Opioids and the Labor Market

Dionissi Aliprantis* Kyle Fee* Mark E. Schweitzer*
Research Department
Federal Reserve Bank of Cleveland

February 28, 2019

Abstract: This paper studies the relationship between local opioid prescription rates and labor market outcomes. We improve the joint measurement of labor market outcomes and prescription rates in the rural areas where nearly thirty percent of the US population lives. We find that increasing the local prescription rate by ten percent decreases the prime age employment rate by 0.50 percentage points for men and 0.17 percentage points for women. This effect is larger for white men with less than a BA (0.70 percentage points) and largest for minority men with less than a BA (1.01 percentage points). Geography is an obstacle to giving a causal interpretation to these results, especially since they were estimated in the midst of a large recession and recovery that generated considerable cross-sectional variation in local economic performance. We show that our results are not sensitive to most approaches to controlling for places experiencing either contemporaneous labor market shocks or persistently-weak labor market conditions. We also present evidence on reverse causality, finding that a short-term unemployment shock did not increase the share of people abusing prescription opioids. Our estimates imply that prescription opioids can account for 44 percent of the realized national decrease in men's labor force participation between 2001 and 2015.

Keywords: Opioid Prescription Rate, Labor Force Participation, Great Recession, Opioid Abuse
JEL Classification Codes: I10, J22, J28, R12

*: *dionissi.aliprantis@clev.frb.org*, *kyle.d.fee@clev.frb.org*, and *mark.schweitzer@clev.frb.org*.

Acknowledgments: We thank Anne Chen and Garrett Borawski for helping us to navigate through the many data complications of this project. We also thank Bruce Fallick, Lawrence Kessler, Roberto Pinheiro, Francisca G.-C. Richter, and Chris Ruhm for helpful comments, as well as seminar participants at the Bureau of Economic Analysis, the Cleveland Fed, Ohio State's Department of Human Sciences, North American Regional Science 2018 Meetings, and the Association of Business and Economics 2018 Fall Meeting.

1 Introduction

Opioids, including prescription pain killers, are widely recognized as the cause of a public health emergency in the United States. Drug overdose has become the leading cause of death for Americans under 50 years old (Katz (2017)), with the increase since 2010 due to opioids like heroin, OxyContin, and fentanyl.[1] Among the many dimensions of the opioid crisis, policy makers have considerable interest in understanding how the crisis is "related to the decline in labor force participation" (Yellen (2017)).

There is growing evidence that the supply of prescription opioids in an area depresses its labor force participation rate, although the full scope of the problem is uncertain. Areas with higher opioid prescription rates have lower labor force participation rates (Krueger (2017)), and quasi-random variation in local prescription rates has been found to decrease labor force participation (Harris et al. (2019), Laird and Nielsen (2016), Deiana and Giua (2018)). But existing results are not uniform: Currie et al. (2018) find that higher opioid prescriptions had a small positive effect on employment-to-population ratios for women.

This paper develops additional evidence on the degree to which the opioid prescription rate in an area affects its labor performance. We examine responses in terms of labor market participation, unemployment, and employment, but focus most of the analysis on employment-to-population ratios. A major obstacle to studying this relationship is simply creating a nationally representative data set that includes both detailed labor market outcomes and opioid prescriptions. Opioid prescription rates can be measured at the county level in the United States from 2006 to 2016. To be able to examine the labor supply response, this prescription data should be geographically-linked to data on labor force participation, unemployment, and employment.

Measuring geographically-specific labor market outcomes is difficult. Individual-level labor market data from the Current Population Survey or the American Community Survey do not report a respondent's county of residence when the county does not meet a minimum population threshold. Approaches to addressing this issue include focusing on high-population areas, aggregating data from low-population areas into state-level observations, or using county-level employment estimates instead of individual-level labor market data. These choices can result in a data set that is not nationally representative, that ignores information on local opioid prescription rates, or that does not report individuals' decisions to participate in the labor market.

The issue of nonidentified counties is potentially consequential, as it affects 29 percent of the US population in our sample from the American Community Survey (ACS). We measure an individual's local area as their county when identified, or else as their Public Use Microdata Area (PUMA), which is a combination of adjacent counties designed to have at least 100,000 residents. The PUMA is the most finely defined geography identified in the ACS along with annual, individual-level labor

[1] According to the National Institute on Drug Abuse (NIDA), "Opioids are a class of drugs that include the illegal drug heroin, synthetic opioids such as fentanyl, and pain relievers available legally by prescription, such as oxycodone (OxyContin), hydrocodone (Vicodin), codeine, morphine, and many others." Quinones (2016) provides a timeline of the crisis.

market outcomes. This geography is labeled the *couma* in Case and Deaton (2017).

Another obstacle to studying the relationship between opioids and the labor market is that geography creates variation in labor market outcomes that might be mistaken for the effects of opioid prescriptions. Consider, for example, that the opioid crisis has occurred simultaneously with the Great Recession and the subsequent recovery. Local labor markets were not equally impacted by the recession and recovery, making controlling for both national and local labor market demand conditions critical. To measure short-term economic shocks, we implement labor demand controls following the approach of Bartik (1991). Similar controls have been widely used in other work (recent examples include Charles et al. (2018) and Maggio and Kermani (2016)) to account for both the effects of trade and technological changes on local labor market demand conditions.

In addition, long-term differences in labor market conditions are very pronounced within the US, and these conditions could be correlated with prescription rates. To account for long-term economic conditions, we include controls for the local labor markets outcomes in the 2000 Census. While factors like these may be picked up in panel controls, the greater specificity of our controls clarifies the role of economic factors in our estimation results.

We have three main empirical results. First, we find that opioid prescription rates and labor force statuses are strongly correlated for both prime-age men and women. We find that increasing the local opioid prescription rate decreases employment rates for both prime-age men and women by lowering labor force participation rates and increasing unemployment. Specifically, increasing the local prescription rate by one-tenth of a log point (roughly 10%) decreases the employment rate of prime age men by 0.50 percent percentage points for men and 0.17 percentage points for prime age women. Examining the growth of national average prescription rates between 2000 to 2016, our estimates imply that opioid prescriptions can account for an important share of the change in labor force participation rates in recent years. For prime age men, they would imply a reduction in the labor force participation rate of 1.3 percentage points, or 44 percent of the decline from 2001 to 2015. For prime age women, our estimates would imply a reduction in the labor force participation rate of 0.5 percentage points, or 17 percent of the decline. Our estimates are larger for men, but smaller for women, than those reported in Krueger (2017).

Second, inspired by Case and Deaton (2017)'s finding of demographic heterogeneity related to mortality, we use our data set to investigate heterogeneity across demographic groups. Consistent with Case and Deaton (2017), we find that opioids have a strong effect on the labor market outcomes of white men with less than a BA. However, we find that coefficients are actually largest for minority men with less than a BA. While NSDUH data indicate lower abuse rates of minority men nationally, we reconcile this seemingly contradictory finding by showing that the exposure to opioids, in terms of local prescription rates, is considerably higher for whites than minorities. The policy implications of this finding are clear: If exposure were to spread and increase for minority men, their labor force participation would likely decline from today's levels.

Third, we also provide evidence in favor of giving our results a causal interpretation. We begin by reviewing other recent literature that tends to provide support for a causal link between opioid

prescription rates and labor market outcomes. We then investigate two concerns against a causal interpretation of our coefficients: reverse causality from economic outcomes to prescriptions and a common cause of local economic outcomes and prescription rates. Looking at data from the National Survey of Drug Use and Health (NSDUH) in the period directly before and after the Great Recession, we find that a short-term unemployment shock does not lead to a greater share of people abusing opioids. To investigate whether those areas impacted by opioids are really just areas that have also experienced short- or long-term labor market challenges, we examine a variety of specifications to control for local labor market conditions. We find that our results are not simply an artifact of the same areas being impacted both by the opioid crisis and by weak labor market conditions; areas with weak or strong labor markets react equivalently to increased exposure to legal opioids.

The remainder of the paper is organized as follows: Section 2 describes the data used in the analysis, Section 3 discusses our empirical specification and identification strategy, and Section 4 presents our results. Section 5 investigates threats to a causal interpretation of our results. Section 6 compares our results to key parts of the literature and Section 7 concludes.

2 Data

2.1 Connecting Local Opioid Prescriptions with Labor Supply

2.1.1 Individual-Level Labor Market Data

When choosing to decide what data source to use for our dependent variable, there are a few options, each with their strengths and weaknesses. Ideally, one would want to use individual-level data representing the entire country and released at a reasonable frequency to conduct this analysis. Individual-level data is preferred because we are studying an individual-level outcome, and aggregating outcomes to larger groups or areas would mean losing important information that can inform how individuals make their labor market decisions given the availability of legal opioids.

The Current Population Survey (CPS) as used in Krueger (2017) is one option. The CPS data set comprises individual-level observations but is limited with respect to the frequency at which the data are available. Krueger (2017) relies on two periods of 3-year pooled data (1999-2001 and 2014-2016) to produce labor market estimates; pooling of the data to boost the sample size limits the frequency at which one has observations. Additionally, geographically-identified CPS data primarily cover large metropolitan areas, and as a result any analysis done with the CPS data has to aggregate most rural areas into state-level remainders.

Another option is the Quarterly Workforce Indicators (QWI) as used in Currie et al. (2018). The QWI is produced annually for most counties in the nation and includes some demographic information (age, sex, educational attainment and race). However, the QWI are the county-level estimated averages of outcomes, which limits the ability to account for the way important characteristics influence individuals' decisions to participate in the labor market and work. Furthermore,

QWI employment figures reflects the location of work rather than the location of residence, which could bias labor market estimates considering that 24 percent of workers work in a county outside their county of residence (authors' calculation using the 1-year 2017 ACS). Additionally, the QWI data is somewhat noisy for counties below 100,000 residents, which causes Currie et al. (2018) to apply additional geographic aggregation to make the data more reliable.

Given the concerns around the CPS and QWI, we decided to use the Integrated Public Use Microdata Series (IPUMS-USA) of the 1% sample of the American Community Survey (ACS) from 2006 to 2016.[2] This annual data set includes detailed information for individuals' labor market status, age, race, sex, and education level, but the county of the individual observations is not always identified. About 80% of counties are not identified because they have an estimated population below 100,000, and this accounts for 29 percent of the US population during our sample period. In those cases the smallest identified geographic unit is a Public Use Microdata Area (PUMA), which by construction has a population over 100,000. Case and Deaton (2017) refer to using the lowest available geographic identifier of counties and PUMAs as coumas, and we adopt this terminology (although not all of our coumas would be identical to those used in Case and Deaton (2017)). We also require the geographic units to be consistent in the ACS over 2006 to 2016, which is challenging in some cases due to PUMA boundary changes in 2010. We use IPUMS-produced identifiers of consistent PUMAs and further aggregation when necessary to reach consistent geographic units, which we refer to as CPUMAs.[3] The scale of the issue raised by nonidentified counties is shown in Figure 1, which displays coumas while distinguishing between counties (tan) and CPUMAs (purple).

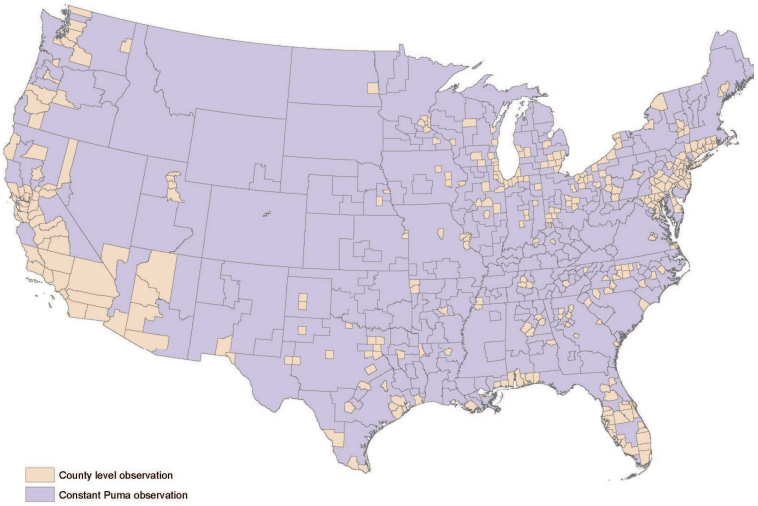

Figure 1: COUMAs
Note: Identified counties between 2006 and 2016 are shown in tan, and nonidentified counties (aggregated into CPUMAs) are shown in purple.

[2]Steven Ruggles, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. IPUMS USA: Version 8.0 [dataset]. Minneapolis, MN: IPUMS, 2018. https://doi.org/10.18128/D010.V8.0

[3]More information on the specifics of the consistent PUMA definition can be found at https://usa.ipums.org/usa/volii/cpuma0010.shtml.

One minor tradeoff we make by drawing individual-level data from the ACS is that this weakens the link to published labor force statistics that are drawn from the CPS. That said, the underlying labor market definitions are conceptually very similar, and the documented differences are mostly the result of different data collection processes (Kromer and Howard (2010)).

2.1.2 County-Level Opioid Prescription Data

We combine the individual-level IPUMS data with the Centers for Disease Control and Prevention's (CDC's) annual county-level data on prescription rates from 2006 to 2016 to measure each individual couma's prescription rate.[4] In cases where the individual's labor market county is not identified or where prescription data in a county are not available, both of which apply only to smaller counties, the individual is assigned the population-weighted average prescription rate of the observed counties within their PUMA. The CDC prescription opioid data set is derived from the records for approximately 59,000 retail (nonhospital) pharmacies, which covers nearly 90 percent of counties and nearly 90 percent of all retail prescriptions in the US. However, the annual CDC prescription data are less specific about the effective quantities of opioids prescribed than the county-level data on the total of morphine milligram equivalents (MME) prescribed (used in Krueger (2017)), which are publicly available only for the year 2015.[5]

2.1.3 Implication of Linking Prescription Rates with Aggregated Labor Market Data

We aimed to link prescription rate data at the lowest level of aggregation feasible to labor market data in order to allow economically-distinct rural areas to influence our empirical results. Figure 2 shows the variation captured when aggregating low-population counties into substate CPUMAs rather than states. The raw 2010 county-level CDC data are presented in Figure 2a. Figure 2b shows these data when nonidentified counties are aggregated into CPUMAs, and Figure 2c shows these data when nonidentified counties are aggregated into states. This alternative aggregation is important because both Currie et al. (2018) and Krueger (2017) aggregate most rural counties up to the state level.

A few states highlight the variation captured by coumas that is lost in state-level aggregation. Starting with Illinois, we can see that when nonidentified counties are aggregated into the state-level average in Figure 2c, they are all around the median prescription rate. However, the coumas in Figure 2b show that northern parts of Illinois have low prescription rates and southern parts of Illinois have high prescription rates. In Michigan, the state-level aggregation assigns a median prescription rate to northern parts of the state. Again we see a contrast with coumas, which assign higher prescription rates to some parts of northern Michigan while assigning the lowest rates to

[4] https://www.cdc.gov/drugoverdose/maps/rxrate-maps.html

[5] The correlation coefficient between a county's number of prescriptions per person and a county's MME prescribed is 0.91 in 2015. Further reassuring us about the appropriateness of county-level prescription counts, the time pattern of national MME quantities is very similar to the time pattern of our average prescription counts between 2006 and 2016 (FDA (2018)).

other parts of northern Michigan. In Indiana, Ohio, and Pennsylvania, we find similar patterns of non-trivial within-state measurement bias when state-level averages are used in place of coumas.

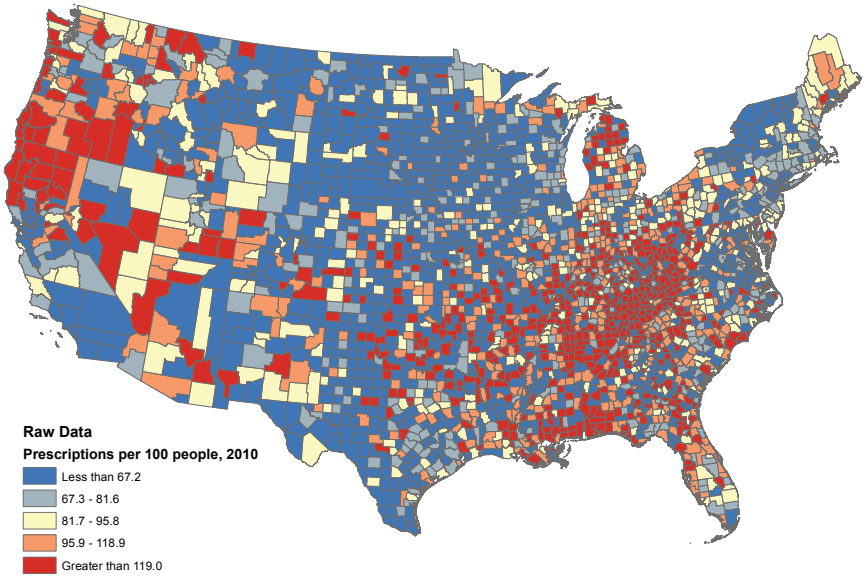

(a) County-Level Data

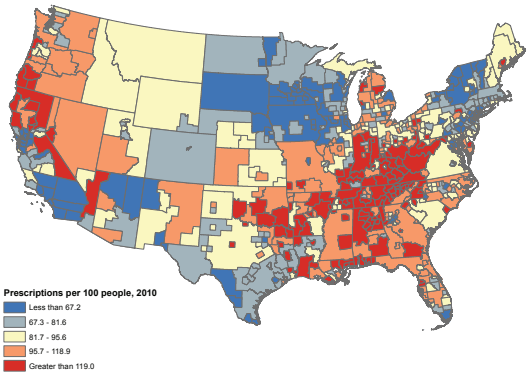

(b) Couma-Based Aggregation

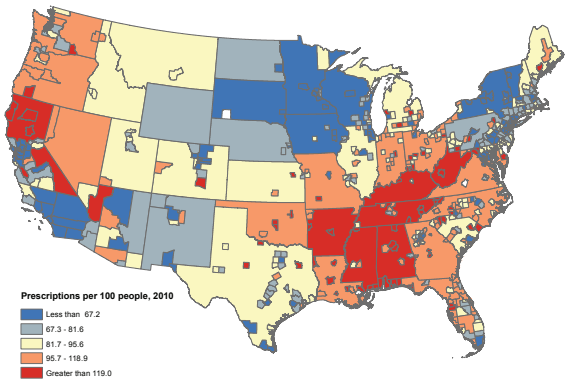

(c) State-Based Aggregation

Figure 2: Geographic Variation in Prescription Rates

A more statistical way of examining the improved measurement of local prescription rates is displayed in Figure 3. Figure 3a shows that the distribution of prescription rates is much more uniform for counties than for states, with the distribution for coumas between the two. We can also look at mismeasurement directly if we consider counties to be the correct scale of measurement. Figure 3b shows mismeasurement of the prescription rates in nonidentified counties when they are assigned, instead of their true prescription rate, the average prescription rate in the couma or

state to which they belong. Twenty percent more of the population in nonidentified counties is less than 0.4 standard deviations from the true county prescription rate when assigned the average prescription rate of their couma rather than the average prescription rate of their state. In the density estimates in Figure 3b this is evident as a considerably larger mass of nonidentified counties (as measured in terms of population) within 0.4 standard deviations when measured at the couma level and a larger mass of nonidentified counties more than 0.4 standard deviations when measured at the state level.

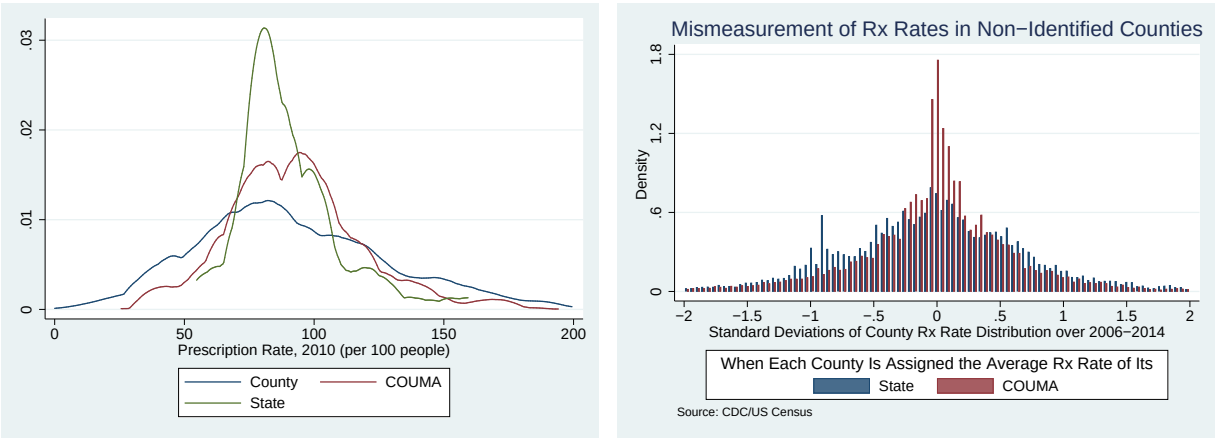

(a) Prescription Rates by Geography (b) Difference Relative to County Measure by Geography

Figure 3: Measurement of Prescription Rates at Different Scales

2.2 Measuring Local Labor Demand

A key challenge to identifying the effects of opioid availability on labor force outcomes is finding appropriate geographic controls that are able to properly account for time and geographically varying economic factors without entirely absorbing the geographic variation in the prescription data. Following Krueger (2017), we start by supplementing individual controls with the manufacturing employment share in the individual's couma and Census division fixed effects. Manufacturing share is calculated annually from the ACS data at the couma level for the prime-age population of interest.

Nonetheless, there may potentially be other reasons why local markets always have a given labor outcome beyond the controls we use. Both Krueger (2017) and Currie et al. (2018) include regressions with detailed geographic fixed effects, and we will consider these models as well. However, it is important to recognize that fixed effects at a finer geography than Census divisions will also absorb the average differences in prescription rates between those finer geographic groups. Figure 4 shows the 10th and 90th percentiles of prescription rates in coumas (Panel 1) and after subtracting the mean levels of prescriptions in each couma's state (Panel 2). This reduced amount of variation could certainly lead to an understatement of the implied effects of opioids. To explore the robustness of our results to regional controls, we examined increasing the level of controls up

to the inclusion of fixed effects for all geographic levels.

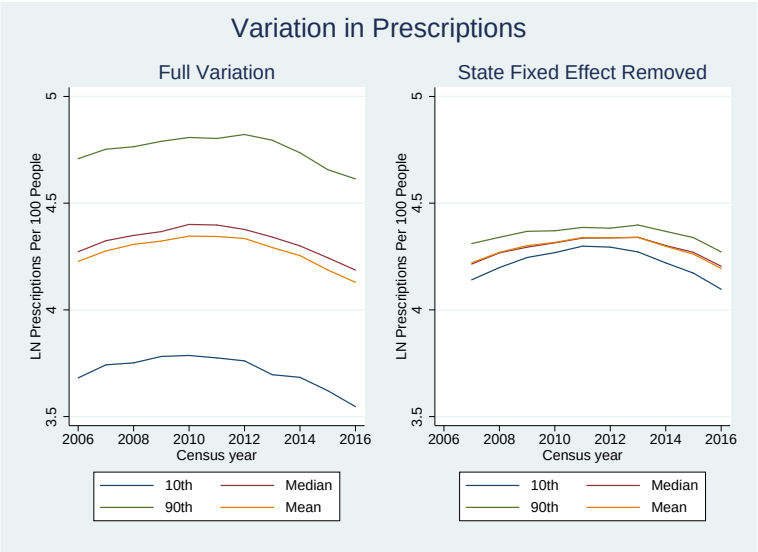

Figure 4: Identifying Variation
Note: This figure shows the variation in the natural logarithm of prescription rates across coumas.

A key contribution of our analysis is the addition of direct measures of short-term (the Bartik local labor demand shock) and long-term (Census measures of labor market conditions in 2000) economic conditions across our geography. Not surprisingly, both of these controls are typically highly important variables in our labor market status regressions.

2.2.1 Short-Term Shocks to the Local Labor Market

To assess local labor market shocks we use geographically and industrially consistent employment data on US counties derived from the US Census' County Business Pattern (CBP) data. CBP data draw on administrative records to estimate annual private non-farm employment figures for all US counties disaggregated by detailed industry following the North American Industrial Classification System (NAICS). Unfortunately, many employment figures are suppressed in the CBP because they might reveal the operations of a single employer. However, in those cases the record is flagged and employment is put into ranges. Isserman and Westervelt (2006) document a process to overcome this suppression to provide consistent point estimates, based on the provided range information and adding up conditions that are applied geographically and industrially. The W. E. Upjohn Institute for Employment Research implements this approach to produce the Wholedata Establishment and Employment Database, which is our source for county-level employment data with 3-digit NAICS coding.

County-level employment figures would still reflect both labor demand and supply conditions, which would include any effects of opioids on the local labor market. To avoid this impact we follow the well-known shift-share approach used by Bartik (1991) and Blanchard and Katz (1992) to isolate local labor demand shocks. Specifically, we follow Maggio and Kermani (2016)'s specifi-

cation to define a demand shock based on the couma's initial industrial composition of employment interacted with national-level changes in employment in narrowly defined industries:

$$B_{j,t} = \sum_{k=1}^{K} \varphi_{j,k,\tau} \frac{\nu_{-j,k,t} - \nu_{-j,k,t-1}}{\nu_{-j,k,t-1}}, \tag{1}$$

where $\varphi_{j,k,\tau}$ is the employment share of industry k in couma j in the base year, 2006, and $\nu_{-j,k,t}$ is the national employment share of industry k excluding couma j in year t. This measure accounts for both business cycle fluctuation impacts on local employment (through differing industry effects) and specific industry-wide shocks, like trade and technology changes, that are differentially experienced depending on the local industry mix.

Figure 5 shows a scatter plot by year of the Bartik local labor demand shock used in our analysis. The variation across places in any given year tends to go down over time. There was considerable variation in each year from 2006 to 2010, with the greatest variation in 2009. But from 2011 to 2015 the variation across space within a given year became quite low. One consequence of these patterns is that most of the shocks in lowest 10 percent of shocks in our sample would be drawn from recession years.

As Figure 5 shows, recessions and recoveries involve important dynamic local labor adjustments which are highly uneven. The labor market recession years of 2008 to 2010 show reliably low local labor market demand shocks, although at the onset of the recession many areas were still experiencing positive local labor demand growth. At the peak of the recession (2009 and 2010) all areas are experiencing negative local labor demand shocks, although 2009 has the largest range of demand shocks of any year. Some recovery years have relatively tight differences in local labor demand, but 2006 and 2016 stand out for the range of experiences across geographies. This variation in local labor market demand is a key reason why it is better to directly control for this important source of variation rather than just using dummy variables, which would only pick up the means for each year.

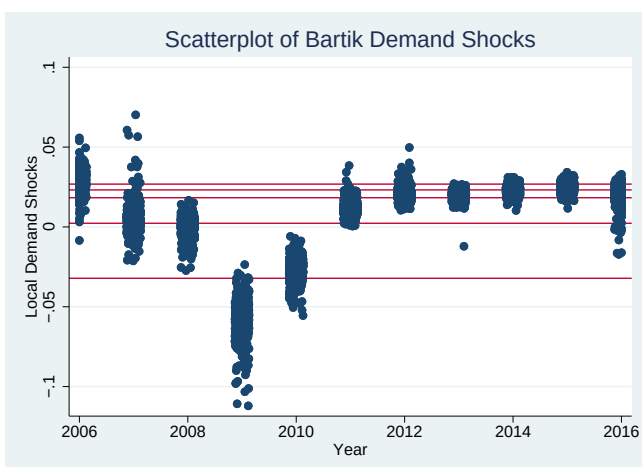

Figure 5: Labor Demand Shocks

2.2.2 Persistent Geographic Patterns in Labor Markets

Geographic regions can also have persistently different labor market outcomes that are not just a function of business cycle shocks. Indeed, there could be a geographic component to the "cumulative disadvantage" of less-educated Americans that contributes to the rising mortality and morbidity found in Case and Deaton (2017). For example, several Appalachian areas have experienced both persistently weak labor markets and high levels of opioid prescriptions and deaths. Determining the direction of causation is particularly hard in such areas. To directly account for this challenge we also include couma-level labor market statuses from the 2000 Census as an additional control.[6] While the 2000 labor market status may have been somewhat impacted by unobserved opioid prescriptions, this control should pick up the persistent labor market weakness seen in some areas of the country. Of course, this also controls for places with persistently strong labor markets which might also have unrelated, persistent patterns in the level of opioid prescriptions.

2.3 Opioid Use

The Census and CDC data let us examine the geographic patterns in prescriptions and the labor market, but they contain limited information on the mechanisms underlying these patterns. When relevant, we turn to other data sets to examine individual-level drug use. The best measurement of drug use among civilians in the United States is the National Survey on Drug Use and Health (NSDUH).[7] Since 2002 the NSDUH has gathered annual individual-level data on drug use by means of in-person interviews with a large national probability sample of the general civilian population aged 12 and older in the United States. Every year, about 70,000 people from the US civilian, noninstitutionalized population age 12 and older are interviewed. The surveys are conducted by the US Department of Health and Human Services (HHS) and use computer-assisted methods to provide respondents with a private and confidential means of responding to questions. Respondents are given $30 for participating in the NSDUH.

We measure opioid abuse as a respondent reporting nonmedical prescription pain reliever use or heroin use in the past 12 months.[8] We focus on abuse over the past 12 months rather than abuse over the past 30 days in the expectation that this longer time period should reduce measurement error, much as results in the literature on the intergenerational elasticity (IGE) in earnings show how transitory fluctuations can attenuate results when outcomes are measured over a short time horizon (Mazumder (2005)).[9] In addition to variables covering drug use in great detail, the NSDUH

[6] County-level estimates from https://www.census.gov/data/datasets/2000/dec/summary-file-3.html aggregated into our couma geography.

[7] While helpful for mechanisms, we do not use the version of these data that includes geographic identifiers. The main reason is that the geographically-identified version of this data set is designed to be representative at the state level, but not below.

[8] The NSDUH questionnaire defines nonmedical use of a prescription drug as using a drug "that was not prescribed for you or that you took only for the experience or feeling it caused." Respondents are explicitly told that nonmedical use does not refer to "over-the-counter" drugs. We refer to the entire class of pain killers (analgesics) in the NSDUH survey as opioids, even though a very small share of such pain killers are nonopioids. See footnote 11 in Carpenter et al. (2017) for a discussion.

[9] The time horizons over which the NSDUH asks respondents about drug use are ever, in the past year, and

has the additional strength for our purposes of having information on demographic characteristics and labor market outcomes such as labor force participation, employment, and hours worked.

2.4 Descriptive Statistics

Table 1 presents summary statistics for our sample broken down by whether a couma is an identified county or a CPUMA. The first detail to notice is that low-population counties, which are consolidated into CPUMAs in our analysis, have opioid prescription rates that are nearly a full standard deviation higher than high-population identified counties. Another detail is that low-population areas tended to have slightly weaker economic performance as measured by their employment-to-population ratios or their labor force participation rates. Interestingly, average high- and low-population areas' labor market outcomes were more similar in 2000 than they were between 2006 and 2016.

Table 1: Sample Summary Statistics

	All Geography		Identified Counties		CPUMAs	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Lagged Rx Rate	78	29	68	25	90	29
Lagged LOG Rx Rate	4.29	0.38	4.16	0.37	4.45	0.33
Emp/Pop	0.59	0.49	0.60	0.49	0.58	0.49
Participation	0.64	0.48	0.65	0.48	0.63	0.48
Unemp/Pop	0.051	0.219	0.053	0.224	0.047	0.212
Demand Shock	0.007	0.025	0.008	0.024	0.006	0.025
Emp/Pop 2000	0.76	0.06	0.75	0.06	0.76	0.05
Participation 2000	0.80	0.05	0.79	0.05	0.80	0.05
Unemp/Pop 2000	0.044	0.017	0.046	0.019	0.042	0.013
Manuf. Share	0.08	0.04	0.07	0.03	0.09	0.04
Age	46	19	45	19	46	19
Male	0.49	0.50	0.49	0.50	0.49	0.50
Less than HS	0.15	0.36	0.15	0.36	0.15	0.36
High School	0.36	0.48	0.34	0.47	0.40	0.49
Some College	0.23	0.42	0.23	0.42	0.23	0.42
College Grad.	0.25	0.44	0.28	0.45	0.22	0.42
White	0.75	0.43	0.70	0.46	0.82	0.38
Married	0.49	0.50	0.48	0.50	0.51	0.50

Note: The lagged prescription (Rx) rate is the number of retail opioid prescriptions per 100 residents. The sample is for the years 2006-2016 and is a combination of data from the IPUMS-USA 1% sample of the ACS and the CDC's annual county-level prescription data.

in the past month. Previous literature has shown that estimates can be sensitive to the reference period used (Carpenter et al. (2017)).

3 Empirical Specification

Our approach to measuring the labor market effects of prescription opioids is a panel version of the specification in Krueger (2017) with additional controls. The primary equation we estimate is a linear probability model on an individual i's labor force status in couma j at time t, Y_{ijt}, based on a combination of their local labor market characteristics, L_{jt}, individual characteristics observed at time t, X_{it}, time effects, δ_t, the natural log of the opioid prescription rate in the individual's couma in the prior year, P_{jt-1}, and unobserved factors, ϵ_{it}:

$$Y_{ijt} = \alpha P_{jt-1} + L'_{jt}\beta + X_i t'\gamma + \delta_t + \epsilon_{it}. \tag{2}$$

The linear probability model summarizes the individual responses into area labor market averages with demographic controls. The results are reliably away from zero and one, making a linear probability model a reasonable approach.

Our baseline approach to controlling for local labor market conditions is to include the Bartik local demand shock, labor market conditions in 2000, a set of Census division dummies, and the manufacturing employment share in the couma. We also include a rich set of individual-level demographic controls: a four-term exponential expansion of age, level of education dummy variables, race dummy variables, and marital status. We also run all regressions separately for men and women, given prior evidence of differential impacts. Finally, we include a time dummy for each year, so that the average effects of national events including but not limited to the Great Recession are absorbed.

It is helpful to see that our estimating equation is easily converted to a panel equation of the effects are prescription rates on labor market averages of interest:

$$\frac{1}{N_j}\sum_{i\in j} Y_{ijt} = \alpha P_{jt-1} + L'_{jt}\beta + \frac{1}{N_j}\sum_{i\in j} X'_{it}\gamma + \delta_t + \frac{1}{N_j}\sum_{i\in j}\epsilon_{it}. \tag{3}$$

Summing within geographies, it becomes clear that the average labor market values (employment-to-population ratio when Y is 1/0 in employment, labor force participation when Y is 1/0 in participation, and unemployment-to-population when Y is 1/0 in unemployment) for area j are estimated, as long as area average demographic effects in area j at time t are accounted for. Not accounting for these effects, which could be empirically large in some areas, would result in a potential omitted variable bias. If after accounting for the effects of the included characteristics ($\mathbf{X'_i}\gamma$) the ϵ_{it} term is mean zero and normally distributed, then geographically informative estimates are easily derived from the individual regressions or vice versa. Time controls and geographic effects above the level of area j are likely to help produce the zero means for the error term but not in all cases. Introducing fixed effects at the level j of geography would result in mean zero errors, but it would also shift the regression to measuring effects on only between-region j variation. To make sure that our standard errors are not understated given that several variables of interest vary only within j and t, we apply Stata's robust cluster variance on clusters within couma clusters (Rogers

(1993)).

3.1 Timing

While we estimate a specification very similar to Krueger (2017), our individual-level labor market and county-level prescription data allow us to estimate regressions on an annual frequency. This improves on the timing in Krueger (2017), which is based on two periods of 3-year pooled CPS data (1999-2001 and 2014-2016) and county-level data from 2015 on opioid prescription rates converted to morphine milligram equivalents (MMEs). As we are interested in the effect of opioid prescription rates on current labor market conditions, we incorporate a 1-year lag in opioid prescription rates in our regressions.

Although it is unclear what the proper lag structure should be, investigating the correlation of prescription opioid rates over time provides some guidance. Table 2 shows that prescription opioid rates are highly correlated over time (the correlation within 2-year windows is always greater than 0.96) and indicates the exact lag structure should be a minor concern. Overall, our data allow us to run panel regressions on individuals' labor force status from 2007 to 2016 with CDC data on average prescriptions per person in 648 coumas.

Table 2: Opioid Prescription Rate Correlation Structure

Time	Correlation
-2	0.970
-1	0.988
0	1.000
+1	0.986
+2	0.962

4 Estimation Results

4.1 Baseline Estimates for the Three Labor Market States

Prime-age individuals (ages 24 to 54) can be sorted into three mutually-exclusive labor market statuses: out of the labor force, employed, or unemployed. Running population-weighted linear probability models on status produces demographically adjusted estimates of the labor force participation rate, the employment-to-population ratio, and the unemployment-to-population ratio for areas, and the marginal impacts of the regressors on these rates.

Table 3 shows the results of these regressions for prime-age men and women using a specification close to the one estimated in Krueger (2017). In the case of both prime-age men and women, the number of opioid prescriptions in their geographic area in the prior year is associated with a lower probability of labor force participation and a lower employment rate with a high level of statistical significance. Consistent with Krueger (2017)'s estimates, the effects on the labor force participation

rate are substantially larger among men than among women (-0.047 versus -0.016). For the fraction of the population that is unemployed, the coefficients on the lagged opioid prescription rate are an order of magnitude smaller and not statistically significant for prime-age women, but also go in the opposite direction (more unemployed), which further boosts the impacts on the employment-to-population ratio.

For both men and women, the results for the employment-to-population ratio and labor force participation are quite similar, which implies that the primary effect that opioid prescription levels have on labor market outcomes appears to be on the individual's labor supply decision with relatively small additional effects associated with individuals searching for work. Recognizing this pattern, we focus our attention on the more economically relevant employment rate going forward although the effects are quite similar for participation.

Table 3 also includes controls for local labor demand shocks and long-term local economic conditions. The second row of Table 3 shows that the local labor market demand shock variable is a highly significant factor for prime-age mens' labor market status. Notably, mens' employment rates are substantially affected by the demand shocks in a pro-cyclical pattern: a negative labor demand shock (like a recession) reduces employment both through lower labor force participation and higher unemployment. In contrast, womens' employment rates are not substantially correlated with the Bartik-style demand shocks, because participation and unemployment rates tend to increase with a negative labor demand shock resulting in offsetting movements. In the period we are considering this could reflect the higher attachment of men to manufacturing and construction, which were nationally impacted, combined with family-level labor supply responses to the shock. While we do not observe the household-level motivations, the highly significant statistical patterns by sex, with room for time and regional variation, make local demand conditions an important source of variation to be controlled for directly, rather than assuming that the fixed effects will absorbe these sources of variation.

The third through fifth rows of Table 3 shows that long-standing local economic conditions are another key predictive variable for individuals' labor market status. As noted earlier, we measure the long-term labor market effects based on a couma's 2000 Census value of the relevant labor market statistics.[10] The positive, statistically significant coefficients indicate that on average individuals in places with a better labor market in 2000 are relatively better off during our sample period. The coefficients are all positive, but the tightness of the relationship is the lowest for the unemployment rate, even when recent labor demand shocks are accounted for. This is consistent with regional differences in unemployment rates converging relatively quickly while participation decisions are relatively persistent.

Despite the change to annual panel data and the inclusion of additional controls, the results in Table 3 are similar to those in Krueger (2017) in the sign and in the pattern of generally stronger effects for men than for women. Given Krueger (2017)'s strategy of estimating over two 3-year

[10] An advantage of using the 2000 labor market values is that they largely represent each location's relative labor market position prior to much of the growth in opioid prescriptions. Of course, while far lower than today, opioid prescriptions were nonzero in 2000, but going to the 1990 Census would stretch the comparison period substantially.

Table 3: Labor Market States of Prime-Age Men and Women with Specific Controls

	Men			Women		
	Participation	Emp/Pop	Unem/Pop	Participation	Emp/Pop	Unem/Pop
Lagged Prescrip.	-0.047***	-0.050***	0.004**	-0.016***	-0.017***	0.002
	(0.005)	(0.005)	(0.001)	(0.003)	(0.004)	(0.001)
Demand Shock	0.346**	0.636***	-0.264***	-0.309**	-0.144	-0.149*
	(0.113)	(0.130)	(0.070)	(0.109)	(0.125)	(0.061)
Particip. 2000	0.593***			0.453***		
	(0.046)			(0.029)		
Emp/Pop 2000		0.519***			0.408***	
		(0.036)			(0.025)	
Unem/Pop 2000			0.287***			0.202***
			(0.035)			(0.030)
Manuf. Share	0.031	0.029	-0.008	-0.016	-0.062	0.029*
	(0.050)	(0.055)	(0.019)	(0.038)	(0.046)	(0.014)
Age	0.070***	0.092***	-0.022**	0.081***	0.085***	-0.003
	(0.013)	(0.015)	(0.008)	(0.020)	(0.020)	(0.007)
Age squared	-0.288***	-0.357***	0.067*	-0.399***	-0.407***	0.006
	(0.048)	(0.056)	(0.033)	(0.076)	(0.077)	(0.028)
Age cubed	0.053***	0.062***	-0.009	0.083***	0.084***	-0.000
	(0.008)	(0.009)	(0.006)	(0.013)	(0.013)	(0.005)
Age to 4th	-0.004***	-0.004***	0.000	-0.006***	-0.006***	-0.000
	(0.000)	(0.001)	(0.000)	(0.001)	(0.001)	(0.000)
Less than HS	-0.193***	-0.243***	0.049***	-0.320***	-0.363***	0.043***
	(0.006)	(0.007)	(0.002)	(0.004)	(0.005)	(0.001)
High School	-0.083***	-0.120***	0.037***	-0.131***	-0.162***	0.031***
	(0.001)	(0.002)	(0.001)	(0.003)	(0.003)	(0.001)
Some College	-0.037***	-0.056***	0.019***	-0.054***	-0.072***	0.019***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)
White	0.019***	0.022***	-0.003**	0.024***	0.031***	-0.007***
	(0.002)	(0.003)	(0.001)	(0.003)	(0.004)	(0.001)
Black	-0.064***	-0.092***	0.028***	0.052***	0.027***	0.025***
	(0.003)	(0.004)	(0.002)	(0.005)	(0.006)	(0.002)
Hispanic	0.061***	0.069***	-0.009***	0.019***	0.015***	0.003**
	(0.003)	(0.004)	(0.001)	(0.003)	(0.003)	(0.001)
Married	0.115***	0.152***	-0.036***	-0.081***	-0.057***	-0.023***
	(0.002)	(0.003)	(0.000)	(0.003)	(0.004)	(0.000)
Constant	0.000	-0.248	0.306***	-0.034	-0.079	0.079
	(0.121)	(0.141)	(0.079)	(0.184)	(0.188)	(0.068)
R^2	0.09	0.11	0.02	0.06	0.06	0.02
N	5,833,789	5,833,789	5,833,789	6,022,049	6,022,049	6,022,049

All regressions include year and Census division fixed effects. Robust standard errors with clustering on coumas.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

periods, the most relevant comparison of his results to ours would be the combination of his "Log Opioids per Capita" and "Log Opioids x Period 2" coefficients. Combining these coefficients from Krueger (2017)'s Table 13, column 6 regressions, which are most similar to our regressions, his results indicate a somewhat smaller effect on participation from a log-point increase in MME of about -0.02 for prime-age men and -0.004 for prime-age women. This latter result (for women) combines a positive impact on labor force participation in the early period with a -0.014 effect of log opioids in the second period. Comparing the scale of our effects with those found in Krueger (2017) is difficult because he uses data on MME where we use data on prescriptions, but our baseline results are qualitatively consistent and we compare our results to Krueger's more extensively in section 6.

4.2 Estimates for Detailed Demographic Groups

Given the number of observations available in the ACS in each couma, it is possible to explore the effects of opioid prescription rates on more narrowly defined subsamples of the population. Given the influential results in Case and Deaton (2015) and Case and Deaton (2017), we explore effects by eight subgroups of sex, education level, and race/ethnicity. For our purposes we split the sample by gender (men and women), non-Hispanic whites (white) and minorities including Hispanics (nonwhite), as well as holding a BA versus some college or lower.

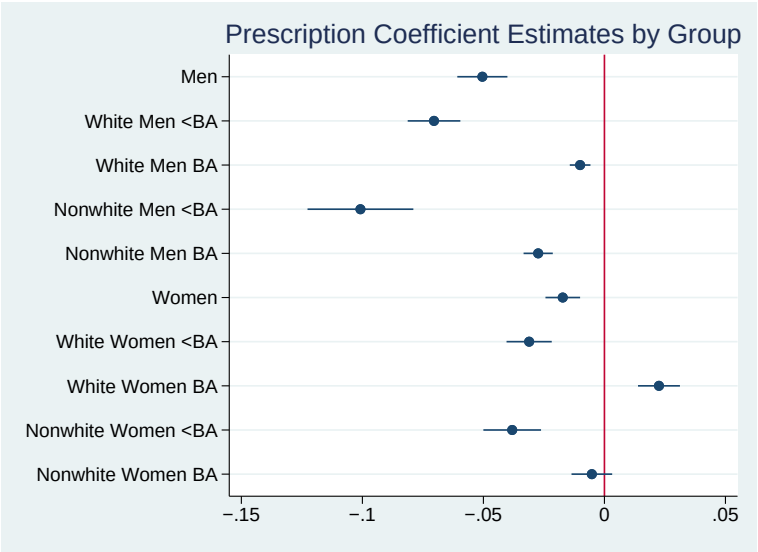

Figure 6: Effects by Demographic Groups

Figure 6 shows the primary results for prime-age men and women, along with the associated 95 percent confidence intervals. While the coefficients for men and women are clearly distinct, there are very large and statistically significant differences between other demographic groups. For men without a college degree, the effect of opioid prescription rates on the employment rate is lower for those with a college degree (see the correlations in the upper half of Figure 6).[11] The coefficient

[11] These categories follow Case and Deaton (2015). We also examined other possible splits and found that measured

for white prime-age men with less than a BA is about seven times higher than the equivalent coefficient for white prime-age men with a BA. The magnitude of the coefficient for nonwhite men without a college degree is even larger than the coefficient for white men without a BA. These results show there are quite large effects for relatively disadvantaged men along the lines suggested in Case and Deaton (2015) and Case and Deaton (2017), even if the mechanism is not identified in this exercise. It is worth emphasizing that this effect is on top of the generally lower participation rate expected for the group of less-educated white men, which is accounted for in the other controls.

Another point worth emphasizing is that while Case and Deaton (2015)'s results focused attention on white households, our results are just as troubling for nonwhite prime-age men. The coefficient for nonwhite men with less than a BA is a startling -0.101, larger than the -0.070 experienced for less-educated white men, although the difference is not quite statistically significant at the 95 percent confidence level. Reinforcing the pattern, nonwhite men with a BA also experience a larger likelihood of not being employed in higher opioid prescription areas than their white counterparts (-0.045 versus -0.021). By our measures it is hard to argue that white prime-age men have been more affected than their minority counterparts.

The lower half of Figure 6 repeats this analysis for groups of prime-age women. For white women with a BA, in contrast to other demographic groups, there is a positive and statistically significant coefficient on being in a higher opioid prescription couma, similar to Currie et al. (2018). Nonwhite women with a BA have a slightly lower and statistically insignificant expected employment rate in higher opioid prescription areas. For both white and nonwhite women without a BA, however, the coefficients are negative and statistically significant, although much smaller than the equivalent male demographic group. While most of the coefficients on lagged log prescription rates continue to be statistically significant for key demographic splits of prime-age women, the coefficients reported here are generally less than half the magnitude of the coefficients for equivalent male populations.

In our analysis, we observe the effects of being in a higher opioid prescription rate location rather than whether an individual received a prescription. This could be an advantage given the documented patterns of prescription opioids being consumed by individuals in the local area other than the individual receiving the prescription (Powell et al. (2015)). On the other hand, our specification does prevent understanding the direct impacts of prescription along the lines of Laird and Nielsen (2016) or Barnett et al. (2017), and there could easily be demographic usage differences within a geographic location. We can, however, examine both the reported usage rates by demographic group from the NSDUH and the relative exposure of demographic groups through their residences within our sample.

Figure 7 shows opioid misuse or heroin use reported in the NSDUH for prime-age (24-49) adults by demographic characteristics over time. The gap in abuse rates across race/ethnicity is larger than the gap across sex and on par with the gap in abuse rates across educational attainment. Considerably lower abuse rates among nonwhite men help to explain why Case and Deaton (2015) focus on the white population to the extent that abuse differences are associated with related

outcomes for individuals with some college were more similar to high school graduates than to college graduates.

mortality differences, a result that may seem at odds with our large estimated labor market effects on nonwhite men.

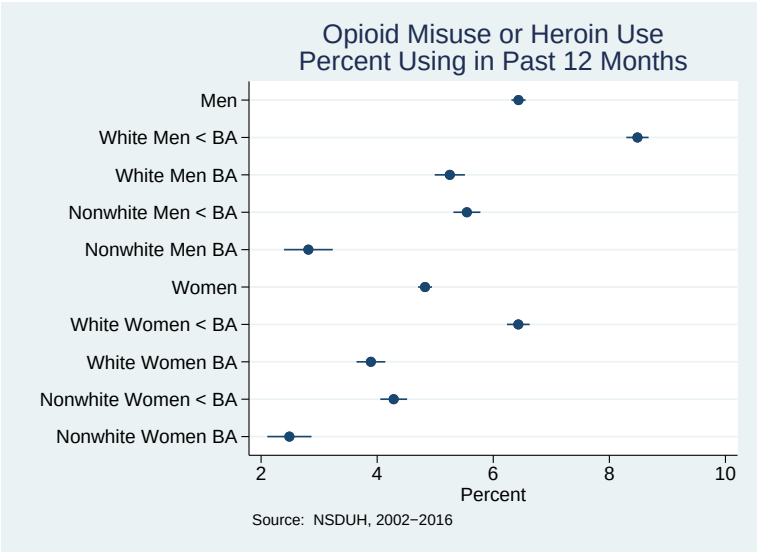

Figure 7: Abuse by Demographic Groups

Is there any way to reconcile our findings for nonwhite men? Our findings are especially puzzling since one might reasonably assume that effects of availability on employment must at some point must pass through abuse. The answer is that yes, these seemingly contradictory results can be explained through differential exposure to local prescription rates. Figure 8a shows that nonwhites are exposed to much lower local prescription rates than their white counterparts. The median white prescription rate is the 73rd percentile of the nonwhite distribution. This compares with relatively similar exposure rates by education and sex.

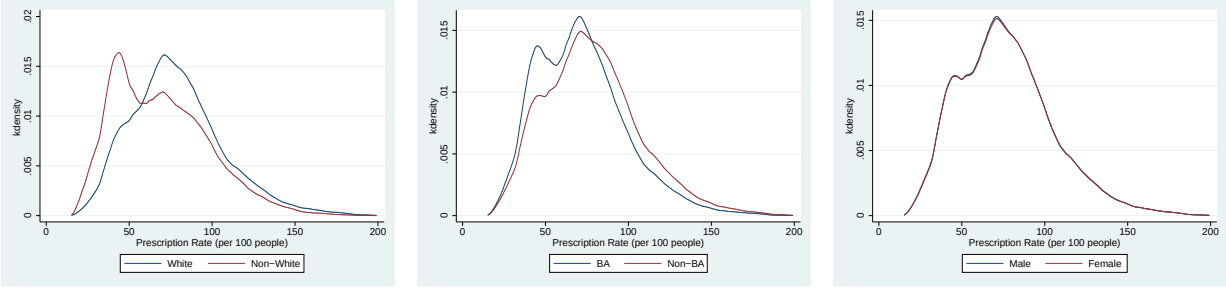

(a) White/Non-White (b) BA/no BA (c) Men/Women

Figure 8: Local Prescription Rates by Demographic Group

Interestingly, a public health review of the literature on pain treatment and race identified a substantial gap in the prescribing rates of opioids for minorities with equivalent symptoms to their white counterparts (Shavers et al. (2010)). This suggests that it could further be the case that minorities are less exposed to opioids even within geographic areas.

4.3 Geographic Identifying Variation and Fixed Effects

As discussed earlier, finding the right level of regional controls is challenging. We have generally sought to use controls for specific regional mechanisms in place of more sweeping fixed effects, but there could still be confounding variation that we are unable to account for.

To account for potential regional variation in labor market outcomes, our baseline regressions include Census division fixed effects. Our first effort to tighten the regional controls uses state fixed effects in the place of the Census divisions. Table 4 shows the baseline and these results. First, it is worth briefly noting that the coefficients on the demographic factors (not shown) are changed only negligibly. Being geographically determined, the coefficients on the demand shock variable, the 2000 employment-to-population ratio and manufacturing share in the couma do change for both men and women, but changes are generally statistically insignificant between the two specifications. Similarly, the coefficients on prescription rates are also altered by the inclusion of state fixed effects, but generally not by a statistically significant difference. For men, the magnitude of the coefficients is a little larger in magnitude, going from -0.050 in the baseline regressions to -0.056. For women, the inclusion of state fixed effects reduces the magnitude of the opioid prescription coefficient to -0.011, but the coefficient is still statistically significant at the 99% level. When we examine the more impacted, lower-education demographic groups, the estimated effects continue to be large and statistically significant. Overall, we conclude that most of the results are generally robust to the inclusion of state-level regional controls.

To complete the range of regional controls, we introduce fixed effects for each county or aggregation of counties. Krueger (2017) includes this level of controls for one regression and identifies a smaller but still significant relationship between prescription rates and participation rates. The third panel of Table 4 shows the results with a full set of couma fixed effects. Not surprisingly, the coefficients of interest are notably smaller, with the standard errors largely unchanged. Notably, all prime-age men in Table 3 had a coefficient of -0.050 but are at a statistically insignificant -0.010 with a full set of geographic fixed effects. The same pattern is true when the results are examined for workers with less than a BA. This indicates that an important amount of the identification of coefficients in Table 3 relied on between-couma differences in (mean) prescription rates. For women, the estimated results in this exercise are more comparable to the estimates without the full set of geographic fixed effects, although the estimates are less precise. At least for women, the estimates do not rely on between-geographic differences.

Once couma fixed effects are included in a specification, the remaining variation in local prescription rates comes across time. Given the very high level of correlations with both leads and lags seen in Table 2, it is not surprising that persistent patterns in prescriptions across coumas are critical to the identification of the effects. On this basis we continue to focus on the baseline results, which combine local controls for specific mechanisms and moderate regional fixed effects. We acknowledge that this approach leaves open questions about whether our estimated relationships reflect two correlated, persistent regional patterns, which are not fully accounted for by the inclusion of 2000 employment rates in the regression.

Table 4: Prime Age Employment Regressions by Geographic Fixed Effects

	Men	Women	White Men <BA	Nonwhite Men <BA	White Women <BA	Nonwhite Women <BA
Baseline: Only Census Division Fixed Effects						
Lagged Prescrip.	-0.050***	-0.017***	-0.070***	-0.101***	-0.031***	-0.038***
	(0.005)	(0.004)	(0.006)	(0.011)	(0.005)	(0.006)
Demand Shock	0.636***	-0.144	0.803***	1.366***	0.037	0.265
	(0.130)	(0.125)	(0.141)	(0.270)	(0.127)	(0.206)
Emp/Pop 2000	0.519***	0.408***	0.539***	0.727***	0.577***	0.535***
	(0.036)	(0.025)	(0.036)	(0.076)	(0.028)	(0.051)
Manuf. Share	0.029	-0.062	0.043	-0.209	-0.056	-0.322***
	(0.055)	(0.046)	(0.049)	(0.132)	(0.037)	(0.073)
State Fixed Effects						
Lagged Prescrip.	-0.056***	-0.011**	-0.073***	-0.118***	-0.023***	-0.039***
	(0.006)	(0.004)	(0.006)	(0.012)	(0.005)	(0.006)
Demand Shock	0.703***	-0.101	0.868***	1.459***	0.092	0.300
	(0.124)	(0.103)	(0.131)	(0.266)	(0.112)	(0.167)
2000 Emp/Pop	0.517***	0.349***	0.548***	0.737***	0.505***	0.505***
	(0.037)	(0.022)	(0.037)	(0.082)	(0.028)	(0.051)
Manuf. Share	0.033	0.029	0.041	-0.252	0.003	-0.181**
	(0.058)	(0.043)	(0.050)	(0.155)	(0.032)	(0.069)
Couma Fixed Effects						
Lagged Prescrip.	-0.010	-0.014**	-0.012	-0.016	-0.011	-0.014
	(0.006)	(0.005)	(0.008)	(0.011)	(0.007)	(0.010)
Demand Shock	0.262***	-0.125	0.376***	0.433*	-0.005	0.170
	(0.076)	(0.084)	(0.090)	(0.178)	(0.097)	(0.133)
Manuf. Share	0.190***	0.202***	0.185***	0.397**	0.208***	0.371**
	(0.042)	(0.043)	(0.048)	(0.120)	(0.054)	(0.124)

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.4 Specific Geographies

Informed scrutiny of the maps in Figure 2 and media accounts of the opioid crisis suggests possible geographic patterns to explore: rural counties, Appalachian counties, and perhaps the Rust Belt region. In each case, it has been be argued that persistently bad economic outcomes might underlay the opioid crisis. To investigate whether these specific, defined regions are critical to the interpretation of our results we ran regressions with interactions on the lagged prescription coefficient to each specified geography. For Appalachian counties we use those listed by the Appalachian Regional Commission.[12] The Rust Belt counties are the counties identified by Schweitzer (2017) as being in the "Industrial Heartland" in 1969. Figure 9 provides a map with the specific boundaries of these regions. The regressions maintain the controls used in prior regressions in order to highlight the particular response of the prescription coefficients to these two interactions.

[12] https://www.arc.gov/appalachian_region/countiesinappalachia.asp

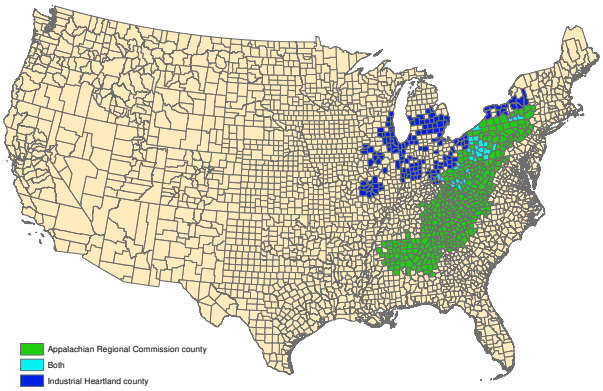

Figure 9: Appalachian Regional Commission and Rust Belt Counties

The primary coefficients on lagged prescriptions are all still statistically significant for both men and women, although in several cases a bit smaller than was seen in Tables 3 and 4. The interactions are generally smaller but in several cases statistically significant, indicating that there are some persistent patterns occurring within these geographies. The regressions indicate that low-population regions (represented by the CPUMA interaction) see a larger reduction in male employment rates but a modestly higher female employment rate than in the high-population identified counties. Appalachian coumas included in the Appalachian Regional Commission (ARC) list of counties also saw a larger employment effect for both men and women in higher prescription counties. Finally, the coefficients for Rust Belt coumas only show a larger coefficient for women. These regressions demonstrate that the correlation between prescription rates and labor market outcomes is not dependent on the inclusion of these challenged geographies.

Table 5: Employment Regressions with Geographic Interactions

	Rural		Appalachia		Rust Belt	
	Men	Women	Men	Women	Men	Women
Lagged Prescrip.	-0.045***	-0.020***	-0.048***	-0.016***	-0.051***	-0.015***
CPUMA*Prescrip	-0.003***	0.001*				
ARC*Prescrip			-0.003**	-0.001*		
RB*Prescrip					0.001	-0.004***
Demand Shock	0.523***	-0.096	0.616***	-0.155	0.616***	-0.079
Emp/Pop 2000	0.501***	0.416***	0.511***	0.403***	0.520***	0.407***
Manuf. Share	0.067	-0.079	0.046	-0.053	0.030	-0.068
R^2	0.11	0.06	0.11	0.06	0.11	0.06
N	5,833,789	6,022,049	5,833,789	6,022,049	5,833,789	6,022,049

All regressions include full set of controls with year and Census division fixed effects.
Robust standard errors with clustering on coumas.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.5 Implications of Increasing Illegal Opioid Usage after 2010

OxyContin was reformulated in 2010 to make it less prone to abuse. Evans et al. (2018) and Alpert et al. (2018) have documented that this reformulation was associated with a rise in heroin deaths that was attributed to consumers substituting to heroin as an inexpensive alternative to OxyContin. In addition, mortality data shows a sharp rise in deaths involving illicit opioids after 2010 (Ruhm (2018). If prescription opioids were less frequently abused after 2010, this could lead our estimates to understate the strength of the relationship pre-2010 and overstate the relationship post-2010.

By design all of our regressions include time dummies to absorb the large cyclical pattern in labor market states associated with the Great Recession and the recovery. These time dummies will also absorb the national average response to opioid prescriptions. That said, evidence of the changing impact of prescriptions after 2010 could still be identified in cross-sectional differences in the couma-level labor market outcomes associated with the levels of prescriptions before and after 2010 (after controlling for economic conditions). In order to assess this effect we interact the prescription rate with a dummy variable that is 1 during the period after 2010 and 0 otherwise, but constrain the time fixed effects to be unchanged from the baseline regressions. Without the constraints the time effects and interaction cannot be independently estimated, but with the constraints the interaction coefficient should show whether there is reliably more or less variation in labor market outcomes associated with opioid prescriptions after 2010.

Table 6 shows the results of these regressions for men and women compared to the baseline. For both sexes the expected effect of opioid prescriptions is essentially unaltered after 2010. This could reflect the tendency for illegal opioid problems to be connected to places that already had higher levels of prescriptions, but this question deserves considerably more analysis than we can do with our data, given the growing prominence of illegal opioids in the mortality data.

Table 6: Employment Regressions with and without Post-2010 Interaction

	Men		Women	
	Baseline	Constrained	Baseline	Constrained
Lagged Prescrip.	-0.050***	-0.050***	-0.017***	-0.017***
Post*Prescrip		-0.000		-0.000
Demand Shock	0.636***	0.651***	-0.144	-0.133***
Emp/Pop 2000	0.519***	0.520***	0.408***	0.408***
Manuf. Share	0.029	0.028	-0.062	-0.062
R-sqr	0.11		0.06	
N	5,833,789	5,833,789	6,022,049	6,022,049

All regressions include year and Census division fixed effects.

Robust standard errors with clustering on coumas.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

5 Evidence on Causality

What differentiates our analysis from those already in the literature is our population, our timeline, and our ability to investigate alternative causal mechanisms. Our approach broadens those already in the literature by generating a data set that improves measurement, covers a longer time horizon, and represents the full US population. In addition, our analysis is able to control for important alternative causes that could drive the correlation between prescription rates and labor market outcomes. That said, we first want to recognize the growing literature working on identifying a causal connection.

5.1 Evidence in the Literature

A growing number of papers have found a large, negative effect of opioid prescriptions on labor market outcomes. These studies have used quasi-random variation for US subpopulations over short time horizons. When used as a source of identifying variation, the location of pill mills (Harris et al. (2019)), moves to new physicians (Laird and Nielsen (2016)), and the timing of state-level opioid laws (Deiana and Giua (2018)) have all provided evidence that increasing local prescription rates decreases participation. Likewise, variation in prescribing behavior across care providers has provided evidence that the duration of disability following an injury increases when long-term opioid treatment is prescribed (Savych et al. (2018)).

The mechanisms creating these labor market effects have also been documented: Increasing the local opioid prescription rate generates diversions to nonmedical uses (Powell et al. (2015)), and increasing the local prescription abuse rate increases an individual's likelihood of abusing prescription opioids (Finkelstein et al. (2018)).[13] At the individual level, increasing a physician's opioid prescribing rate increases the long-term opioid use of their patients (Barnett et al. (2017)).[14]

In the opposing causal direction, both "poisoning deaths," which in recent data are dominated by drug overdoses, and the intensity of clinical disorders associated with opioid abuse have been found to be mildly countercyclical in the US (Ruhm (2015), Carpenter et al. (2017)). Similarly, Hollingsworth et al. (2017) find that short-term fluctuations in the local unemployment rate increase opioid abuse and deaths from opioids.[15] This result suggests a potential increase in drug use in recessions and after labor demand shocks.

5.2 Reverse Causality: The Impact of Unemployment Shocks on Opioid Use

During our sample period, the easiest alternative explanation for the strong correlation between opioid prescriptions and lower labor force outcomes is simple reverse causality: poor labor market

[13] In addition to the evidence related to the reformulation of OxyContin cited earlier, there is also evidence on the relationship between the abuse of legal and illegal opioids. The majority of heroin users today are typically found to have abused prescription opioids before ever using heroin (Compton et al. (2016)), a reverse of the ordering of abuse from the 1960s (Cicero et al. (2014)).

[14] Barnett et al. (2017) document wide variation in rates of opioid prescribing among physicians practicing within the same emergency department. They use this within-hospital variation as their identifying variation.

[15] Ruhm (2018) finds that medium-term economic conditions do not appear to be the main driver opioid deaths.

outcomes might drive opioid misuse. While we cannot directly assess this possibility in terms of prescription rates and labor market statuses, we can look at the relationship between labor market outcomes and opioid misuse in the NSDUH, especially just before and after the onset of the Great Recession. From the perspective of opioid research, the Great Recession is a large source of exogenous variation. Figure 10 shows that while there was a large change in labor market outcomes over the course of the Great Recession, this was not accompanied by any discernible increase in the rate of opioid misuse. If we look in terms of linear trends as in Figure 11, we might even conclude that the Great Recession caused an inflection point that decreased opioid misuse.

We interpret these time series patterns in the data to be evidence against the role of reverse causality driving our estimation results, at least in terms of short-term economic conditions. Much like Case and Deaton (2017), who found that short-term economic conditions do not seem to explain the secular increase in mortality among midlife non-Hispanic whites in the US, we suspect that any labor market impacts on the share of the population abusing opioids would need to be long term in nature.

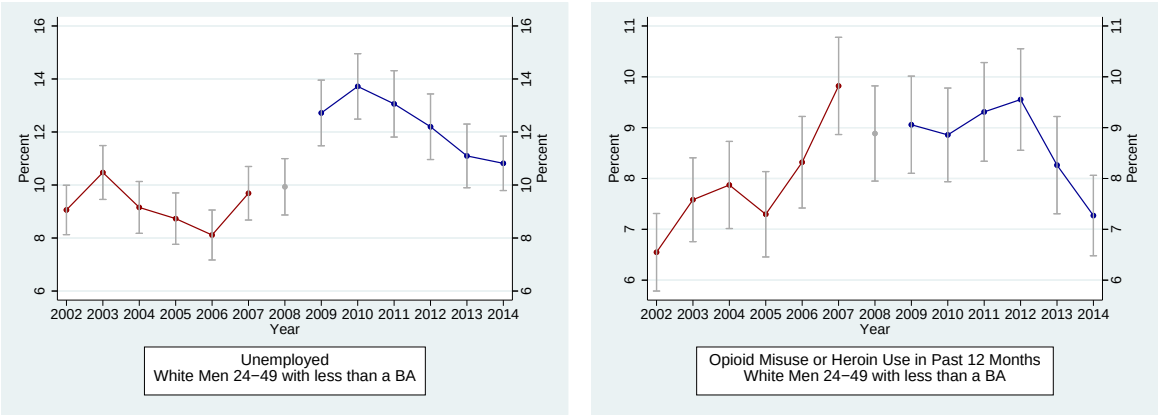

(a) Unemployed (b) Opioid Misuse or Heroin Use

Figure 10: Labor Market Outcomes and Opioid Use over Time

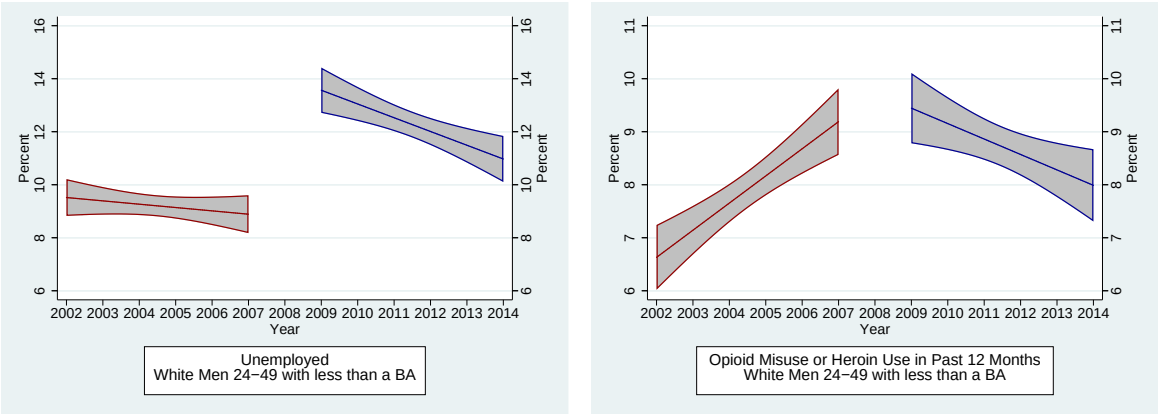

(a) Unemployed (b) Opioid Misuse or Heroin Use

Figure 11: Labor Market Outcomes and Opioid Use over Time

5.3 Local Economic Conditions and Geographic Sorting

A more subtle alternative explanation for our results is that there is a common cause of local labor market conditions and local opioid availability. This concern is especially salient because some areas with notably poor local labor market conditions tend to also have higher prescription rates (Appalachia for example) than areas with good local labor market conditions.

The Directed Acyclic Graph (DAG) in Figure 12 helps to illustrate this concern. The issue of a common cause is that the same unobserved local factors U_j that drive economic performance might also drive local prescription rates. Alternatively, individuals with "low" unobservables U_i with respect to their labor market outcomes Y_i might choose to reside in locations with poor local labor market conditions L_j and high prescription rates P_j.

Throughout our analysis we have directly controlled for long-term and short-term economic conditions (L_j) that influence labor market outcomes in the coumas. We have also partially controlled for individuals' geographic sorting by controlling for a rich set of demographic variables. Since these earlier specifications might have relied too heavily on a linearity assumption (Imbens (2015)), here we examine more flexible controls for long-term economic conditions and short-term labor market shocks. To be concrete, we allow the coefficient on opioid prescription rates to vary with the economic conditions of the place.

$U \equiv$ Unobserved Characteristics in Location j or of Individual i
$X \equiv$ Observed Characteristics of Individual i
$L \equiv$ Labor Market Conditions in Location j
$P \equiv$ Prescription Rate in Location j
$Y \equiv$ Labor Market Outcome of Individual i
$O \equiv$ Opioid Use

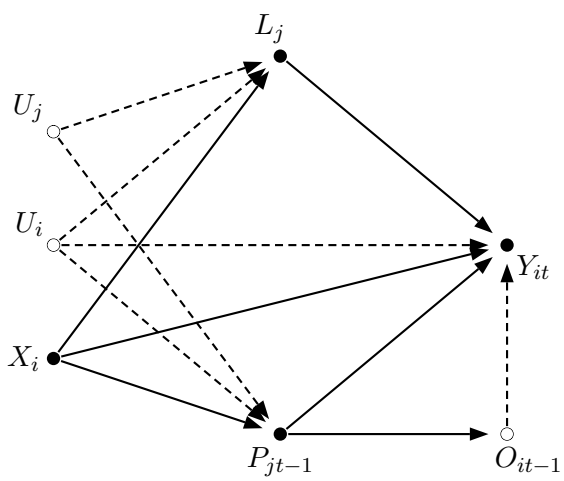

Figure 12: Directed Acyclic Graph of Opioids Affecting Labor Force Participation
Note: This figure follows the convention from Pearl (2009) of communicating that a variable is observed by drawing a solid line to its descendants and communicating that a variable is unobserved by drawing a dashed line to its descendants.

5.3.1 Initial Local Labor Market Conditions

We first consider a variation of our baseline specification in Equation 2 where we estimate the coefficient on local prescription rates separately for areas depending on where they are located in the distribution of initial labor market conditions, L_{j0}, as measured by the local employment rate in the year 2000:

$$Y_{ijt} = \alpha_1 \mathbf{1}\{0 <= L_{j0} < 10\} P_{jt-1} + \cdots + \alpha_6 \mathbf{1}\{90 < L_{j0} <= 100\} P_{jt-1} + X_i' \gamma + \delta_t + \epsilon_{it}. \quad (4)$$

If the observed correlation between prescription rates were being primarily driven by individual selection on unobservables U_i or some common local factor U_j that our previous specification did not adequately control for, then we would expect to see variation in coefficient estimates across areas based on their local labor market conditions. In contrast, if the pattern is driven by variation in prescription practices (exposure) independent of initial economic conditions then there is no reason to expect the estimated coefficients to vary.

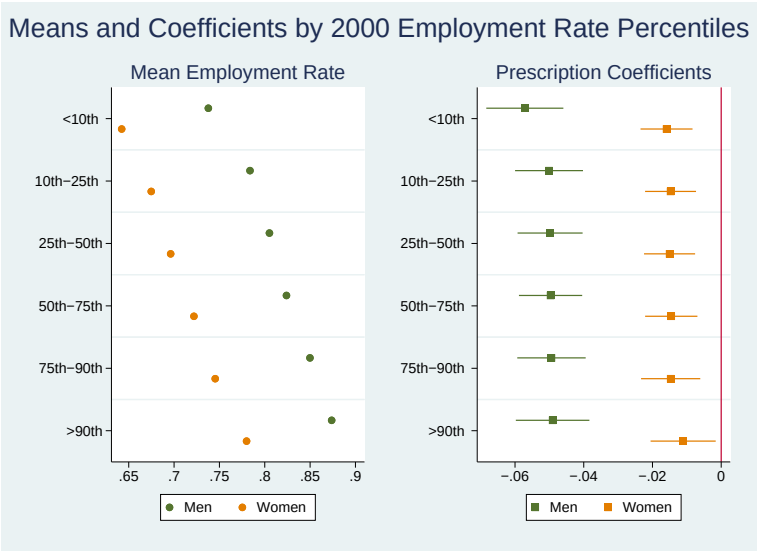

Figure 13: Coefficients by Initial Local Labor Market Conditions

The left-hand panel of Figure 13 plots the means for both men and women according to the distribution of 2000 labor market status, showing that 2000 economic conditions substantially alter the mean labor market status during the sample period. This result is neither new nor surprising given that the coefficients on 2000 status variables have been uniformly positive and significant. The right-hand panel shows the coefficient estimates when they are interacted with the locations' 2000 employment-to-population ratio. Variation in these coefficients allows us to observe subtle non-linearity in the response that could point to a deeper causal factor. For example, if individuals residing in long-term poor labor markets were also inherently more susceptible to addiction and therefore weaker labor market prospects going forward, we should anticipate that the interacted coefficient should be higher for those labor markets even after we control for the average effect of long-term market status on all labor markets.

The pattern indicating an omitted common cause is not the result that we see Figure 13: the coefficients are remarkably consistent across weak and strong (in 2000) labor markets for both men and women.

5.3.2 Shocks to Local Labor Market Conditions

We also estimate the coefficient on local prescription rates separately depending on where in the national distribution the labor demand shock of an area falls into. This would represent the

possibility that residents of places which get a particularly bad economic shock relative to the nation are more prone in those periods to opioid addiction, potentially resulting in a weaker labor market outcome for any given opioid prescription rate. One case for this possibility would be the literature that suggests recessions make people more generally prone to addiction (Carpenter et al. (2017)) and opioids in particular (Hollingsworth et al. (2017)). In terms of our sample, most of the shocks in the lower 10 percent would be drawn from recession years (see Figure 5 again).

The left-hand panel of Figure 14 shows that again worse local labor demand shocks are associated with lower employment rates for men. In addition, men's employment rates are also positively impacted by the top 25 percent of labor demand shocks. While these differences may look small when scaled to prime-age population, they are roughly equivalent to a 5 percentage point change in the conventional unemployment rate. There are relatively limited differences for female employment rates across labor demand shocks, which is consistent with the weak coefficients for women in the prior regression tables. Turning to the prescription coefficients, there are no statistically significant differences for men or women, as the results show stable coefficients across considerably different local labor markets. These results would not be expected if recessionary periods significantly increased opioid abuse and led to weaker labor market outcomes; however, this is consistent with employment effects being primarily identified by non-cyclical opioid prescription patterns.

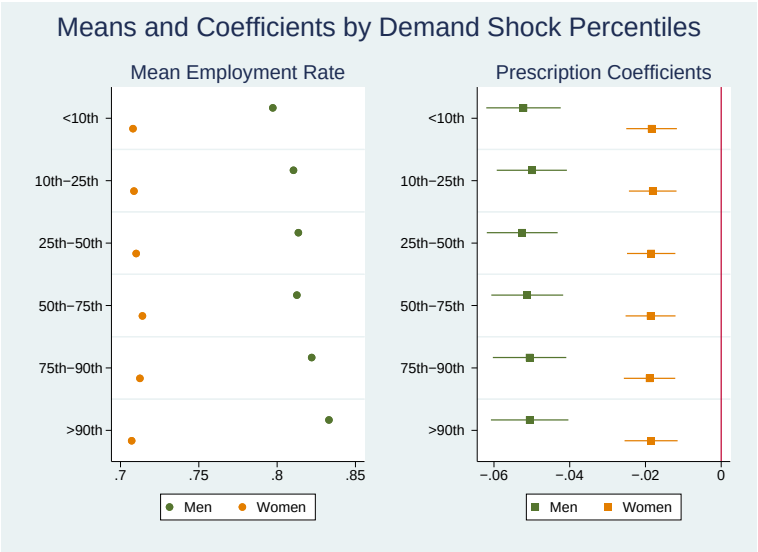

Figure 14: Coefficients by Contemporaneous Local Labor Market Conditions

6 Comparing Our Results to Key Parts of the Literature

6.1 Krueger and the Implied Scale of the Effects

By design, this analysis has used variation between areas in their opioid prescription rate and their labor market conditions to measure the impact of opioids on the labor market. Using this design, it easy to say what the difference in employment rate is expected for a high opioid prescrip-

tion area relative to a low prescription area. As a rough measure, the difference between the 10th and 90th percentiles is roughly one log point. This implies the coefficients can be directly interpreted as the expected impact of being in a high-opioid-prescription-rate area. These are clearly substantial values, particularly for men, but differences in the measures (opioid prescription rates versus morphine milligram equivalents), data sources (American Community Survey versus Current Population Survey), and structure of timing (10-year panel versus a cross-sectional comparison of two 3-year windows) makes direct comparisons of coefficients across our studies challenging.

However, from Krueger (2017) onward there has been an interest in associating the rise of the opioid crisis with declines in labor market outcomes, especially those seen for male participation. As Equation 3 implies, we should be able to apply the coefficients to the average prescription rate even if the coefficients are identified off of cross-sectional variation. During our sample period (2006-2016), the net changes in participation due to opioid prescriptions would be negligible, since nationally the prescription rate rises and then falls back.

The national rise of opioid prescriptions largely predates 2006, and much of the impact of opioid prescriptions on participation would accrue during these years. Fortunately, we were able to find a source for national opioid prescription rates from 2000 to 2006, Kenan et al. (2012), that is derived from a very similar set of pharmacy records used to generate our county-level data. Using that 2001-2006 data, we can infer the effects of prescription rates on labor market outcomes over the entire period from 2001 to 2017. Importantly, the change in national prescription rates is well within the variation seen between low- and high-opioid-prescription-rate areas within our sample.

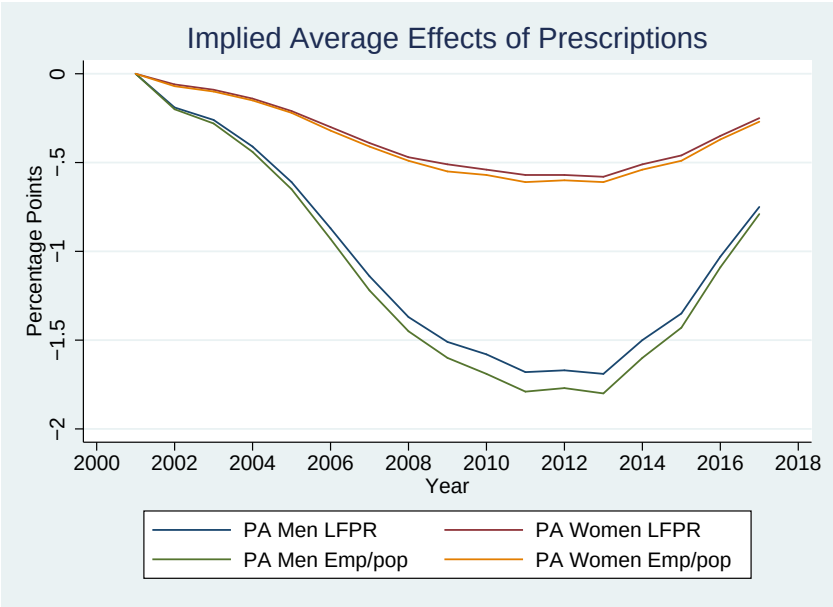

Figure 15: Coefficients by Initial Local Labor Market Conditions

Figure 15 shows the implied average effects for male and female labor force participation and the employment-to-population rate. As our discussion of Table 3 suggested, there is little difference between the implied time patterns between the labor force participation rate and the employment-

to-population ratio. Indeed, the gap (most notable for men) is the much smaller implied unemployment effect.

We can use these results to compare the expected decline in participation rates implied by our estimates with the declines implied by Krueger (2017)'s estimates. Krueger (2017)'s results imply a 0.6 percentage point decline in participation for men between 1999 and 2015 and a 0.8 percentage point decline for women over the same period. Our results imply a larger decline for prime-age men and a smaller decline prime-age women. For men we find that the decline in participation associated with rising opioid prescriptions is 1.3 percentage points between 2001 and 2015, but that the decline is only 0.5 percentage points for women. In our estimation, then, opioid prescriptions account for 44 percent of the realized decline in the prime-age male participation rate and 17 percent for prime-age women.

6.2 Consistency with Harris, Kessler, Murray, and Green (2019)

Harris et al. (2019) use county-level regressions on data from 10 states with Prescription Drug Monitoring Programs (PDMPs) primarily between 2013 and 2015.[16] They use the plausibly exogenous variation in the concentration of high-volume prescribers as an instrument to achieve causal identification, which are not available for most of the US. Their empirical specification is similar to ours, but the opioid prescription is measured in per capita rates rather than log rates per hundred residents. They use county-level labor market outcomes whereas we use individual-level labor market outcomes, but as Equation 3 shows, our measures can be related to one another. Loose comparisons of the coefficients can be made by considering a one-standard deviation shift in the relevant prescription variables: one standard deviation is 0.42 in Harris et al. (2019) and 0.39 in our analysis.

Harris et al. (2019)'s results for the 10 states in their sample are considerably larger than our estimates for the entire US. Starting with employment-to-population results, they have a coefficient of -0.070 for all adults in their preferred IV specification compared to our estimates of -0.050 for prime-age men and -0.017 for prime-age women. A one standard deviation shift in prescriptions in Harris et al. (2019) produces a 2.9 percentage point decline in employment rates. A one standard deviation increase in log prescriptions produces a 1.9 percentage for prime age men and 0.7 percentage point decline for prime age women. Taking the average of male and female results in a combined decline of 1.3 percentage points. The difference in these results is unlikely due to the inclusion of younger and older demographic groups in Harris et al. (2019). While not reported, our estimates for both younger (age less than 24) and older (age greater than 54) have smaller coefficients and larger standard errors.

Overall, our estimates are roughly half of those in Harris et al. (2019). There are many differences in the analyses, but key differences include our longer sample period, broader set of states, and controls for individual and labor-market conditions controls. While we cannot identify the relevant margins that produce the differences in the estimate effects, we view Harris et al. (2019) as

[16] See Buchmueller and Carey (2018) and Horwitz et al. (2018) for related evidence on PDMPs.

focused on and strongly supportive of the causal impacts of opioid prescriptions. Thus Harris et al. (2019) complement our results, which focus on improving the geography covered and controls used in order to better estimate the economy-wide impacts of opioid prescriptions on prime-age adults.

6.3 Sources of Differences with Currie, Jin, and Schnell (2018)

Conversely, Currie et al. (2018) find that opioid prescriptions have a positive and significant effect on the employment-to-population ratio for women and no effect for men. Given the substantive differences in their estimates and ours, we undertook a more thorough comparison of the two sets of results. There are key differences in data used, geographic groupings, age-sex groupings, and regression specifications. In terms of data, Currie et al. (2018) use quarterly prescription data that are specific to age-sex groups along with aggregate labor market measures from the Quarterly Workforce Indicators (QWI) for the same age-sex groups. Our prescription rate data are broader, limited to an annual frequency with no age or sex information, but our labor market data are at the individual level, allowing for controls from the ACS. Currie et al. (2018) also rely on a different approach to geographically grouping less populous counties. Their approach uses a county-state grouping that consolidates all of the counties in a state with fewer than 100,000 people into one observation. Section 2.1 describes our rationale for using coumas. Finally, Currie et al. (2018) rely more on fixed effects rather than controlling for individual-level demographics and area-level economic conditions that can be time-varying.

Table 7 shows several regressions on annual data, which we estimated to better understand the differences in results. The first row of Table 7 reports the relevant results from Currie et al. (2018) (Table 3). Our first regression (row 2) is designed to "replicate" this result from Currie et al. (2018), although we do use a different specification as well as annual data for both the labor market state and the opioid prescription rates. Our "replication" results are less precise, but still produce positive coefficients for some, but not all, of the groups. Row 3 replaces geographic fixed effects with our controls for county economic conditions (demand shocks and 2000 emp-pop), which shrinks or makes negative two out of four of the coefficients. Overall, our regressions with county-state geographic groupings and QWI labor market data are a bit imprecise but still not in line with the results in Currie et al. (2018).

The fourth row swaps out the aggregate QWI data with individual-level ACS data to determine if differing labor market data is behind our conflicting result. Just exchanging the ACS data for the QWI data shifts the results closer to the results of our analysis. The significant positive coefficient for older women is now negative and significant, along with both male groupings' coefficients. The data source is different at this stage, and we suspect what is most important is our individual-level data reflecting the connection between residency and pharmacy data. Using the QWI to measure the dependent variable results in places with employment concentrations above their population (think Manhattan) that are linked in the analysis to urban prescriptions rates. Work-place based employment-to-population ratios are cyclically-sensitive both to general decline of unemployment economy-wide and to the relative recovery in urban versus commuting versus rural counties. This

issue is difficult to address with fixed effects, which was our motivation for including specific labor market controls that vary both over time and over the geography. The fifth row incorporates the FRBC detailed specification, which produces results that are qualitatively similar but with smaller magnitudes than row four. For comparability purposes, rows six and seven use our couma geographic grouping with the ACS but for the demographic groups used in Currie et al. (2018); the results in these last two lines are similar to our baseline results. One interesting feature of this table is that it shows, in rows 5 and 7, that the introduction of our controls tends to moderate the estimated coefficients while maintaining the size pattern of the coefficients.

Overall, we believe our results differ from Currie et al. (2018) in several ways, but the key to the sign difference is the different geographic assignment of labor market outcomes in our data sets. In our view, an individual's labor market decisions and the local opioid prescription rate they face are better linked using the individual's residence rather than their place of work.

Table 7: Employment to Population Ratio and Log Prescription Rate: Model and Data Comparison

	Women 19-44	Women 45-64	Men 19-44	Men 45-64
Currie et al. (2018) Table 3, Column 1	0.024***	0.014*	0.020**	0.005
QWI with County-State Geography				
Currie et al. (2018) "Replication"	0.026	-0.007	0.063*	0.032
FRBC Controls Rather than FEs	0.056	-0.068	0.021	-0.143*
ACS with County-State Geography				
No Controls, No Fixed Effects	-0.010	-0.049***	-0.052***	-0.076***
FRBC Baseline Specification	0.016**	-0.021***	-0.034***	-0.040***
ACS with Couma Geography				
No Controls, No Fixed Effects	-0.010	-0.053***	-0.053***	-0.082***
FRBC Baseline Specification	-0.001	-0.043***	-0.049***	-0.066***

All regressions include year robust standard errors with clustering on geographic units. Currie et al. (2018) "Replication" includes geographic fixed effects. FRBC controls include demand shock and 2000 emp-pop ratio. FRBC baseline specification from Table 3.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

7 Conclusion

This paper makes a contribution to our understanding of the scale and scope of the opioid crisis' effect on the labor market. We constructed a data set allowing us to accurately and jointly measure individuals' labor market outcomes along with the opioid prescription rate in their area. Our use of coumas to link a nationally representative data set on individuals' labor market outcomes with data on local opioid prescription rates allowed us to: (i) improve measurement of prescription rates, particularly in the rural areas that are a critical part of the opioid crisis, (ii) investigate demographic heterogeneity in the relationship between prescription rates and labor market outcomes, and (iii)

account for geographic variation in short-term labor demand shocks, long-term economic conditions, and residential sorting.

The scale of the opioid crisis makes it likely that the crisis would generate labor market impacts. Our results confirm that these effects have been substantial, depressing economic outcomes in counties that have had high rates of opioid prescriptions even within the same state. To see how large these effects could be, compare a couma roughly one standard deviation above the average prescription rate in 2010 (Macon County Illinois, a relatively small, identified county in central Illinois) to one roughly one standard deviation below the mean (Ogle and Stephenson counties, which form a CPUMA in northwest Illinois). Macon county had an opioid prescription rate of 104.7 per hundred residents, while Ogle and Stephenson had an average prescription rate of 50.5 per hundred residents. This 0.73 log point differential would imply that, ceteris paribus, we would predict the prime-age employment rate of men to be 3.6 percentage points lower in Macon County than in Ogle and Stephenson counties. The prime-age employment rate of women would be expected to be 1.3 percentage points lower in Macon County.

Although we have confidence in the measured impacts of higher prescriptions rates, we have little evidence on the reversibility of these effects. That would require carefully examining places that had effectively reduced their prescription rates (and likely illegal opioid availability as well) through specific measures. At this point, the data does not include any viable candidate areas, although prescription rates have gradually declined in most areas of the United States since 2010 [17].

Our results paint a picture of widespread impacts within the group most affected by the opioid crisis: less-educated men. While abuse and mortality rates have rightly focused attention on white men with less than a BA, we found that the labor market outcomes of minority men with less than a BA are even more impacted when exposed to high prescription rates. And while there is a correlation between economic conditions and opioid prescription rates, we showed a variety of evidence indicating that the labor market effects of opioid prescriptions are consistent across areas, regardless of their long- or short-term economic conditions.

While many relevant policy issues are outside the scope of this paper, our work serves to show the scale of the impact of the opioid crisis on the labor market. In our view, the impact of the opioid crisis on regional and national labor markets looks to be large and statistically robust.

References

Alpert, A., D. Powell, and R. L. Pacula (2018). Supply-side drug policy in the presence of substitutes: Evidence from the introduction of abuse-deterrent opioids. *American Economic Journal: Economic Policy 10*(4).

Barnett, M. L., A. R. Olenski, and A. B. Jena (2017). Opioid-prescribing patterns of emergency

[17] Two CPUMAs (Oklahoma pan handle and eastern Texas) actually experience large declines in prescription opioid rates over our sample time frame but the reasons for those declines are unkown.

physicians and risk of long-term use. *New England Journal of Medicine 376*(7), 663–673. PMID: 28199807.

Bartik, T. J. (1991). Who benefits from state and local economic development policies?

Blanchard, O. and L. Katz (1992). Regional evolutions. *Brookings Papers on Economic Activity 1*, 1–76.

Buchmueller, T. C. and C. Carey (2018). The effect of prescription drug monitoring programs on opioid utilization in Medicare. *American Economic Journal: Economic Policy 10*(1), 77–112.

Carpenter, C. S., C. B. McClellan, and D. I. Rees (2017). Economic conditions, illicit drug use, and substance use disorders in the United States. *Journal of Health Economics 52*, 63 – 73.

Case, A. and A. Deaton (2015). Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century. *Proceedings of the National Academy of Sciences 112*(49), 15078–15083.

Case, A. and A. Deaton (2017). Mortality and morbidity in the 21st century. *Brookings Papers on Economic Activity Spring*, 397.

Charles, K. K., E. Hurst, and M. Schwartz (2018). The transformation of manufacturing and the decline in US employment. *NBER WP 24468*.

Cicero, T. J., M. S. Ellis, H. L. Surratt, and S. P. Kurtz (2014). The changing face of heroin use in the United States: A retrospective analysis of the past 50 years. *JAMA Psychiatry 71*(7), 821–826.

Compton, W. M., C. M. Jones, and G. T. Baldwin (2016). Relationship between nonmedical prescription-opioid use and heroin use. *New England Journal of Medicine 374*(2), 154–163. PMID: 26760086.

Currie, J., J. Y. Jin, and M. Schnell (2018). US employment and opioids: Is there a connection? *NBER WP 24440*.

Deiana, C. and L. Giua (2018). The US opioid epidemic: Prescription opioids, labour market conditions and crime. *MPRA Paper 85712, University Library of Munich, Germany*.

Evans, W. N., E. Lieber, and P. Power (2018). How the reformulation of OxyContin ignited the heroin epidemic. *The Review of Economics and Statistics*. Forthcoming.

FDA (2018, March 1). *FDA Analysis of Long-Term Trends in Prescription Opioid Analgesic Products: Quantity, Sales, and Price Trends*. US Food and Drug Administration.

Finkelstein, A., M. Gentzkow, and H. Williams (2018). What drives prescription opioid abuse? Evidence from migration. *Stanford Institute for Economic Policy Research Working Paper 18-028*.

Harris, M. C., L. M. Kesslery, M. N. Murray, and M. E. Glenn (2019). Prescription opioids and labor market pains: The effect of schedule II opioids on labor force participation and unemployment. *Journal of Human Resources*. Forthcoming.

Hollingsworth, A., C. J. Ruhm, and K. Simon (2017). Macroeconomic conditions and opioid abuse. *Journal of Health Economics 56*, 222–233.

Horwitz, J., C. S. Davis, L. S. McClelland, R. S. Fordon, and E. Meara (2018). The problem of data quality in analyses of opioid regulation: The case of prescription drug monitoring programs. *NBER Working Paper 24947*.

Imbens, G. W. (2015). Matching methods in practice: Three examples. *Journal of Human Resources 50*(2), 373–419.

Isserman, A. M. and J. Westervelt (2006). 1.5 million missing numbers: Overcoming employment suppression in County Business Patterns data. *International Regional Science Review 29*(3), 311–335.

Katz, J. (2017, August 3). Short answers to hard questions about the opioid crisis. *The New York Times*.

Kenan, K., K. Mack, and L. Paulozzi (2012). Trends in prescriptions for oxycodone and other commonly used opioids in the United States, 2000–2010. *Open Medicine 6*(2), e41.

Kromer, B. K. and D. J. Howard (2010). Comparison of ACS and CPS data on employment status.

Krueger, A. B. (2017). Where have all the workers gone? An inquiry into the decline of the US labor force participation rate. *Brookings Papers on Economic Activity*.

Laird, J. and T. Nielsen (2016). The effects of physician prescribing behaviors on prescription drug use and labor supply: Evidence from movers in Denmark. *Mimeo., Harvard University*.

Maggio, M. D. and A. Kermani (2016). The importance of unemployment insurance as an automatic stabilizer. *NBER WP 22625*.

Mazumder, B. (2005). Fortunate sons: New estimates of intergenerational mobility in the United States using social security earnings data. *Review of Economics and Statistics 87*(2), 235–255.

Pearl, J. (2009). *Causality: Models, Reasoning and Inference* (2nd ed.). Cambridge University Press.

Powell, D., R. L. Pacula, and E. Taylor (2015). How increasing medical access to opioids contributes to the opioid epidemic: Evidence from Medicare Part D.

Quinones, S. (2016). *Dreamland: The True Tale of America's Opiate Epidemic*. Bloomsbury Press.

Rogers, W. (1993). Regression standard errors in clustered samples. *Stata technical bulletin 13*, 19–23.

Ruhm, C. J. (2015). Recessions, healthy no more? *Journal of Health Economics 42*, 17 – 28.

Ruhm, C. J. (2018). Deaths of despair or drug problems? *NBER WP 24188*.

Savych, B., D. Neumark, and R. Lea (2018). The impact of opioid prescriptions on duration of temporary disability.

Schroeder, J. and D. V. Riper (2016). Harmonizing PUMAs across time: An aggregation algorithm to achieve minimally acceptable consistency. *Mimeo., University of Minnesota*.

Schweitzer, M. E. (2017). Manufacturing employment losses and the economic performance of the Industrial Heartland. *FRB of Cleveland Working Paper 17-12*.

Shavers, V. L., A. Bakos, and V. B. Sheppard (2010). Race, ethnicity, and pain among the US adult population. *Journal of Health Care for the Poor and Underserved 21*(1), 177–220.

Yellen, J. L. (2017, July 13). *Federal Reserve's Second Monetary Policy Report for 2017* (S. HRG. 115-108 ed.). Washington, DC: US Senate Committee on Banking, Housing, and Urban Affairs.

A Couma Data Appendix

As discussed in the text in Section 2.1.1, the geographic area we use in our analysis is the couma, a mix of counties and Public Use Microdata Area (PUMAs). PUMAs have the desirable characteristics that they generally aggregate adjoining counties that can be matched to county-level data sources. An undesirable characteristic of PUMAs, at least for our purposes, is that their boundaries change during our sample period. For this reason we actually use the IPUMS-provided consistent-PUMA (CPUMA) geography, which "is an aggregation of one or more 2010 US Census PUMAs that, in combination, align closely with a corresponding set of 2000 PUMAs" (Schroeder and Riper (2016)). In a small number of cases where CPUMA boundaries cross county lines we have to be further aggregate to identified consistent areas in both individual-level data and prescription rates.

To clarify the relationship between counties, CPUMAs, and coumas, Figure 16 show maps for two example using the state of Nebraska and Franklin County, Ohio (the most populous county in Ohio). In Nebraska, sets of adjoining counties are aggregated into 7 CPUMAs while two counties are directly identified. These 9 units of observation are the coumas in Nebraska. Franklin County includes several CPUMAs, but the entire county is a single observation in our coumas since prescription data is not available below the county level.

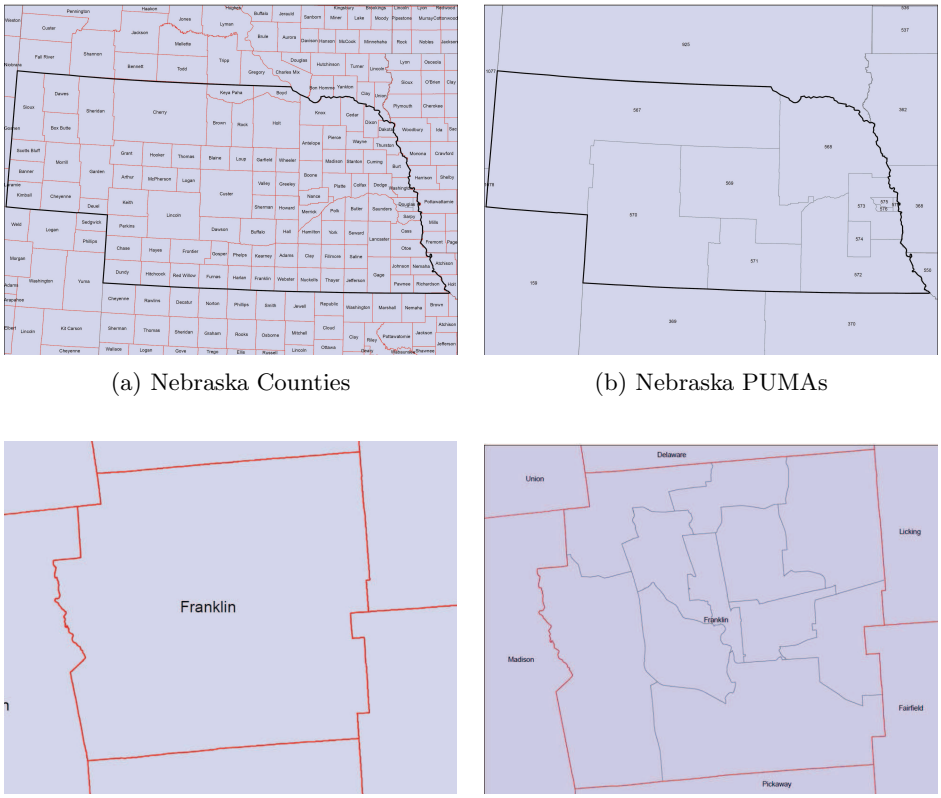

(a) Nebraska Counties

(b) Nebraska PUMAs

(c) Franklin County

(d) Franklin County PUMAs

Figure 16: Nebraska and Franklin County, Ohio