Opioids and the Labor Market

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

August 11, 2020

Abstract: This paper studies the relationship between local opioid prescription rates and labor market outcomes in the United States between 2006 and 2016. To understand this relationship at the national level, we assemble a data set that allows us to both include rural areas and to estimate the relationship at a disaggregated level. We control for geographic variation in both short-term and long-term economic conditions. In our preferred specification, a 10 percent higher local prescription rate is associated with a lower prime-age labor force participation rate of 0.46 percentage points for men and 0.15 percentage points for women. Multiple robustness checks generate results that are inconsistent with theories of reverse causation that prescriptions rates rose in response to weakness in the labor market. These checks include instrumental variable estimation on a subset of areas; estimates of the effect of unemployment on prescription opioid misuse; and estimation within 2000 labor market quintiles.

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 with 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, the Upjohn Institute, Ohio State's Department of Human Sciences, North American Regional Science 2018 Meeting, the Association of University Business and Economics Research 2018 Meeting, the 2019 ACS Data Users Conference, and the International Association of Applied Econometrics 2019 Meetings.

1 Introduction

Opioids, including prescription pain killers, are widely recognized as the cause of a public health emergency in the United States. By 2016 drug overdose had 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]

The interaction between opioids and the labor market is an important dimension of the opioid crisis (Case and Deaton (2017), Krueger (2017)). There is growing evidence that the supply of prescription opioids in an area depresses its labor force participation rate. Areas with higher opioid prescription rates have been shown to 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)). However, results are not uniform: Currie et al. (2019) find that higher opioid prescription rates had a small positive effect on employment-to-population ratios for women and negligible effects for men.

This paper focuses on understanding the national impacts of prescription opioids on US labor markets. While deaths and direct costs to state and local governments have been a focus of attention, the potential for general labor market impacts is another important channel for costs of the opioid crisis. Labor market impacts allow us to understand the breadth of the affected population, which could help with scaling an appropriate response. Finally, improved estimates on this critical issue can help to inform policy made in response to the state of the labor market, such as that of the Federal Reserve (Yellen (2017), Powell (2019)). Our primary contributions to understanding the impacts of prescription opioids on US labor markets are (i) to improve the joint measurement of labor market outcomes and local prescription rates and (ii) to improve the quality of controls for short- and long-term local labor market conditions.

One reason for the literature's focus on sub-national geographies is the search for credible identification. Another reason, though, is simply that it is difficult to create 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, these prescription data should be geographically linked to data on labor force participation, unemployment, and employment.

[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.

Creating a geographically-linked data set with prescription and labor market information typically forces researchers to choose between constructing 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. 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.

Our improvement to these approaches for handling nonidentified counties is consequential, as 29 percent of the US population in our sample from the 2006 to 2016 American Community Surveys (ACS) resides in a nonidentified county. We measure an individual's local area as his/her county when identified, or else as his/her 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 market outcomes. This geography is labeled the *couma* in Case and Deaton (2017).

Turning back to the issue of credible identification, a central concern is that geography could create variation in labor market outcomes that might be mistaken for the effects of opioid prescriptions. To illustrate this issue over a short time horizon, consider 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. Over a longer time horizon, differences in labor market conditions are very pronounced within the US, and these conditions could be correlated with prescription rates. We address this issue by estimating a panel model with specific controls for business cycles and local labor market performance. We also estimate a difference-in-differences specification of our model, and perform robustness checks by estimating models where identifying variation comes from an instrumental variable and from prescription rates within groups of similar local labor markets.

We have three main empirical results. First, 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 (sometimes by) increasing unemployment. In our preferred specification a 10 percent higher local prescription rate is associated with a lower prime-age labor force participation rate of 0.46 percentage points for men and 0.15 percentage points for women. Since Krueger (2017), there has been interest in how these effects might explain the decline labor force participation after 2000. Using these coefficients and national data on prescription rates our preferred results suggest that a decline of 1.5 percentages points of prime age male participation is associated with prescription growth from 2001 to 2015. The figure for prime age women is 0.5 percentage points, which is smaller but still a large fraction of the realized change in participation rates.

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 the National Survey of Drug Use and Health (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 finds opioid prescription rates cause labor market outcomes. We then estimate an instrumental variables specification for a subset of our sample using the state-based triplicate prescriptions programs documented by Alpert et al. (2019) as an instrument for prescription rates. Our IV results are qualitatively identical to those from our preferred specification, both overall and for specific demographic groups. We also investigate reverse causality from economic conditions to opioid prescription rates. We find no increase in opioid misuse following the short-term unemployment shock that was the Great Recession in data from the NSDUH. Finally, we estimate large coefficients of opioid prescription rates on participation when restricting attention to areas with weak labor markets in 2000. It turns out there is substantial variation in prescription rates within areas with similar local labor market conditions. If the direction of causality went from labor market outcomes to opioid prescription rates, then we would expect to find both less variation in prescription rates within weak labor markets and a small estimated coefficient within these weak labor markets.

The remainder of the paper is organized as follows: Section 2 describes the data used in the analysis, Section 3 discusses our empirical specifications and identification strategy, and Section 4 presents our results. Section 5 investigates threats to a causal interpretation of our results. Section 6 discusses the implied scale of our estimates for national outcomes and Section 7 concludes. Appendix A provides detail and estimates from additional specifications, and Section C compares our results to key parts of the literature.

2 Data

2.1 Connecting Local Opioid Prescriptions with Labor Supply

2.1.1 Individual-Level Labor Market Data

When choosing which data source to use for our dependent variable, there are a few options, each with its 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 are 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 three-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. (2019). The QWI are 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 individual-level characteristics influence decisions to participate in the labor market and work. Furthermore, QWI employment figures reflect 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 one-year 2017 ACS). Additionally, the QWI data are somewhat noisy for counties below 100,000 residents, which causes Currie et al. (2019) to apply additional geographic aggregation to make the data more reliable.

Given the concerns around the CPS and the QWI, we decided to use the Integrated Public Use Microdata Series (IPUMS-USA) of the 1 percent sample of the American Community Survey (ACS) from 2006 to 2016.[2] This annual data set includes detailed information for

[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

individuals' labor market status, age, race, sex, and education level, but the county of the individual observations is not always identified. About 80 percent 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 identified counties (tan) and CPUMAs (purple).

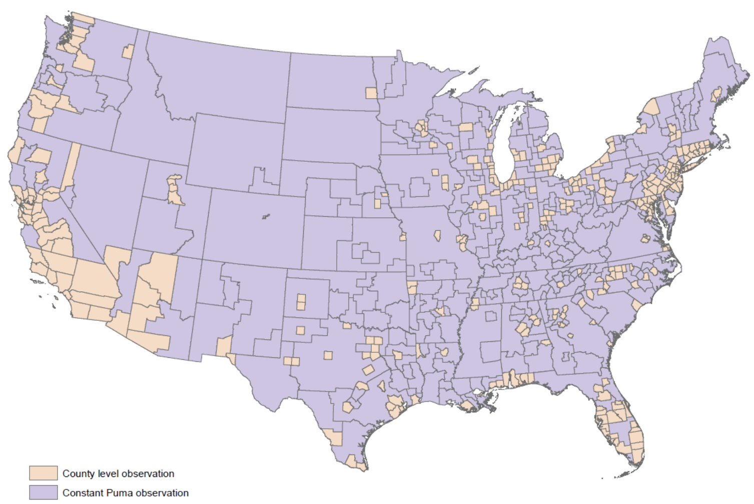

Figure 1: COUMAs
Note: Identified counties between 2006 and 2016 are shown in tan, and nonidentified counties (aggregated into CPUMAs) are shown in purple.

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 (2011)).

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

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 his/her PUMA. The CDC prescription opioid data set is derived from the records for approximately 59,000 retail (non-hospital) pharmacies, which cover nearly 90 percent of counties and nearly 90 percent of all retail prescriptions in the US. While the precise morphine milligram equivalents (MME) prescribed would be preferable to the number of prescriptions, which we use, these data are only publicly-available for 2015. Moreover, these variables appear to provide similar measures, as 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 using prescription counts, the time series of national MME quantities is very similar to the time series of our average prescription counts between 2006 and 2016 (FDA (2018)).

2.1.3 Implication of Aggregating Prescription Rates to Larger Geographies

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. (2019) and Krueger (2017) aggregate most rural counties up to the state level, and if one takes the perspective that closer measures are better than more aggregated ones, we can interpret aggregated measures in terms of measurement error.

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

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

a contrast with coumas, which assign higher prescription rates to some parts of northern Michigan while assigning the lowest rates to 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.

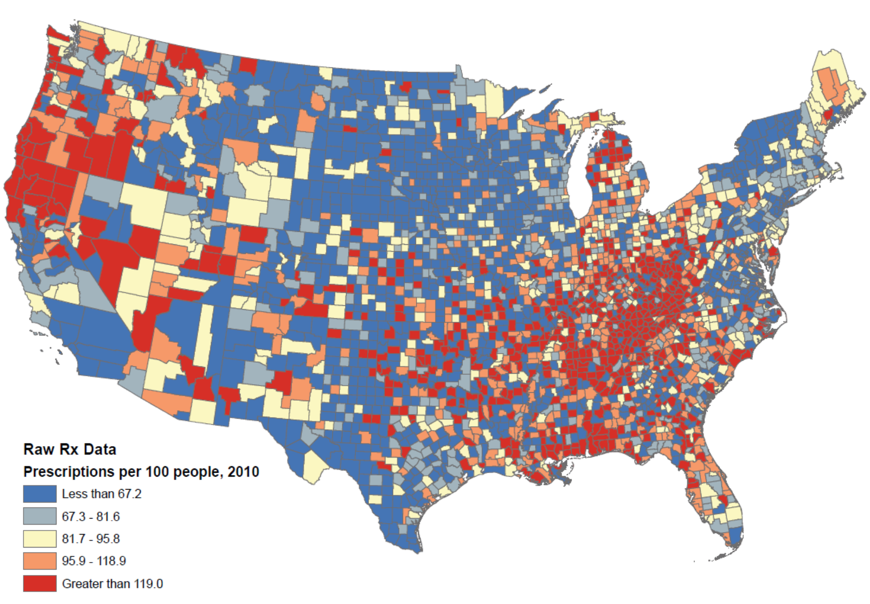

(a) County-Level Data

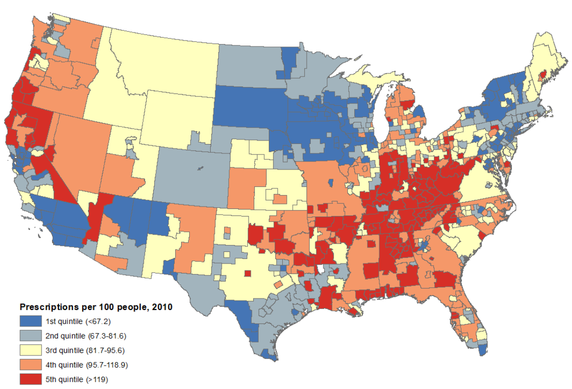

(b) Couma-Based Aggregation

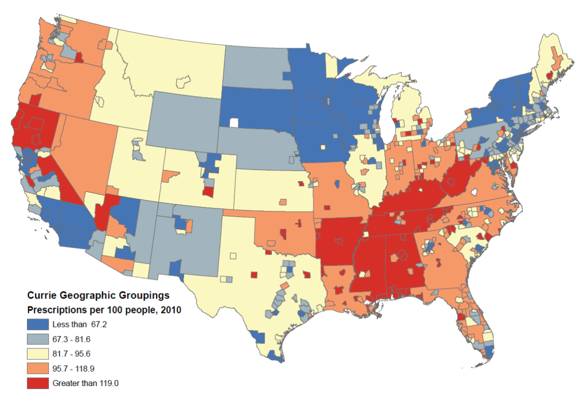

(c) State-Based Aggregation

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.

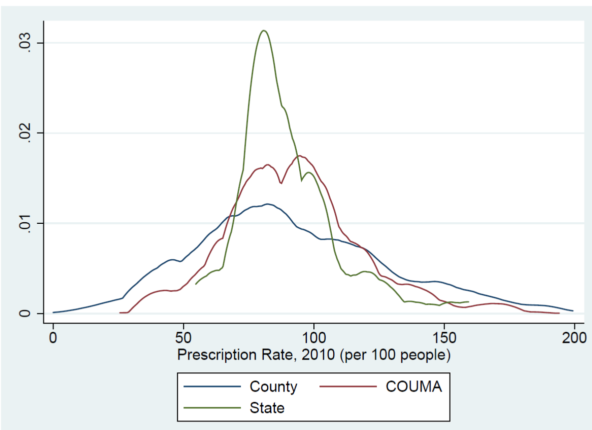

(a) Prescription Rates by Geography

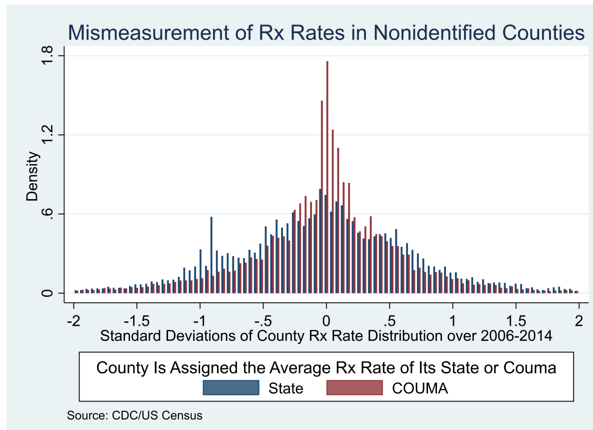

(b) Difference Relative to County Measure by Geography

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.

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 Patterns (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 three-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 the Maggio and Kermani (2016) specification 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:

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

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, such as trade and technology changes, that are differentially experienced depending on the local industry mix.

Figure 4 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 the lowest 10 percent of shocks in our sample would be drawn from recession years.

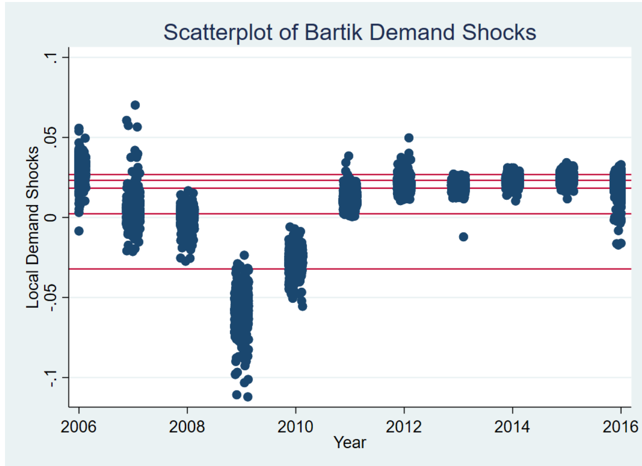

Figure 4: Labor Demand Shocks

As Figure 4 shows, recessions and recoveries involve important dynamic local labor adjustments that 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.

2.2.2 Persistent Geographic Patterns in Labor Markets

While the labor market outcomes we have in mind are the focus of both cyclical and trend analyses, geographic regions can also have persistently different labor market outcomes. 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). To directly account for this challenge we also include couma-level labor market statuses from the 2000 Census as an additional control .[5] The 2000 labor market status should be relatively less impacted by opioid prescriptions.

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

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).[6] 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 age 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.[7] 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).[8] In addition to variables covering drug use in great detail, the NSDUH 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

[6]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.

[7]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.

[8]The time horizons over which the NSDUH asks respondents about drug use are ever, in the past year, and in the past month. Previous literature has shown that estimates can be sensitive to the reference period used (Carpenter et al. (2017)).

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.
Rx Rate	78	29	68	25	90	29
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 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.

3 Empirical Specifications

Conceptually, the experiment that we are interested in observing is how individuals' labor market outcomes respond to randomly-assigned levels of opioid prescriptions in their area. Because we do not observe this experiment, we are instead faced with the options of searching for quasi-random variation in prescription rates or else trying to control for confounders, or common causes, of local prescription rates and individuals' labor market outcomes. Since we are interested in understanding the magnitude of the effect of opioids on labor market outcomes at a national scale, we focus on the latter approach.

Our approach follows individuals' labor market outcomes as they respond to the conditions in their area in terms of both the level of prescription opioids and evolving labor market conditions. The equations we estimate are all a form of 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 the natural log of the opioid prescription rate in the individual's couma, P_{jt}, a function of the current local economic conditions (LECs) facing the individual, $f(LEC_{jt})$, individual characteristics observed at time t, X_{it}, and a term to represent any unobserved factors, ϵ_{it}:

$$Y_{ijt} = \alpha P_{jt} + f(LEC_{jt}) + \beta X_{it} + \epsilon_{it}. \tag{2}$$

While this specification is similar to the specifications used in 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 three-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). This leaves the proper lag structure unclear, but prescription opioid rates are highly correlated over time (the correlation within two-year windows is always greater than 0.96). This indicates that the exact lag structure should result in only minor estimation differences so we use contemporaneous log prescriptions rates. Our data allow us to run panel regressions on individuals' labor force status from 2006 to 2016 with CDC data on average prescriptions per person in 648 coumas. We run all regressions separately for men and women, given prior evidence of differences in labor market attachment and differential impacts of opioid prescriptions.

LECs are almost certainly the most important confounding variable of local opioid prescription rates and an individual's labor market outcomes, so we consider three approaches to controlling for this potential confounder.[9] 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. The samples in each local area are reasonably large, but in any even given period might over- or under-represent demographic groups that could impact the measured labor market performance of the location. Summing the dependent variables within geographies can generate average labor market values (employment-to-population ratio when Y is 1/0 in employment, labor force participation

[9]Each of these models represents an extension of the specification used in Krueger (2017) that is possible due to our use of panel data.

when Y is 1/0 in participation, and unemployment-to-population when Y is 1/0 in unemployment) for area j. Equivalent models could be estimated using average labor market outcomes if the individual characteristics could be controlled for or were assumed to be unchanged.

Our preferred approach is to use specific controls that directly measure LEC_{jt}. This alternative controls for both cyclical and long-term economic conditions. The cyclical control, D_{jt}, is a Bartik measure of local impacts of national economic changes that is observable in each location over time. Average labor market outcomes in the location in a time prior to most of the growth in opioid prescriptions, $\overline{Y}_{j,2000}$ is a simple, and highly effective, approach to measuring the effects of long-term economic differences. In addition, we still include time effects and a location fixed effect, at a level of aggregation k above j, yielding:

$$f(LEC_{jt}) = \eta D_{jt} + \theta \overline{Y}_{j,2000} + \gamma_k + \delta_t. \tag{3}$$

This approach maintains panel elements to account for unobserved sources of variation, but allows some of the variation between places to identify opioid effects once the cyclical and long-term differences are accounted for.

The reason that specific controls is our preferred specification is that it includes controls for both cyclical and longer-term local labor market outcomes that are critical controls for estimating the relationship between opioids and the labor market, while not ignoring the remaining cross-sectional variation in prescriptions which helps in the identification of results on an inherently time-limited panel. The employment to population ratio is quite cyclical and, as we noted earlier when discussing Figure 4, cyclical impacts were uneven across the country during the Great Recession and its aftermath. Our estimates for employment to population ratios, particularly for men, are closer to zero in the difference-in-differences specification presented in Appendix A. This result is very likely due to the included fixed effects absorbing identifying variation.[10]

4 Estimation Results

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

[10] While there are also differences that arise from measurement of place of work and place of residence in the Quarterly Workforce Indicators (QWI), geographic groupings, age-sex groupings, and available controls, we believe that the differences between our results and those in Currie et al. (2019) are most likely due to their use of models that rely on fixed effects rather than controls for individual-level demographics and area-level economic conditions that can be time-varying.

population-weighted linear probability models of 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 changes associated with the regressors on these rates.

4.1 Specific Cyclical and Local Labor Market Controls

For a specific control strategy to be effective it must control for factors that are both consequential sources of variation in the difference-in-differences model while allowing some cross-sectional variation absorbed by fixed effects to be used in identification. The controls that we add are a measure of cyclical demand (Bartik shocks) and a measure of the longer-run performance of the the local labor market (2000 Census labor prime-age labor force statistics). Importantly, both of these factors vary across coumas, so that if these were the factors that we aimed to account for with the couma fixed effects, then we could look to a higher aggregation of regional fixed effects.

The specific control specification estimates in Table 2 report the effects on the labor force participation rate and the employment to population ratio. Participation effects for men are estimated at -0.046, with the effect being approximately three times larger in magnitude for men than for women (-0.046 versus -0.015).

Table 2: Labor Market States of Prime-Age Men and Women

	Men Participate	Men Emp/Pop	Men Unem/Pop	Women Participate	Women Emp/Pop	Women Unem/Pop
Prescrip. Rate	-0.046***	-0.049***	0.004**	-0.015***	-0.017***	0.003
	(0.005)	(0.005)	(0.001)	(0.003)	(0.004)	(0.001)
Demand Shock	0.309**	0.606***	-0.251***	-0.346**	-0.102	-0.199***
	(0.118)	(0.132)	(0.068)	(0.109)	(0.118)	(0.057)
2000 Particip.	0.593***			0.450***		
	(0.044)			(0.027)		
2000 Emp/Pop		0.519***			0.396***	
		(0.034)			(0.023)	
2000 Unem/Pop			0.292***			0.188***
			(0.032)			(0.029)
R-sqr	0.09	0.11	0.02	0.06	0.06	0.02
N	6410469	6410469	6410469	6625455	6625455	6625455

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

The second row of Table 2 shows that the local labor market demand shock variable is a highly significant factor for prime-age men's labor market status. Notably, men's employment rates are substantially affected by the demand shocks in a pro-cyclical pattern: A negative labor demand shock (such as a recession) reduces employment both through lower labor force participation and higher unemployment. In contrast, women's employment rates are not strongly correlated with the Bartik-style demand shocks, because both 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 make local demand conditions an important source of variation to be controlled for directly, rather than assuming that time or location fixed effects will absorb these sources of variation.

The third through fifth rows of Table 2 show 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. 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 from Krueger (2017)'s use of a cross-sectional model to our use of annual panel data and the inclusion of additional controls, the results in Table 2 are similar to those in Krueger (2017) in the sign and in the pattern of generally stronger effects for men than for women. Given the Krueger (2017) strategy of estimating over two three-year 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 the 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. While the opioid prescription variables are different, which makes the coefficient comparisons less direct, these results are qualitatively consistent. We will further compare our results to Krueger's 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. The influential results in Case and Deaton (2015) and Case and Deaton (2017) suggest exploring effects by subgroups of sex, education level, and race/ethnicity. For our purposes we examine eight subgroups by splitting 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.

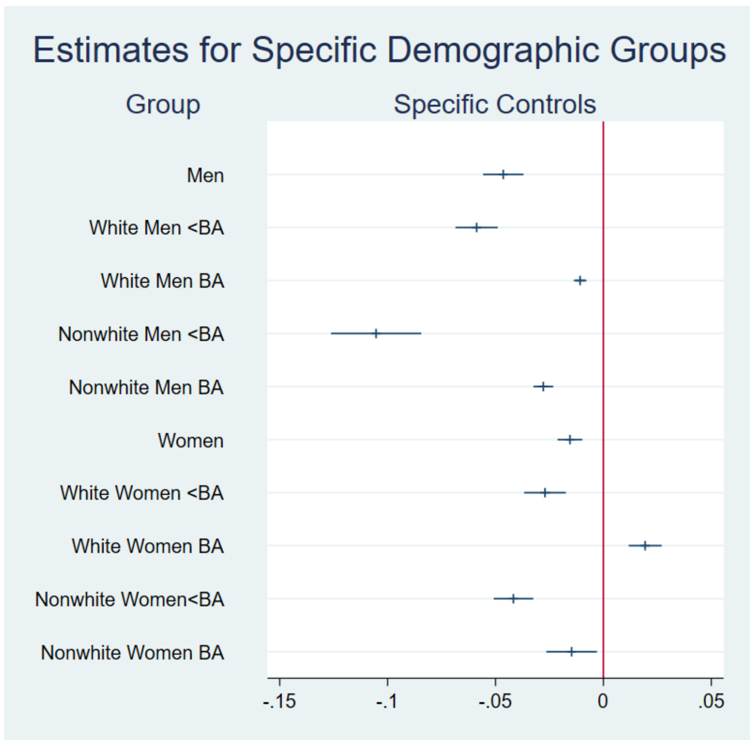

Figure 5: Labor Force Participation Effects by Demographic Groups

Figure 5 shows the estimated coefficients for the demographic groups in each of the models, along with the associated 95 percent confidence intervals. For men without a college degree, the effect of opioid prescription rates on the labor force participation rate is larger in magnitude than for those with a college degree (see the correlations in the upper half of Figure 5).[11] The coefficient 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

[11]These categories follow Case and Deaton (2015). We also examined other possible splits and found that measured outcomes for individuals with some college were more similar to those of high school graduates than to those of college graduates.

than the coefficient for white men without a BA. These results show that 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 men, which predated the growth of opioid prescriptions. That effect is accounted for in the other controls.

Another point worth emphasizing is that while Case and Deaton (2015) focus 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 are more vulnerable to opioids than their minority counterparts.

The lower half of Figure 5 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. (2019). Nonwhite women with a BA have a slightly lower and statistically insignificant expected participation 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 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 location with a higher opioid prescription rate 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 prescriptions 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 6 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 mortality differences, a result that may seem at odds with our large estimated labor market effects on nonwhite men.

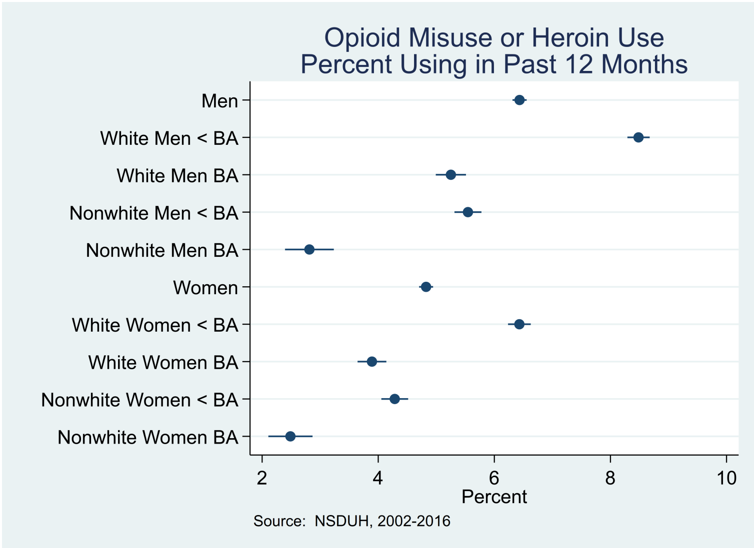

Figure 6: Abuse by Demographic Groups

Is there any way to reconcile our findings for minority men? Our findings are especially puzzling since one might reasonably assume that the effects of availability on employment 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 7a 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.

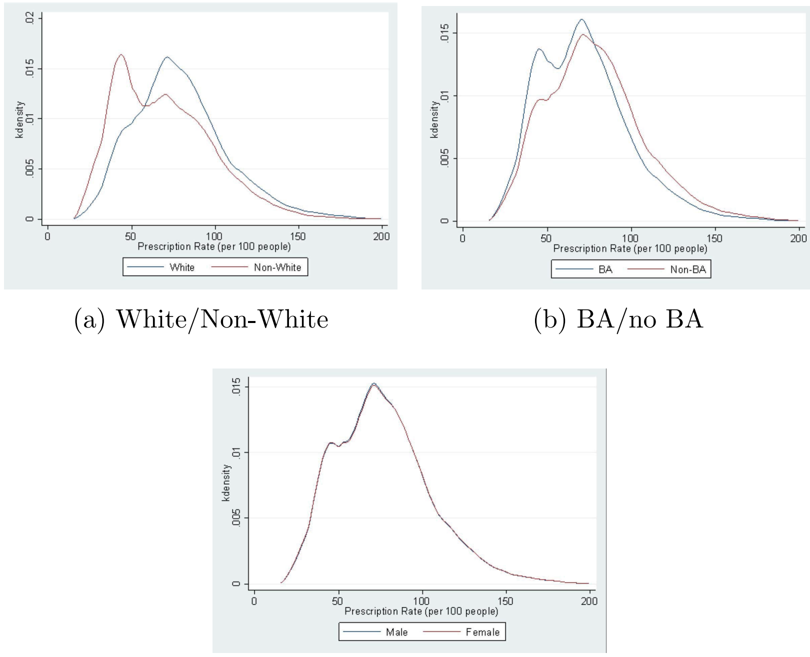

(a) White/Non-White

(b) BA/no BA

(c) Men/Women

Figure 7: 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 symptoms equivalent to those of 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 Specific Geographies

Informed scrutiny of the maps in Figure 2 and media accounts of the opioid crisis suggest possible geographic patterns to explore: rural counties, Appalachian counties, and perhaps the Rust Belt region. In each case, it has been argued that persistently bad economic outcomes might underlie 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 prescription coefficient to each specified geography. For Appalachian counties we use those listed by the Appalachian Regional Commission (ARC).[12] The Rust Belt counties are the counties identified by Schweitzer (2017) as being in the "Industrial Heartland" in 1969. Figure 8 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

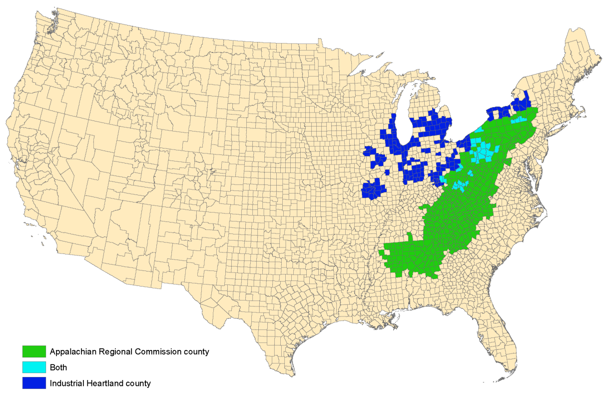

Figure 8: Appalachian Regional Commission and Rust Belt Counties

The primary coefficients on opioid prescriptions are all still statistically significant for both men and women, although in several cases a bit smaller than was seen in Table 2. 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 participation rates but a modestly (and statistically insignificant) higher female participation rate than in the higher-population identified counties. Appalachian coumas included in the ARC list of counties also saw a larger effect for both men and women in higher prescription counties. Finally, the coefficients for Rust Belt coumas show a smaller effect on men and a larger effect on women in Industrial Heartland counties. While there are modestly larger impacts in some of these geographies, these regressions demonstrate that the correlation between prescription rates and labor market outcomes is not dependent on the inclusion of these challenged geographies.

Table 3: Labor Force Regressions with Geographic Interactions

	CPUMA		Appalachia		Industrial Heartland	
	Men	Women	Men	Women	Men	Women
Prescrip. Rate	-0.038***	-0.017***	-0.044***	-0.014***	-0.047***	-0.014***
CPUMA*Prescrip	-0.004***	0.001				
ARC*Prescrip			-0.002*	-0.002**		
IH*Prescrip					0.003*	-0.003***
Demand Shock	0.045	-0.309**	0.268*	-0.381***	0.254*	-0.279**
2000 LFPR	0.573***	0.452***	0.586***	0.442***	0.595***	0.447***
R-sqr	0.09	0.06	0.09	0.06	0.09	0.06
N	6424995	6641288	6424995	6641288	6424995	6641288

All regressions include demographic variables, 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.

This section provides additional evidence on how strongly we should interpret our estimated coefficients as representing local prescription rates causing labor market outcomes. We first survey the growing literature working on identifying the causal effect of prescription rates on labor market outcomes. We then present results from an instrumental variables specification that may be deemed more credible for identifying causal effects if one believes our preferred specification does not adequately control for confounders. We finish our focus on causality by looking for evidence of reverse causality, both in terms of opioid abuse responding to the Great Recession's labor market shock, as well as in terms of variation of and correlations with prescription rates within areas that have similar long-term labor market conditions.

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, while Charles et al. (2018) show that declining local manufacturing employment is associated with increased local opioid abuse. Together, these results suggest a potential increase in drug use in recessions and after labor demand shocks.

Other research discounts the purely economic motivations for opioid abuse. Glei and Weinstein (2019) observe that drug abuse may be related to perceived economic distress that is not captured by headline statistics. Ruhm (2019) finds that drug death rates increased more in places experiencing relative economic decline than in places with more robust economic growth, but he attributes this weak relationship to confounding factors related to the demographic composition of places. The lack of social capital (Zoorob and Salemi (2017)) and social determinants of health (physical and psychological trauma, concentrated disadvantage, isolation, and hopelessness) have also been linked as root causes to drug and opioid

[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.

abuse (Dasgupta et al. (2018)).

The analysis in the literature we consider to be nearest to our own is Harris et al. (2019). Harris et al. (2019) use plausibly exogenous variation in the concentration of high-volume prescribers to estimate county-level regressions on data from 10 states with prescription drug monitoring programs (PDMPs) primarily between 2013 and 2015.[15] Overall, our estimates are roughly half of those in Harris et al. (2019). The key differences in our analyses include our longer sample period, broader set of states, and controls for individual and labor-market conditions. 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.

5.2 Applying the Triplicate Prescription Instrumental Variable

Alpert et al. (2019) examine the effects of the 1996 introduction and marketing of OxyContin, exploiting recently unsealed court documents to show that state-based triplicate prescriptions programs were an obstacle to Purdue Pharma's marketing efforts. Triplicate prescription programs, essentially a precursor to more recent prescription drug monitoring programs that required the use of special state-issued prescription forms for opioids like OxyContin, resulted in about 50 percent lower OxyContin distributions in states where they were required. Alpert et al. (2019) provide evidence that the prior-existence of these triplicate laws represents an exogenously-determined source of state-level differences in opioid prescription rates, that is independent of labor market performance. They specify their estimates as a state-level difference-in-differences analysis of the effects of the introduction of OxyContin to the market.

We apply the "triplicate-state" identifier of Alpert et al. (2019) as an instrument in an IV regression on the labor market outcomes examined in our analysis. Specifically, as a first stage we estimate:

$$P_{jt} = \sum_{t=1996}^{2017} \beta_t \times 1(\{\text{Non-triplicate} = j\} \times 1\{\text{Year} = t\} + \delta_t + \epsilon_{jt} \tag{4}$$

where the regional variation in couma-level prescriptions is limited to state-level variation by the triplicate instrument's inherent variation.

The first stage estimation results on our prescription rate data are very strong, paralleling the state-level drug distribution results for OxyContin and other other prescription opioids reported in Alpert et al. (2019). This approach focuses directly on exogenous variation to

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

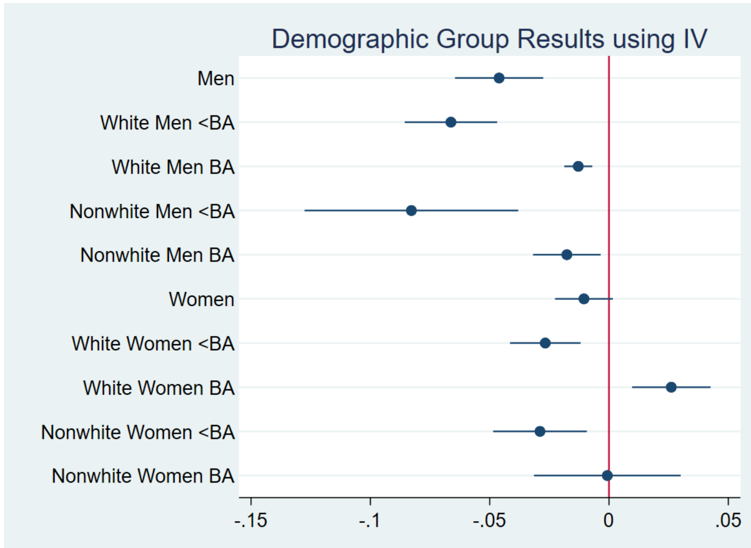

Figure 9: Triplicate-State IV Results by Demographic Groups

avoid potential endogenous interactions between the labor market and prescriptions, but its focus on limited state differences makes this a less effective tool for examining nationwide effects because there is important variation below the state level and between triplicate and non-triplicate states.

The resulting IV regression results largely confirm our existing results, although the variation in prescription rates is more limited resulting in larger standard errors. The IV regression point estimates and patterns between demographic groups shown in Figure 9 are quite similar to the estimates shown in Figure 5. We are encouraged by the fact that both the point estimates and the patterns between demographic groups are closely aligned when the variation used to explore the labor market effects is only from differences in OxyContin marketing due to the triplicate prescription rules across states.

5.3 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 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. 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 10, 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.

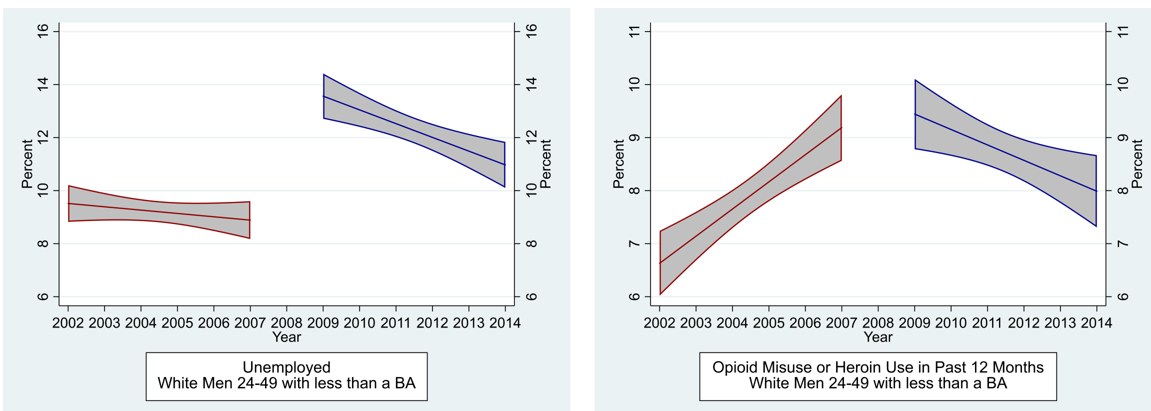

(a) Unemployed

(b) Opioid Misuse or Heroin Use

Figure 10: Labor Market Outcomes and Opioid Use over Time

5.4 Local Economic Conditions and Causal Paths

The simplest path to reverse causality would be that prescriptions simply flowed more frequently into depressed labor markets. The medical literature and Krueger (2017) both suggest a surprising degree of randomness in prescription frequencies across the United States. Usefully, we have direct evidence in the form of the cross-sectional variation between 2000 prime-age labor market participation rates and early opioid prescription rates. Figure 11 offers some immediate conclusions on the association between labor market status in 2000 and prescription rates. First, there are both high and low prescription rate places in at least the first four quintiles of 2000 participation rates. So we would have to be concerned if our identification of prescription effects was overly reliant on the fifth quintile. Second, a simple regression line (shown in red) shows a near-zero upward slope of prescription rates to improving participation rates, with a low correlation of 0.05. Third, the remarkable range in prime-age participation rates (from 56.9 to 89.9 percent) is largely a result of weaker performing places, as 60 percent of coumas have participation rates between 76.1 and 83.2 percent. Overall, there does seem to be substantial variation in prescription rates across places that is not associated with the relative performance of their labor markets prior to the opioid crisis.

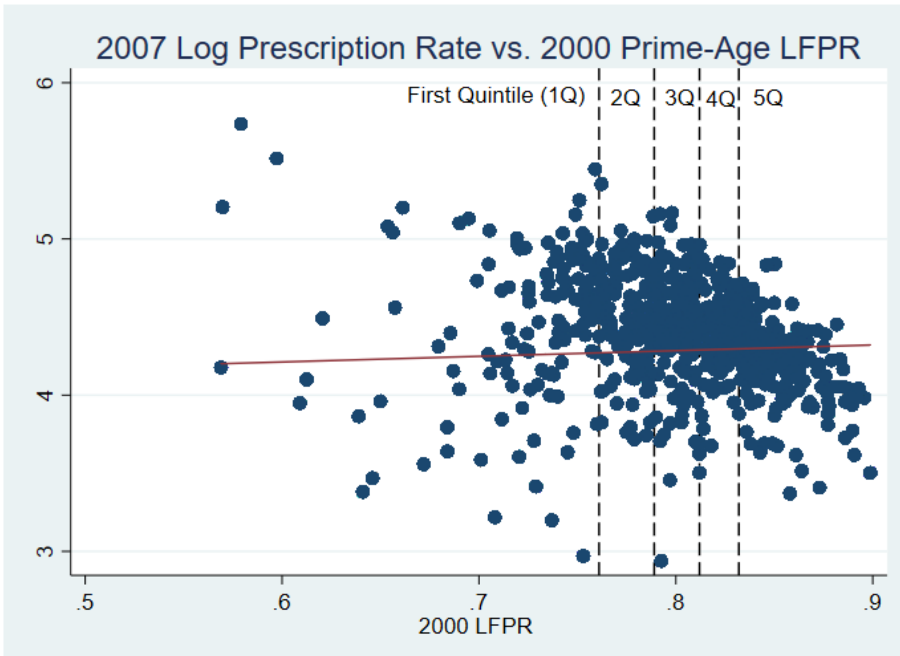

Figure 11: Association between 2000 LFPR and 2007 Prescription Rates

A more subtle alternative explanation for a noncausal interpretation of our results is that there might be 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 LEC_j and high prescription rates P_j.

Throughout our analysis we have directly controlled for LECs (LEC_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 prior economic conditions of the place.

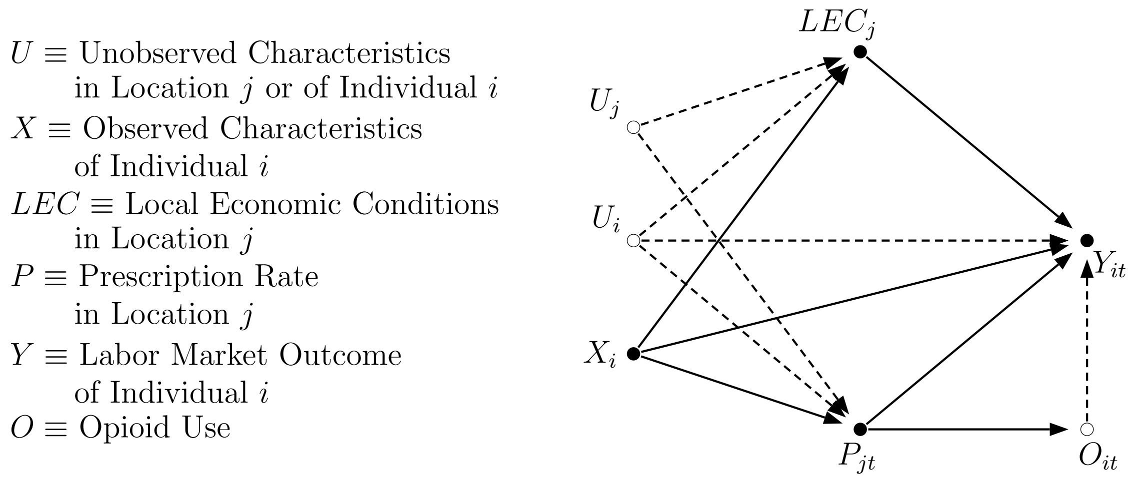

Figure 12: Directed Acyclic Graph of Opioids Affecting Labor Force Participation
Note: This figure follows the convention from Pearl (2009) of communicating that a variable is observed by drawing a solid line to its descendants and communicating that a variable is unobserved by drawing a dashed line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

Specifically, we estimate the coefficients on prescription rates separately for areas depending on their position in the distribution of initial labor market conditions, LEC_{j0}, as measured by the labor force participation rate in 2000. The specifications are unchanged from those described in Section 3, but constrained to only include observations in a given quintile. 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 near zero coefficient estimates within the quintiles 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 the estimated coefficients should be similar to the panel results shown in Tables 2 and 4. For this comparison we again limit ourselves to labor force participation.

Figure 13 shows the results, which provide evidence against the selection hypothesis. Recall that our preferred specification using specific controls relies on the 2000 labor market participation as its key control variable, so estimating it within quintiles of the same variable has the potential to weaken key identifying restrictions on the effects of prior labor market information on the current value. Still, the men's results all remain statistically significant and suggest relatively large impacts, even if the better preforming quintiles are estimated to have smaller estimated coefficients. The women's estimates for the specific controls, in contrast, show that only the first and fifth quintile samples produce statistically significant

Figure 13: Coefficeint Estimates by Quintile of 2000 Labor Market Conditions

estimates. The estimated effects for women are always smaller, so these weaker results are not surprising given the reduced precision of these estimates. We interpret the more uniform results of the difference-in-difference specification for men as evidence that the specific controls specification is able to control for important short-term cyclical variation in labor market conditions.

Overall, we think the fact that these coefficients can generally be estimated within quintiles of labor market performance argues against versions of the reverse causality story that rely on unobserved factors persistently affecting both prescription rates and local labor market.

6 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. It is easy to say what the difference in labor force participation rate is expected to be for a high opioid prescription area relative to a low prescription area. As a rough measure, the difference between the 10th and 90th percentiles is roughly 1 log point. This implies that the coefficients can be directly interpreted as the expected impact of being in a high-opioid-prescription-rate area.

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. While our estimates are clearly substantial values, particularly for men, direct comparison of coefficients with Krueger (2017) is made difficult by 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 three-year windows). One approach to comparison is to use the implied aggregate effects over time. As Equation 4 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, with the model uncertainty partially revealed by comparing the range of estimates. 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.

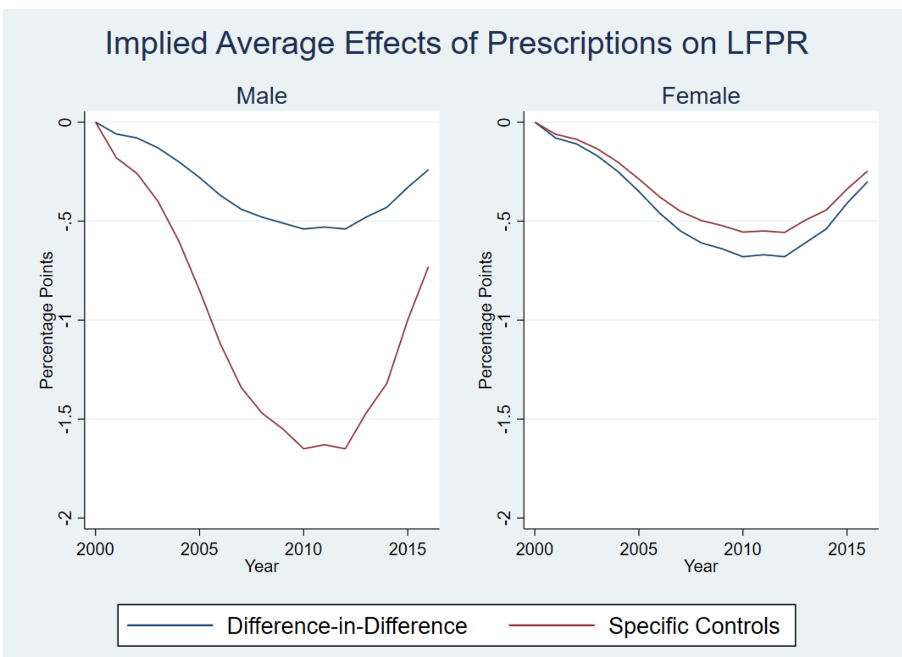

Figure 14: National Labor Market Effects Estimates

Figure 14 shows the implied average effects for male and female labor force participation for each of our estimated models.[16] Model uncertainty is present in the male estimates, but both of these effects estimates would have to to be viewed as economically important. The difference-in-difference results still suggest a peak reduction in the prime-age male labor participation rate of a half percentage point. The estimates based on specific controls imply a more than 1.5 percentage point decline in participation of prime-age men. The women's estimates are all clustered around a peak effect of 0.5 percentage points.

We can use these results to compare the expected decline in participation rates implied by our estimates with the declines estimated in Krueger (2017). Krueger (2017)'s results imply a decline of 0.6 percentage points in participation for men between 1999 and 2015 and a decline of 0.8 percentage points for women over the same period. Our results generally imply a larger decline for prime-age men and a smaller decline for prime-age women, with our preferred specific-controls model showing substantially larger effects for men. For men we find that the decline in the labor force participation rate associated with rising opioid prescriptions is between 0.6 and 1.7 percentage points between 2000 and 2016 (depending on the model specification), but that the decline is between 0.5 and 0.7 percentage points for women. In our estimation, then, opioid prescriptions account for between 8 and 24 percent of the realized decline in the prime-age male participation rate and 10 to 13 percent for prime-age women from 2000 to 2016. If instead of measuring the effects in 2016 after a considerable rebound associated with lower prescription rates we compared the period from 2000 to 2010, opioid prescriptions would account for between 22 and 69 percent of the decline in prime-age male labor force participation and between 36 and 46 percent of the decline in prime-age female labor force participation.

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

[16]Estimates could be calculated for other labor market statuses, including the employment to population ratio. As our discussion of Table 2 suggested, there is little difference between the implied time patterns between the labor force participation rate and the employment-to-population ratio.

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. 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. 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.

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 do not include many viable candidate areas, although prescription rates have gradually declined in most areas of the United States since 2010.[17]

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, Abby E., William N. Evans, Ethan M.J. Lieber, and David Powell (2019). "Origins of the opioid crisis and its enduring impacts." *NBER Working Paper 26500*. doi:10.3386/w26500.

Angrist, Joshua D. and Jörn-Steffen Pischke (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Barnett, Michael L., Andrew R. Olenski, and Anupam B. Jena (2017). "Opioid-prescribing patterns of emergency physicians and risk of long-term use." *New England Journal of Medicine*, 376(7), pp. 663–673. doi:10.1056/NEJMsa1610524.

[17]Two CPUMAs (the 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 unknown.

Bartik, Timothy J. (1991). *Who Benefits from State and Local Economic Development Policies?* WE Upjohn Institute for Employment Research. doi:10.17848/9780585223940.

Blanchard, Olivier and Lawrence Katz (1992). "Regional evolutions." *Brookings Papers on Economic Activity*, 1, pp. 1–76. doi:10.2307/2534556.

Buchmueller, Thomas C. and Colleen Carey (2018). "The effect of prescription drug monitoring programs on opioid utilization in Medicare." *American Economic Journal: Economic Policy*, 10(1), pp. 77–112. doi:10.1257/pol.20160094.

Carpenter, Christopher S., Chandler B. McClellan, and Daniel I. Rees (2017). "Economic conditions, illicit drug use, and substance use disorders in the United States." *Journal of Health Economics*, 52, pp. 63–73. doi:10.1016/j.jhealeco.2016.12.009.

Case, Anne and Angus 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), pp. 15,078–15,083. doi:10.1073/pnas.1518393112.

Case, Anne and Angus Deaton (2017). "Mortality and Morbidity in the 21st century." *Brookings Papers on Economic Activity*, Spring, p. 397. doi:10.1353/eca.2017.0005.

Charles, Kerwin Kofi, Erik Hurst, and Mariel Schwartz (2018). "The transformation of manufacturing and the decline in U.S. employment." *NBER WP 24468*. doi:10.3386/w24468.

Cicero, Theodore J., Matthew S. Ellis, Hilary L. Surratt, and Steven 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), pp. 821–826. doi:10.1001/jamapsychiatry.2014.366.

Compton, Wilson M., Christopher M. Jones, and Grant T. Baldwin (2016). "Relationship between nonmedical prescription-opioid use and heroin use." *New England Journal of Medicine*, 374(2), pp. 154–163. doi:10.1056/NEJMra1508490.

Currie, Janet, Jonas Jin, and Molly Schnell (2019). "US employment and opioids: Is there a connection?" In *Health and Labor Markets*, pp. 253–280. Emerald Publishing Limited.

Currie, Janet and Molly Schnell (2018). "Addressing the opioid epidemic: Is there a role for physician education?" *American Journal of Health Economics*. doi:10.1162/ajhe_a_00113. Forthcoming.

Dasgupta, Nabarun, Leo Beletsky, and Daniel Ciccarone (2018). "Opioid crisis: No easy fix to its social and economic determinants." *American Journal of Public Health*, 108(2), pp. 182–186. doi:10.2105/AJPH.2017.304187.

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

Evans, William N., Ethan Lieber, and Patrick Power (2019). "How the reformulation of OxyContin ignited the heroin epidemic." *The Review of Economics and Statistics*, 101(1), pp. 1–15. doi:10.1162/rest_a_00755.

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

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

Glei, Dana A. and Maxine Weinstein (2019). "Drug and alcohol abuse: The role of economic insecurity." *American Journal of Health Behavior*, 43(4), pp. 838–853. doi:10.5993/AJHB.43.4.16.

Harris, Matthew C., Lawrence M. Kesslery, Matthew N. Murray, and M. Elizabeth Glenn (2019). "Prescription opioids and labor market pains: The effect of schedule II opioids on labor force participation and unemployment." *Journal of Human Resources*. doi:10.3368/jhr.55.4.1017-9093R2.

Hollingsworth, Alex, Christopher J. Ruhm, and Kosali Simon (2017). "Macroeconomic conditions and opioid abuse." *Journal of Health Economics*, 56, pp. 222–233. doi:10.3386/w23192.

Horwitz, Jill, Corey S. Davis, Lynn S. McClelland, Rebecca S. Fordon, and Ellen Meara (2018). "The problem of data quality in analyses of opioid regulation: The case of prescription drug monitoring programs." *NBER WP 24947*. doi:10.3386/w24947.

Imbens, Guido W. (2015). "Matching methods in practice: Three examples." *Journal of Human Resources*, 50(2), pp. 373–419. doi:10.3368/jhr.50.2.373.

Imbens, Guido W. and Jeffrey M. Wooldridge (2009). "Recent developments in the econometrics of program evaluation." *Journal of Economic Literature*, 47(1), pp. 5–86. doi:10.1257/jel.47.1.5.

Isserman, Andrew M. and James Westervelt (2006). "1.5 Million missing numbers: Overcoming employment suppression in County Business Patterns data." *International Regional Science Review*, 29(3), pp. 311–335. doi:10.1177/0160017606290359.

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

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

Kromer, Bracdyn K. and David J. Howard (2011). "Comparison of ACS and CPS data on employment status." *US Census Bureau WP SEHSD-WP2011-31*.

Krueger, Alan 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*. doi:10.1353/eca.2017.0012.

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

Maggio, Marco Di and Amir Kermani (2016). "The importance of unemployment insurance as an automatic stabilizer." *NBER WP 22625*. doi:10.3386/w22625.

Mazumder, Bhashkar (2005). "Fortunate sons: New estimates of intergenerational mobility in the United States using social security earnings data." *Review of Economics and Statistics*, 87(2), pp. 235–255. doi:10.1162/0034653053970249.

Pearl, Judea (2009). *Causality: Models, Reasoning and Inference*. Cambridge University Press, second edition. doi:10.1017/CBO9780511803161.

Powell, David, Rosalie Liccardo Pacula, and Erin Taylor (2015). "How increasing medical access to opioids contributes to the opioid epidemic: Evidence from Medicare Part D." *NBER WP 21072*. doi:10.3386/w21072.

Powell, Jerome (2019). *Federal Reserve's Second Monetary Policy Report for 2019*. US Senate Committee on Banking, Housing, and Urban Affairs, Washington, DC.

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

Ruhm, Christopher J. (2015). "Recessions, healthy no more?" *Journal of Health Economics*, 42, pp. 17–28. doi:10.1016/j.jhealeco.2015.03.004.

Ruhm, Christopher J (2016). “Taking the measure of a fatal drug epidemic.” *NBER WP 22504*. doi:10.3386/w22504.

Ruhm, Christopher J. (2018). “Deaths of despair or drug problems?” *NBER WP 24188*. doi:10.3386/w24188.

Ruhm, Christopher J. (2019). “Drivers of the fatal drug epidemic.” *Journal of Health Economics*, 64, pp. 25–42. doi:10.1016/j.jhealeco.2019.01.001.

Savych, Bogdan, David Neumark, and R. Lea (2018). “The impact of opioid prescriptions on duration of temporary disability.”

Schweitzer, Mark E. (2017). “Manufacturing employment losses and the economic performance of the Industrial Heartland.” *FRB of Cleveland WP 17-12*. doi:10.26509/frbc-wp-201712.

Shavers, Vickie L., Alexis Bakos, and Vanessa B. Sheppard (2010). “Race, ethnicity, and pain among the US adult population.” *Journal of Health Care for the Poor and Underserved*, 21(1), pp. 177–220. doi:10.1353/hpu.0.0255.

Yellen, Janet L. (2017). *Federal Reserve's Second Monetary Policy Report for 2017*. US Senate Committee on Banking, Housing, and Urban Affairs, Washington, DC, s. hrg. 115-108 edition.

Zoorob, Michael J. and Jason L. Salemi (2017). “Bowling alone, dying together: The role of social capital in mitigating the drug overdose epidemic in the United States.” *Drug and Alcohol Dependence*, 173, pp. 1–9. doi:10.1016/j.drugalcdep.2016.12.011.

A Difference-in-Differences

A.1 Specification

Recall that our baseline specification is

$$Y_{ijt} = \alpha P_{jt} + f(LEC_{jt}) + \beta X_{it} + \epsilon_{it}. \tag{5}$$

In addition to the approach of using specific variables to control for local economic conditions, LEC, we also use the panel structure of our data along with linearity and separability assumptions to account for the LEC term using fixed effects. In this case,

$$f(LEC_{jt}) = \gamma_j + \delta_t. \tag{6}$$

With both a time and geographic fixed effect, this creates a difference-in-differences specification where the parameter α is identified based on how the difference-in-differences in labor market outcomes across locations relate to their difference-in-differences in prescription rates.[18] This model is well identified when t is large so that the measured differences in LECs across time can be substantial.

$$f(LEC_{jt}) = \eta D_{jt} + \theta \widehat{Y}_{gt} + \gamma_k + \delta_t. \tag{7}$$

This specification has the advantages of Equation 3, but allows for historically weak/strong labor markets to have better/worse labor market developments over the sample period.

A.2 Estimates

Table 4 shows the results of difference-in-difference regressions for prime-age men and women using the specification that is widely used in the literature. For these estimates we rely on year fixed effects to absorb the business cycle and other general time patterns in participation. The couma-level fixed effects pick up the average local differences in the period, leaving the coefficient on prescription rates to be identified by the time variation in localities.

For both prime-age men and women, the number of opioid prescriptions in their geographic area is associated with a lower probability of labor force participation and a lower employment rate with a high level of statistical significance. The results for the employment-to-population ratio and labor force participation are similar to the participation effects. Sta-

[18] Related discussions can be found in Imbens and Wooldridge (2009) and Angrist and Pischke (2009).

tistically weak unemployment effects (that suggest increased employment) combined with statistically stronger participation effects indicate that opioid prescription levels appear to primarily affect labor markets through the individual's labor market participation decision. These estimates, which end up representing the lower bound on the magnitude of our estimated effects, are statistically significant and indicate that opioid prescription rates have economically relevant impacts on labor market outcomes. The difference in predicted participation rates between high and low prescription areas (about 1 log point) is a 1.5 percentage point difference for men and 1.9 for women. Given the results of Krueger (2017), one surprising result is that estimates are lower in magnitude for prime-age men than for women.

Table 4: Labor Market States of Prime-Age Men and Women, Difference-in-Differences

	Men			Women		
	Participate	Emp/Pop	Unem/Pop	Participate	Emp/Pop	Unem/Pop
Prescrip. Rate	-0.015***	-0.010*	-0.005	-0.019***	-0.015***	-0.004
	(0.003)	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)
Age	0.067***	0.090***	-0.023**	0.080***	0.084***	-0.004
	(0.011)	(0.013)	(0.008)	(0.014)	(0.015)	(0.007)
Age squared	-0.278***	-0.352***	0.073*	-0.397***	-0.405***	0.008
	(0.042)	(0.049)	(0.031)	(0.055)	(0.058)	(0.027)
Age cubed	0.051***	0.061***	-0.010	0.084***	0.084***	-0.001
	(0.007)	(0.008)	(0.005)	(0.009)	(0.010)	(0.005)
Age to 4th	-0.004***	-0.004***	0.001	-0.006***	-0.006***	0.000
	(0.000)	(0.001)	(0.000)	(0.001)	(0.001)	(0.000)
Less than HS	-0.190***	-0.239***	0.049***	-0.320***	-0.363***	0.044***
	(0.002)	(0.003)	(0.001)	(0.002)	(0.002)	(0.001)
High School	-0.078***	-0.115***	0.037***	-0.130***	-0.161***	0.030***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)
Some College	-0.034***	-0.053***	0.019***	-0.053***	-0.072***	0.018***
	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.000)
White	0.025***	0.028***	-0.003***	0.025***	0.031***	-0.006***
	(0.001)	(0.001)	(0.000)	(0.001)	(0.002)	(0.000)
Black	-0.069***	-0.096***	0.027***	0.051***	0.026***	0.024***
	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)	(0.001)
Hispanic	0.055***	0.067***	-0.011***	0.017***	0.016***	0.001**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Married	0.115***	0.151***	-0.036***	-0.080***	-0.058***	-0.023***
	(0.001)	(0.001)	(0.000)	(0.001)	(0.002)	(0.000)
R-sqr	0.10	0.12	0.02	0.06	0.06	0.02
N	6440267	6440267	6440267	6655308	6655308	6655308

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

The other coefficient estimates shown in Table 4 indicate, at least when the model is estimated in this multilevel approach, that demographic factors should be accounted for. If the regression were estimated as a panel of geographic averages, many of these demographic differences would likely average out, but the highly significant coefficients raise the possibility that sampling variation by demographic group within small geographic areas could make the estimation less efficient. Demographic controls are included in all of our analyses, but to aid readability of the tables, we exclude these coefficients in the remaining tables.

While the difference-in-differences estimator controls for a wide variety of potential confounding factors, it does narrow the range of variation used to identify effects. In addition, we use a limited sample of years and both labor force status and drug use are known to be persistent. These problems suggest looking for additional identifying variation from the cross-section. Table 5 shows results for the same specification, Equation 4, for labor force participation regressions, with no additional controls, when the fixed effects are at a higher level of aggregation. When at the state level, this allows differences in prescription rates within the state to identify effects. At the Census division level, cross-state differences are also generally relevant.

Table 5: Participation of Prime-Age Men and Women, Varying Fixed Effects

	Men			Women		
	Census Division	State	Couma	Census Division	State	Couma
Prescrip. Rate	-0.041***	-0.050***	-0.015***	-0.010***	-0.006**	-0.019***
	(0.002)	(0.003)	(0.003)	(0.003)	(0.002)	(0.004)
R-sqr	0.09	0.09	0.10	0.06	0.06	0.06
N	6440267	6440267	6440267	6655308	6655308	6655308

All regressions include demographic variables and year fixed effects.
Robust standard errors with clustering on couma x year.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5 indicates that the results are sensitive to the level of fixed effects used in the regression. Notably, the coefficient estimates become more negative for men with more aggregated fixed effects and fall somewhat for women. While it could be argued that state legal codes are a good reason to have fixed effects at the state level, we treat Table 5 primarily as evidence that pursuing alternative control group strategies could yield substantially different results.

A.3 Non-nested Model Comparisons with Akaike Information

The two models we have estimated are tightly related, but they are not formally nested. We use Akaike information to measure the role of the predictive factors in the alternative models of labor force participation. Specifically, we consider the incremental gain in Akaike information stemming from the addition of groups of controls:

$$\Delta AIC(X_2, X_1) = -2\ln(L|X_2) + k_2 - (2\ln(L|X_1) + k_1), \tag{8}$$

where $\ln(L|X_2)$ represents the log likelihood of model 2 with X_2 data in the model and k_2 number of parameters compared to nested model 1, which must have fewer parameters. This formula is designed for nested models but it also yields informative comparisons across models when the included data sample is the same.

Clearly this measure is order dependent. In each case, we remove the demographic variation before the figure is plotted, because in our multilevel estimation strategy the demographic variables are large explanatory variables at the individual level, even though they are not the focus of the analysis. We then add time-varying factors (specific and then general) and then local labor market factors (specific and then general from most aggregated to least), before finally adding the prescription rate variable. This in no way invalidates or alters the significance tests in the individual regressions, but it can help clarify factors identifying variation in the models.

Figure 15 shows the analysis of the two models for female labor force participation. Starting with the difference-in-differences model, it is clearly evident that the variation picked up by the year fixed effects is far smaller than the cross-sectional variation identified by Census divisions, states and coumas. This is despite the fact that the sample period included a large recession and recovery. In the other model, this variation is largely picked up by the combination of the Bartik demand shock and the year fixed effects, although that combination accounts for more Akaike information because it includes some cross-sectional variation.

The limited impact of time variation leads us to focus on a specification that can identify more-specific forms of cross-sectional variation. While the 2000 prime-age labor force participation rate is just a single added variable in the specific controls model, it still absorbs a substantial fraction of the variance accounted for by the far more numerous fixed effects in the difference-in-difference model. Allowing the labor force participation rate to evolve with a related comparison group expands the information increment associated with the local labor market factor, but other patterns are similar. In our view this is a strong indication that pre-opioid labor market performance is a critical control that explains much of the persistence in differences between places.

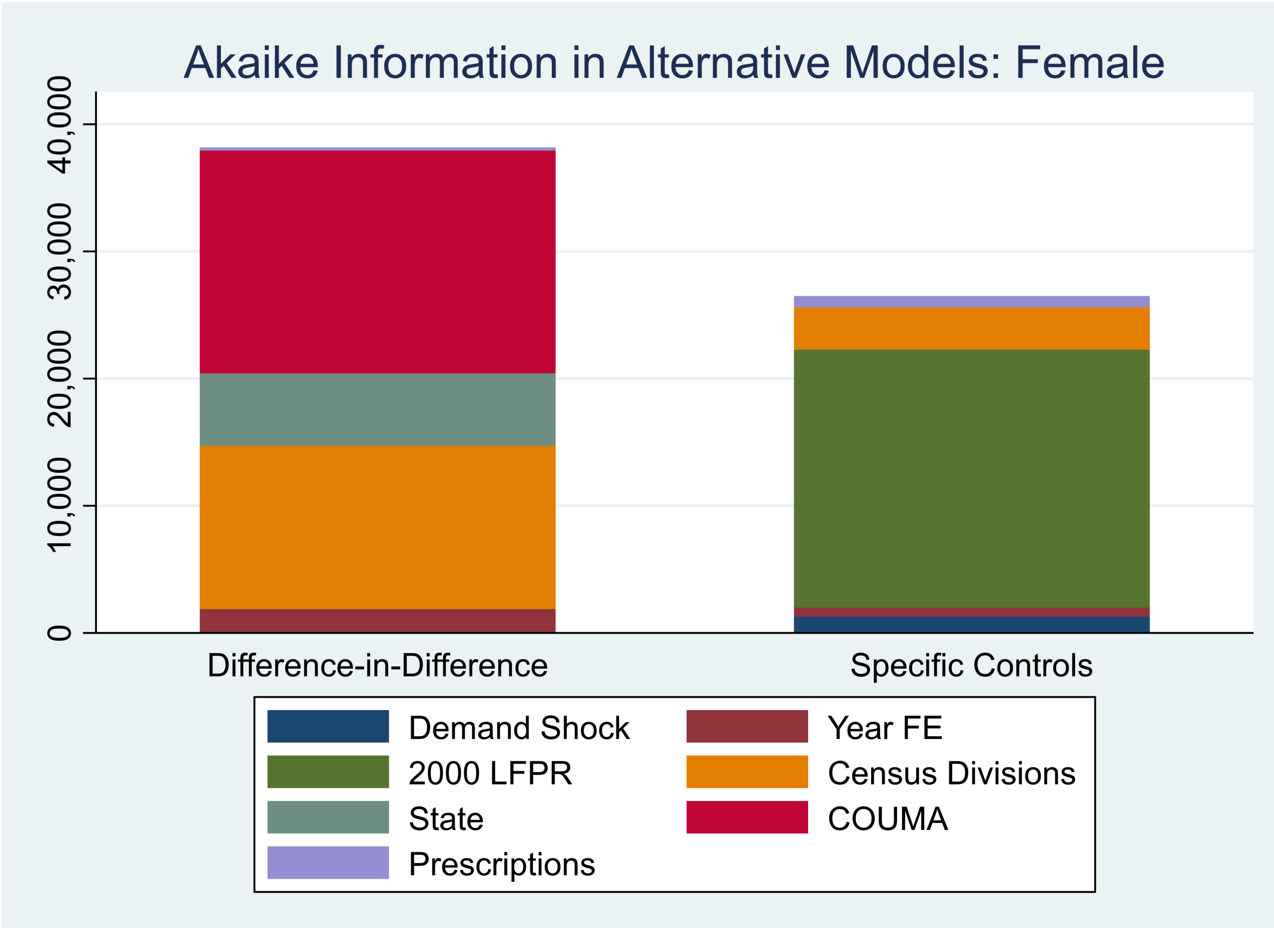

Figure 15: Role of Explanatory Variables in the Alternative Models

Prescriptions are added last and the explained variation is more limited due the inclusion of prior variables, but the increment in the AIC still varies between the models shown in Figure 15. The identified coefficients are quite similar for women across the two models: ranging form -0.015 to -0.019. Despite this agreement in the coefficient estimates, the AIC value for the difference-in-differences model is particularly small, reflecting the limited variation in prescription rates across coumas.

It is reassuring that the identification restrictions of the difference-in-difference model still imply a similar coefficient for women, but the other model's identification restrictions seem to absorb important differences in coumas without absorbing all of the differences picked up by couma-level fixed effects. From a simple Akaike information perspective, the difference-in -differences model would be preferred, but the other models offer more clarity on the specific comparisons that are being made while still explaining significant sources of confounding variation.

Figure 16 shows that the same qualitative features are evident in the same decomposition for males.

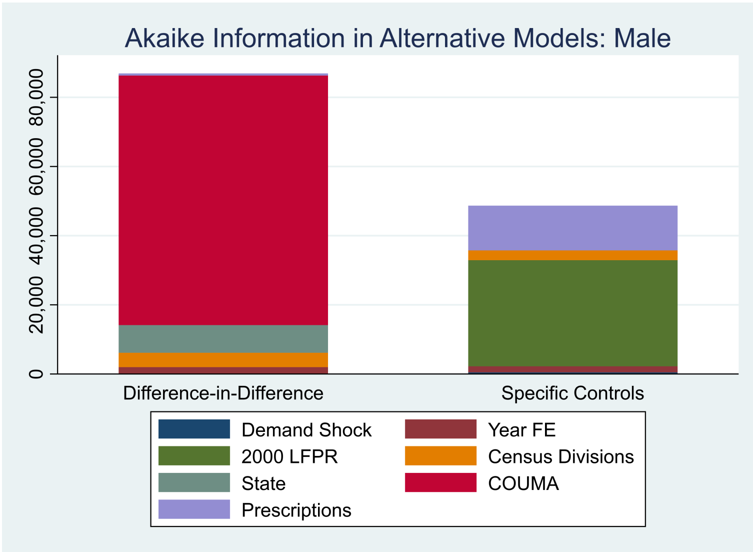

Figure 16: Role of Explanatory Variables in the Alternative Models

B Time Effects and the Implications of Increasing Illegal Opioid Usage after 2010

OxyContin was reformulated in 2010 to make it less prone to abuse. Evans et al. (2019) and Alpert et al. (2019) have documented that this reformulation was associated with a rise in heroin deaths that was attributed to consumers substituting heroin as an inexpensive alternative to OxyContin. In addition, mortality data show 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 our panel regression models all 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 and any national turn from opioid prescriptions to illegal opioids. 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 whether coefficients have shown a time pattern, we run a variant of the specific control model as a sequence of cross-sectional regressions.[19]

$$Y_{ijt} = \alpha P_{jt} + f(LEC_{jt}) + \beta X_{it} + \epsilon_i, \qquad \text{for each possible year } t \tag{9}$$

Of course, $f(LEC_{jt})$ can no longer include fixed effects on years, which eliminates the difference-in-differences model.

Applying a panel regression as in Table 2 does increase the available variation for identification while imposing constant coefficients over time. Figure 17 shows the results of these sequential regressions for prime-age male labor force participation.

[19] In our view, it does not make sense to consider the difference-in-difference model because its only cyclical controls are the time effects.

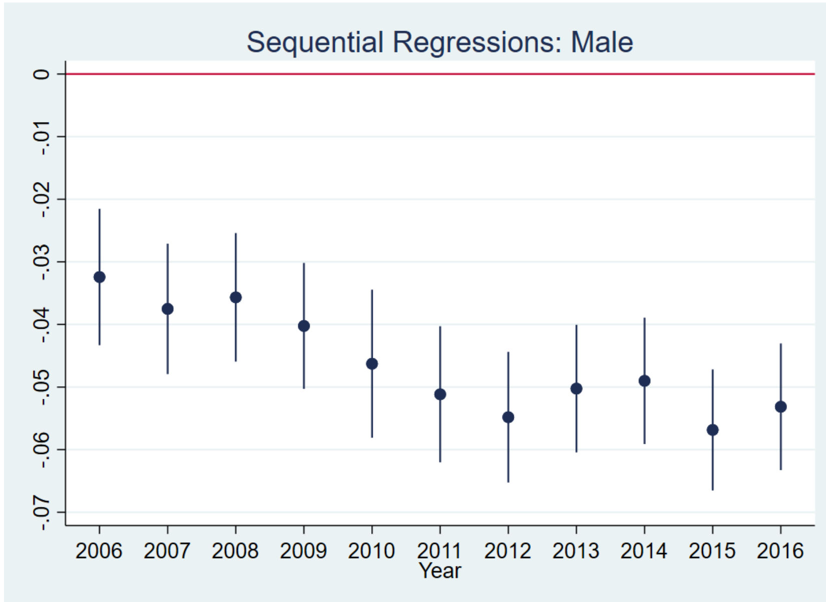

Figure 17: Coefficient Estimates from Sequential Models: Male

First, it is worth noting that the coefficient estimates of the sequential models are on average very similar to the panel model results in Table 2, but with larger confidence intervals. Next, it is immediately evident that cross-sectional variation in the post-2010 years for labor market performance continues to be strongly correlated with prescription rates. This could reflect the tendency for illegal opioid problems to be higher in 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; see Ruhm (2016).

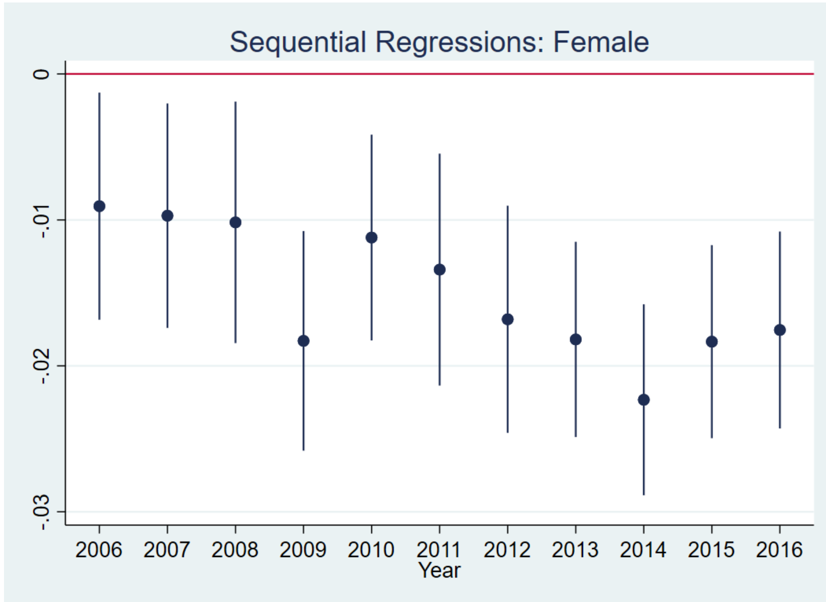

Figure 18: Coefficient Estimates from Sequential Models: Female

The same analysis is repeated for women in Figure 18 with both of the conclusions from the men's analysis repeated: Similar coefficient estimates with wider confidence bands and effects get modestly stronger in later years. The one additional conclusion that can be drawn from this is that there may be some cyclical effects showing up in the sequential results, notably in the larger negative coefficient for 2009. This is the worst year for labor force participation, and female participation is more cyclical than male participation. This is another reason to favor the panel estimates, even though this approach does not allow for an easy way of assessing whether the estimated effects of opioid prescriptions are declining in recent years. For both sexes and both models the expected effect of opioid prescriptions is modestly larger after 2010.

C Comparing Our Results to Key Parts of the Literature

C.1 Consistency with Harris, Kessler, Murray, and Glenn

Harris et al. (2019) use county-level regressions on data from 10 states with prescription drug monitoring programs (PDMPs) primarily between 2013 and 2015.[20] They use the plausibly exogenous variation in the concentration of high-volume prescribers, which are not available for most of the US, as an instrument to achieve causal identification. 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 4 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) show results for the 10 states in their sample that are considerably larger than our estimates for the entire US. Focusing on their labor force participation results, they have a coefficient of -0.065 for all adults in their preferred IV specification compared to our largest estimates (from the specific controls model) of -0.046 for prime-age men and -0.015 for prime-age women. A one standard deviation shift in prescriptions in Harris et al. (2019) produces a 2.7 percentage point decline in the labor force participation rate. Using our estimates, a one standard deviation increase in log prescriptions produces a 1.8 percentage point decline for prime-age men and a 0.6 percentage point decline for prime-age women. Taking the weighted average of male and female results in a combined decline of 1.2 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) groups 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. While we cannot identify the relevant margins that produce the differences in the estimate effects, we view Harris et al. (2019) as 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

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

prescriptions on prime-age adults.

C.2 Sources of Differences with Currie, Jin, and Schnell

Currie et al. (2019) examine directions of causation between quarterly prescription data that are specific to age-sex groups and quarterly employment to population ratios from the Quarterly Workforce Indicators (QWI) for the same age-sex groups. They find that opioid prescriptions have a positive and significant effect on the employment-to-population ratio for women and no effect for men. They find ambiguous effects when they explore the effects of employment on opioid prescriptions rates. Given the substantive differences in their estimates and ours, we undertook a more thorough comparison of the two sets of results, even though we cannot directly estimate their model because we do not have access to the prescription data used in their analysis. There are key differences in geographic groupings, age-sex groupings, the available controls, and regression specifications. 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 specific demographic controls available in the ACS. Currie et al. (2019) 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. (2019) use models that rely on fixed effects rather than controlling for individual-level demographics and area-level economic conditions that can be time-varying.

While our methods and results differ from those in Currie et al. (2019) in several ways that cannot be fully compared, differences in our results by model specification (particularly Table 4) can provide some suggestive evidence on key differences. Importantly, the employment to population ratio is quite cyclical and, as we noted earlier, cyclical impacts were uneven across the country during the Great Recession and its aftermath. Our estimates for employment to population ratios, particularly for men, are closer to zero in the difference-in-differences specification, where the included fixed effects may be absorbing relevant variation. As this paper has tried to emphasize, the specific nature of the controls for both cyclical and longer-term local labor market outcomes is critical for estimating this relationship. This is further complicated in Currie et al. (2019) because the QWI-based employment to population ratio includes commuters in the numerator but not the denominator, which may exaggerate the urban and suburban cyclical changes. Having examined the three possible labor market outcomes, we find that most of the action in the result is in the labor market participation equation, which is better measured using prescriptions and outcomes at the individual's residence rather than his/her place of work. Currie and Schnell (2018) find important non-economic factors that drive prescription rates, which is consistent with the patterns that we observe for prescription rates, which underlies our analysis of causality. Overall, we would emphasize in the comparison of results that our labor force participation estimates are economically substantial and statistically robust to a range of model specifications. In addition, our results are consistent with demographic patterns evident in mortality statistics (Ruhm (2019)).