Federal Reserve Bank of Cleveland Working Paper Series

Opioids and the Labor Market

Dionissi Aliprantis, Kyle D. Fee, and Mark E. Schweitzer

Working Paper No. 18-07R3

July 2022

Suggested citation: Aliprantis, Dionissi, Kyle D. Fee, and Mark E. Schweitzer. 2022. "Opioids and the Labor Market." Working Paper No. 18-07R3. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-201807r3.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Opioids and the Labor Market

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

July 12, 2022

Abstract: This paper quantifies 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 both to 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 0.53 percentage points for men and 0.10 percentage points for women. We focus on measuring the impact of opioid prescriptions on labor markets, so we evaluate the robustness of our estimates to an alternative causal path, unobserved selection, and an instrumental variable from the literature.

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 OxyContin, heroin, and fentanyl.[1]

The interaction between opioid prescriptions 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. Krueger (2017) instigated research on this topic by showing that areas with higher opioid prescription rates have lower labor force participation rates. Newer research using quasi-random variation in local prescription rates has also found a decrease in labor force participation (Harris et al. (2019), Laird and Nielsen (2016), Deiana and Giua (2018), and Beheshti (2022)). This research is helpful in assessing the causality of opioids but necessarily focuses on deviations around a particular policy or geography.[2] Abraham and Kearney (2018) summarize the literature on opioids and the labor market as inconclusive based on the arrows of causality potentially running in both directions.

This paper focuses on understanding the scale of the national impacts of prescription opioids on US labor markets. We begin by more rigorously estimating the strong correlation between higher opioid prescription rates and labor market outcomes for prime-age adults and key sub-groups. 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)). Whether opioids caused an individual to leave the labor market or complicated a return to employment, our results indicate that areas with high opioid prescription rates experience weaker labor market outcomes, even after after controlling for other employment shocks and the persistent component of weak or strong labor markets.

Our primary enhancements to the understanding of the general impacts of prescription opioids on US labor markets relative to Krueger (2017) result from (i) improving the joint measurement of labor market outcomes and local prescription rates and (ii) improving the quality of controls for short- and long-term local labor market conditions. We assemble a nationally representative data set that includes both individual labor market outcomes and local opioid prescription rates. Creating such a geographically linked data set with

[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.

[2]Results are also not uniform: Currie et al. (2019) find that higher opioid prescription rates had a small positive effect on employment-to-population ratios for women.

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. For example, individual-level labor market data from the Current Population Survey or the American Community Survey (ACS) do not report a respondent's county of residence when the county does not meet a minimum population threshold.

Our approach improves measurement precisely in the low-population areas most affected by the opioid crisis. Our approach is to measure an individual's local area in the ACS as his/her couma, which is his/her county when identified, or else as his/her Public Use Microdata Area (PUMA) for the 29 percent of the US population residing in nonidentified counties.[3] When PUMAs are a combination of adjacent counties, they are designed to have at least 100,000 residents and are the most finely defined geography identified in the ACS along with annual individual-level labor market outcomes. We aggregate the Centers for Disease Control Prevention's (CDC) county-level data on opioid prescriptions to coumas to then connect individual-level labor market outcomes in the ACS with local prescription rates. In our data from 2006 to 2016, prescription rates are 32 percent higher in low-population PUMAs than in identified counties. And we show that aggregating data from low-population counties into coumas, rather than states, allows for more accurate measurement of local prescription rates. Thus, our approach to addressing measurement in nonidentified counties is an improvement over previous approaches that either focus entirely on high-population areas, aggregate data from low-population areas into state-level observations, or use county-level employment estimates instead of individual-level labor market data.

We use this data set to estimate regressions characterizing the relationship between local opioid prescription rates and labor market outcomes. A common approach to estimation in this environment is a difference-in-differences specification. This modeling approach relies on only a small fraction of the variation in the data, an approach that has consequences for the efficiency of the estimation in a short panel. In order to assess the sensitivity of the estimation to short panel concerns, we estimate a panel model with specific controls for business cycles and local labor market performance with a range of location fixed effects covering a spectrum from Census division-level to state to complete (the difference-in-differences model). We then use specifications to explore possible causal paths; this includes a careful examination of the potential for geographic selection; more general patterns of selection on unobservables Oster (2019); and an instrumental variable for geographic variation in opioid prescriptions.

We have three main empirical results. First, regardless of the modeling approach, we find that opioid prescription rates and labor force statuses are strongly correlated for both

[3] This geography, which is a mix of counties and PUMAs, is labeled a *couma* in Case and Deaton (2017).

prime-age men and women. We find that an increase in the local opioid prescription rate is associated with a decrease in employment rates for both prime-age men and women, with an accompanying decrease in labor force participation rates and (sometimes) an increase in 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 an interest in how these effects might explain the decline in 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 in the participation rate of prime-age males 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 association with 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 among 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 for 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 apply three approaches to assess the appropriateness of our measurement of the correlation between opioid prescriptions and labor market conditions to a causal interpretation at the level of the labor market. First, to address the role of labor market conditions in causing higher prescription rates, we estimate the model restricting attention to areas with similar labor markets in 2000. There is substantial variation in prescription rates within areas with similar local labor market conditions and the coefficients within similar labor markets are consistent with our primary results. 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. Secondly, we apply methods used in Oster (2019) to test the potential implications of unobserved selection for our estimated coefficients. We find that our preferred estimates would require very high levels of selection on unobservables to explain away our results. Finally, we estimate an instrumental variables specification that can be applied to the full sample using the state-based triplicate prescriptions programs documented by Alpert et al. (2019) as an instrument for prescription rates. This is not as efficient an estimator, but our IV results are consistent with those from our preferred specifications, both overall and for specific demographic groups.

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, we had 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 include 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, a fact that 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, since 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 (Ruggles et al. (2018)). This annual data set includes detailed information for individuals' labor market status, age, race, sex, and education level, but the county of the individual observations is not always identified. About 80 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.[4] 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).

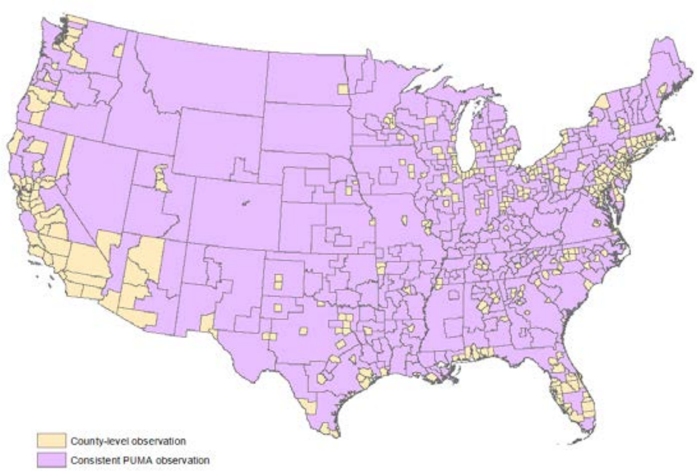

Figure 1: Coumas
Note: Identified counties between 2006 and 2016 are shown in tan, and nonidentified counties (aggregated into CPUMAs) are shown in purple.

One minor tradeoff we make by drawing individual-level data from the ACS is that this

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

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)).

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.[5] 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 (nonhospital) 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 Implications 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.

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

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 the northern parts of Illinois have low prescription rates and the southern parts of Illinois have high prescription rates. In Michigan, the state-level aggregation assigns a median prescription rate to the northern parts of the state. Again we see a contrast with coumas, which assign higher prescription rates to some parts of northern Michigan while assigning the lowest rates to other parts of northern Michigan. In Indiana, Ohio, and Pennsylvania, we find similar patterns of nontrivial within-state measurement bias when state-level averages are used in place of coumas.

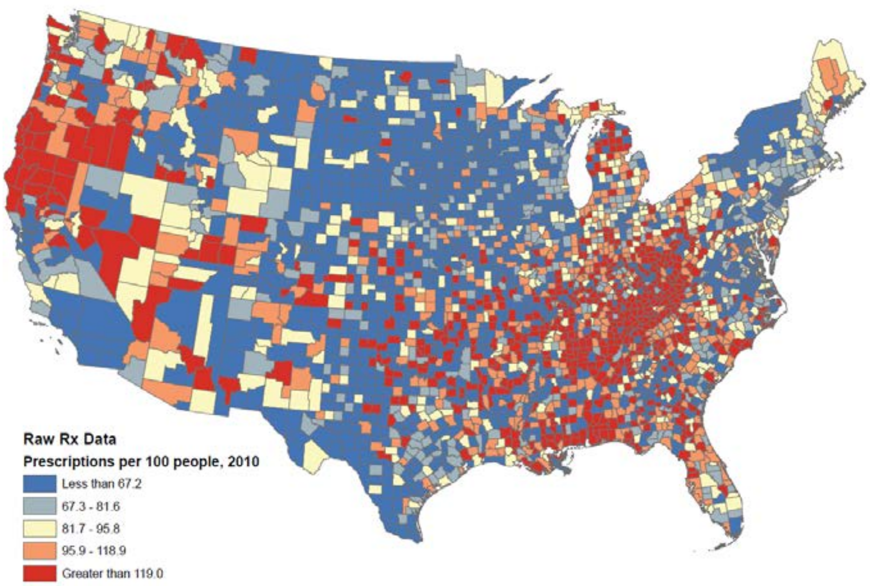

(a) County-Level Data

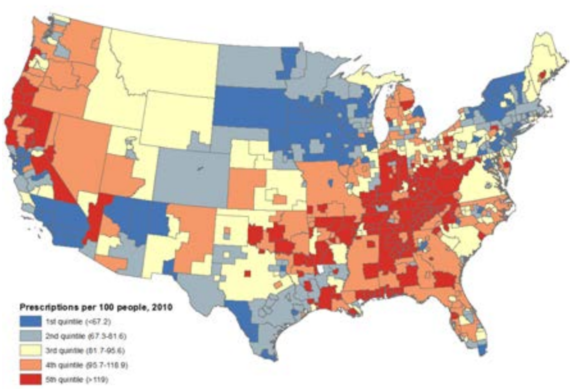

(b) Couma-Based Aggregation

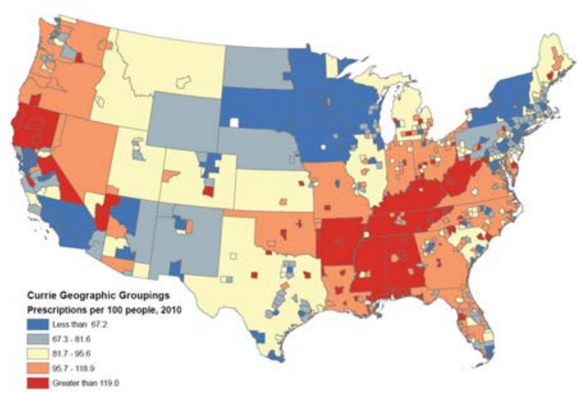

(c) State-Based Aggregation

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.

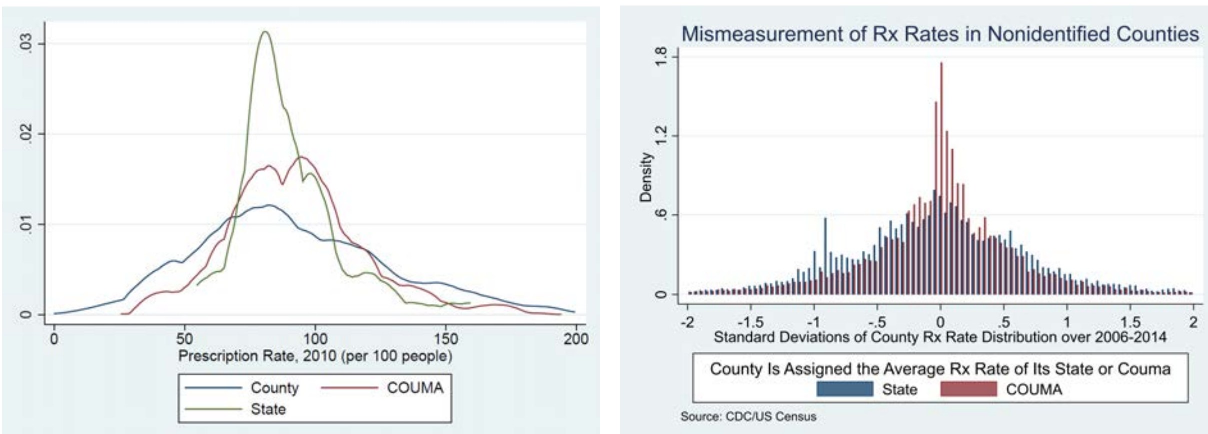

(a) Prescription Rates by Geography

(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 (Bartik et al. (2018)).

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 Di 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.

Goldsmith-Pinkham et al. (2020) show that the Bartik estimator as we have applied it is the equivalent to using industry shares as instruments. For some applications the exogeneity of these instruments is questionable. In our case, both the local shares of the NAICS industries prior to our sample period are likely exogenous to any subsequent labor supply impacts associated with opioid prescriptions, and opioid-induced labor supply changes will be mostly excluded in the national-level industry employment changes. To the extent that some component of industry share reflects the opioid-induced labor supply, our estimates controlling for demand shocks would be overstated and the estimates associated with opioid prescription rates would be understated.

The first panel of Figure 4 shows a scatter plot by year of the Bartik local labor demand shock used in our analysis. Recessions and recoveries involve important dynamic local labor adjustments that are highly uneven geographically. 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 unemployment peak following the recession (2009 and 2010) all areas are experiencing negative local labor demand shocks, although 2009 has the largest range of demand shocks of any year. Some recovery years have relatively tight differences in local labor demand, but 2006 and 2016 stand out for the range of experiences across geographies. This variation in local labor market demand is a key reason why it is better to directly control for this important source of variation rather than just using dummy variables, which would only pick up the means for each year.

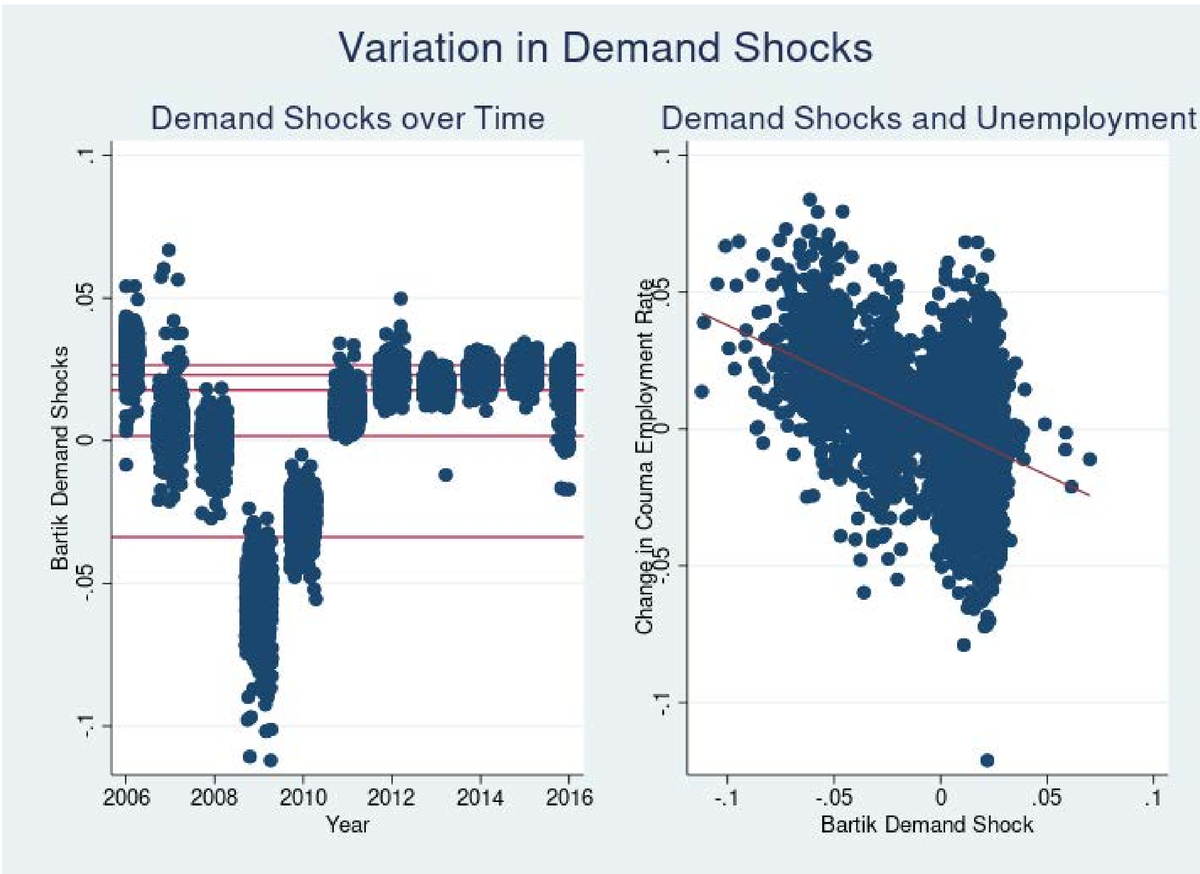

Figure 4: Labor Demand Shocks

The second panel of Figure 4 shows the Bartik demand shock compared to the change in the prime-age unemployment rate in the couma. There is a strong relationship between local labor market outcomes and Bartik demand shocks. Demand shocks explain over 25 percent of the variation in local unemployment rate changes. While the effects on the participation rate are muted, this signals the importance of introducing strong cyclical controls that allow for variation in local effects. Overall, the Bartik demand shock variable should account for local labor market shocks while not picking up local labor market supply conditions that could be impacted by opioid prescriptions.

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 when specifications allow for couma-level controls.[6] The 2000 labor market status should be relatively less impacted by opioid prescriptions.

2.3 Descriptive Statistics

Table 1 presents summary statistics for our sample broken down by whether a couma is an identified county or a CPUMA. The first detail to notice is that low-population counties, which are consolidated into CPUMAs in our analysis, have opioid prescription rates that are nearly a full standard deviation higher than those of 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.

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

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 must decide between searching for quasi-random variation in prescription rates or trying to control for confounding factors that alter local prescription rates and/or individuals' labor market outcomes. The disadvantage of pursuing quasi-random variation is that these shocks are often limited to narrow geographies, populations, or time periods. Since we are interested in understanding the magnitude of the effect of opioids on labor market outcomes on 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 the 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.[7] 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 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

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

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, but excludes any local labor supply changes. It 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 include time effects and a location fixed effect, at a level of aggregation k above the geographic units of observation 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 are 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 in an inherently time-limited panel. The labor market outcomes of interest are cyclical, and, as we noted earlier when discussing Figure 4, cyclical impacts were uneven across the country during the Great Recession and its aftermath. Variation across locations in cyclical impacts argues for having direct controls for the business cycle rather than letting time dummies absorb the average impact.

It is helpful to see that our estimating equation is easily converted to a location-oriented panel equation with estimated terms for prescription rates, local economic demand conditions, the demographic composition of the sample, local persistent effects, and panel controls. 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 when Y is 1/0 in participation, and unemployment-to-population when Y is 1/0 in unemployment) for area j:

$$\frac{1}{N_j}\sum_{i\in j} Y_{ijt} = \alpha P_{jt} + L'_{jt}\beta + \frac{1}{N_j}\sum_{i\in j} X'_{it}\gamma + \theta \overline{Y}_{jt_0} + \delta_t + \frac{1}{N_j}\sum_{i\in j}\epsilon_{it}. \tag{4}$$

If $\theta \overline{Y}_{jt_0}$ were estimated with a fixed effect at the couma level, this follows the standard

FE approach. However, this regression can also be interpreted as a difference-in-differences type of estimator. This alternative specification is useful when we turn back to considering aggregate impacts.

4 Main 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 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.

The specific control specification estimates in Table 2 report the effects on the labor force participation rate, the employment-to-population ratio, and unemployment as a share of the working age population. Participation effects for men are estimated at -0.053, with the effect being approximately five times larger in magnitude for men than for women (-0.053 versus -0.010).

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

	Men			Women		
	Participate	Emp/Pop	Unem/Pop	Participate	Emp/Pop	Unem/Pop
Prescrip. Rate	-0.053***	-0.055***	0.003*	-0.010***	-0.011**	0.002
	(0.005)	(0.006)	(0.002)	(0.003)	(0.003)	(0.001)
Demand Shock	0.317**	0.689***	-0.311***	-0.456***	-0.202*	-0.210***
	(0.105)	(0.117)	(0.057)	(0.093)	(0.098)	(0.051)
2000 Particip.	0.637***			0.409***		
	(0.050)			(0.029)		
2000 Emp/Pop		0.517***			0.349***	
		(0.036)			(0.021)	
2000 Unem/Pop			0.263***			0.169***
			(0.027)			(0.021)
R-sqr	0.09	0.11	0.02	0.06	0.06	0.02
N	6424995	6424995	6424995	6641288	6641288	6641288

All regressions include demographic variables, year, and state 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 through both 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 remain 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.

4.1 Heterogeneous Effects by 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.

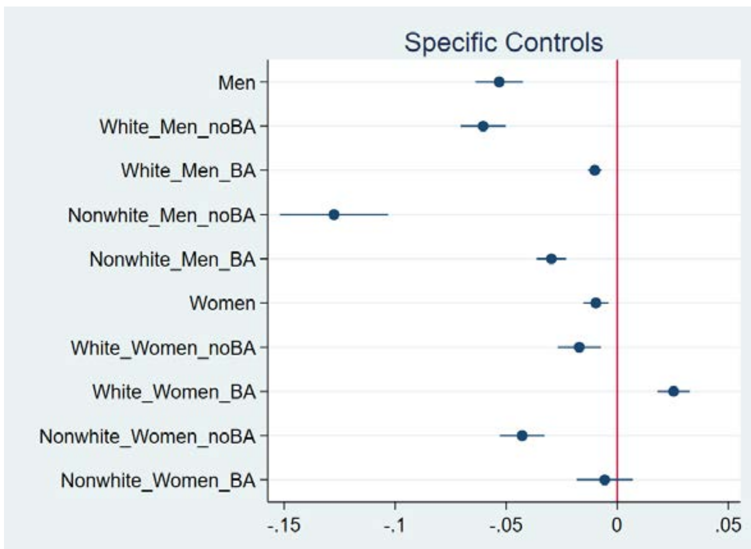

Figure 5: Labor Force Participation Effects by Demographic Groups

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 point estimates in the upper half of Figure 5).[8] 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 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

[8]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.

for the group of less-educated men, which predated the growth of opioid prescriptions. That effect is accounted for in the other controls.

Another result 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 the results of 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.

4.2 Alternative Difference-in-Differences Specification

Recall that our baseline specification (equation 6) leaves open how local economic conditions are controlled for. In this section we estimate a specification that uses the panel structure of our data along with linearity and separability assumptions to account for the LEC term using fixed effects. In this case,

$$f(LEC_{jt}) = \gamma_j + \delta_t. \tag{5}$$

With both a time and geographic fixed effect, this creates a standard 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.[9] The year fixed effects absorb the national 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

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

within localities.This model is well identified when t is large so that the measured differences in LECs across time can be substantial.

Table 3 shows the results of difference-in-differences regressions for prime-age men and women using the specification that is widely used in the literature. 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 difference-in-differences results for the employment-to-population ratio and labor force participation are similar to the estimates shown in table 2. Statistically 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 mostly 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) and Currie et al. (2019), one surprising result is that estimates are lower in magnitude for prime-age men than for women.

Table 3: 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)
R-sqr	0.10	0.12	0.02	0.06	0.06	0.02
N	6424995	6424995	6424995	6641288	6641288	6641288

All regressions include demographic controls, year, and couma fixed effects.

Robust standard errors with clustering on couma x year.

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

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. The time period of elevated opioid prescriptions limits the available years, and persistence in both labor force status and drug use complicates the short time horizon. These problems suggest looking for additional identifying variation from the cross-section. Table 4 shows results for the same specification, equation 6, 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 4: Participation of Prime-Age Men and Women, Varying Fixed Effects

	Men			Women		
	Census Division	State	Couma	Census Division	State	Couma
Prescrip. Rate	-0.046***	-0.053***	-0.014***	-0.016***	-0.010***	-0.017***
	(0.005)	(0.005)	(0.003)	(0.003)	(0.003)	(0.004)
R-sqr	0.09	0.09	0.10	0.06	0.06	0.06
N	6424995	6424995	6424995	6641288	6641288	6641288

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 4 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 state legal codes are certainly a reason to have fixed effects at the state level, we treat Table 4 primarily as evidence that modeling approaches that include some cross-sectional information could yield substantially different results.

To further explore how alternative sub-state level control groups yield different results, we employ a K-means approach to cluster or group all coumas found within a state based upon their geographic location (latitude and longitude) and long-term labor force participation rate, repeating this process using a progressively larger number of clusters (1, 3, 5, 10, 15, 20, and 25). Fixed effects based on these clusters are introduced to the model shown in equation 2, including the 2000 labor force status and Bartik controls. In this approach, one cluster is essentially a state fixed effect, and as the number of clusters increases, we move toward couma fixed effects.[10]

Figure 6 shows the coefficients and 95 percent confidence intervals on opioid prescription rates for prime-age men and women. The results remain stable across 1 to 5 clusters of coumas within states but gradually shift and become more uncertain as the number of clusters rises. Figure 6a shows that for males the impact of opioid prescriptions on participation rates is only materially different from state fixed effects at the couma fixed effect level, while Figure 6b shows that female participation is less sensitive to various levels of geographic fixed effects. This illustrates that while the difference-in-differences model does alter coefficient estimates, narrower geographic controls can be introduced without substantively altering the preferred model's results.

[10]In our dataset, 42 states have 25 or fewer coumas. California, Florida, Michigan, New York, North Carolina, Ohio, Pennsylvania, and Texas have more than 25 coumas.

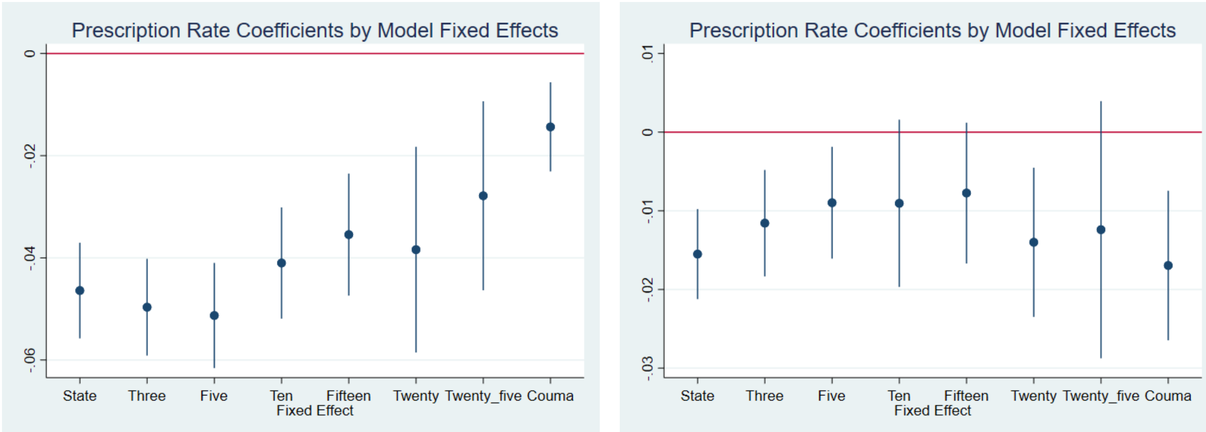

(a) Men

(b) Women

Figure 6: Sub-state Level Fixed Effects

4.3 Estimates for 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. Charles et al. (2018) argue that areas that experienced manufacturing declines experienced higher opioid prescription rates. While the causality issues are not solved in this analysis, the pattern would suggest a Rust Belt-centric 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).[11] The Rust Belt counties are the counties identified by Schweitzer (2017) as being in the "Industrial Heartland" in 1969. Figure 7 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.

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

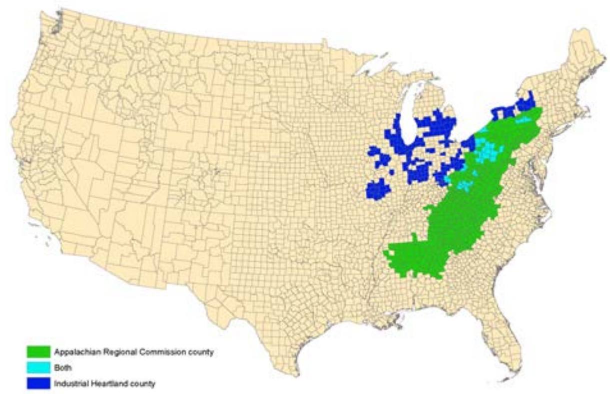

Figure 7: Appalachian Regional Commission and Rust Belt Counties

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 5: Labor Force Regressions with Geographic Interactions

	CPUMA		Appalachia		Industrial Heartland	
	Men	Women	Men	Women	Men	Women
Prescrip. Rate	-0.045***	-0.011***	-0.051***	-0.008**	-0.054***	-0.009**
CPUMA*Prescrip	-0.005***	0.001				
ARC*Prescrip			-0.003**	-0.002**		
IH*Prescrip					0.003*	-0.002***
Demand Shock	0.069	-0.425***	0.275**	-0.482***	0.275**	-0.422***
2000 LFPR	0.602***	0.414***	0.626***	0.402***	0.633***	0.413***
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 Measurement and Causal Paths

5.1 Selection into Local Labor Markets

The simplest path to reverse causality would be that prescriptions simply flowed more frequently into depressed labor markets, as suggested by Charles et al. (2018) and Hollingsworth et al. (2017) 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 8 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.

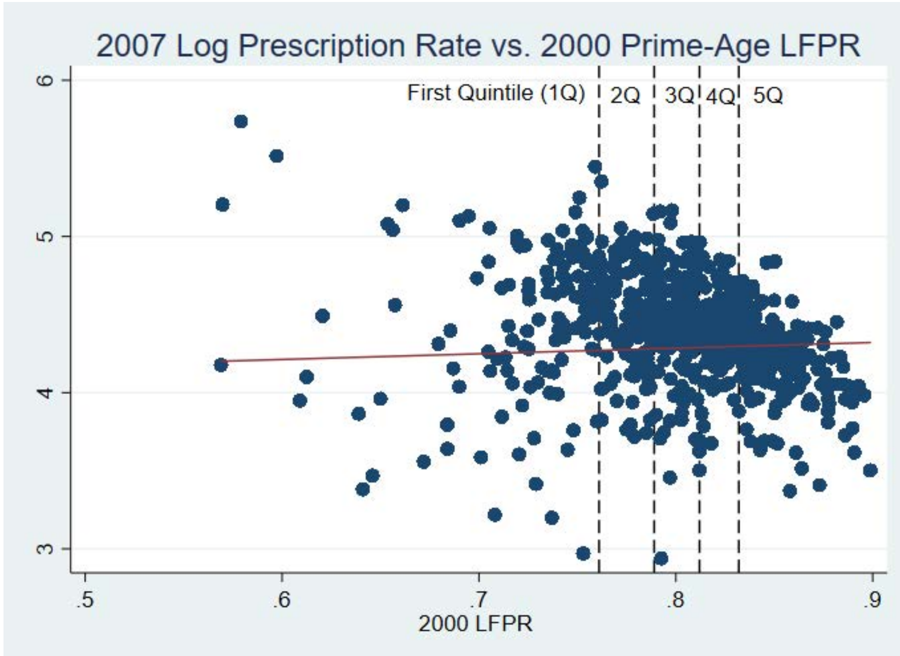

Figure 8: Association between 2000 LFPR and 2007 Prescription Rates

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 9 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.

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

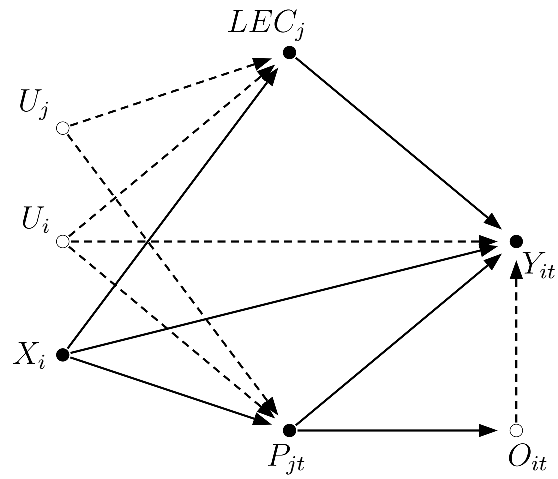

Figure 9: Directed Acyclic Graph of Opioids Affecting Labor Force Participation
Note: This figure follows the convention from Pearl (2009) of communicating that a variable is observed by drawing a solid line to its descendants and communicating that a variable is unobserved by drawing a dashed line to 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 3. For this comparison we again limit ourselves to labor force participation.

Figure 10 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 rate 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 estimates. The estimated effects for women are always smaller, so these weaker results are not surprising given the reduced precision of this estimation strategy. We interpret the more uniform results of the difference-in-differences specification for men as evidence that

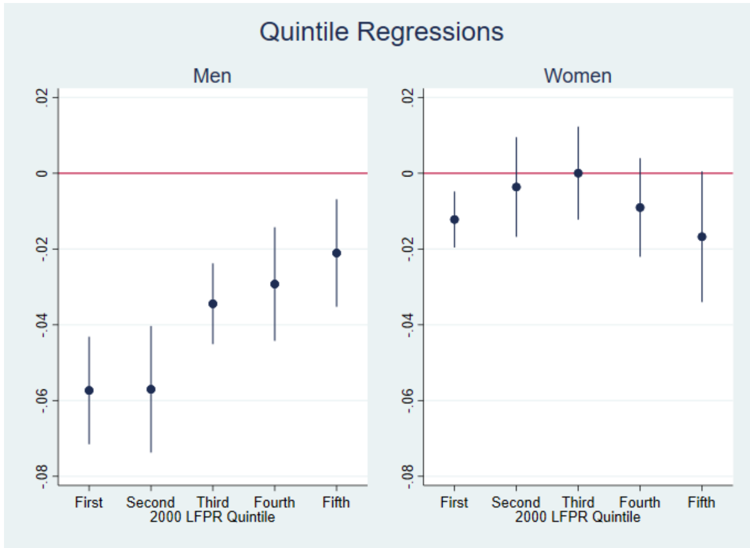

Figure 10: Coefficeint Estimates by Quintile of 2000 Labor Market Conditions

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 markets.

5.2 Selection on Unobserved Characteristics: Oster Sensitivity Analysis

A concern for both the measurement and causality implications for our results is the potential for unobserved selection. We apply the Oster (2019) model of unobserved selection and coefficient stability to estimate the potential for omitted variables to bias our results. Our model has introduced strong controls for local economic shocks (the labor demand measure based on Bartik (1991) and longer-run labor market differences (labor market statistics from the 2000 Census), but these measures are potentially incomplete. Nonetheless the most relevant omitted variables (local labor demand shocks and persistent labor challenges) are likely to be positively correlated with the full set of controls included in our model.

We implement Oster (2019) using the PSACALC Stata model Oster (2013). For our estimates the variables potentially correlated with unobserved selection in coumas are local labor demand shocks, 2000 labor market outcomes, state fixed effects, and individual characteristics, which in our two-level design control for couma average characteristics over time.

The national time pattern is treated as uninformative about unobserved variation within a couma. A key parameter for the Oster model is a maximum R-squared. For the purposes of our estimates, which include both geographic and individual variation in labor market outcomes, we set the maximum R-squared using a panel regression with couma-year fixed effects. The most that couma-level opioid prescription data might explain is couma average labor market outcomes.

Oster (2019) recommends beginning with a δ test for the degree of selection on unobservables relative to observables that would be necessary to explain away the result. Using this δ criteria, we get 9.2 for prime-age men and 5.0 for women. These are strong results, indicating that selection on unobservables would need to be several times as strong as the observed effects for the coefficient on opioid prescriptions to go to zero.

Oster (2019)'s techniques can also be used to estimate a bias-adjusted β under the assumption that the degree of unobserved selection is equal to that of the observed selection. Oster argues that this measure can be treated as a reasonable upper bound on selection effects and shows that these estimates are more robust estimators of the true treatment effects in an experiment where known variables are omitted. Because these are empirical estimates we collect 500 couma-clustered bootstrap repetitions to estimate standard errors. These standard errors augment the sampling variation in the coefficient estimates for uncertainty in the bias correction. For prime-age men the estimate is little changed at -0.0541 with a bootstrap standard error of 0.0077, versus a β of -0.0532 with a standard error of 0.0054 in the baseline estimates. For prime-age women, the point estimate shrinks to -0.0082 with a bootstrap standard error of 0.0042 from -0.0093 with a standard error of 0.0029. The 95 percent confidence interval for women still excludes zero, but the Oster's bounding exercise highlights that the result for women is more sensitive to selection.

5.3 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 prescription 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 OxyContin distributions that were about 50 percent lower in states where these programs 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.

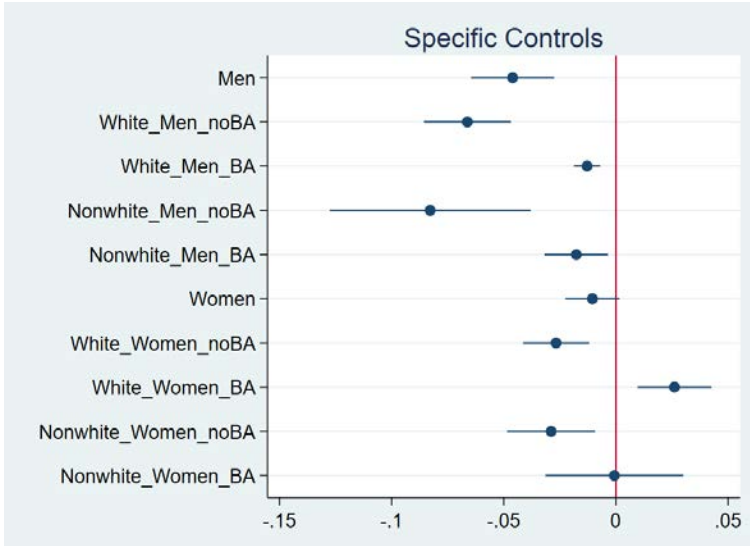

Figure 11: Triplicate-State IV Results by Demographic Groups

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:

$$P_{jt} = \sum_{t=1996}^{2017} \beta_t \times 1(\{\text{Non-triplicate} = j\} \times 1\{\text{Year} = t\} + \delta_t + \epsilon_{jt} \tag{6}$$

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 prescription opioids reported in Alpert et al. (2019). This approach focuses directly on exogenous variation to 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 nontriplicate states.

The subsequent IV regression findings 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 11 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.

6 The Implied Scale of the Effects

By design, this analysis has used the 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.

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. However, Abraham and Kearney (2018) conclude that the effects are still uncertain. 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 the 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 partially identified off of cross-sectional variation. During our sample period (2006-2016), the net changes in participation due to opioid prescriptions would be negligible, since nationally the prescription rate rises and then falls back.

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

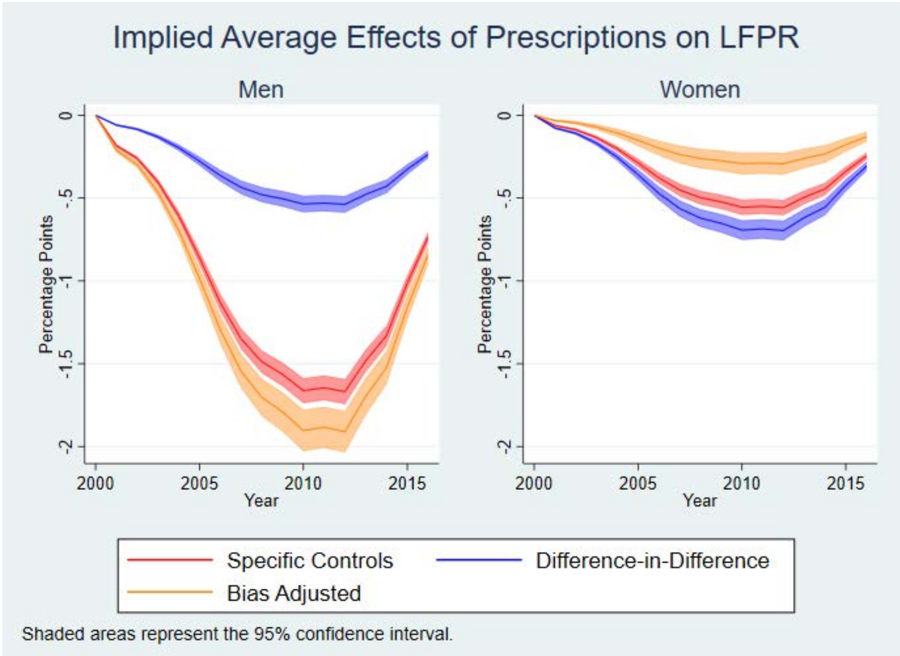

Figure 12: National Labor Market Effects Estimates

Figure 12 shows the implied average effects for male and female labor force participation for each of our estimated models.[12] Our primary estimates along with 95 percent confidence intervals are shown in red, and are the middle set of estimates for both men and women. Model uncertainty is present in these estimates, motivated generally by unobserved selection concerns. The orange estimates are the Oster bound for the case where unobserved selection is assumed to be as large as the observed selection. This assumes a large amount of unobserved selection (equal to the observed variation) and should be viewed as a selection-adjusted upper (or lower) bound to our primary estimates. The difference-in-differences results (shown in blue) are an alternative method for controlling for cross-sectional selection, which can induce a bias when the time period is short. These results are the product of excluding all cross-sectional variation in the estimates: As we showed earlier, including even limited cross-sectional variation with controls for local economic conditions produced higher estimated effects for men and lower estimates for women.

Importantly, all of the estimates show meaningful peak reductions in participation. Our baseline estimates based on specific controls imply a more than 1.7 percentage point decline in the participation of prime-age men. The baseline women's estimates peak at 0.6 percentage points. Generally, model uncertainty is a larger factor than the estimation standard errors

[12]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.

in differences in the estimates. The respective peak declines for men and women in the Oster bias-adjusted coefficient estimates are 1.9 and 0.3 percentage points. The peak reductions as estimated in the difference-in-differences model are sharply lower for men at just 0.5 percentage points. In contrast, the difference-in-differences estimates are a bit larger for women.

We can use the baseline 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 0.7 and 0.8 percentage points (based on the Oster bounds) between 2000 and 2016, but that the decline is between 0.2 and 0.1 percentage points for women. In our estimation, then, opioid prescriptions account for between 25 and 29 percent of the realized decline in the prime-age male participation rate and 7 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 64 and 73 percent of the decline in prime-age male labor force participation and between 23 and 43 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 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 (partially offset by more frequently residing in low-prescription counties).

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 prescription rates, we have little evidence on the reversibility of these effects. That would require carefully examining places that had effectively reduced their prescription rates without increasing illegal opioid availability and use. 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.[13] In addition, the nature of the opioid crisis has shifted from legal prescriptions to the widespread illegal use of opioids. Our results are not designed to identify the effects of fentanyl and other illegal opioids on the labor market.[14] Our data cannot pick up that important shift.

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

Abraham, Katharine G. and Melissa S. Kearney (2018). "Explaining the decline in the US employment-to-population ratio: A review of the evidence." *NBER WP 24333*. doi:10.3386/w24333.

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

Angrist, Joshua D. and Jörn-Steffen Pischke (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press. doi:10.1515/9781400829828.

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

[13]Two CPUMAs (the Oklahoma panhandle and eastern Texas) actually experienced large declines in prescription opioid rates over our sample time frame, but the reasons for those declines are unknown.

[14]See Park and Powell (2021) for an examination of the effects of illicit opioids on labor markets.

Bartik, Timothy J, Stephen CY Biddle, Brad J Hershbein, and Nathan D Sotherland (2018). "Wholedata: unsuppressed county business patterns data: version 1.0 [dataset]." *Kalamazoo: WE Upjohn Institute for Employment Research*.

Beheshti, David (2022). "The impact of opioids on the labor market: Evidence from drug rescheduling." *Journal of Human Resources*, p. 0320. doi:10.3368/jhr.0320-10762r1.

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

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.

Currie, Janet, Jonas Jin, and Molly Schnell (2019). "US employment and opioids: Is there a connection?" In *Health and Labor Markets*, volume 47 of *Research in Labor Economics*, pp. 253–280. Emerald Publishing Limited. doi:10.1108/S0147-912120190000047009.

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.

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

FDA (2018). "FDA analysis of long-term trends in prescription opioid analgesic products: Quantity, sales, and price trends." White paper, US Food and Drug Administration.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift (2020). "Bartik instruments: What, when, why, and how." *American Economic Review*, 110(8), pp. 2586–2624.

Harris, Matthew C., Lawrence M. Kessler, 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.

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*. URL `https://www.nytimes.com/interactive/2017/08/03/upshot/opioid-drug-overdose-epidemic.html`.

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–e47. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3659213/`.

Kromer, Braedyn K. and David J. Howard (2011). "Comparison of ACS and CPS data on employment status." Working paper SEHSD-WP2011-31, US Census Bureau. URL `https://www.census.gov/library/working-papers/2011/demo/SEHSD-WP2011-31.html`.

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." Job market paper, Harvard University. URL `https://scholar.harvard.edu/lairdja/publications/effects-physician-prescribing-behaviors-prescription-drug-use-and-labor-supply`.

Oster, Emily (2013). "Psacalc: Stata module to calculate treatment effects and relative degree of selection under proportional selection of observables and unobservables."

Oster, Emily (2019). "Unobservable selection and coefficient stability: Theory and evidence." *Journal of Business & Economic Statistics*, 37(2), pp. 187–204.

Park, Sujeong and David Powell (2021). "Is the rise in illicit opioids affecting labor supply and disability claiming rates?" *Journal of Health Economics*, 76, p. 102,430.

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

Powell, Jerome H. (2019). "The semiannual monetary policy report to the Congress." URL `https://www.govinfo.gov/app/details/CHRG-116shrg37911`.

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

Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek (2018). "Ipums usa: Version 8.0 [dataset]. minneapolis, mn: Ipums, 2018."

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.

Yellen, Janet L. (2017). *The Semiannual Monetary Policy Report to the Congress*. URL `https://www.govinfo.gov/app/details/CHRG-116shrg37911`.