Wildfires and local labor markets

Raphaelle G. Coulombe*   Akhil Rao†

November 15, 2022

Preliminary and incomplete - Please do not cite without authors' permission

Abstract

We analyze the dynamic effects of wildfires on employment in US counties. We construct a novel spatially-detailed dataset of wildfire exposure between 2000-2021 using hourly satellite imagery linked with monthly county-level economic statistics. We find that on average an increase in burn area leads to a decrease in employment for about two and a half years. These effects vary greatly across counties. In particular, employment decreases only in counties with a relatively low level of education, a relatively high degree of income inequality, or with a relatively small number of industries. These effects do not vary systematically by average county income or population size and are not well-explained by out-migration. We find some positive spillover effects on employment in neighboring counties, particularly over longer horizons, which are economically but not statistically significant. The state of the economy matters: we find that the negative effect of wildfires on employment is much larger in periods of slack *i.e.,* when a county's unemployment rate is high compared to its average unemployment rate.

JEL classification: R11, Q54, Q58, E24, H84.

Keywords: wildfires, business cycles, employment, spatial analysis.

*Department of Economics, Middlebury College. E-mail: rgauvincoulombe@middlebury.edu

†Department of Economics, Middlebury College. E-mail: akhilr@middlebury.edu.

1 Introduction

Technological progress has advanced human capacity to withstand and recover from natural disasters. Yet as climate change progresses, these disasters are becoming increasingly frequent and severe (Westerling, Hidalgo, Cayan, and Swetnam (2006); Jolly, Cochrane, Freeborn, Holden, Brown, Williamson, and Bowman (2015); Coronese, Lamperti, Keller, Chiaromonte, and Roventini (2019); Bhatia, Vecchi, Knutson, Murakami, Kossin, and Dixon (2019); Kossin, Knapp, Olander, and Velden (2020); Vecchi, Landsea, Zhang, Villarini, and Knutson (2021)). Wildfires, in particular, are becoming increasingly common in the United States. These disasters can spread rapidly and inflict concentrated damages in counties and smaller communities. How have wildfires tended to affect economic activity in US counties? Are they short, sharp shocks with quick recoveries? Or slower, more drawn-out affairs?

In this paper we analyze the effects of wildfires on employment in the US using a novel spatially-detailed monthly dataset of county-level fire exposure and employment over 2000-2021. Figure 1 shows changes in total burned area in US counties over this period. Though fires are generally concentrated in the West and South, the long-run trend is for more counties in these regions to experience more significant fire exposures.

Recent analyses of the effects of wildfires on labor market outcomes in the US have found mixed effects across time and space. Nielsen-Pincus, Moseley, and Gebert (2013) examine the effect of large wildfires (based on USDA Forest Service classifications of suppression spending) on wage growth and volatility in the US West, finding that suppression spending can have positive effects on wage growth and employment while also increasing their volatility. More recently, Tran and Wilson (2022) study the local economic impact of natural disasters over 1980-2017 using Federal Emergency Management Agency (FEMA) administrative data on disaster declarations and assistance. They find a temporary employment boost and increase in hours though negative effects and spatial spillovers in the longer-run at the state level. Borgschulte, Molitor, and Zou (2022) focus on the effects of wildfire smoke on annual earnings, labor force participation, and Social Security claims. They find that smoke exposure has negative effects on earnings and participation and increases Social Security claims. Their analysis sheds light on an important biological channel through which wildfires may impact economic activity. We contribute to this literature by (1) constructing a novel geophysical measure of fire exposure that facilitates analysis of all fire events (regardless of suppression or aid funding involved) in the entire US over 2000-2021, (2) identifying monthly economic effects of a marginal increase in fire area, and (3) identifying communities that are especially

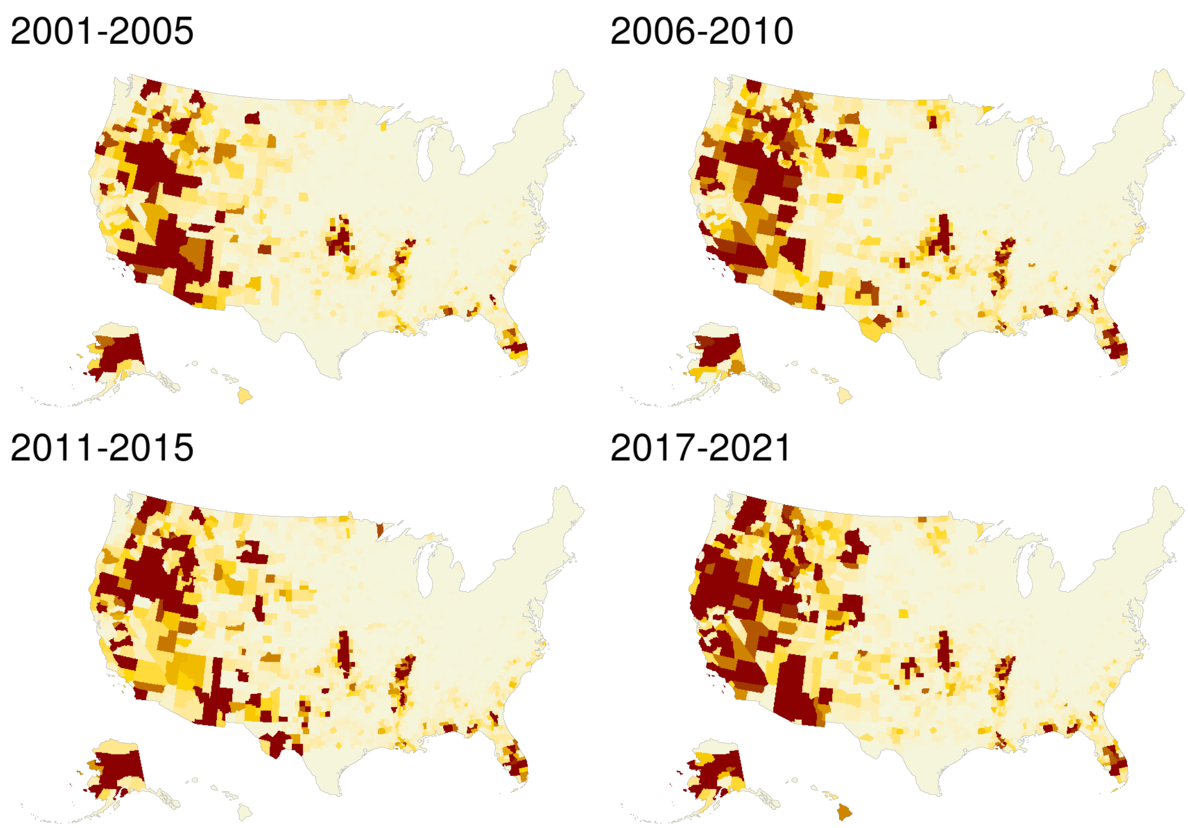

Figure 1: Total burned area in US counties over time. Darker colors indicate larger burned areas within counties. *Data source:* NASA Earthdata LPDAAC MCD64A1 product, (Giglio et al., 2018).

vulnerable to the economic effects of wildfires.

There are two key challenges in studying monthly changes in employment growth induced by wildfires. First, though they may have indirect effects at the national level through trade and migration linkages, the direct effects of wildfires may be confined to relatively small regions of a large country like the US. This limits the extent to which their impacts can be measured using aggregate economic statistics at the national or even regional level. Second, employment is often studied using quarterly data, but wildfires tend to be short and sharp disruptions which may last mere days or weeks. While some disasters can have long-lasting effects (Long and Siu (2018); Boustan, Kahn, Rhode, and Yanguas (2020)), the brevity of wildfires themselves limits the extent to which their effects on employment can be characterized using quarterly data.

To overcome these challenges, we combine satellite imagery from the National Aeronautics and Space Administration (NASA) of the universe of wildfires affecting the continental US since November of 2001, compiled and hosted by the National Oceanic and Atmospheric Administration (NOAA) and US Geological Survey (USGS) with monthly unemployment rate data from the Bureau of Labor Statistics (BLS). Following the literature on real-time economic activity measurement, we use the employment rate—a leading indicator of economic activity in US counties (Sahm, 2019; Dupraz, Nakamura, and Steinsson, 2019)—as our dependent variable. Our final dataset contains county-month observations with county characteristics, the employment rate, and an estimate of the area burned in wildfires in that county-month. Our dataset also allows us to parse out the heterogeneous employment effects of wildfires by county characteristics like racial and age composition, median income, and degree of urban development. We use local projections (Jordà, 2005) to estimate impulse responses of employment to fire impulses. One key advantage of this method is its flexibility in allowing for state dependence where the dynamic response of the unemployment rate can vary across space and time and according to different scenarios.

Our analysis contributes to the growing literature on wildfires and economic activity, as well as the growing literature on the economic impacts of natural disasters, in three ways. First, we focus specifically on wildfires using a geophysical measure of fire exposure. Both official disaster declarations and dollar value measures may be endogenous to economic conditions, an issue which satellite measurements of burn area avoid.[1] To our knowledge this analysis is the first to both use a geophysical measure of fire exposure as well as cover the entire United States for the 2000-2021 period. Our use of geophysical measures of disaster activity mirrors Strobl (2011) and Felbermayr and Gröschl (2014), albeit for wildfires instead of hurricanes or other disaster types. Our estimates are thus both highly informative about the distribution of impacts over space and time, and can be used to simulate future impacts under different fire conditions or infer historical effects of fire exposure on labor market outcomes. We conduct such an exercise to determine the historical variation in regional employment attributable to fire exposure over the sample period.

Second, while there is much evidence regarding the long-term growth effects of natural disasters broadly (*e.g.,* Noy (2009); Cavallo, Galiani, Noy, and Pantano (2013); Lackner (2018); Boustan et al. (2020)), there is less evidence regarding county-level impacts of wildfires (or disasters more generally) at the monthly frequency. Borgschulte et al. (2022) provides the

[1]For example, dollar-denominated damages may be systematically higher in counties with greater economic capacity to recover from a fire event, which would bias the estimated impulse response towards projecting faster recoveries.

closest match to our study in these respects. They conduct a county-level analysis using high-frequency geophysical measures of wind flows and smoke exposure, albeit at a quarterly frequency. Our analysis complements theirs, as we focus on fires rather than smoke, and jobs (*i.e.,* extensive margin) rather than hours worked (*i.e.,* intensive margin). Our results are consistent with their study (*i.e.,* fire exposure tends to reduce employment in affected counties for nearly three years after the event), and we show that the extensive margin effects can be substantial—a monthly fire impulse in the 90th percentile can cost on the order of 10% of the monthly average employment growth rate for a county in the US West.

Third, we provide novel evidence on the heterogeneous impacts of wildfires across county-level characteristics such as average education levels, industrial concentration, and income inequality, as well as across slack states of the local labor market. We document economically and statistically significant differences in the effects of wildfires on labor market outcomes across counties with different characteristics, as well as increased sensitivity to wildfires during high-slack periods. These results can help policymakers better target relief efforts over space and time to minimize the overall cost of wildfires.

2 Data

Fire Exposure Measurements We obtain fire exposure data from the NASA Earthdata LPDAAC MCD64A1 product, hosted by USGS (Giglio, 2015). The dataset divides the Earth into a set of tiles which are further split into 500 m^2 grid cells. Each cell contains information about burn status when sampled, the date of the detected fire, and burn measurement uncertainty. The underlying data comes from the MODIS satellite products, hosted by NOAA, and contains hourly 1 km^2-pixel measurements of fire activity captured by the MODIS satellite (NOAA, NOAA). The LPDAAC MCD64A1 data product identifies burn status within 500 m^2 grid cells by using information on cell characteristics such as burn-sensitive vegetation and reflectance (Giglio et al., 2018). We obtain the latitude-longitude coordinates of each grid cell using the inverse mapping described in the LPDAAC MCD64A1 user guide (Giglio, Boschetti, Roy, Hoffmann, Humber, and Hall, 2020), then link each coordinate to a county FIPS code using the FCC's Census Block Conversions API (FCC, 2022), and finally aggregate from hourly to monthly frequency. The dataset covers the period from November of 2000 to December of 2021.

Table 1 shows some summary statistics of burn areas for the US as a whole and within

Census regions over the 2000-2021 period. Most counties do not experience fires in most months, but those which do experience an average monthly burn area of roughly 26,100 km^2. This varies widely across regions: counties in the West which experience fires observe roughly 48,900 km^2 of burn area per month on average, while counties in the Northeast which have fires see burn areas of roughly 4,300 km^2 per month.

Table 1: Summary statistics for county-level burn areas in thousands of square kilometers per month over 2000-2021

	Mean	Mean | burn	95th percentile	Obs
US	1.9	26.1	2	824,589
Midwest	0.9	21.5	0	268,419
Northeast	0	4.3	0	55,120
South	1.4	15.6	2.5	362,719
West	7.3	48.9	12	115,217

The first column shows the unconditional mean, the second column shows the mean for counties which have fires, and the third column shows the 95th percentile across all counties. The final column shows the unconditional sample size for each region.

Labor Force Statistics We obtain monthly data on the number of people employed at the county level from the Quarterly Census of Employment and Wages (QCEW) program available from the Bureau of Labor Statistics (BLS). The number of people employed in a county is constructed from employer reports. Workers are therefore included in the county where they work—not where they live. Monthly employment data is available since January 1990 and includes both part-time and full-time workers as well as workers who are on paid vacations. We use the X-12 algorithm for seasonal adjustment as the number of employed workers shows strong seasonal patterns. While we use the full labor force statistics dataset for seasonal adjustment, we restrict our sample of labor force statistics to the period after November of 2000 to match our fire data.

Table 2 shows some summary statistics of employment growth for the US as a whole and within Census regions over the 2000-2021 period. The mean employment growth rates nationally and across regions are close to zero, but counties in the South and West with fire exposure appear to have higher and more widely varied employment growth. Counties in the Midwest and Northeast which did and did not experience fires over the sample period experienced employment declines on roughly the same order of magnitude, while in the South and West there are small differences. The differences in employment growth between counties with and without fire exposure in the US as a whole, the South, and the West are statistically significant at the 5% level.

Table 2: Summary statistics for county-level monthly employment growth in percentage change units over 2000-2021

	Mean	Mean | burn	5th percentile	Median	95th percentile	Obs
US	-1e-05	3e-04	-0.0173	4e-04	0.0174	824,589
Midwest	-2e-04	-7e-04	-0.0175	1e-04	0.0171	268,419
Northeast	-2e-04	-8e-04	-0.0097	1e-04	0.0101	55,120
South	4e-05	3e-04	-0.0163	6e-04	0.0163	362,719
West	4e-04	1e-03	-0.0249	9e-04	0.0256	115,217

The first column shows the unconditional mean, the second column shows the mean for counties which have fires, and the third, fourth, and fifth columns show the 5th, 50th, and 95th percentiles across all counties. The final column shows the unconditional sample size for each region.

County-to-County Migration Data on county-to-county migration comes from the IRS Tax Stats Database and is computed from year-to-year address changes as reported on individual income tax returns filed with IRS. Migration data at the county level is only available at an annual frequency. Our sample covers the 2003-2019 period. Since migration data is constructed from tax returns, this dataset does not include individuals who do not file a tax return and those who file late.

Table 3 shows some summary statistics of percentage changes in net out-migration (out-migration − in-migration) over the 2003-2019 period. The mean changes in net out-migration tend to be small and fairly similar for counties with and without fire exposure, though some counties do experience larger declines or increases in net out-migration. The differences in out-migration between counties with and without fire exposure in this sample are not statistically significant at the 5% level for any region or the US as a whole.

3 Empirical Framework

Our estimation strategy follows Jordà (2005)'s local projection method, modified for panel data. Let t index months and c index counties. Let $y_{c,t+h}$ denote the natural logarithm of a labor market outcome of interest, such as the number of people employed or the labor force participation rate, and $D_{c,t}$ denote county-level wildfire exposure. Let $X'_{c,t}$ denote a vector

Table 3: Summary statistics for percentage changes in county-level annual out-migration over 2003-2019

	Mean	Mean | burn	5th percentile	Median	95th percentile	Obs
US	0.002	0.005	-0.325	0.002	0.328	53,348
Midwest	-0.002	0.004	-0.366	-0.002	0.361	17,901
Northeast	0.004	0.032	-0.295	-0.005	0.325	3,689
South	0.005	0.005	-0.291	0.007	0.297	24,162
West	0.002	0.004	-0.341	0.003	0.337	7,596

The first column shows the unconditional mean, the second column shows the mean for counties which have fires, and the third, fourth, and fifth columns show the 5th, 50th, and 95th percentiles across all counties. The final column shows the unconditional sample size for each region.

of control variables. We estimate the following equation for a series of horizons $h \geq 0$:

$$y_{c,t+h} - y_{c,t-1} = \alpha_{c,h} + \beta_h D_{c,t} + X'_{c,t}\gamma_h + \epsilon_{c,t+h}. \tag{1}$$

The dependent variable represents the percent change in the labor market outcome of interest at horizon h relative to before the fire occurs. The county-specific intercept, $\alpha_{c,h}$, may reflect differences in labor force attachment across counties or other time-invariant county characteristics. The parameter β_h measures the response of the labor market outcome at horizon h to an increase in wildfire exposure as measured by the number of fire pixels detected in month t. γ_h measures the impact of the exogenous variables at horizon h. In the benchmark model, we control for the county's fire exposure in the previous twelve months ($\sum_{s=1}^{12} D_{c,t-s}$) as fire exposure tends to be highly persistent and may have an effect on the labor market over several months. We also include twelve lags of the county's labor market outcome ($\sum_{s=1}^{12} y_{c,t-s}$) to account for serial correlation in labor market outcomes. To account for factors that are common across counties that may have an impact on the labor market, such as the financial crisis or the COVID-19 pandemic, we include month fixed effects in the baseline model.

3.1 Identification

We rely on variation in fire exposure across counties and over time to identify the causal effect of marginal wildfire exposure at month t on a county's labor market outcome in month $t+h$, *i.e.*, β_h. There are two key identifying assumptions. First, conditional on the control variables, fire exposure at time t must be exogenous to county-level labor market outcomes at times $t, ..., t+h$. This is a "no anticipation" assumption. Since wildfires are generally unpredictable events—though the probability of a fire at any location-time pair can be calculated, the occurrence of a fire itself is random—and changing jobs or leaving the workforce is often a costly action, it seems unlikely that people adjust their employment or workforce status in anticipation of fire events over short horizons. The reverse causality channel—where people anticipating certain labor market outcomes deliberately create or prevent fires—similarly seems unlikely.[2]

Second, conditional on the control variables, counties exposed to wildfires should have similar labor market trends as counties not exposed to wildfires. This is a "parallel trends" assumption. Though counties are unique in many dimensions, our control variables and fixed effects likely hold relevant differences constant across counties. In particular, lagged labor market outcomes account for heterogeneous labor market trends, while lagged fire exposures account for heterogeneous fire propensity (*e.g.,* availability of fuel following recent fires, fire seasonality). Controlling for these lags also addresses a backwards-looking Stable Unit Treatment Value Assumption (SUTVA) condition that counties experiencing fires at month t are unaffected by fires in the same county at months prior to t conditional on controls.

However, we still require a spatial SUTVA condition: that counties other than c are unaffected by fire exposure in county c, conditional on controls. This condition is unlikely to hold even with temporal lags of fire exposure and labor market outcomes, as migration between counties in response to fire activity may induce labor market changes in counties not exposed to fires (at t or at all). We conduct two robustness checks to measure the degree to which such effects may bias our results. First, using IRS tax returns data, we use local projections to measure the impulse response of county-level out-migration to wildfire activity. We discuss these results further in section 4.1. Second, to measure potential spatial spillovers more granularly while remaining agnostic to the source of such spillovers, we augment our baseline specification in equation 1 with spatial lags (Halleck Vega and Elhorst, 2015). We discuss these results further in section 4.2.

While our fire exposure measure allows us to study the effects of wildfires in unprecedented granularity, measurement error is a concern. The data-generating process is informative about the type of measurement error we might expect. The MODIS satellite captures images of the Earth's surface, which are then fed into a contextual classification algorithm to determine which pixels contain fires (Giglio et al., 2020). Though the algorithm con-

[2]Employment at local fire departments may respond to labor market outcomes or fire activity, *e.g.,* workers who are marginally attached to the labor force may join a volunteer fire department to assist in firefighting activities during a wildfire event. Such actions are part of the effect we are studying.

tains processing steps to remove false positives due to factors like sun glints and water reflections, it may still miss small or low-intensity fires (incorrectly labeling such pixels as "unburned"). Such small fires may not be representative of wildfire activity and so their exclusion may be appropriate. A related source of measurement error in our fire exposure measure comes from the discreteness of fire detection. The use of discrete measurement levels (*i.e.,* burned/unburned at the 500 m^2) level induces further measurement error when burned areas are not divisible by 500, *e.g.,* a "true" burn area of 2750 m^2 may at best be recorded as 2500 m^2 or 3000 m^2. Such measurement error is likely to attenuate our estimates toward zero.

Finally, our estimation framework has similarities with difference-in-differences (DiD) frameworks, in that we consider the causal impact of wildfires on US counties and identify these effects through comparison to counties experiencing less or no wildfire exposure. One may thus wonder about potential bias arising from previously treated units (*i.e.,* counties which have experienced wildfires) being used as controls. The use of such controls has recently been shown to create "negative weights" in traditional two-way fixed-effects (TWFE) estimators, particularly when treatment effects change over time and vary across treated units (De Chaisemartin and d'Haultfoeuille, 2020; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021). Though we do not rely on a TWFE estimator, our treatment-control comparisons are similar to those implied by TWFE—*i.e.,* we are implementing a local projections difference-in-difference (LP-DiD) estimator with a continuous treatment variable (Dube, Girardi, Jorda, and Taylor, 2022).

To see if negative weighting may present issues we run our analysis restricted the sample to treated counties and "clean controls" (Cengiz, Dube, Lindner, and Zipperer, 2019; Dube et al., 2022). To account for the fact that counties may experience wildfires multiple times over our sample period, we construct a sample with a clean control group using only counties which have not experienced a wildfire for the past 3 years. Our results comparing counties experiencing fires to only clean controls closely match our results from the full sample (Appendix B, Figure B.1).

4 Economic Effects of Wildfires

We first investigate the effect of an increase in fire activity on counties' employment. We estimate the model based on expression 1 with the percent change in the total number of people employed as the dependent variable. The impulse response is computed from the series of estimated $\{\beta_h\}_0^{36}$. Figure 4 shows the impulse response of employment following a

60.91 km^2 impulse (the size of July 2021's average fire impulse across counties experiencing fires) of additional fire exposure with 95% confidence intervals. Among counties which experienced wildfires in Julys of 2000-2021, the mean burn area was about 44.7 km^2. Even among these counties however fire exposure is highly unequal: the median exposure was 8 km^2 but the maximum was 7,505 km^2. July 2021 represents an 89th-percentile event in terms of monthly fire exposures among counties experiencing fires—more severe than what most counties experience, but considerably less severe than the most that some counties experience.

Figure 2: Response of employment to an increase in fire exposure

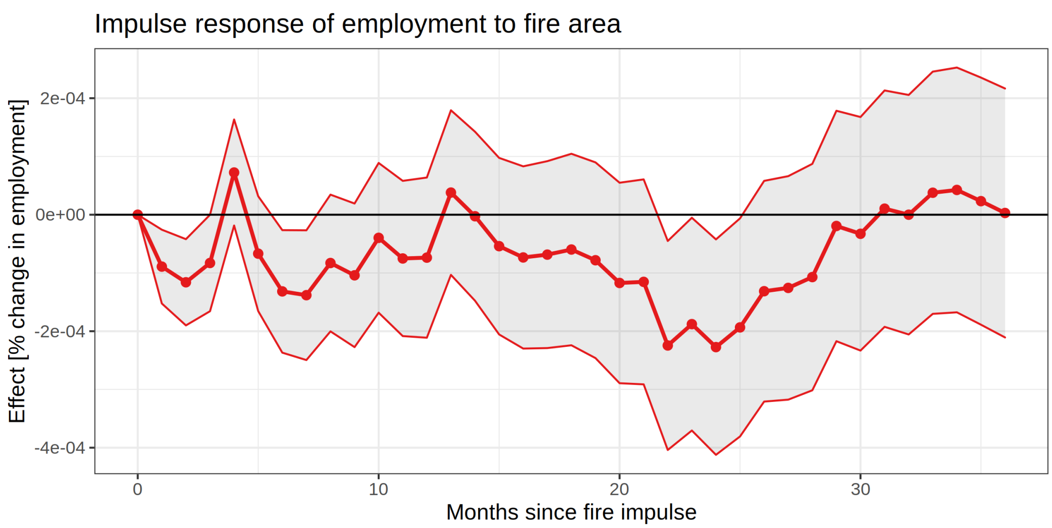

The y axis shows percentage changes in response to a burn area impulse of about 61 km^2—the mean burn area in counties that experienced fires in July of 2021. Shaded areas indicate 95% CIs computed assuming IID errors conditional on covariates, which include county and month-year fixed effects and 12 monthly lags of county employment and burn area.

The initial decline in employment in the months following an increase in fire activity is about 0.01%, growing to about 0.02% two years after the fire occurs and fading a little around the one year mark. By three years the response is consistently noisy and centered on zero. To put effect magnitude in context, Table 2 shows that the average monthly employment growth rate in Western counties which experience fires—the region with the highest growth rate in this period—is around 1e-03×100%=0.1%. One month after a fire, a decline of 1e-04×100%=0.01% (the impact of the average July 2021 fire impulse one month later) implies a loss of around 10% of the monthly employment growth.

4.1 Net Out-migration

The observed decrease in employment following a fire exposure impulse might reflect job destruction or net out-migration. To assess the degree to which net out-migration may drive our results in Figure 4, we estimate equation 1 with the percent change in net out-migration as the dependent variable. Figure 4.1 shows the impulse response of out-migration to an increase in fire burn area with H set at 10 (*i.e.,* up to 10 years since the fire impulse) since county-to-county migration data are only available at the annual frequency. The figure shows that it takes approximately 3 years for fire exposure to have a statistically detectable effect on net out-migration. The increase in out-migration after two years might explains the greater decrease in employment around that time.

Figure 3: Response of out-migration to an increase in fire exposure

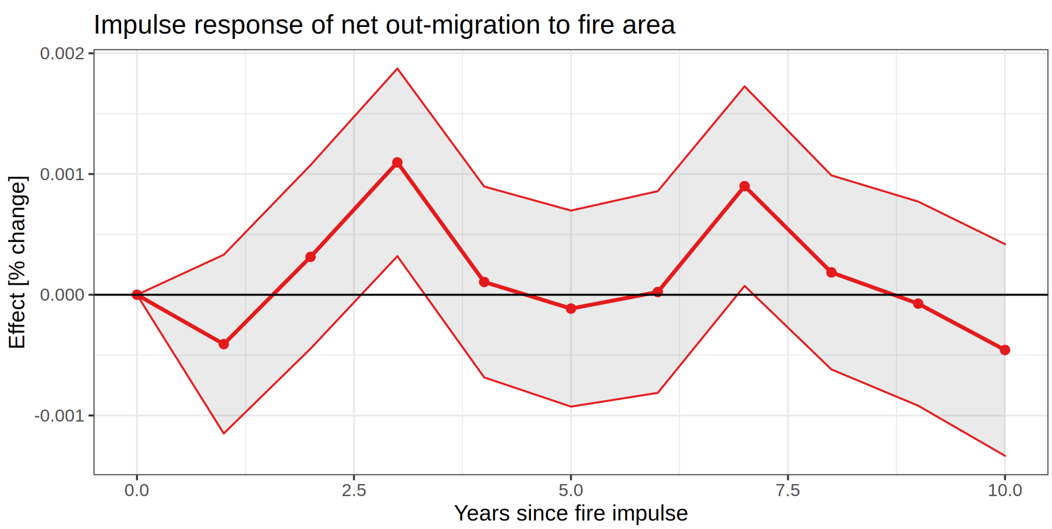

The y axis shows percentage changes in response to a burn area impulse of about 61 km^2—the mean burn area in counties that experienced fires in July of 2021. Shaded areas indicate 95% CIs.

Our results indicate no significant out-migration effects of wildfires until 3 years after the fire impulse, though the data are only available at the annual frequency and do not measure out-migration by individuals who did not file tax returns. The peak in out-migration at 3 years after the fire impulse is consistent with the trough in the impulse response shown in Figure 4 at about 2 years after the fire impulse, given that tax returns are a lagging indicator of migration.

4.2 Spatial Spillovers

As Figure 1 shows, fire exposure is spatially correlated. In addition, there may be employment spillover effects as people and jobs move across county lines in response to fire damages and suppression or recovery efforts. We model these spillover effects by augmenting equation 1 with variables for fire exposure in counties neighboring c and counties neighboring c's neighbors (*i.e.,* c's first two spatial lags) and their temporal lags. These variables are the product of a spatial weight matrix W with county fire exposures $D_{c,t}$. We use county adjacency relationships to construct W and its lags. A typical element in W is 1 if two counties share a border and 0 otherwise, so that $WD_{c,t}$ measures fire exposure in counties neighboring c. Our specification becomes

$$y_{c,t+h} - y_{c,t-1} = \alpha_{c,h} + \beta_{0,h} D_{c,t} + W D_{c,t} \beta_{1,h} + W^2 D_{c,t} \beta_{2,h} + X'_{c,t} \gamma_h + \epsilon_{c,t+h}, \tag{2}$$

where W^2 is the lagged weight matrix reflecting second-degree adjacencies (elements are 1 if two counties share a neighbor and 0 otherwise, with the diagonal normalized to 0).[3] $\beta_{0,h}$ is the effect of fires in c on employment growth in c (the "own effect"), $\beta_{1,h}$ is the effect of fires in c's neighbors on employment growth in c (the "first-degree neighbor effect"), and $\beta_{2,h}$ is the effect of fires in c's neighbors' neighbors on employment growth in c (the "second-degree neighbor effect"). The set of controls $X_{c,t}$ in equation 2 includes twelve lags of fires in first- and second-degree neighbors to control for spatiotemporal correlations in fire activity. The impulse response coefficients $\beta_{k,h}$ from the spatially-lagged regression are more likely to satisfy the necessary spatial SUTVA condition.

With spatial lags included, $\beta_{0,h}$ now represents the effect of fire impulses only in county c holding fire exposure in c's first- and second-degree neighbors constant. Figure 4 shows the own effect with and without controlling for fires in neighboring counties. Comparison reveals that spillovers from fires in neighboring counties appear to slightly offset the direct losses of fires in a given county.

Figure 5 shows the effects of fires in c on monthly employment growth c's first- and second-degree neighbors. Though not statistically significant at the 5% level over a three-year post-fire horizon and noisier over time, the effect is economically significant in first-degree

[3] Formally,

$$W^2 = W \times W.$$

The diagonal may contain values larger than 1 due to loops from c to its neighbors and back to c, so is removed to leave only the off-diagonal elements reflecting second-degree adjacencies.

Figure 4: Response of employment to an increase in fire exposure with and without spatial lags

Impulse responses of fires in a county on its employment with and without controls for fires in neighboring counties. The y axis shows percentage changes in response to a burn area impulse of about 61 km^2—the mean burn area in counties that experienced fires in July of 2021. Shaded areas indicate 95% CIs.

neighbors and growing. Three years later, the fire impulse in c induces employment growth in c's first-degree neighbors on the order of 1e-4$\times$100%=0.01%, or roughly half the largest decline caused in c (-0.02%, roughly two years after the impulse) and about the magnitude of the immediate decline a month after the impulse (-0.01%, lasting about three months after the impulse). The positive first-degree neighbor effect appears to persist, with the point estimate growing over the full three-year horizon. The second-degree neighbor effects appear to be more precisely estimated "zeros": though they trend negative over time, they are usually not statistically significant at the 5% level.

Overall, the immediate change in the own-effect relative to the baseline model is negligible and the effect appears to dissipate over space. In the longer run counties neighboring those experiencing fires may see some employment growth. The effect at three years after the impulse in c's neighbors will offset the effects of a new impulse in c one month later, but will only partially offset the medium-run ($\approx$2 year) effects of a new impulse in c.

Figure 5: Response of employment to an increase in fire exposure in first- and second-degree neighbors

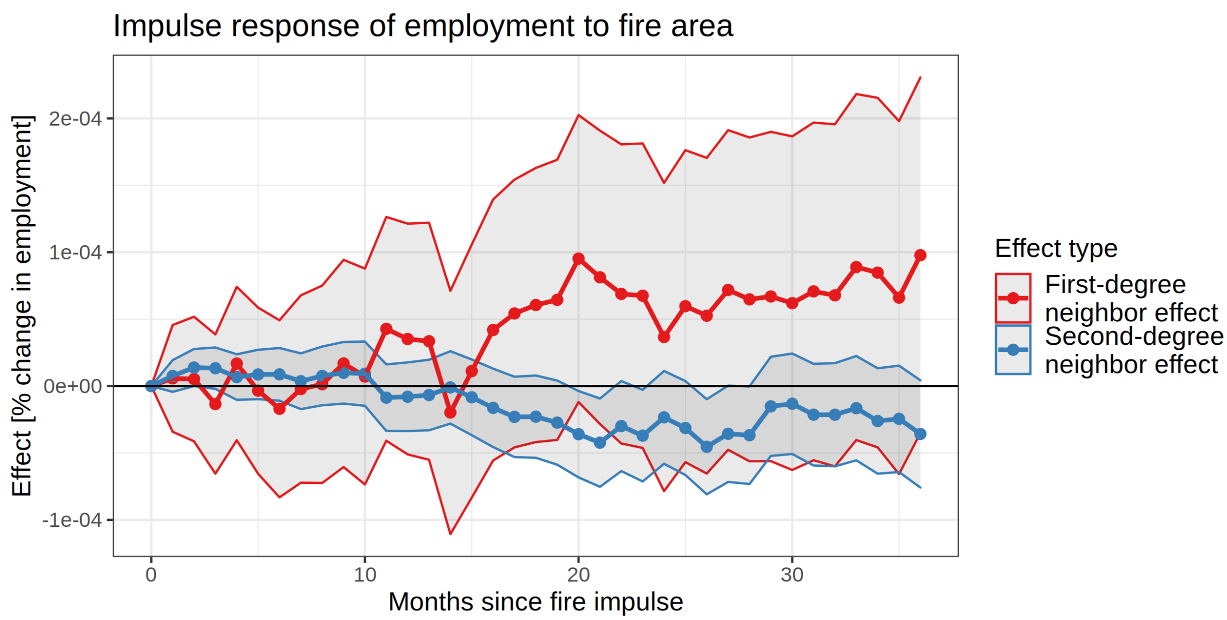

Impulse responses of fires in a county on employment in its neighbors (first-degree) and neighbors' neighbors (second-degree). The y axis shows percentage changes in response to a burn area impulse of about 61 km^2—the mean burn area in counties that experienced fires in July of 2021. Shaded areas indicate 95% CIs.

4.3 County Characteristics

Some communities might be especially vulnerable to economic impacts of wildfires. The large number of US counties gives us substantial spatial variations allowing us to analyze how key county characteristics, such as the level of education or the degree of market concentration, matter for the economic effects of wildfires.

4.3.1 Education

We first investigate how the level of education of a county's population might affect the response of employment to wildfires. Less-educated workers are known to have a lower labor force attachment, as shown by a higher unemployment rate and lower employment-to-population ratio, and might be more likely to work in certain industries that are more affected by fires, such as the tourism or agricultural industry. More-educated workers might be more likely to be insured against fires and might have a greater capacity to adapt (*e.g.,* due to higher savings or mobility). For each county, we calculate the average percentage of the population with a high school diploma since 2003 using data from the Census Bureau. We then estimate our model in equation 1 for counties that are above and below the median level of education. Figure 6 shows the impulse response of employment for counties that are below the median level of education (right panel) and above the median level of education (left panel).

Interestingly, only least educated counties see a decrease in employment following wildfires. The effect lasts for about two years after the shock. For counties with more educated workers, there is a positive effect of wildfires on employment after one year, which could be explained by reconstruction efforts. While least educated counties tend to be in rural areas and have smaller population, we do not find any difference in economic impacts when we split counties based on population size or the share of population living in urban/rural areas.

4.3.2 Industrial Concentration

As a second cut at the data, we split our counties based on their degree of industrial concentration. We measure industrial concentration using the Herfindahl index, which captures the distribution of employment across a set of industries:

$$H_c = \sum_{s=1}^{S_c} \left(\frac{e_{s,c}}{e_c} \right)^2 \tag{3}$$

Figure 6: Response of employment to an increase in fire exposure split by county education level

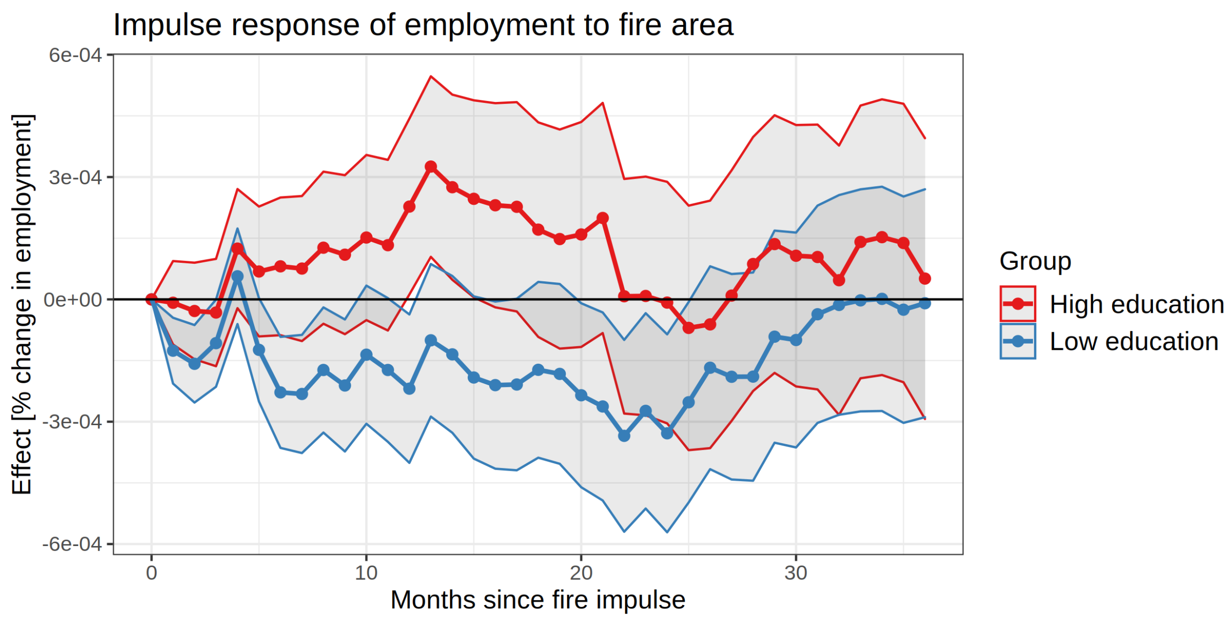

The level of education of a county is calculated from the percentage of the population without a high school diploma. The y axis shows percentage changes in response to a burn area impulse of about 61 km^2—the mean burn area in counties that experienced fires in July of 2021. Shaded areas indicate 95% CIs.

where S_c is the total number of industries in county c, $e_{s,c}$ is employment in industry s in county c, and e_c is the total employment in county c. High Herfindahl indexes can be interpreted as high industrial concentration such that a relatively small number of industries employs a large proportion of the population. In contrast, a county where employment is dispersed across a large number of industries would have a low Herfindahl index. We use 2-digit NAICS codes as our industry classification. Data on the share of the population employed in each industry is from the American Community Survey.

Figure 7 shows that counties that rely on a greater number of industries do not experience a decrease in employment after wildfires. Counties that are less economically diverse are more vulnerable to the economic effects of wildfires. This finding is consistent with the literature in other settings showing that more-diversified economies are more robust to economic shocks (*e.g.,* Kluge (2018); Coulson, McCoy, and McDonough (2020)).

4.3.3 Income Inequality

Next, we divide our sample of counties based on their degree of income inequality, as measured from the Gini index. Figure 8 shows that counties that have a relatively high level of

Figure 7: Response of employment to an increase in fire exposure split by industrial concentration level

The level of industrial concentration of a county is calculated from the Herfindahl-Hirschman index. The y axis shows percentage changes in response to a burn area impulse of about 61 km^2—the mean burn area in counties that experienced fires in July of 2021. Shaded areas indicate 95% CIs.

income inequality, *i.e.,* high Gini index, are negatively affected by wildfires both on impact and after three years. The effect is statistically significant at the 5% level. Workers at the bottom of the income distribution might be less attached to where they live in counties with high income inequality, which could explain the large decline in employment. It is worth noting that the group of counties with high income inequality is very different than the counties with a lower level of education. Table 4 in Appendix A shows that regions with high income inequality tend to have large populations and to be located in urban areas.

Overall, these results show that counties characteristics matter a great deal. Communities especially vulnerable to the economic effects of wildfires tend to have lower levels of education, be more specialized, and be more unequal. Of course, counties might differ across other characteristics. We don't find any evidence that the economic effects of wildfires depend on counties' size, their income, or the share of population living in urban vs rural areas.

Figure 8: Response of employment to an increase in fire exposure split by income inequality

The level of income inequality of a county is calculated from the Gini index. The y axis shows percentage changes in response to a burn area impulse of about 61 km^2—the mean burn area in counties that experienced fires in July of 2021. Shaded areas indicate 95% CIs.

4.4 The State of the Economy

The effect of wildfires on employment might depend on the state of the economy when the fire starts. For example, during a recession more businesses might be on the verge of bankruptcy, which could exacerbate the economic impacts of wildfires. Following the approach taken in Auerbach and Gorodnichenko (2012), Owyang, Ramey, and Zubairy (2013), and Ramey and Zubairy (2018) in the context of fiscal multipliers, we modify the baseline model to allow for the effect of marginal fire exposures to depend on the state of the business cycle of the region before the fire occurs. We use county-specific unemployment rate data to measure the level of slack of a region. Let $I_{c,t}$ denote an indicator variable that equals one if the unemployment rate in county c at month t is greater than the 70^{th} percentile of county c's unemployment rate and 0 otherwise.[4] We estimate the following model

$$\begin{aligned} y_{c,t+h} - y_{c,t-1} =& I_{c,t} \left[\alpha^H_{c,h} + \beta^H_h D_{c,t} + X'_{c,t}\gamma^H_h\right] \\ &+ (1 - I_{c,t}) \left[\alpha^L_{c,h} + \beta^L_h D_{c,t} + X'_{c,t}\gamma^L_h\right] + \epsilon_{c,t+h} \end{aligned} \tag{4}$$

[4]We find very similar results when we vary this threshold.

where H denotes the high slack state and L the low slack state. All coefficient estimates are allowed to vary across states. Figure 9 presents the results based on expression 4. The effect of fires on employment in a low slack state, *i.e.*, when the county's unemployment rate is below the 70th percentile before the fire, is depicted in red, while blue denotes the effect during a high slack state. The state of the economy appears to matter a great deal. During periods of high unemployment, fires lead to a much more pronounced decline in employment in the two years following the shock. In periods of low slack states, however, the effect is not statistically significant.

Figure 9: Response of employment to an increase in fire exposure split by level of slack

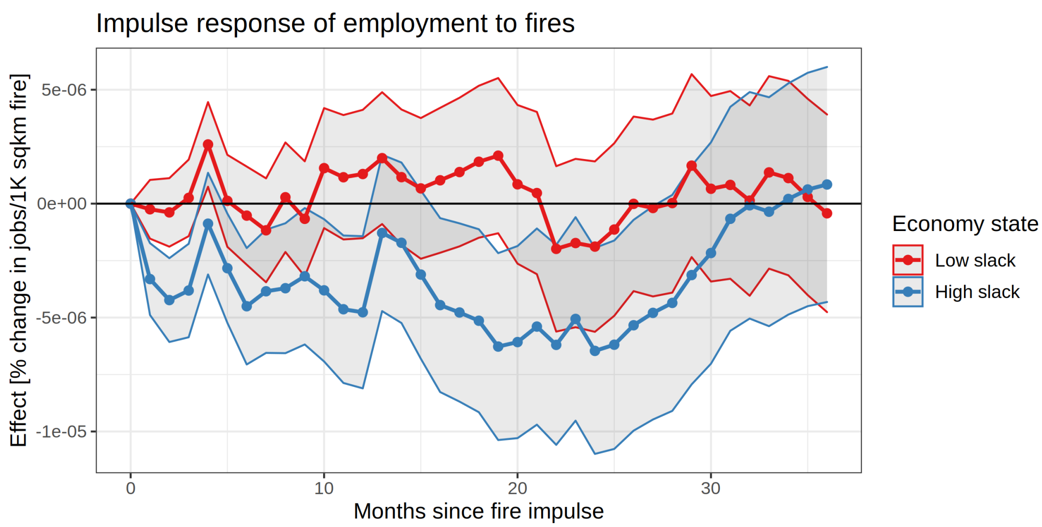

The level of slack before the fire occurs is based on the county's unemployment rate. We consider a high slack state if the unemployment rate in county c is greater than the 70^{th} percentile of county c's unemployment rate. The y axis shows percentage changes in response to a burn area impulse of about 61 km^2—the mean burn area in counties that experienced fires in July of 2021. Shaded areas indicate 95% CIs.

4.5 Cumulative Employment Impacts

As noted earlier, one advantage of using geophysical measures of fire exposure is the ability to simulate the implied cumulative employment impacts historically and for potential future fire scenarios. To illustrate this we simulate the historical cumulative employment impact of fires in each US Census Region. The cumulative employment impact of a burn sequence $(a_{c0}, \ldots, a_{ch})$ is the summation of present and lingering employment effects of fires which have occurred in county c over the previous 36 months. The cumulative employment impact,

$CE_c(h)$, is shown in equation 5.

$$CE_c(h) = \sum_{j=0}^{h} \beta_{h-j} a_{cj}, \quad \beta_k = 0 \ \ \forall k \notin (0, 36] \tag{5}$$

We compute the cumulative employment impacts for each county, then sum across counties within Census regions. The resulting cumulative fire-driven regional employment fluctuations are shown in Figure 10. The cumulative employment impacts reveal the total cost of recurring fire exposure. As might be expected, the US West region is most severely impacted on average, followed by the South, the Midwest, and finally the Northeast. The magnitude of the effect in all regions except the Northeast is substantial: an ongoing reduction in employment of roughly 0.03 (in the Midwest) to 0.15 (in the West) percent compared to the no-fire counterfactual. This is a large effect relative to the monthly regional employment growth over this period, which ranges from roughly 0.3 percent in the South to 0.34 percent in the Midwest, and is net of fire impacts. Recurring fire exposure in the West is thus reducing regional employment growth by roughly 7.5 times the observed monthly growth rate in the 95th percentile, while in the Midwest the monthly employment growth losses are "only" roughly 3 times the observed monthly growth rate in the 95th percentile (see Table 2).

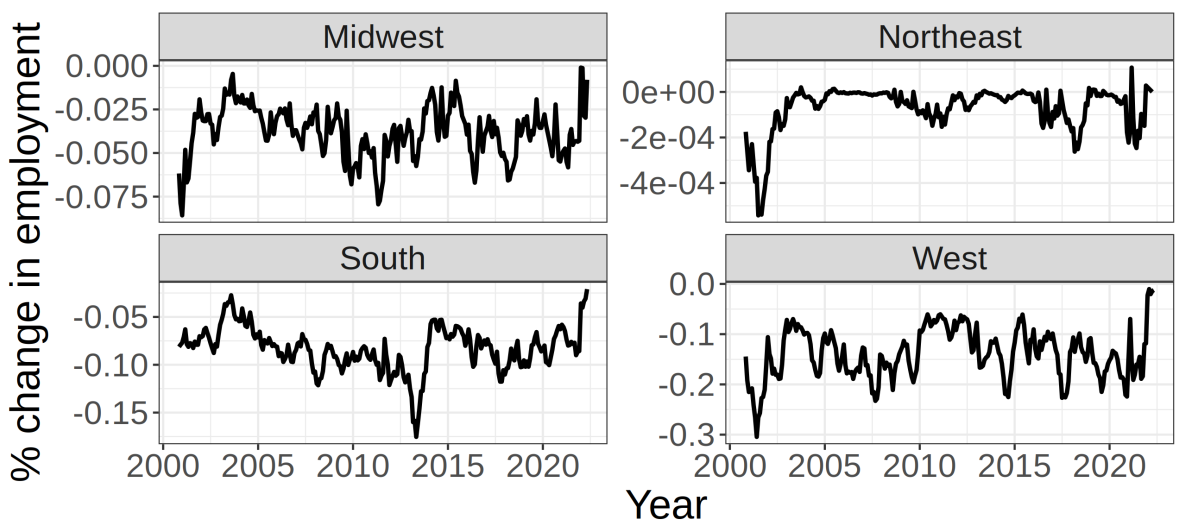

Figure 10: Cumulative fire-driven regional employment fluctuations 2000–2021

This technique can be extended to include spatial spillover effects by augmenting equation 5 with spatial impulse response coefficients (*e.g.,* $\beta_{1,h}$ and $\beta_{2,h}$ from equation 2) and burn sequences from neighboring counties. It can also be used to forecast fire-driven employment losses in a set of counties given a projected burn area sequence $(\hat{a}_{c0}, \ldots, \hat{a}_{ch})$.

5 Conclusion

In this paper we measure the dynamic monthly impact of wildfires on county-level employment growth using a novel geophysical measure of fire exposure at the county-month level for the entire US. We show that employment decreases immediately for around a quarter after a fire impulse, and decreases again beginning around a year after the impulse and bottoming out roughly two years later. The magnitude of the effect can be significant: a large (roughly 90th percentile) fire impulse can erase around 10% of monthly employment growth in an average county experiencing fires in the US West—the Census region with the highest average monthly employment growth.

The dynamics of this effect may be explained by both short-term (≈1-6 months after the impulse) net job destruction and medium-term (≈1-2 years after the impulse) net out-migration. In the long term (≈3 years after the impulse), some counties may experience modest positive growth effects from fires in neighboring counties, though perhaps not enough to fully offset the effects of their own fires. These may be driven by reversion of housing prices after an initial post-fire decline—*e.g.,* McCoy and Walsh (2018) find ≈3-year-long declines in housing prices in Colorado following wildfires. These patterns vary systematically by county education levels, industrial concentration, and inequality, with less educated, more concentrated, and more unequal counties experiencing more significant negative impacts. Highly educated, less concentrated, and less unequal counties may even experience net positive impacts from fires, perhaps suggesting greater economic resilience or greater receipts of recovery and relief funding. The effects overall are more markedly negative during periods of labor market slack.

Though our effects appear robust to some likely threats to identification, our results shed little light on the mechanisms through which these changes to employment occur. While the patterns in annual net out-migration are suggestive of medium-term flight from fires and short-term job destruction, they are not conclusive. Establishment-level data on jobs and vacancies or higher-frequency migration data (ideally covering individuals who do not submit tax returns as well) would help in determining when and the degree to which job destruction or out-migration are responsible for the employment effects of fires. Additionally, though our fire exposure data is very spatially detailed, we lack similarly-detailed population concentration measures. Such measures could be used to measure the per-capita effects of exposure to fire in regions smaller than counties, complementing existing analyses of the effects of smoke exposure on labor market outcomes (Borgschulte et al., 2022). Lastly, fiscal aid to states in response to fire exposure varies over time. The effects we measure are net of such funding flows; they do not isolate the effects of wildfires on employment in the absence of such spending. This packaging of effects limits the extent to which these results can predict the effects of marginal fire aid allocations. Incorporating data on federal aid flows could address this.

Our results are consistent with the growing economic literature on the dynamic short- and longer-term impacts of wildfires (Nielsen-Pincus et al., 2013; Tran and Wilson, 2022; Borgschulte et al., 2022). The effect magnitudes suggest that government spending on fire prevention, suppression, and recovery efforts could have large fiscal multipliers, particularly when targeted across space and time to counties. Our analysis identifies counties with lower education, higher industrial concentration, greater income inequality, and slack labor markets as being more likely to experience negative impacts. Fire-related fiscal support to such counties is likely to be particularly effective. The use of a geophysical fire measure facilitates using our results with fire exposure forecasts (*e.g.,* from coupled weather-fire models) to forecast regional economic activity and project how it will respond to different fire patterns. Applied to the 2000-2021 period, the results suggest that wildfires may have appreciably reduced overall employment growth in particularly fire-stricken regions, such as the West and South.

Measuring economic activity is among the oldest and most-central issues in economics, and wildfires are among the oldest types of shocks to human civilization. Climate change threatens to make wildfires both more frequent and more severe for ever-greater fractions of the population, making them an increasingly important source of economic fluctuations. Understanding how wildfires affect economic activity is an important step in adapting to a warming planet.

References

Auerbach, A. and Y. Gorodnichenko (2012). *Fiscal multipliers in recession and expansion.* University of Chicago Press.

Bhatia, K. T., G. A. Vecchi, T. R. Knutson, H. Murakami, J. Kossin, and C. E. Dixon, K. W. Whitlock (2019). Recent increases in tropical cyclone intensification rates. *Nature Communications 10*(635).

Borgschulte, M., D. Molitor, and E. Zou (2022). Air pollution and the labor market: Evidence from wildfire smoke. Technical report, National Bureau of Economic Research.

Boustan, L., M. Kahn, P. Rhode, and M. L. Yanguas (2020). The effect of natural disasters on economic activity in US counties: A century of data. *Journal of Urban Economics 118*, 1032–1057.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Cavallo, E., S. Galiani, I. Noy, and J. Pantano (2013). Catastrophic natural disasters and economic growth. *The Review of Economics and Statistics 95*(5), 1549–1561.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics 134*(3), 1405–1454.

Coronese, M., F. Lamperti, K. Keller, F. Chiaromonte, and A. Roventini (2019). Evidence for sharp increase in the economic damages of extreme natural disasters. *Proceedings of the National Academy of Sciences 116*(43), 21450–21455.

Coulson, N. E., S. J. McCoy, and I. K. McDonough (2020). Economic diversification and the resiliency hypothesis: Evidence from the impact of natural disasters on regional housing values. *Regional science and urban economics 85*, 103581.

De Chaisemartin, C. and X. d'Haultfoeuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–96.

Dube, A., D. Girardi, O. Jorda, and A. Taylor (2022). A local projections approach to difference-in-differences event studies. *NBER Working Paper*.

Dupraz, S., E. Nakamura, and J. Steinsson (2019). A plucking model of business cycles. Technical report, National Bureau of Economic Research.

FCC (2022). Census block conversions api. https://www.fcc.gov/general/census-block-conversions-api-v100.Version1.0.0.

Felbermayr, G. and J. Gröschl (2014). Naturally negative: the growth effects of natural disasters. *Journal of Development Economics 111*, 92–106.

Giglio, L., C. J. (2015). Mod14a2 modis/terra thermal anomalies/fire 8-day l3 global 1km sin grid v006. Distributed by NASA EOSDIS Land Processes DAAC, https://doi.org/10.5067/MODIS/MOD14A2.006.

Giglio, L., L. Boschetti, D. Roy, A. A. Hoffmann, M. Humber, and J. V. Hall (2020). Collection 6 modis burned area product user's guide version 1.0. *NASA EOSDIS Land Processes DAAC: Sioux Falls, SD, USA*.

Giglio, L., L. Boschetti, D. P. Roy, M. L. Humber, and C. O. Justice (2018). The collection 6 modis burned area mapping algorithm and product. *Remote sensing of environment 217*, 72–85.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

Halleck Vega, S. and J. P. Elhorst (2015). The slx model. *Journal of Regional Science 55*(3), 339–363.

Jolly, W. M., M. A. Cochrane, P. H. Freeborn, Z. A. Holden, T. J. Brown, G. J. Williamson, and D. M. Bowman (2015). Climate-induced variations in global wildfire danger from 1979 to 2013. *Nature Communications 6*(7537).

Jordà, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review 95*(1), 161–182.

Kluge, J. (2018). Sectoral diversification as insurance against economic instability. *Journal of Regional Science 58*(1), 204–223.

Kossin, J. P., K. R. Knapp, T. L. Olander, and C. S. Velden (2020). Global increase in major tropical cyclone exceedance probability over the past four decades. *Proceedings of the National Academy of Sciences 117*(22), 11975–11980.

Lackner, S. (2018). Earthquakes and economic growth. *FIW Working Paper, No. 190*.

Long, J. and H. Siu (2018). Refugees from dust and shrinking land: Tracking the dust bowl migrants. *Journal of Economic History 78*(4), 1001–1033.

McCoy, S. J. and R. P. Walsh (2018). Wildfire risk, salience & housing demand. *Journal of Environmental Economics and Management 91*, 203–228.

Nielsen-Pincus, M., C. Moseley, and K. Gebert (2013). The effects of large wildfires on employment and wage growth and volatility in the western united states. *Journal of Forestry 111*(6), 404–411.

NOAA. Hazard mapping system fire and smoke product. `https://www.ospo.noaa.gov/Products/land/hms.html#data`.

Noy, I. (2009). The macroeconomic consequences of disasters. *Journal of Development Economics 88*(2), 221–231.

Owyang, M. T., V. A. Ramey, and S. Zubairy (2013). Are government spending multipliers greater during periods of slack? evidence from twentieth-century historical data. *American Economic Review: Papers & Proceedings 103*(3), 129–134.

Ramey, V. A. and S. Zubairy (2018). Government spending multipliers in good times and in bad: Evidence from us historical data. *Journal of Political Economy 126*(2), 850–901.

Sahm, C. (2019). Direct stimulus payments to individuals. *Recession ready: Fiscal policies to stabilize the American economy*, 67–92.

Strobl, E. (2011). The economic growth impact of hurricanes: Evidence from U.S. coastal counties. *The Review of Economics and Statistics 93*(2), 575–589.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

Tran, B. R. and D. Wilson (2022). The local economic impact of natural disasters. Federal Reserve Bank of San Francisco.

Vecchi, G. A., C. Landsea, W. Zhang, G. Villarini, and T. Knutson (2021). Changes in atlantic major hurricane frequency since the late-19th century. *Nature communications 12*(1), 1–9.

Westerling, A., H. Hidalgo, D. Cayan, and T. Swetnam (2006). Warming and earlier spring increase western U.S. forest wildfire activity. *Science 313*(5789), 940–943.

Appendix A Summary statistics

Table 4 shows county-level summary statistics across the three characteristics described in Section 4.3. While these counties can be quite different in terms of their covariates, they all have significant fire exposure.

Table 4: Average outcomes of interest by county characteristics

Sample	Fire [500 m^2]	Emp-to-pop [%]	Population [people]	Income [$]	Urban [%]
Full sample	3,156,251	45.5%	92,639	46,726.7	42.5%
	(84.43)	(7.1%)	(299,025)	(98,885.7)	(32.0%)
Share w/o HS degree					
above median	1,635,809	42.2%	66,200	35,402.8	36.2%
below median	1,520,442	48.8%	119,637	58,873.5	49.2%
Industrial concentration					
above median	1,387,383	45.4%	73,947	39,261.1	39.0%
below median	1,768,868	45.6%	111,201	54,206.3	46.0%
Income inequality					
above median	1,878,826	43.6%	132,398	57,875.3	46.8%
below median	1,277,425	47.3%	54,826	35,612.73	38.2%

The table shows the total fire burn area (in 500 m^2) and average employment-to-population ratio, number of people (rounded to the nearest integer), GDP per capita, and percentage of population living in an urban area by county characteristics. Standard errors are in parentheses. Industrial concentration is measured using the Herfindahl-Hirschman Index, and income inequality is measured using the Gini index.

Appendix B Robustness checks

B.1 Clean controls

To determine whether negative weighting issues are biasing our estimates, we construct a "clean control" group from counties which have not experienced a fire in the past 36 months. We then compare the estimated IRF using only clean controls to the IRF estimated using the full sample in Figure B.1. The results are broadly similar with either type of identification strategy.

Figure 11: Response of employment to an increase in fire exposure

The y axis shows percentage changes in response to a burn area impulse of about 61 km^2—the mean burn area in counties that experienced fires in July of 2021. Shaded areas indicate 95% CIs.