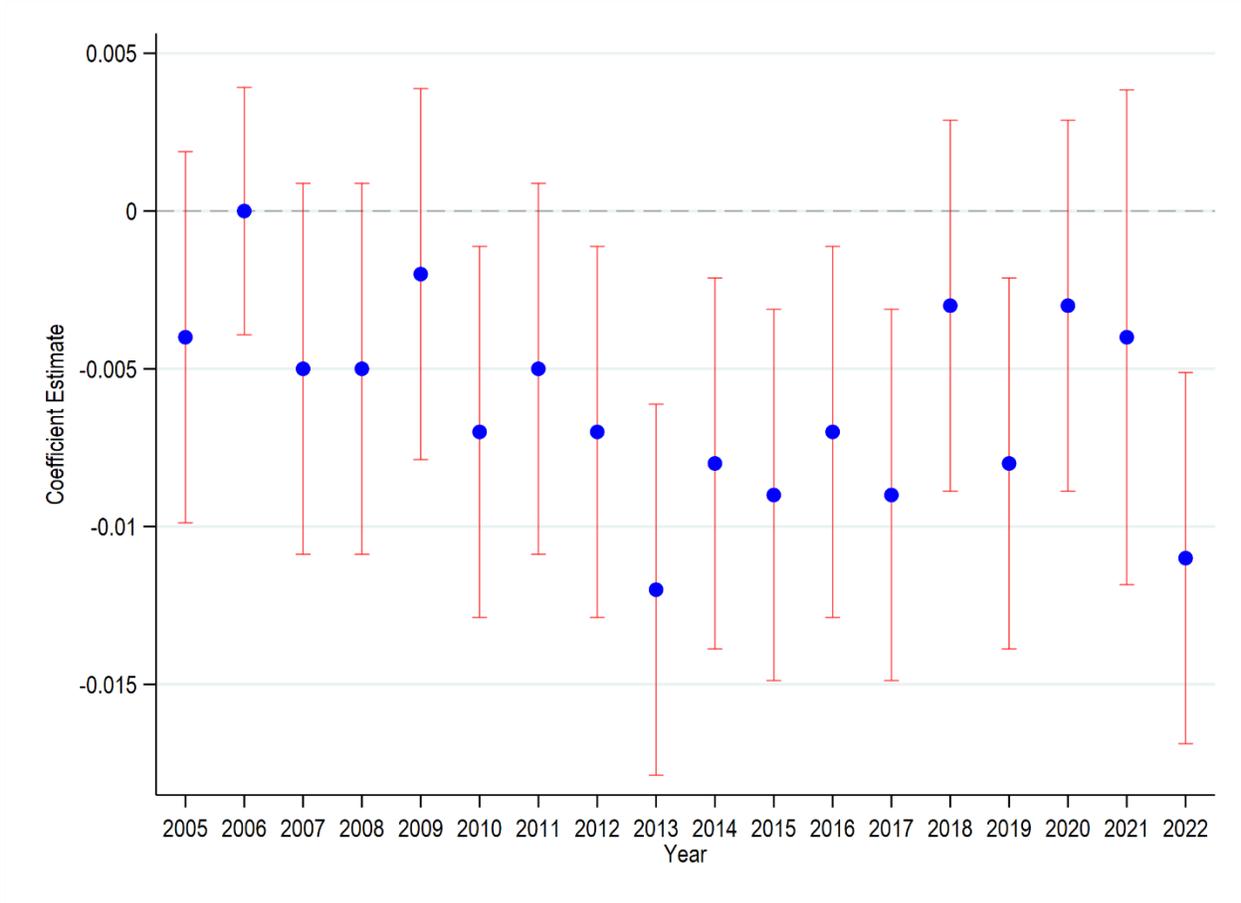
Notes: All fracking counties are excluded from the analysis. The treatment group is defined as non-fracking “neighboring” counties that are within 100 miles of a high fracking county. Sample and variable definitions are otherwise similar to Figure 4. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 5,588,000 observations (rounded) in the sample.

Figure A8: Log Earnings Results for Restaurant and Retail Industry Workers



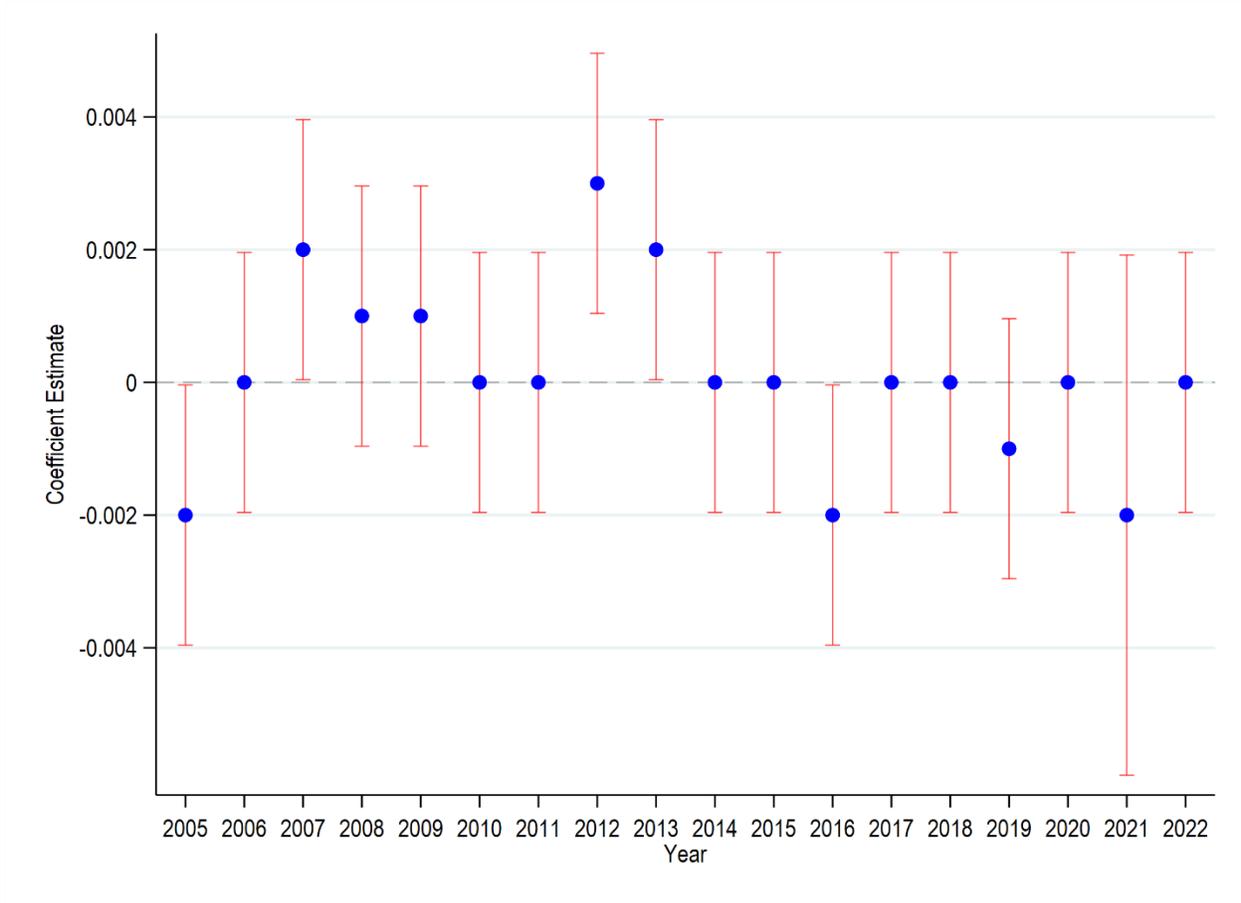
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to workers employed in the restaurant and retail industries. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 586,000 observations (rounded) in the sample.

Figure A9: Fracking Impacts on Restaurant and Retail Employment for Prior Residents



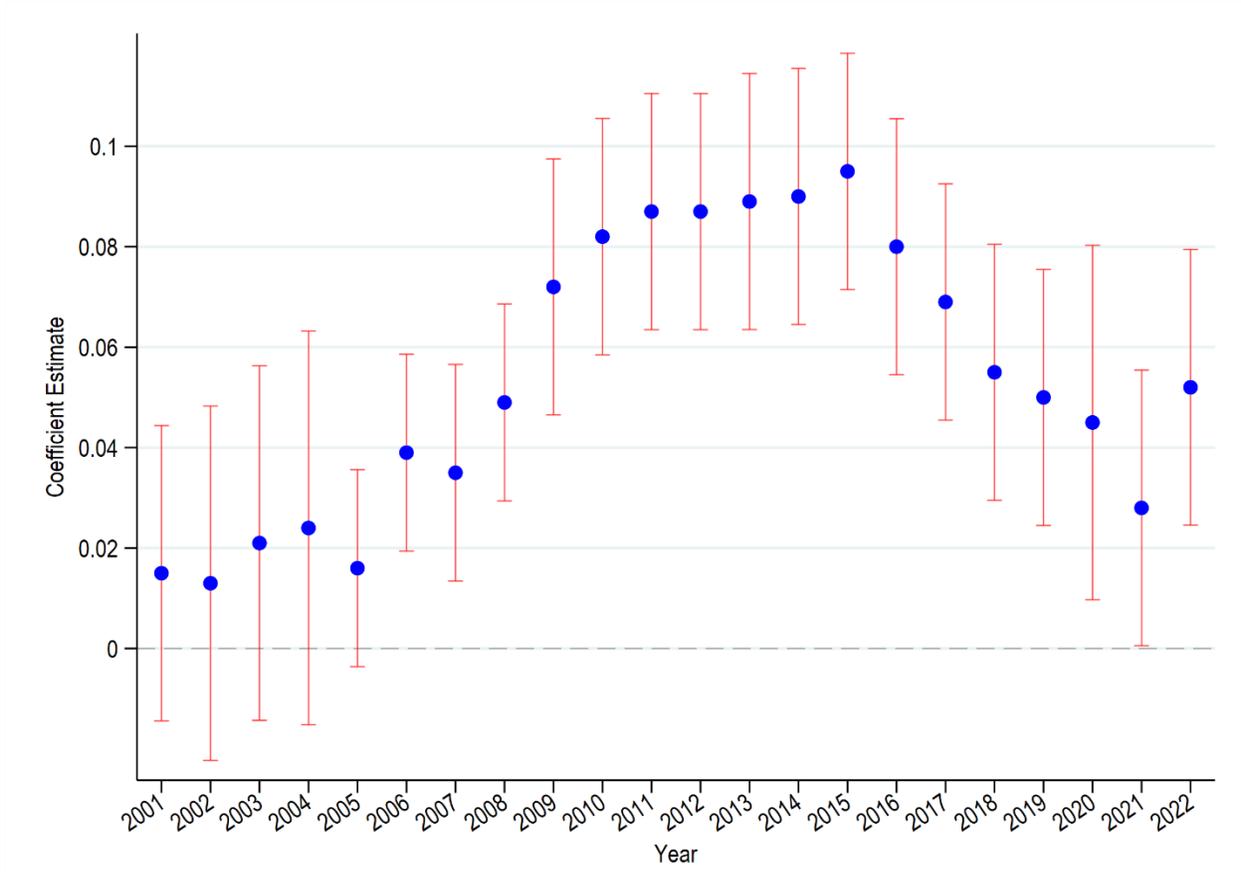
Notes: The dependent variable is an indicator for being employed in the restaurant or retail industries. The sample and variables are otherwise the same as the main analysis. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure A10: Fracking Impacts on Wholesale Trade Employment for Prior Residents



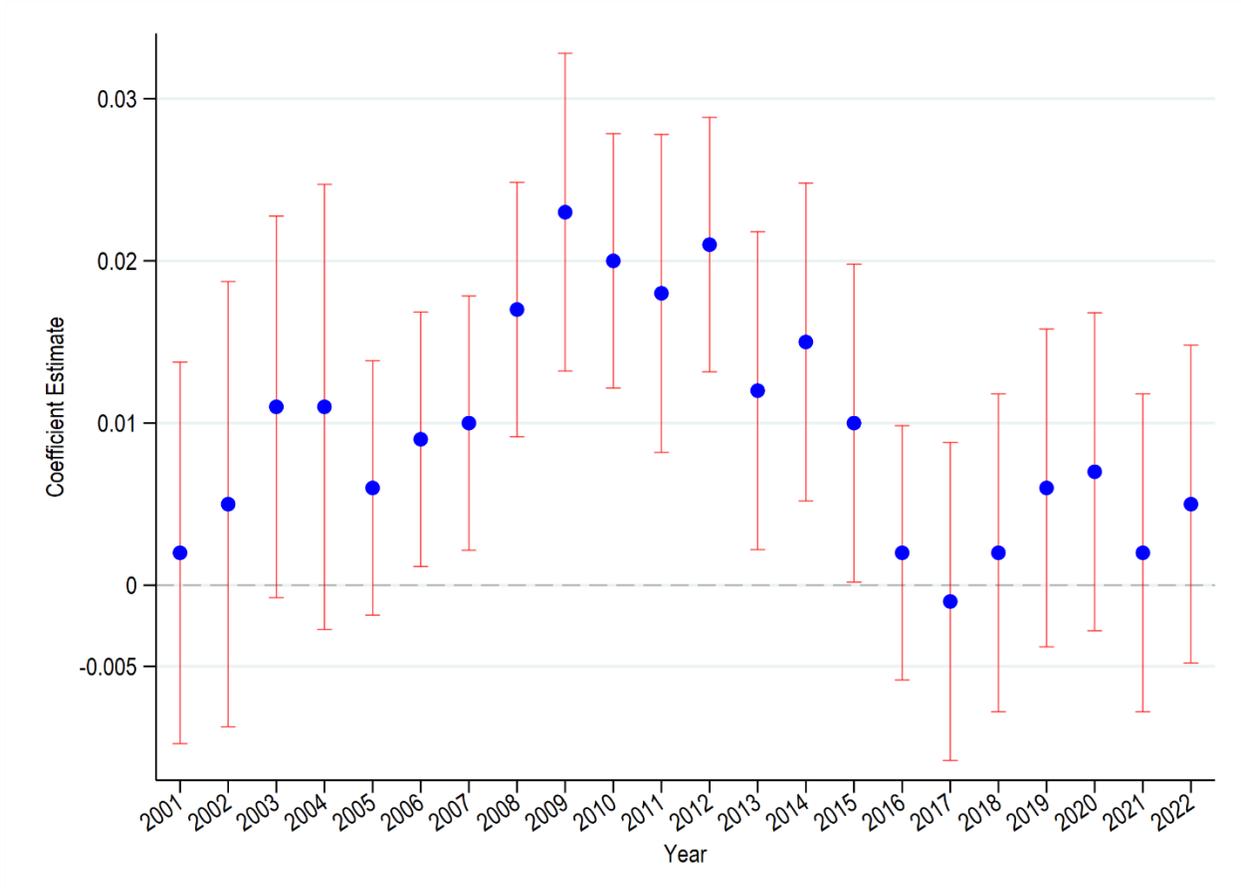
Notes: The dependent variable is an indicator for being employed in the wholesale trade industry. The sample and variables are otherwise the same as the main analysis. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure A11: Log Earnings Results Including the 2001-2004 ACS



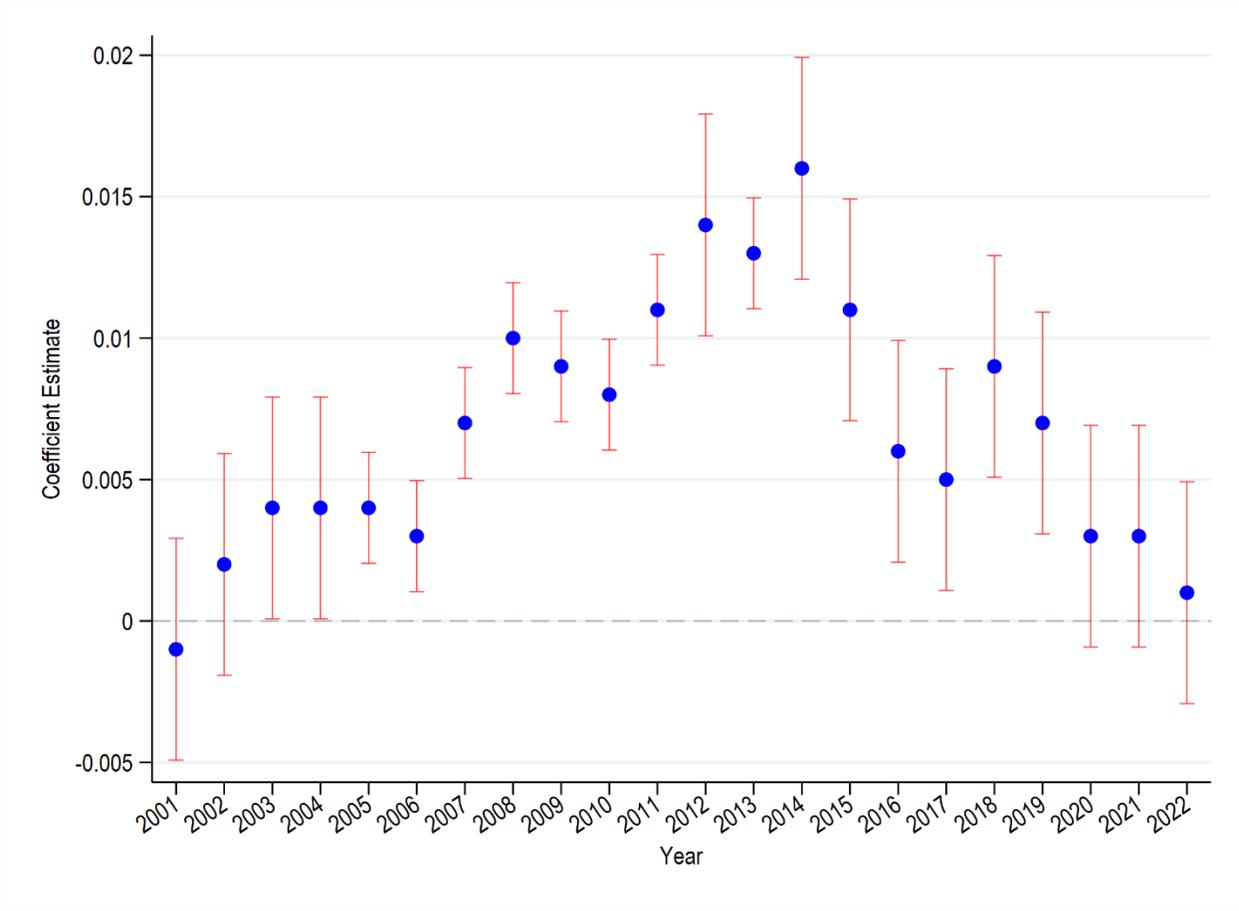
Notes: Sample and measurement are similar to Figure 4 except the sample also includes the 2001-2004 ACS. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 5,444,000 observations (rounded) in the sample.

Figure A12: Employment Indicator Results Including the 2001-2004 ACS



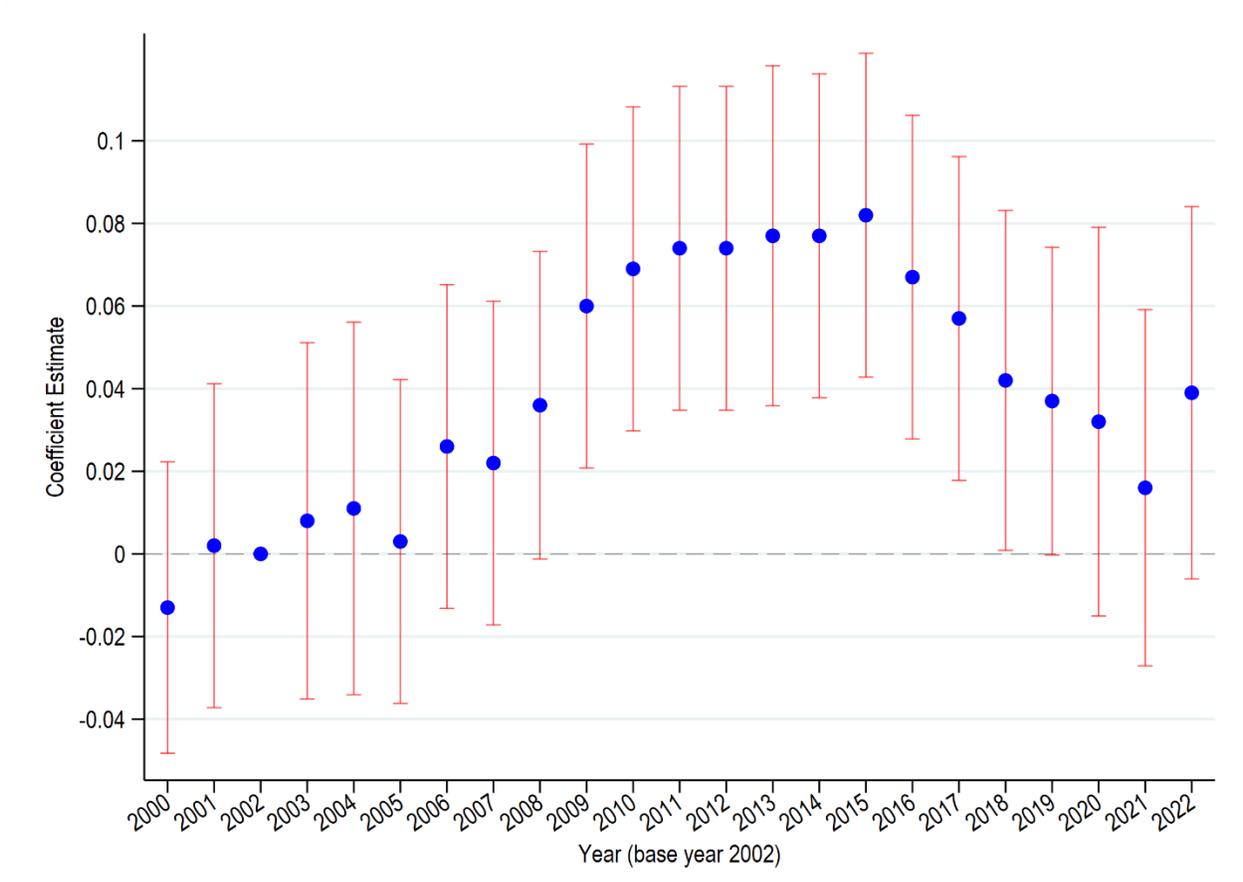
Notes: Sample and measurement are similar to Figure 5 except the sample also includes the 2001-2004 ACS. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,312,000 observations (rounded) in the sample.

Figure A13: Mining Employment Indicator Results Including the 2001-2004 ACS



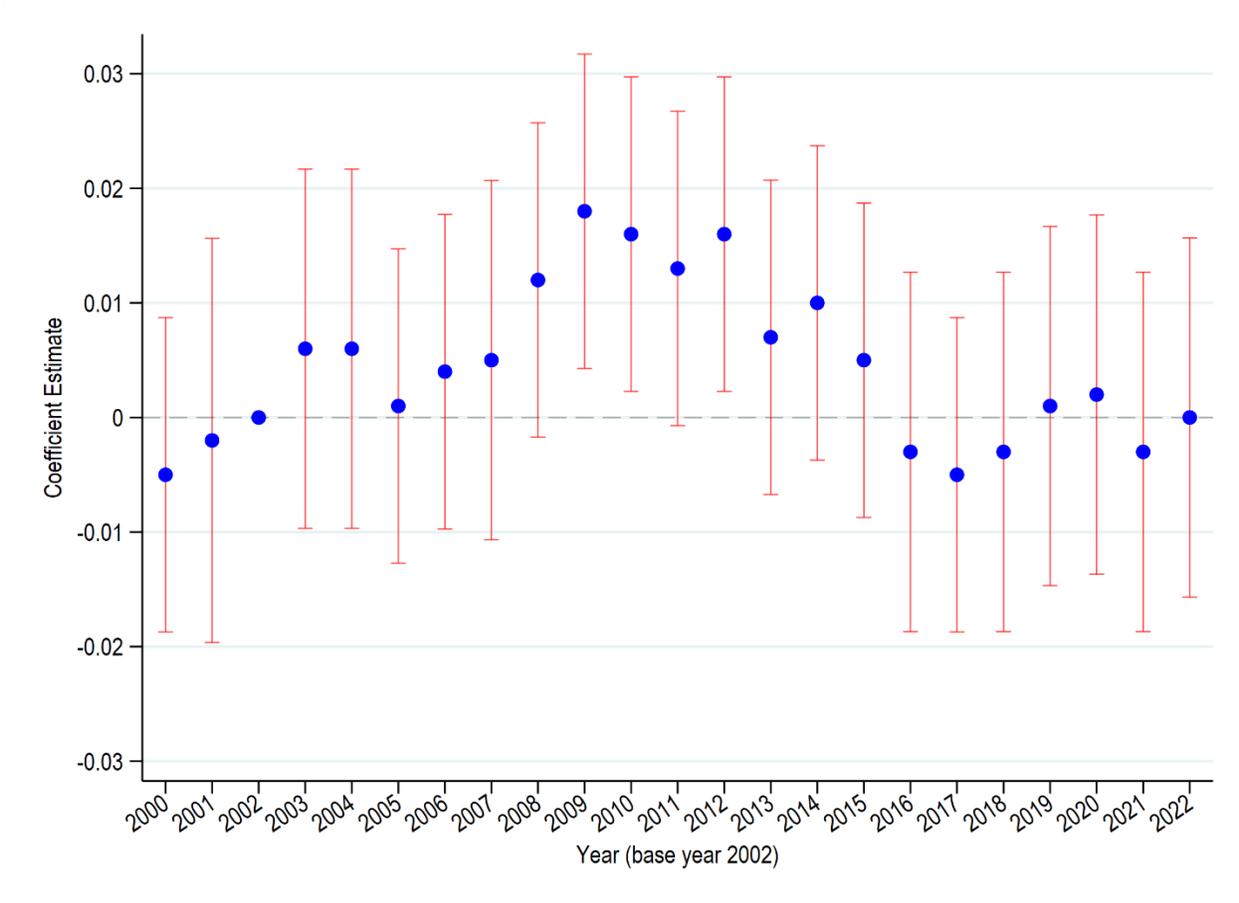
Notes: Sample and measurement are similar to Figure 6 except the sample also includes the 2001-2004 ACS. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,312,000 observations (rounded) in the sample.

Figure A14: Log Earnings Results with 2002 as the Base Year



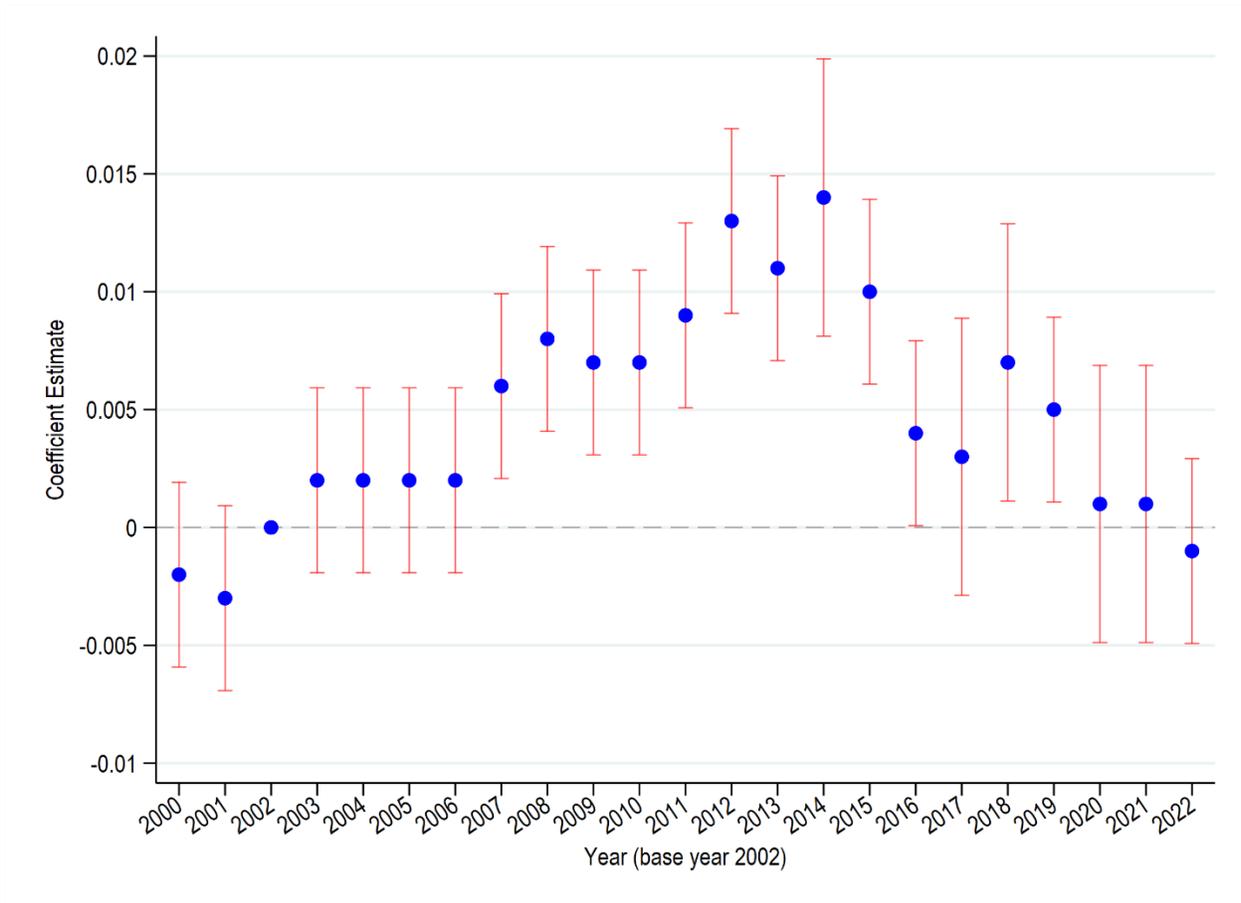
Notes: Sample and measurement are similar to Figure 4 except the sample also includes the 2001-2004 ACS and 2002 is the base year. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2002), and the vertical error bars represent 95% confidence intervals. There are 5,444,000 observations (rounded) in the sample.

Figure A15: Employment Indicator Results with 2002 as the Base Year



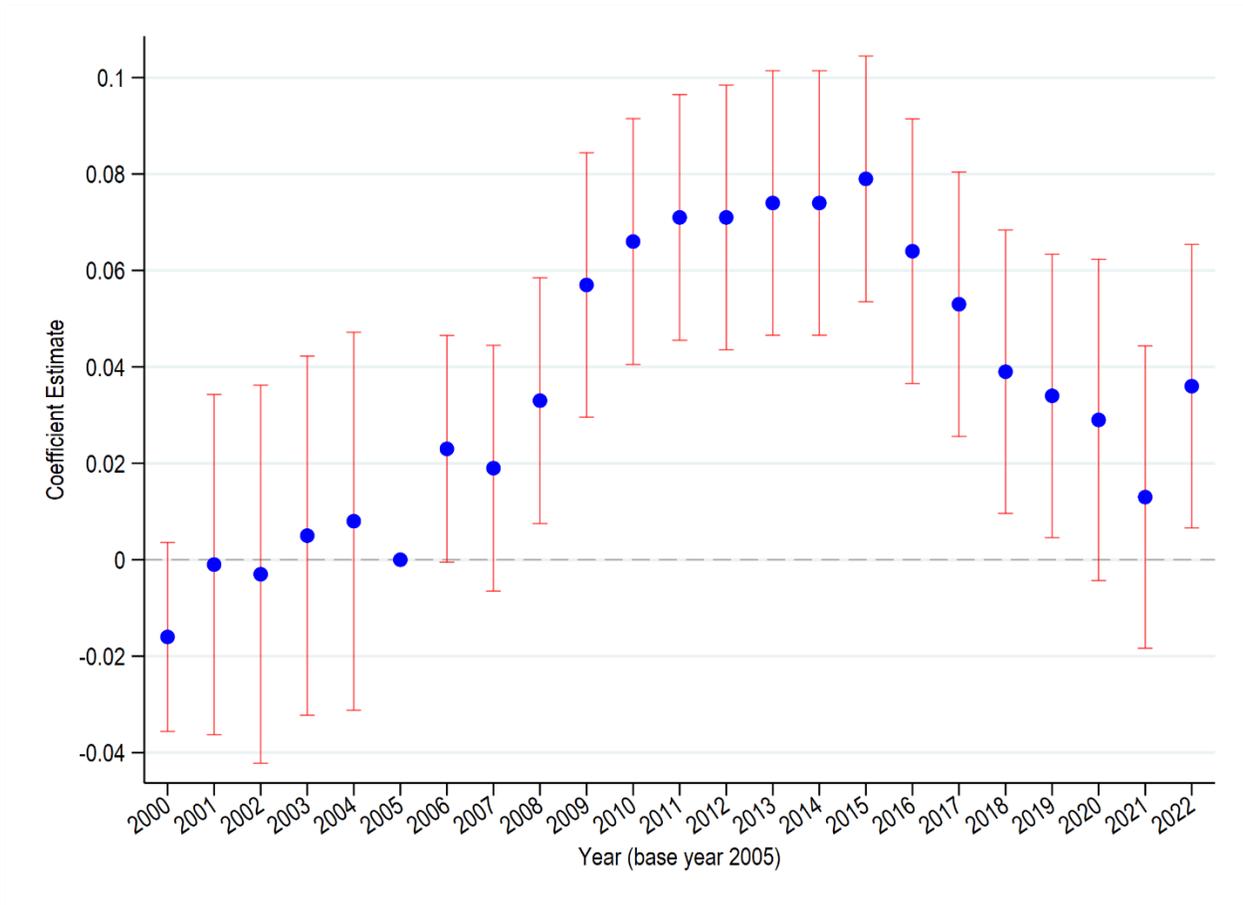
Notes: Sample and measurement are similar to Figure 5 except the sample also includes the 2001-2004 ACS and 2002 is the base year. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2002), and the vertical error bars represent 95% confidence intervals. There are 6,312,000 observations (rounded) in the sample.

Figure A16: Mining Employment Indicator Results with 2002 as the Base Year



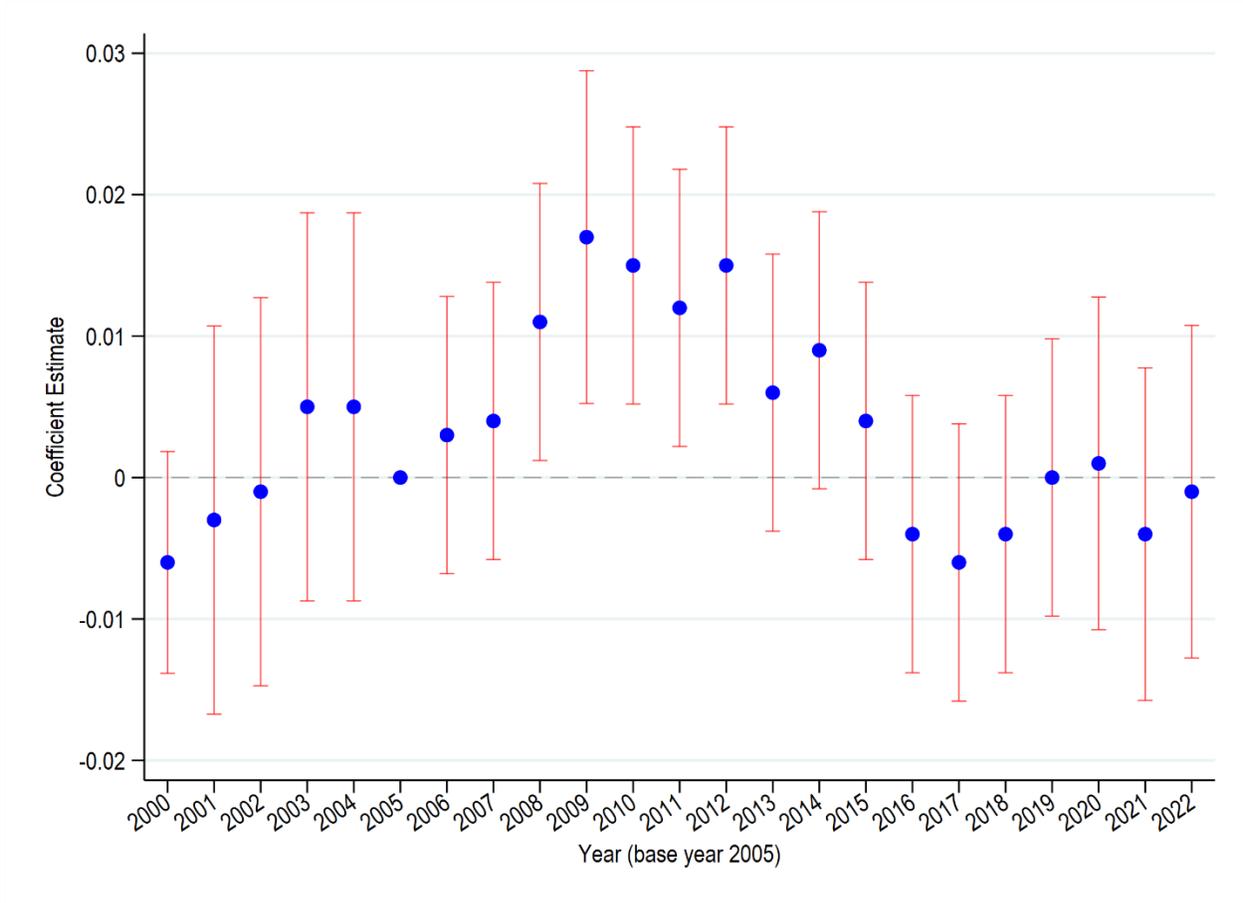
Notes: Sample and measurement are similar to Figure 6 except the sample also includes the 2001-2004 ACS and 2002 is the base year. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2002), and the vertical error bars represent 95% confidence intervals. There are 6,312,000 observations (rounded) in the sample.

Figure A17: Log Earnings Results with 2005 as the Base Year



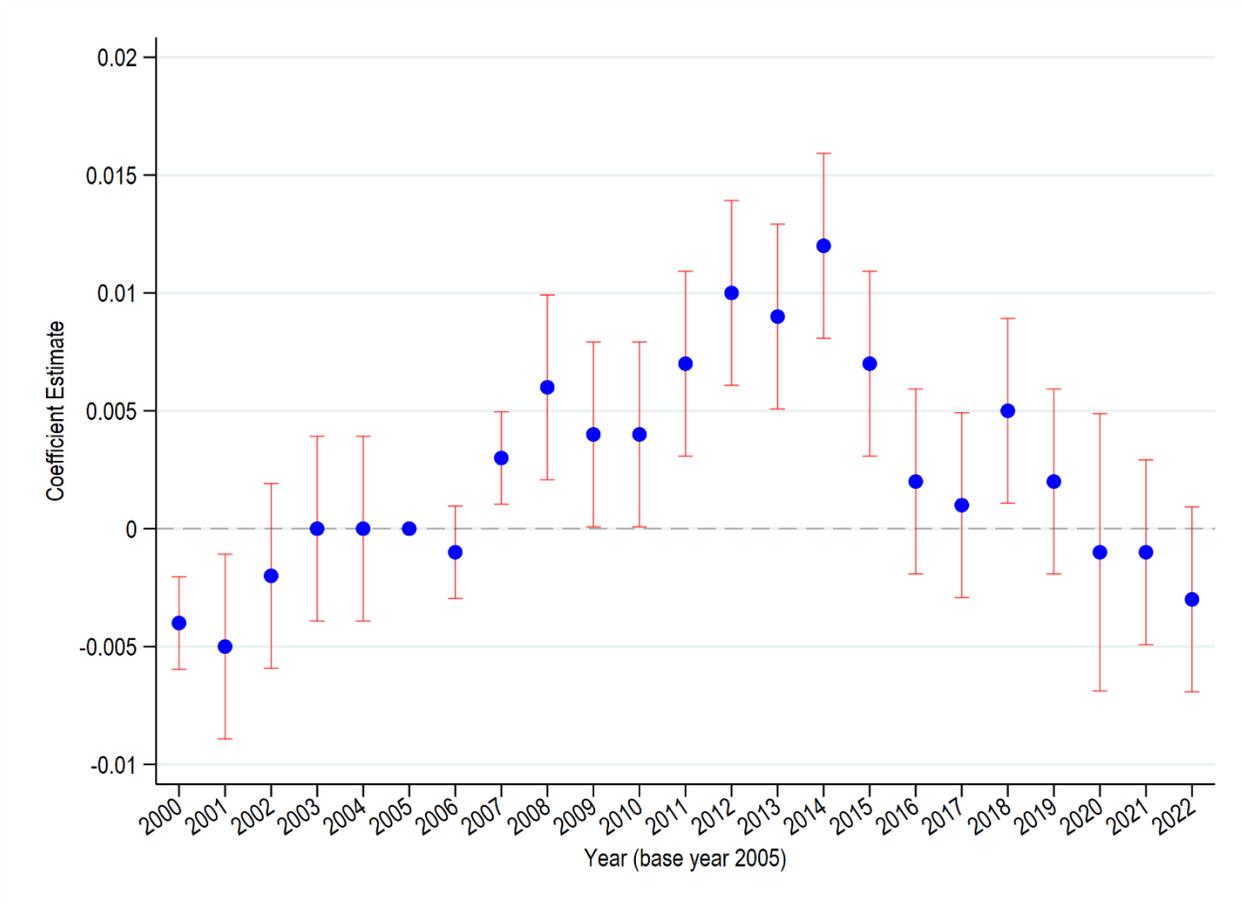
Notes: Sample and measurement are similar to Figure 4 except the sample also includes the 2001-2004 ACS and 2005 is the base year. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2005), and the vertical error bars represent 95% confidence intervals. There are 5,444,000 observations (rounded) in the sample.

Figure A18: Employment Indicator Results with 2005 as the Base Year



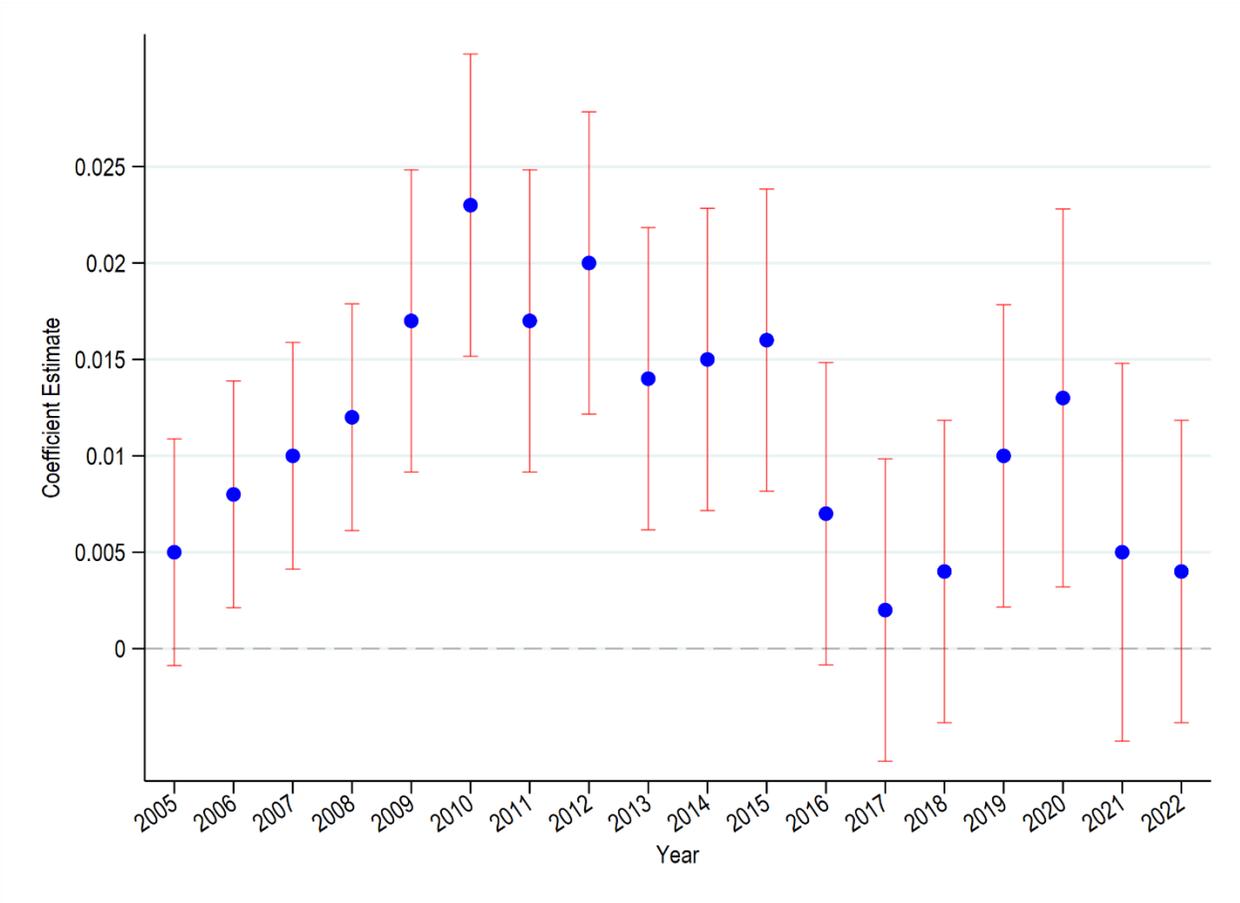
Notes: Sample and measurement are similar to Figure 5 except the sample also includes the 2001-2004 ACS and 2005 is the base year. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2005), and the vertical error bars represent 95% confidence intervals. There are 6,312,000 observations (rounded) in the sample.

Figure A19: Mining Employment Indicator Results with 2005 as the Base Year



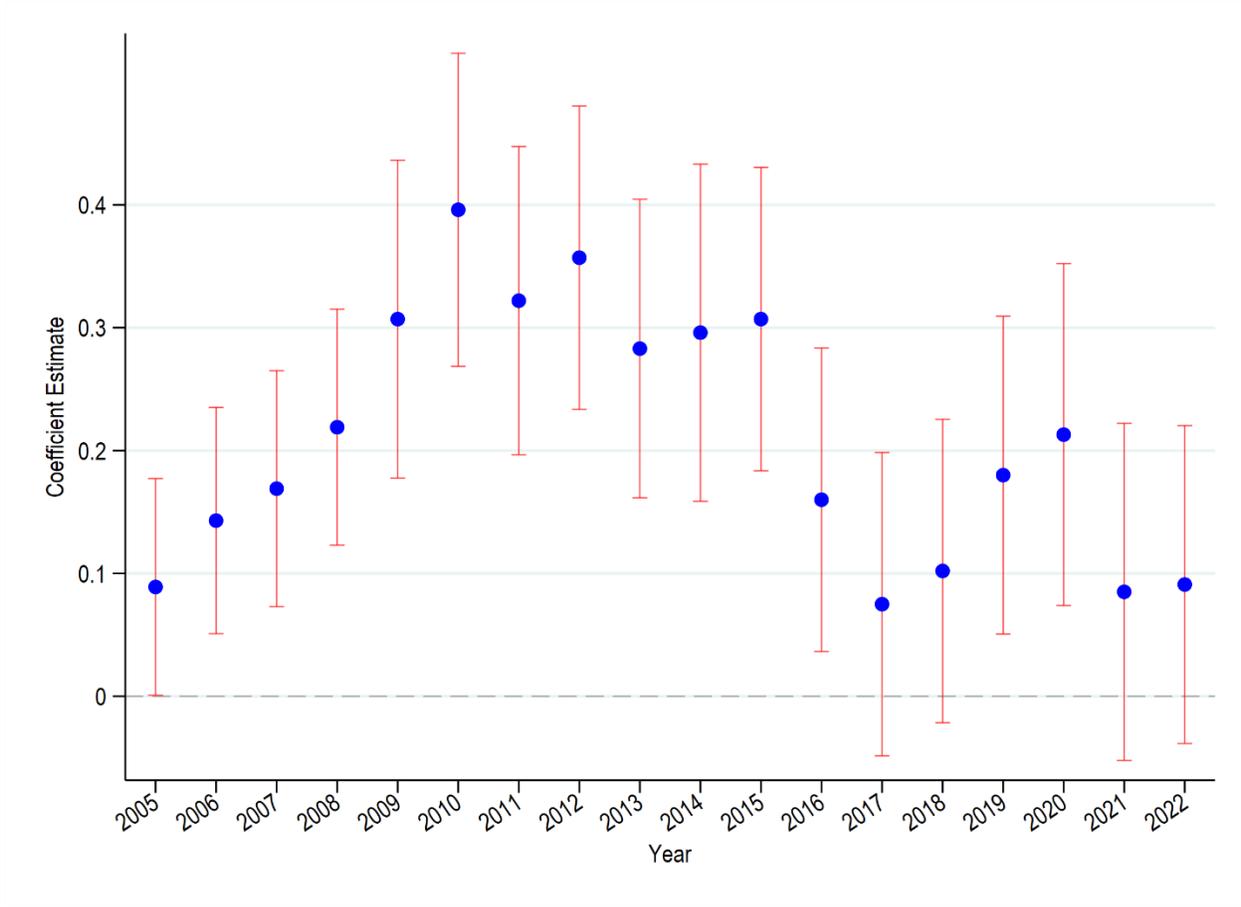
Notes: Sample and measurement are similar to Figure 6 except the sample also includes the 2001-2004 ACS and 2005 is the base year. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2005), and the vertical error bars represent 95% confidence intervals. There are 6,312,000 observations (rounded) in the sample.

Figure A20: Fracking Impacts on Having Positive Annual Earnings



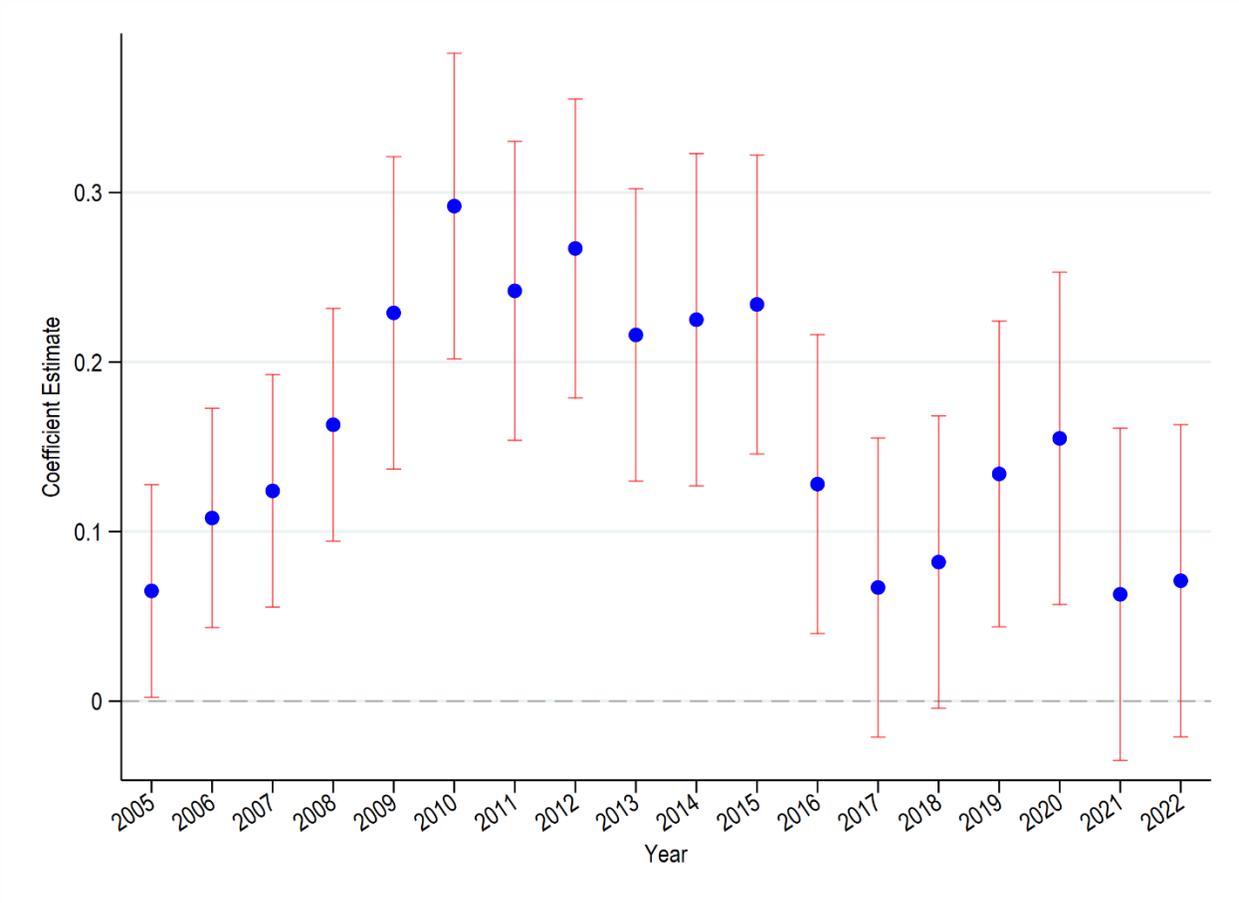
Notes: Sample and measurement are similar to Figure 5 except the dependent variable is an indicator for having positive annual earnings. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure A21: Fracking Impacts on Log Earnings with Non-Positive Earnings Set to \$0.01



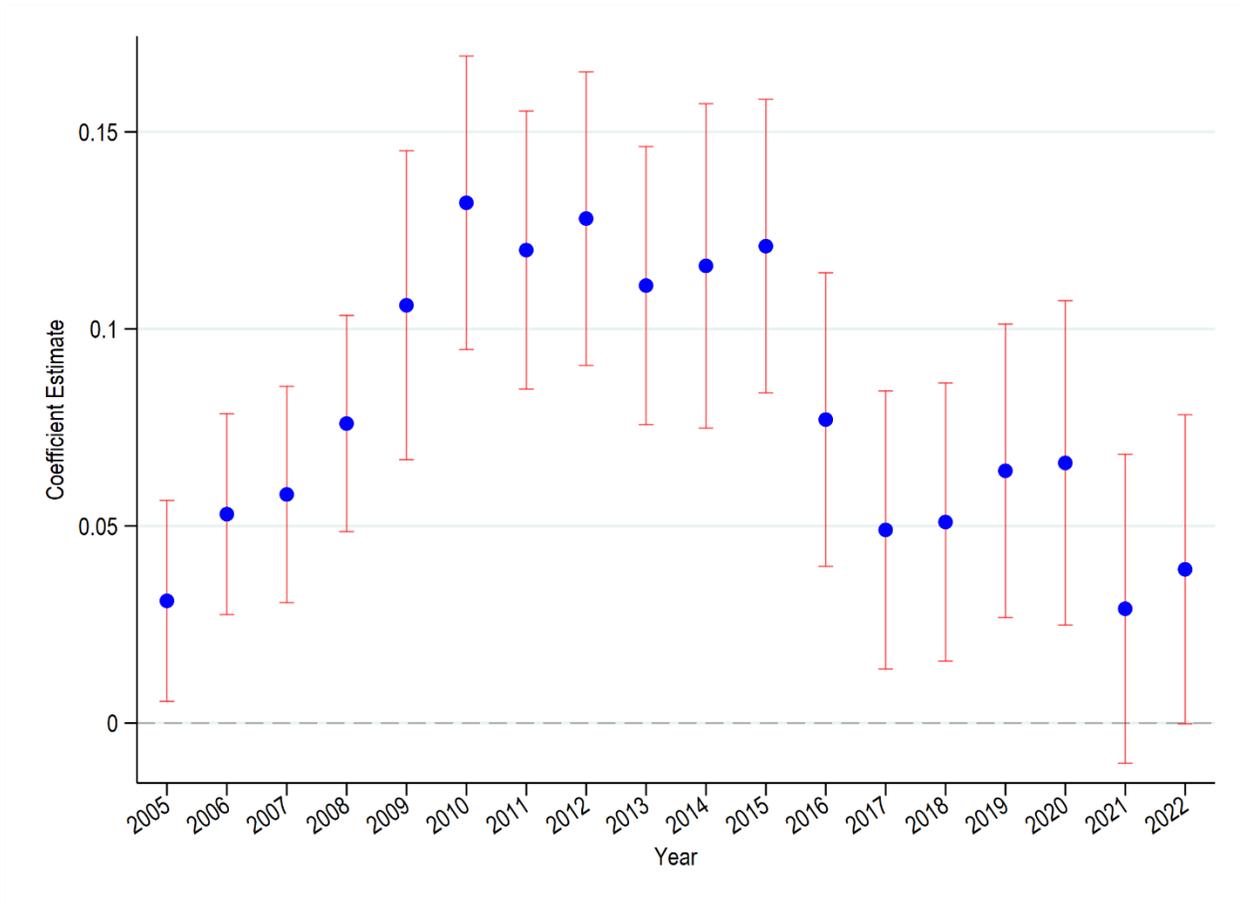
Notes: Sample and measurement are similar to Figure 4 except non-positive earnings are set to \$0.01 before taking logs. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure A22: Fracking Impacts on Log Earnings with Non-Positive Earnings Set to \$1



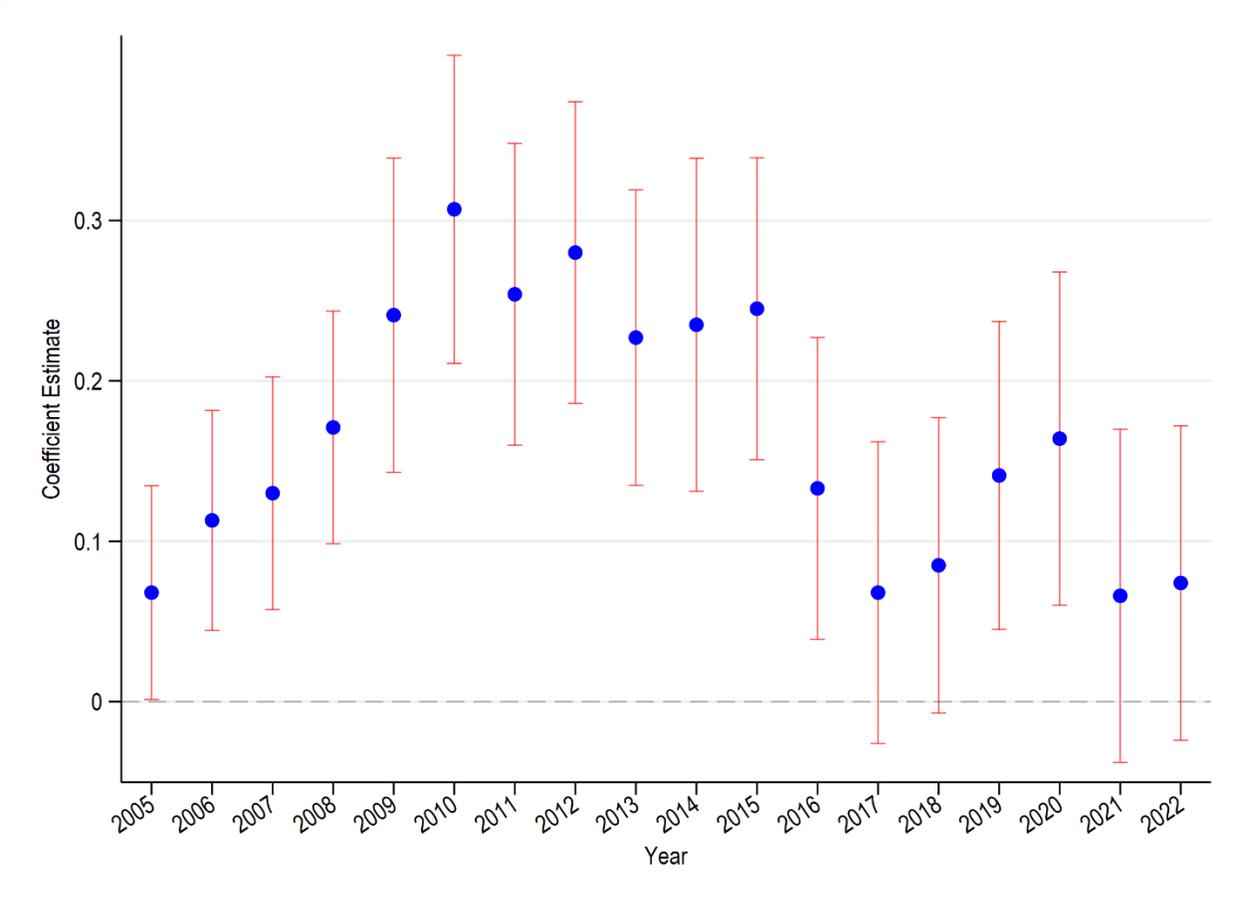
Notes: Sample and measurement are similar to Figure 4 except non-positive earnings are set to \$1 before taking logs. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure A23: Fracking Impacts on Log Earnings with Non-Positive Earnings Set to \$1000



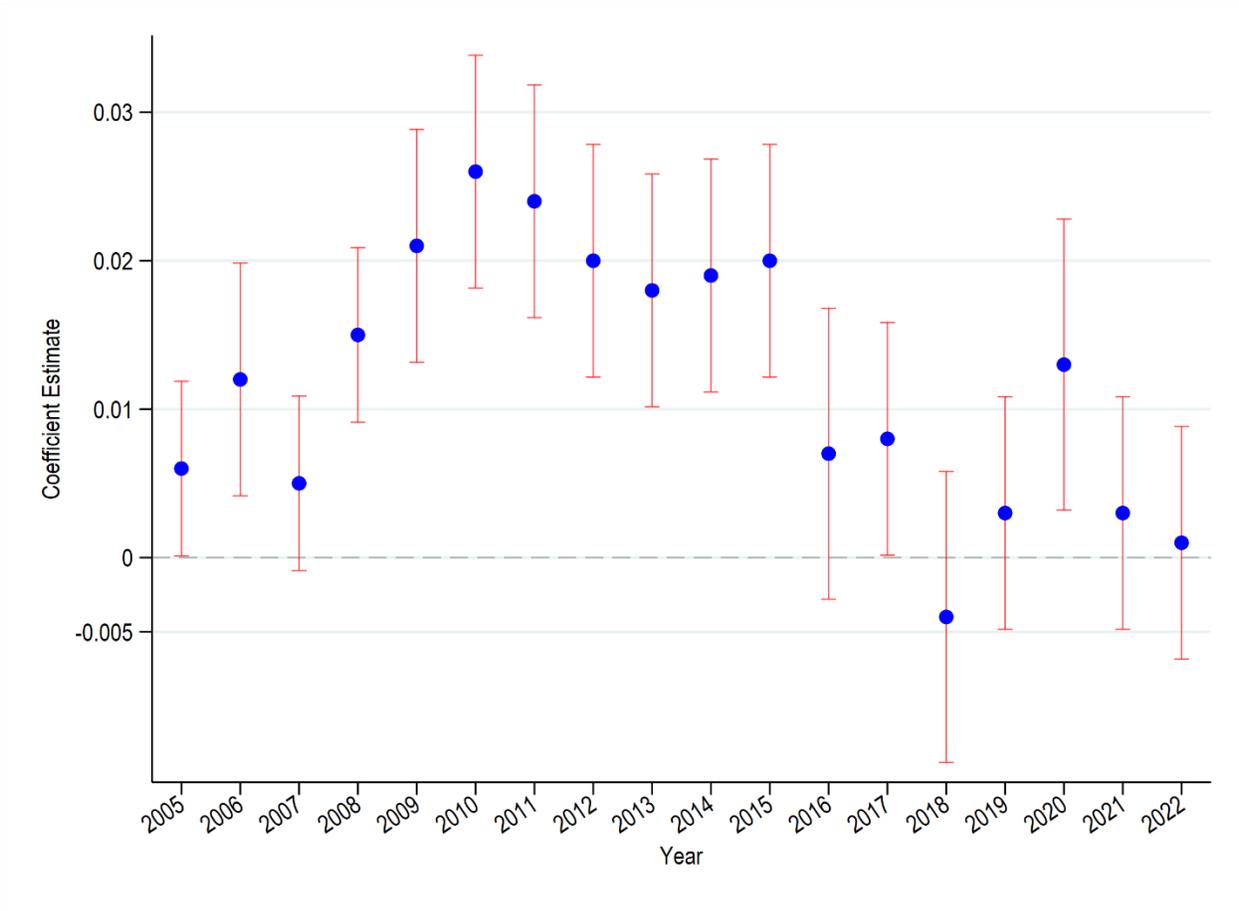
Notes: Sample and measurement are similar to Figure 4 except non-positive earnings are set to \$1000 before taking logs. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure A24: Fracking Impacts on Inverse Hyperbolic Sine Transformation of Annual Earnings



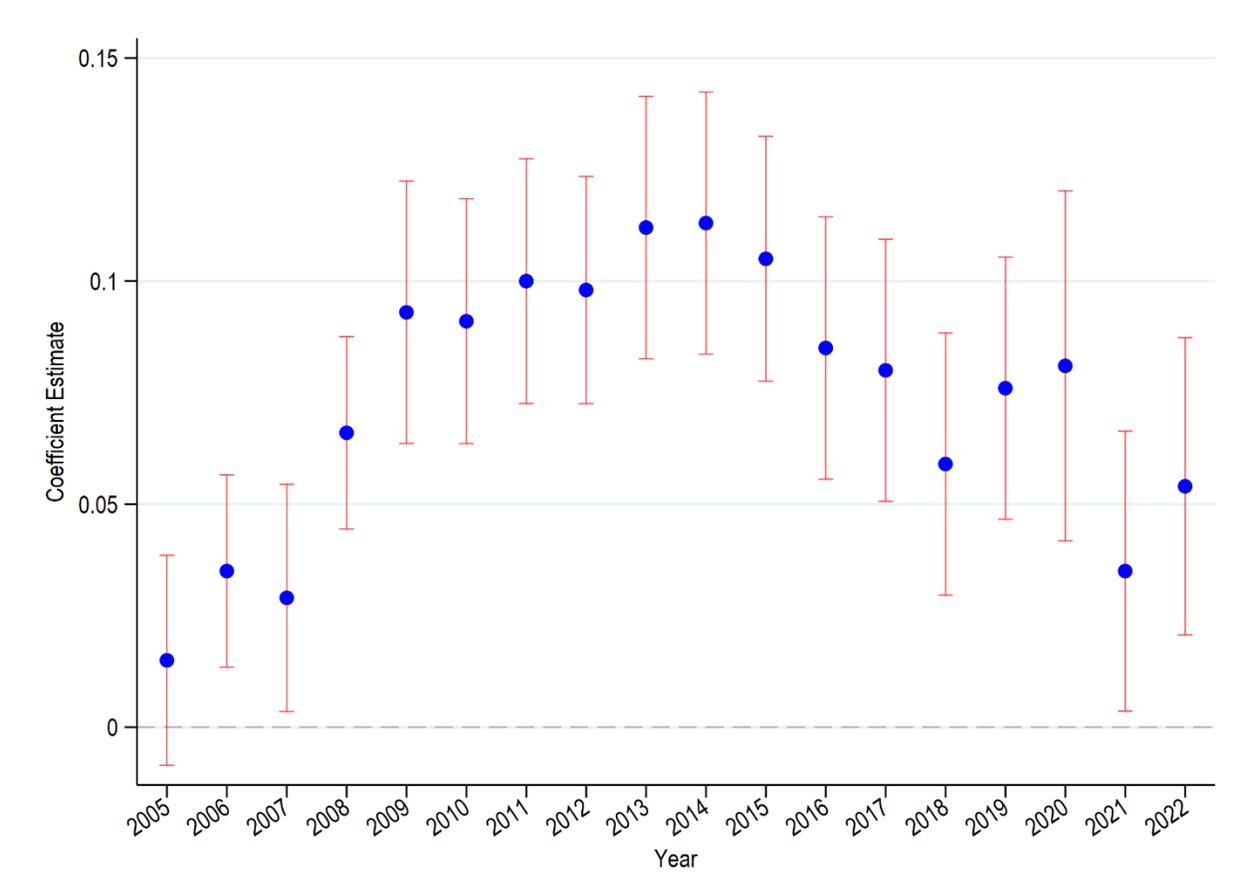
Notes: Sample and measurement are similar to Figure 4 except the dependent variable is the inverse hyperbolic sine transformation of annual earnings. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure A25: Fracking Impacts on Household Income above Poverty Threshold for Ages 25-54



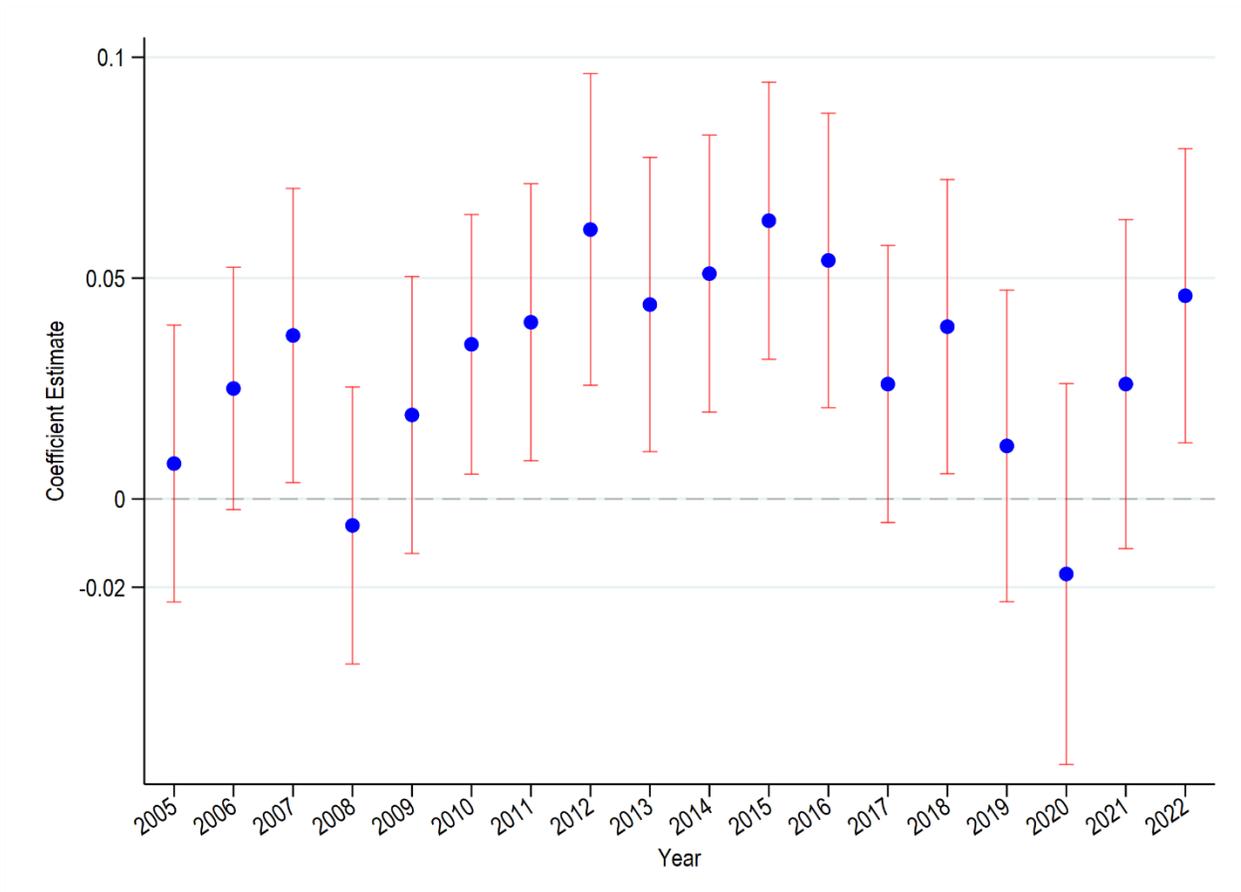
Notes: Sample and measurement are the same as indicated in Figure 4 notes. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 6,044,000 observations (rounded) in the sample.

Figure A26: Fracking Impacts on Log Annual Earnings for Non-College Graduate Workers



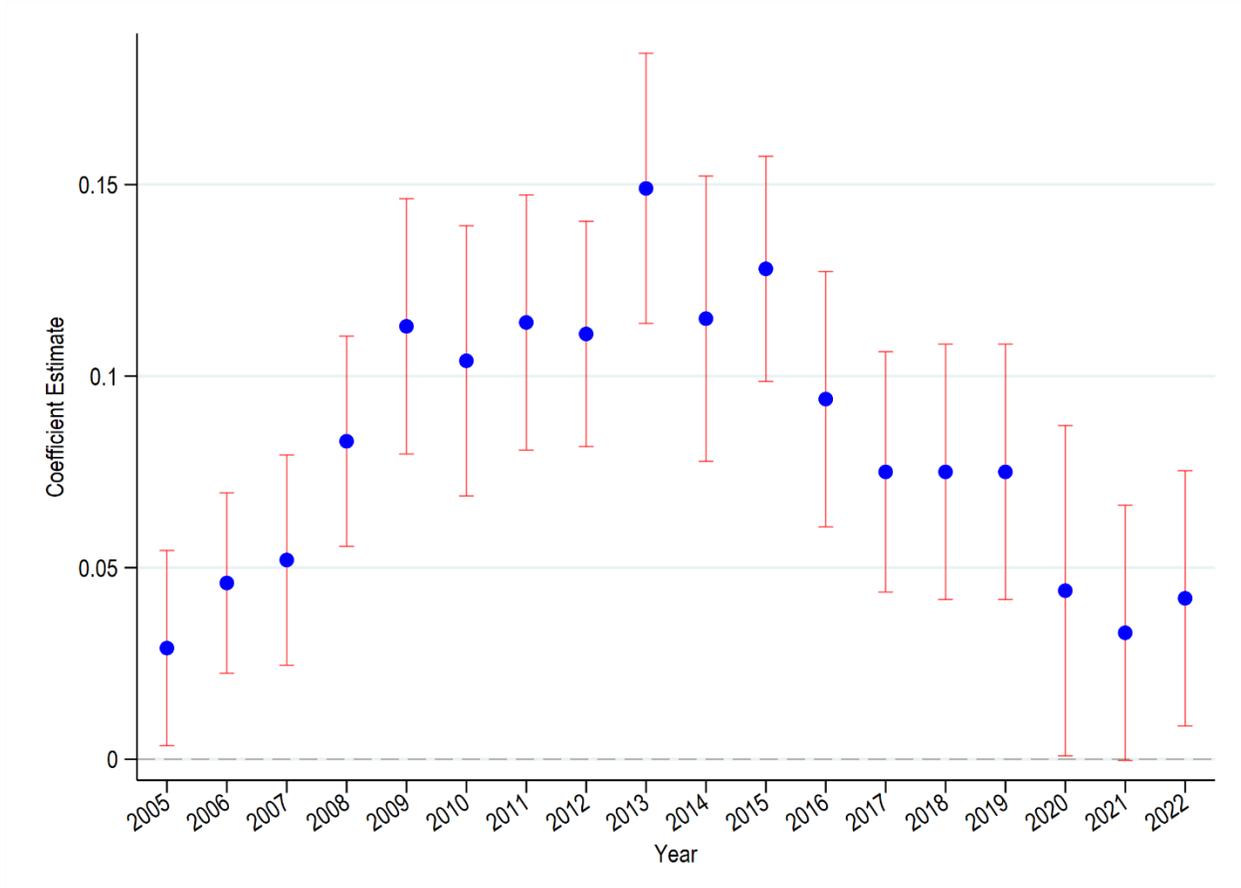
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to workers with less than a bachelor's degree. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 3,788,000 observations (rounded) in the sample.

Figure A27: Fracking Impacts on Log Annual Earnings for College Graduate Workers



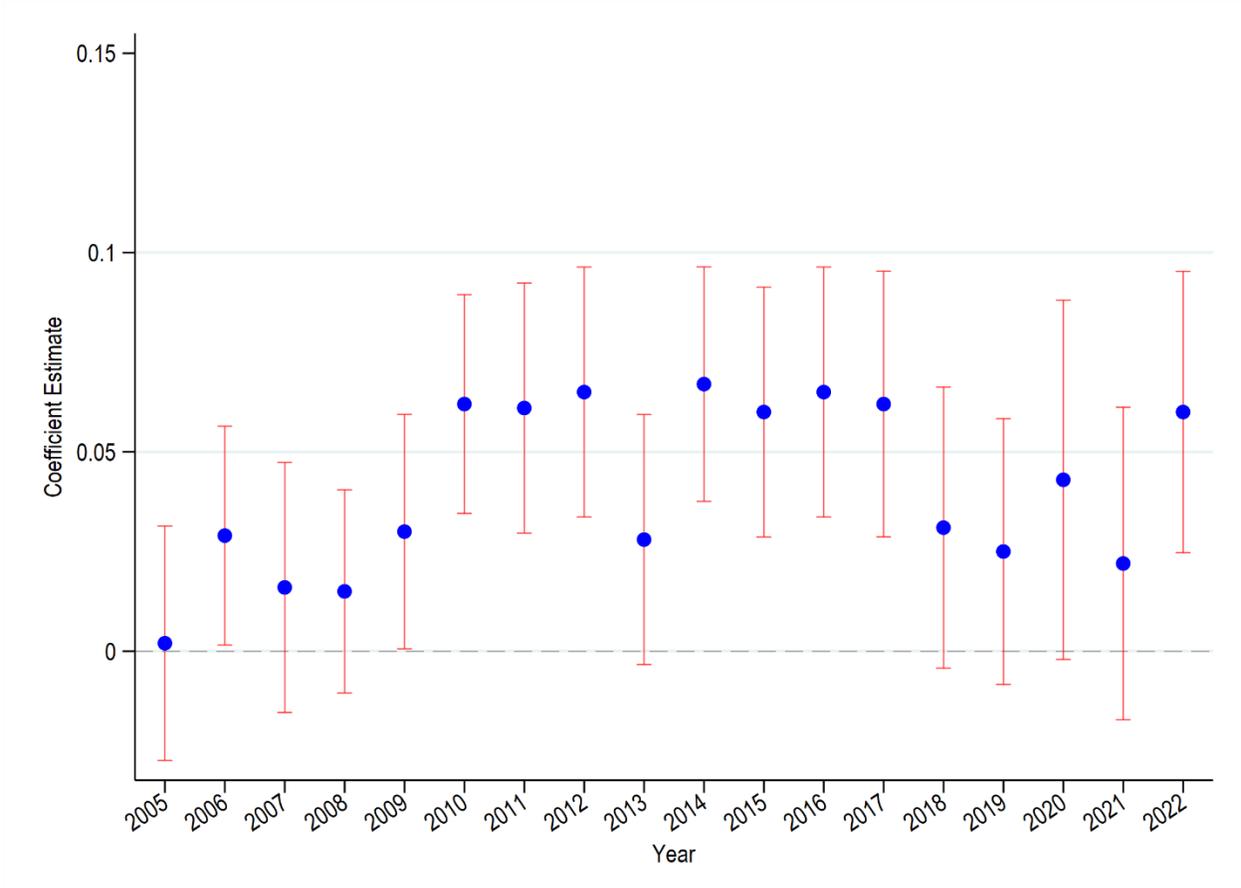
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to workers with a bachelor's degree or higher education. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 1,419,000 observations (rounded) in the sample.

Figure A28: Fracking Impacts on Log Annual Earnings for Men



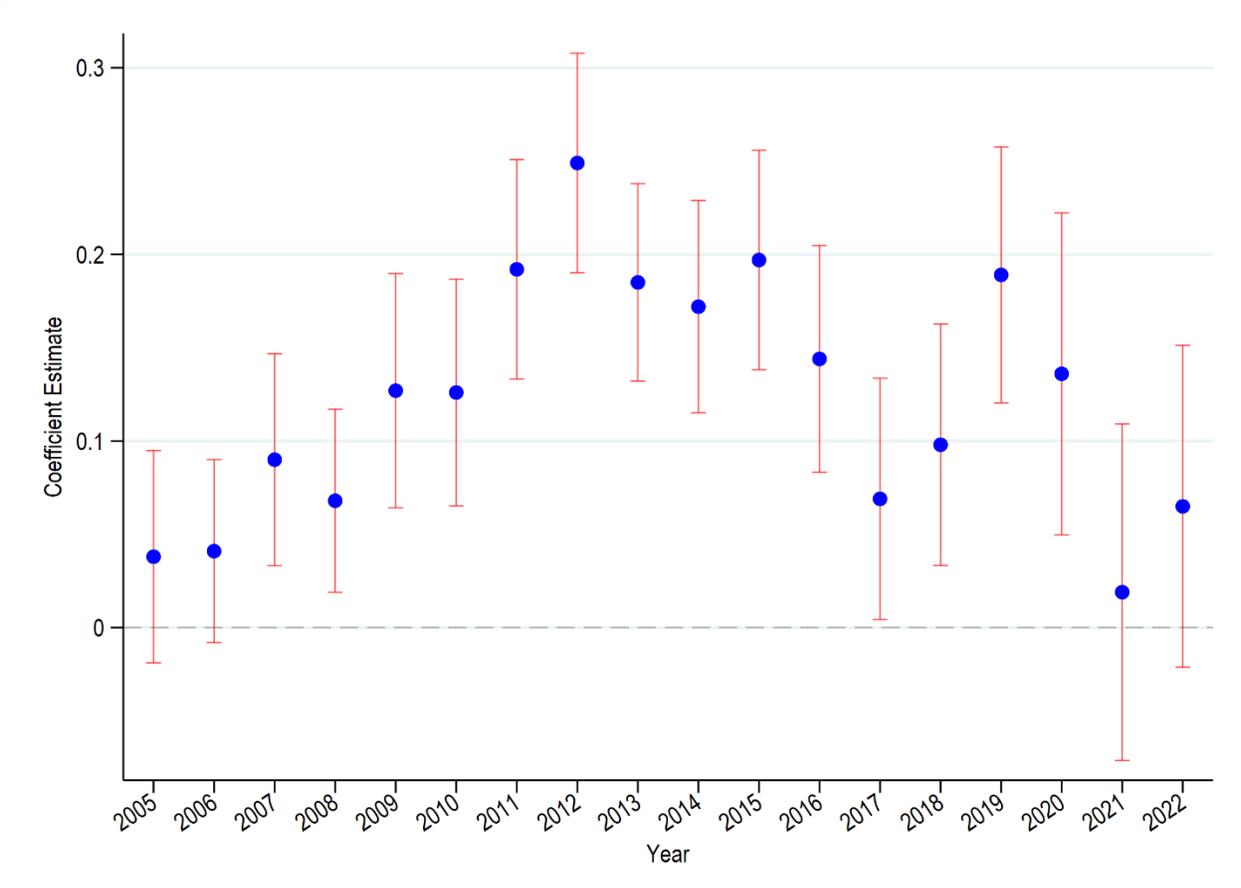
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to male workers. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 2,713,000 observations (rounded) in the sample.

Figure A29: Fracking Impacts on Log Annual Earnings for Women



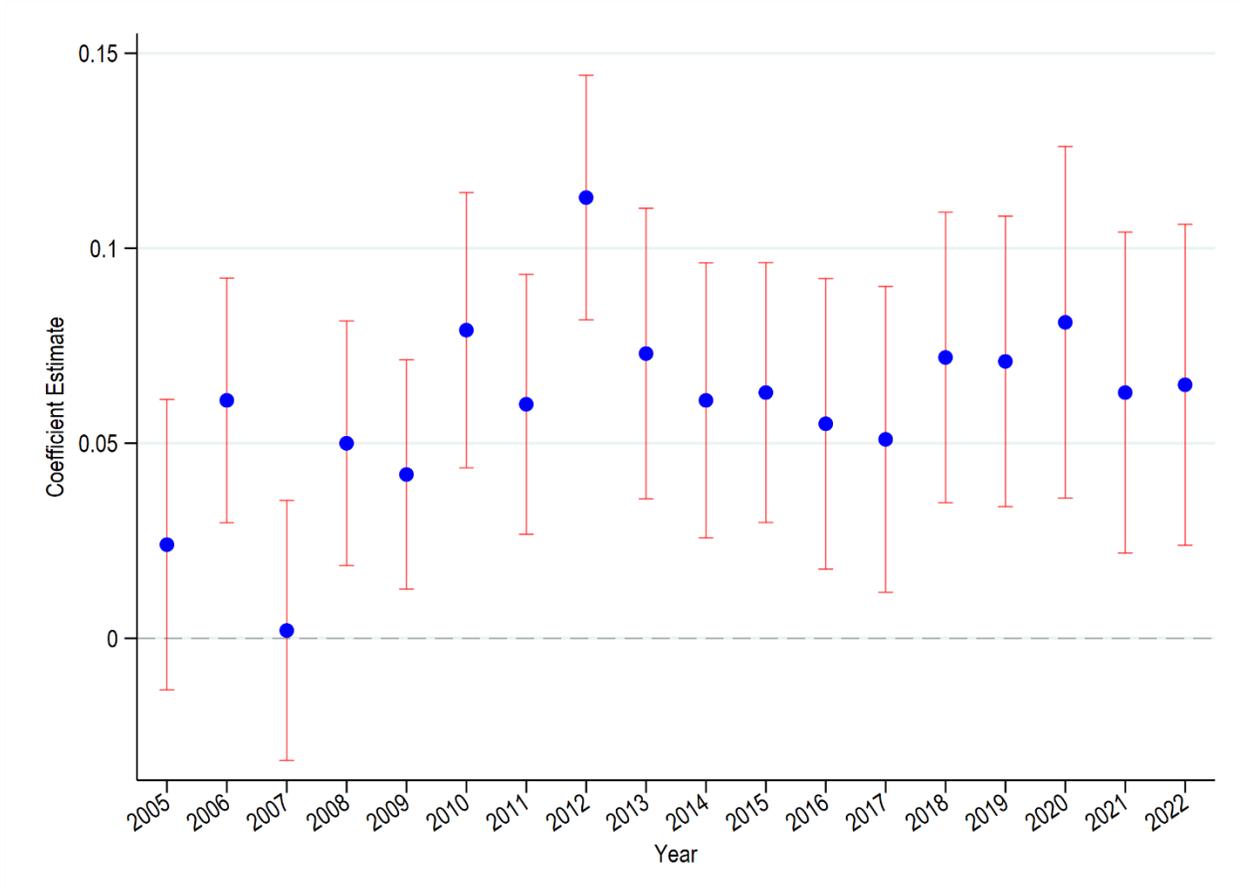
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to female workers. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 2,495,000 observations (rounded) in the sample.

Figure A30: Fracking Impacts on Log Annual Earnings for Ages 18-24



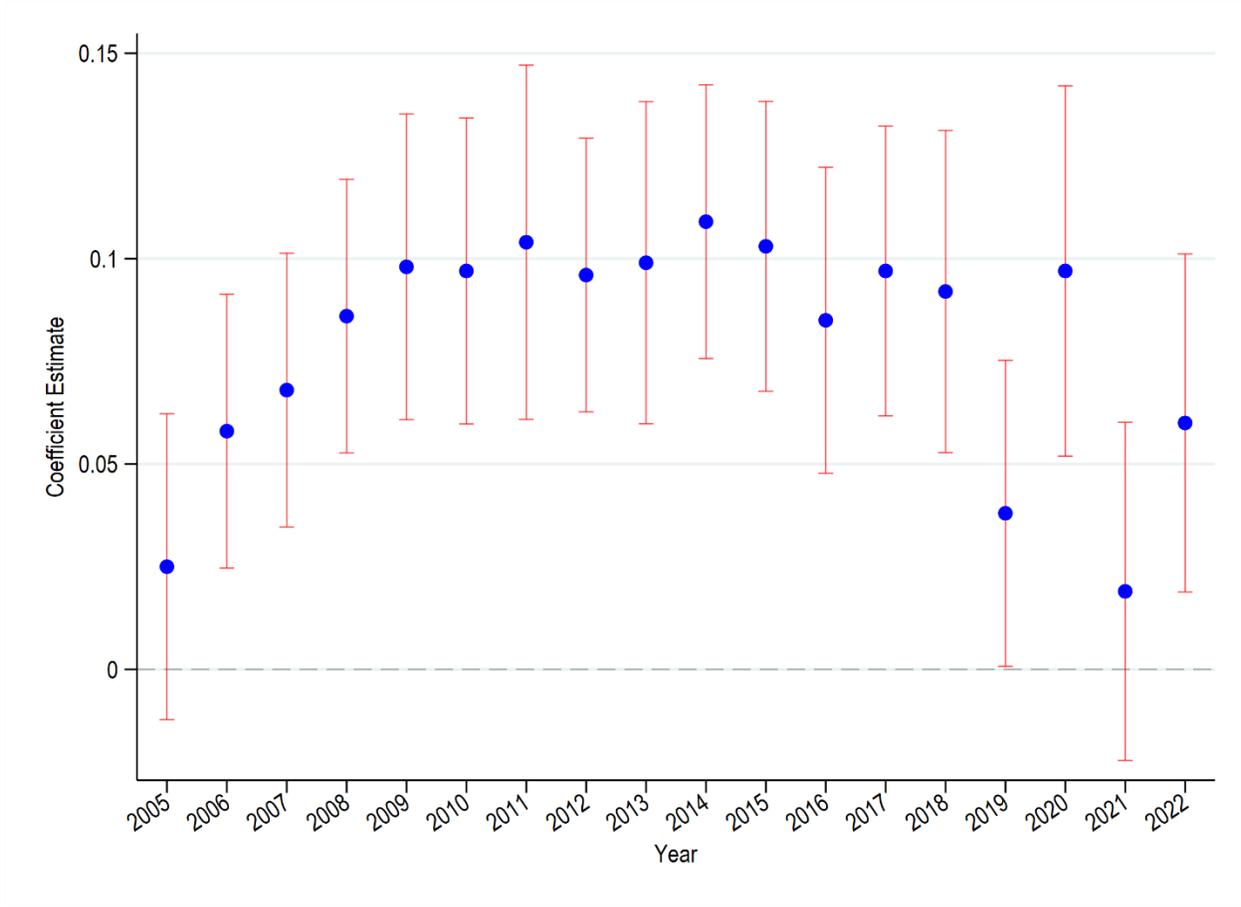
Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to ages 18-24. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 809,000 observations (rounded) in the sample.

Figure A31: Fracking Impacts on Log Annual Earnings for Ages 55-64



Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to ages 55-64. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 1,665,000 observations (rounded) in the sample.

Figure A32: Fracking Impacts on Log Annual Earnings for Cohorts Born in 1966-1975



Notes: Sample and measurement are the same as indicated in Figure 4 notes except the sample is limited to cohorts born in 1966-1975. These cohorts were roughly ages 25-34 in year 2000 and ages 47-56 by 2022. Each blue dot denotes the coefficient estimate for a given year relative to the omitted baseline year (2000), and the vertical error bars represent 95% confidence intervals. There are 1,533,000 observations (rounded) in the sample.

Figure A33: Oil Productivity in the United States

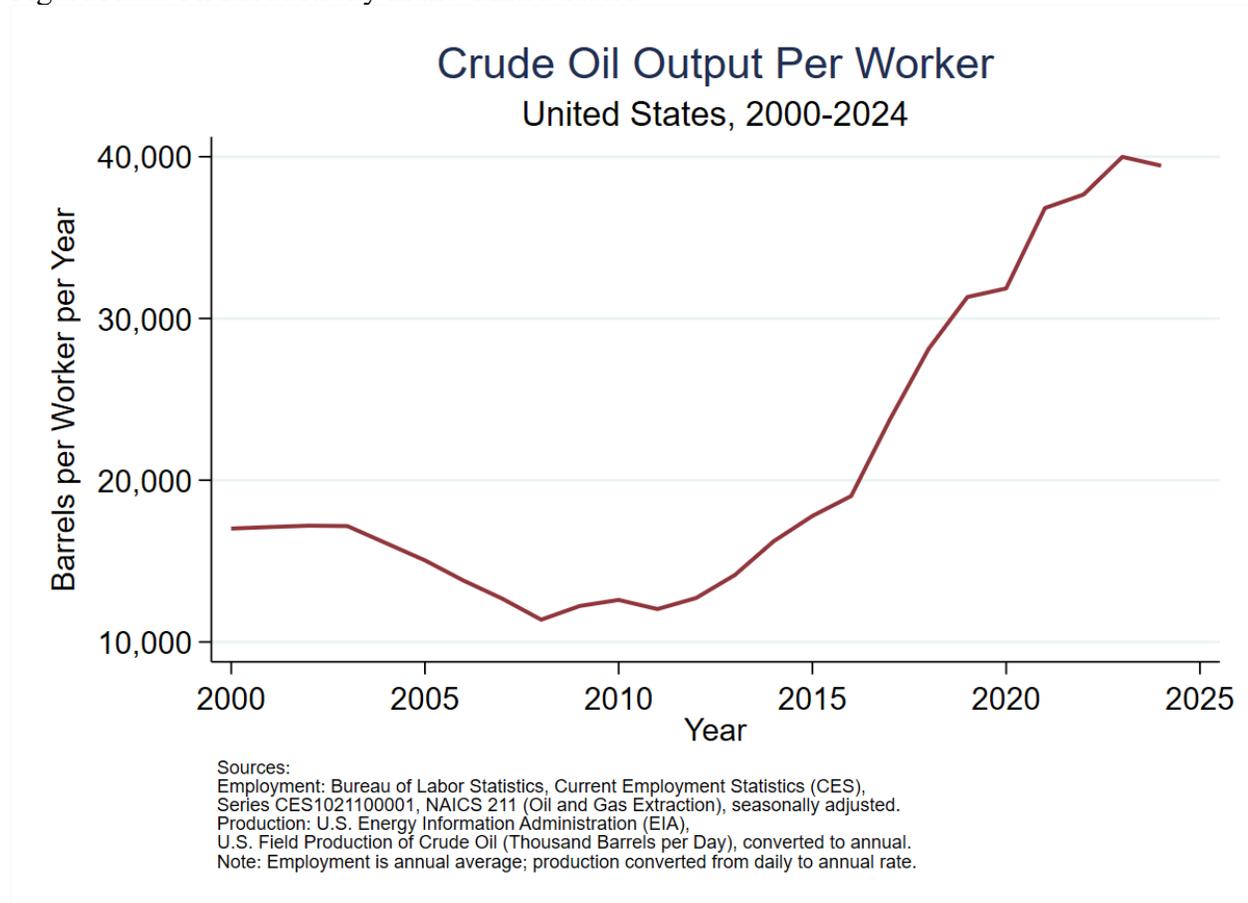


Figure A34: Natural Gas Productivity in the United States

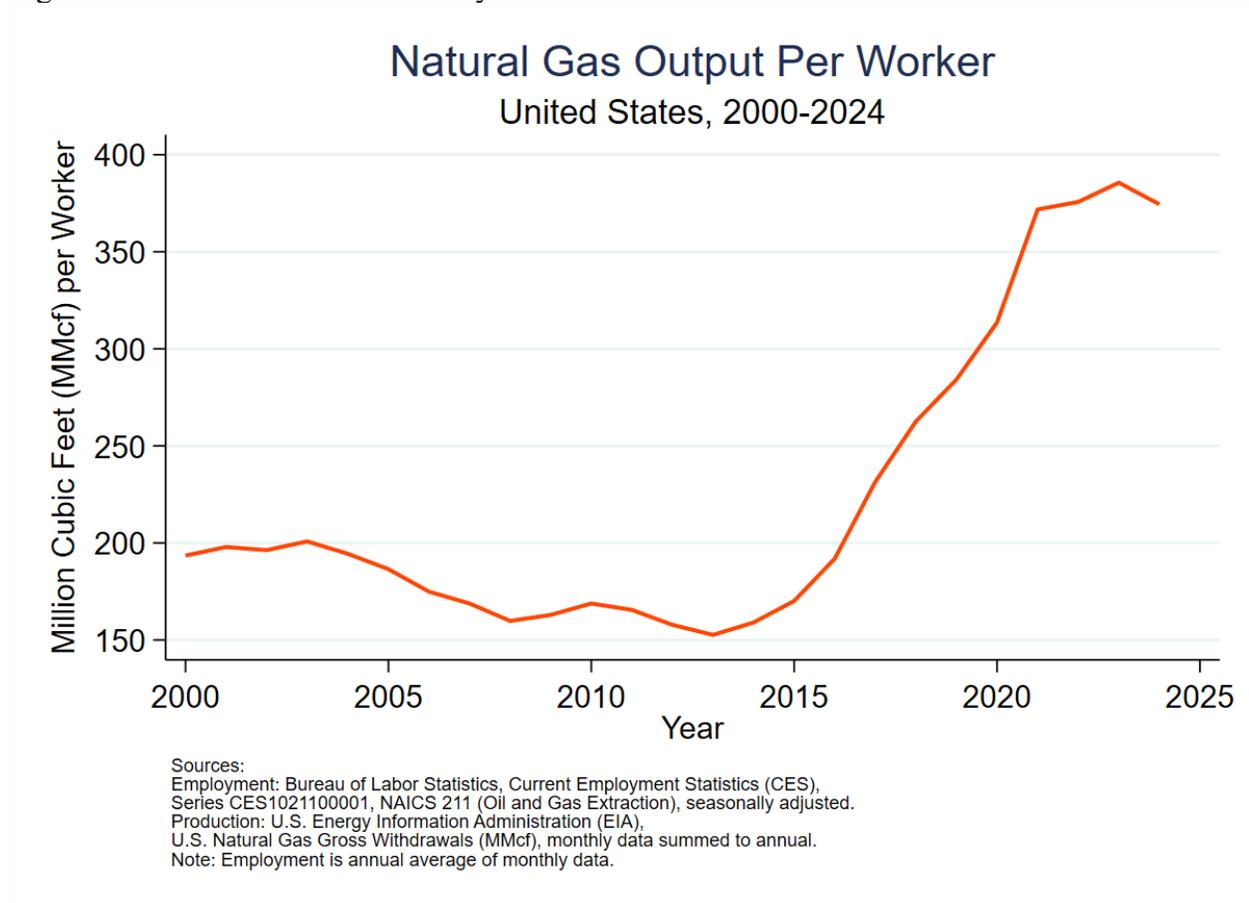


Figure A35: Oil and Gas Employment Trends in the United States

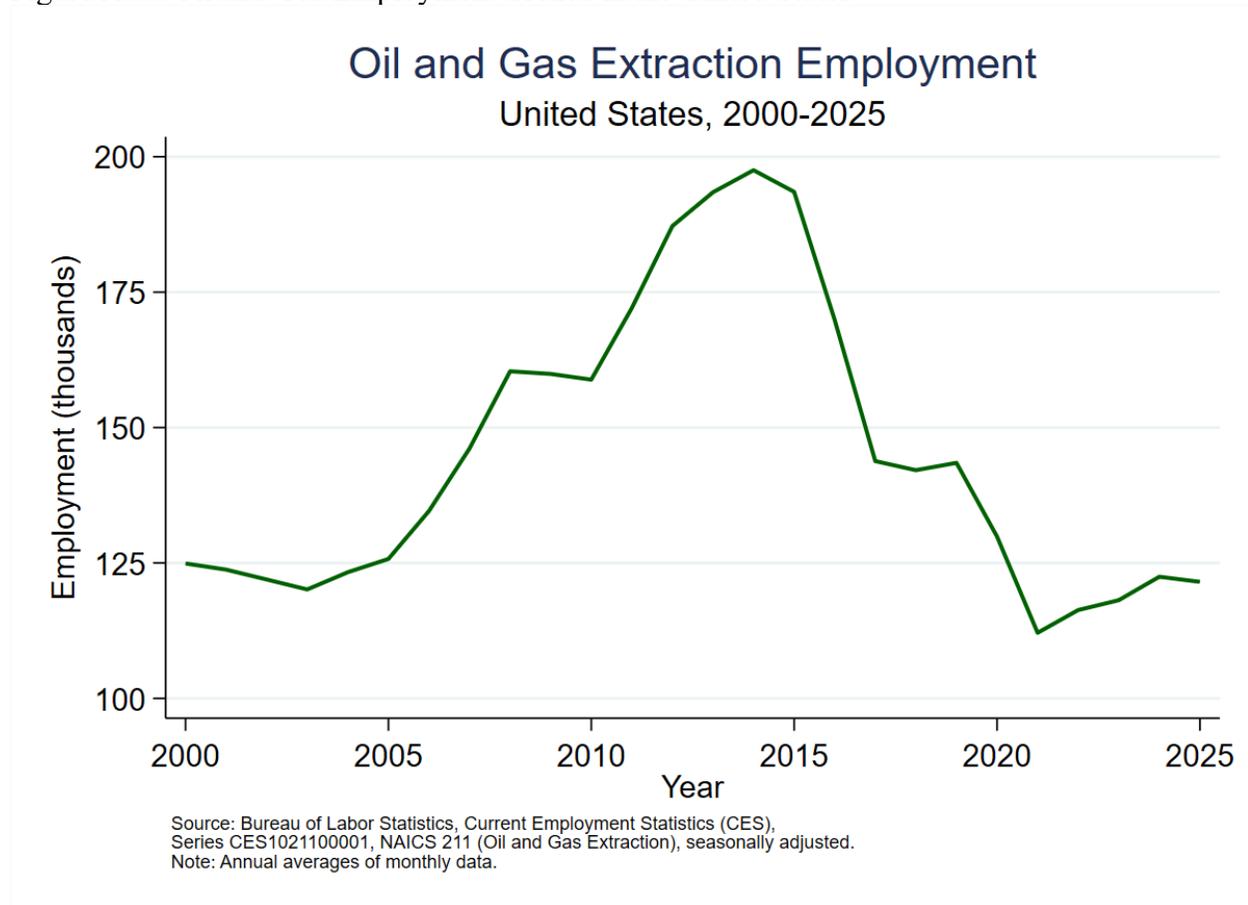
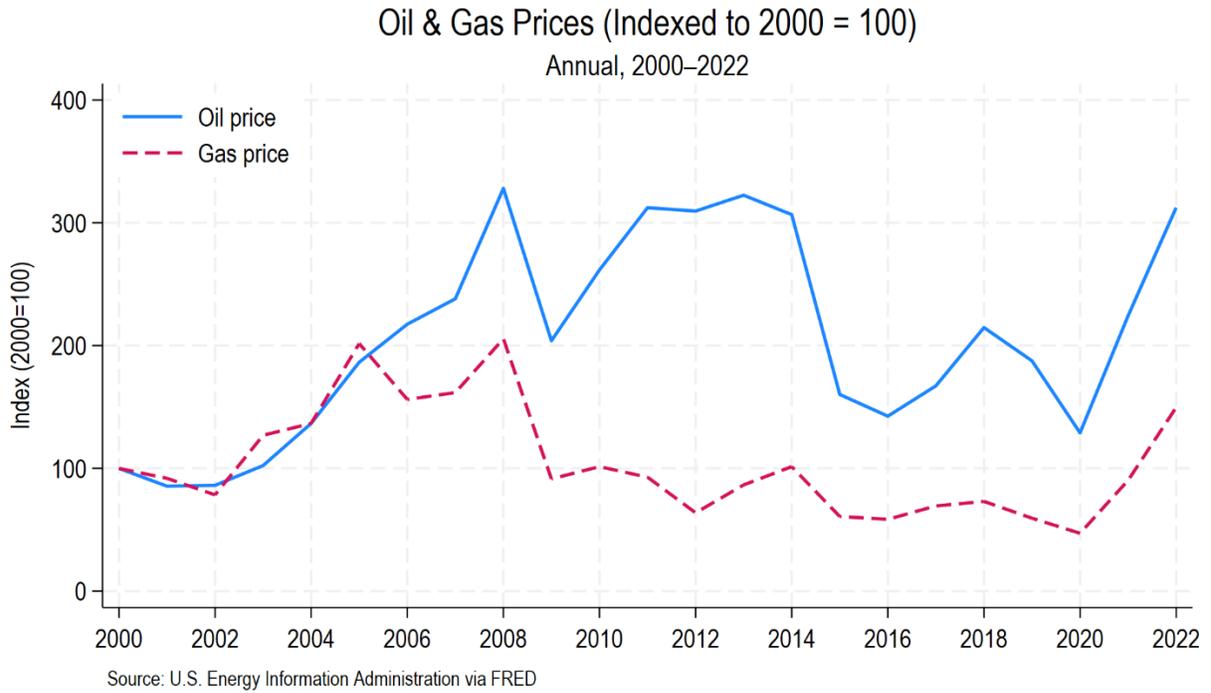


Figure A36: Oil and Natural Gas Price Trends



Notes: Oil prices are annual averages for West Texas Intermediate (WTI) - Cushing, Oklahoma. Natural gas prices are annual averages for Henry Hub Natural Gas Spot Price.