A APPENDIX: Counterfactuals and Robustness

A.1 Counterfactuals

In order to further contextualize the magnitude of the labor market results, I conduct a simple simulation to estimate the aggregate impact of the rescheduling on the labor market. Using the estimates from equation 2, I predict the paths of the aggregate LFP rate and EPR under the observed levels of pre-period hydrocodone exposure as well as the under the counterfactual that every zip3 had zero pre-period exposure. This counterfactual can be thought of as an estimate of what may have happened had hydrocodone not been rescheduled, since if every zip3 had no pre-period hydrocodone exposure there would have been no scope for the rescheduling to affect the outcome variables. Mathematically, I estimate

$$Y_{zt}^{cf} = \hat{\lambda}_z + \hat{\psi}_t + \sum_Z \hat{\phi}_Z \cdot 1(z = Z) \cdot t + \sum_{T>\text{Sept. '14}} \hat{\theta}_T \cdot 1(t = T) \times \overline{\text{Hydro}}_z^{pre} + \hat{\eta} X_{zt} \tag{3}$$

where $\overline{\text{Hydro}}_z^{pre}$ is either the actual level of pre-period hydrocodone consumption or zero. The results of this exercise are shown in panels (A) and (B) of Figure A10 for the LFP rate and EPR, respectively. In both panels the solid black line shows the predicted outcome variable under the observed level of hydrocodone exposure, while the dashed blue line shows the counterfactual estimate. The estimates in panel (A) indicate that the LFP rate would have declined at a steeper rate absent the rescheduling, resulting in about a half of a percentage point (0.8 percent) lower LFP rate two and a half years after the rescheduling than what was actually observed. In the case of the EPR, this exercise suggests that the rescheduling can account for an approximately 0.61 percentage point (1 percent) increase in the EPR in the two and a half years following the rescheduling. This suggests that high levels of opioid prescribing may depress local labor market conditions, and reducing unnecessary prescriptions could have positive spillover effects on the labor market.

For completeness, I show the results for the UPR in Figure A11. Consistent with the discussion in Section 5.2, this figure shows no real effects of the rescheduling on the UPR. Given the small response of the UPR, I also show the 95 percent confidence interval for the counterfactual to check whether I can rule out large effects of the rescheduling. These confidence intervals are tightly concentrated around the counterfactual UPR, which is itself very close to the observed UPR. Quantitatively, I can rule out changes in the UPR of more than about one twentieth of a percentage point.

A.2 Robustness

The primary explanatory variable measuring the intensity of treatment used throughout section 5 is drawn from ARCOS. This has the advantage of being compiled from administrative reports filed by drug manufacturers directly to the DEA, measuring the entire legal supply of hydrocodone in

the US. However, these data do not directly measure consumption of hydrocodone, but rather shipments of the drug. Furthermore, they are only available at the zip3 level, necessitating an imperfect transformation of the county level labor force statistics. In light of these potential issues, I investigate the robustness of the main regression results to the use of an alternate treatment variable, namely, per capita hydrocodone consumption from Medicare Part D. While these data only measure hydrocodone consumption among the Medicare Part D population, it can be measured at the county level and therefore does not require any changes to the labor force data. The regression results using this treatment variable are shown in Figure A12. Panels (A) and (B) show the non-parametric and parametric results for the LFP rate, respectively, while Panels (C) and (D) present the results for the EPR. Overall, the results look very similar to those with the ARCOS treatment variable shown in Figure 6, albeit slightly larger in magnitude. This provides some assurance that measurement error introduced from the transition from county to zip3 labor force statistics or differential stockpiling of hydrocodone across zip3s are unlikely to significantly influence the regression results.

Next, I investigate the sensitivity of the results to the inclusion of controls for other non-hydrocodone opioids. Specifically, I convert several of the other most commonly prescribed opioids into morphine equivalent dosages and then aggregate them together in an 'other' opioid variable, pooled over the same period as the primary treatment variable.[61] I then include interaction terms for both the pre-period hydrocodone treatment variable as well as the measure of pre-period use of other opioids with time fixed effects in the regressions. The purpose of this exercise is that in principle it is possible that the results presented in the previous section were not driven by the rescheduling of hydrocodone, but rather some omitted factor which changed coincident with the rescheduling. In order for this to be the case the omitted factor would need to have differentially affected areas with higher pre-period consumption of hydrocodone. If areas which consume large amounts of hydrocodone also tend to consume larger amounts of other opioids, it is possible that this omitted factor was not directly related to hydrocodone consumption, but rather opioid consumption more generally.[62] For example, some set of policies enacted in late 2014 in areas with high levels of opioid consumption which changed labor market conditions could have generated a similar set of results to those reported above. If this was the case we would expect the hydrocodone variable to have little additional explanatory power once we control for other opioids. However, Figure A13 demonstrates that the inclusion of these additional controls has almost no impact whatsoever on the regression results. The top panels (A and B) show the coefficients on the interaction of the hydrocodone variable with month fixed effects with the LFP rate as the dependent variable, while the bottom panels (C and D) show the same coefficients with the EPR as the dependent variables. Panels (A) and (C) show the non-parametric results for each outcome, while panels (B) and (D) show the coefficients from the parametric model. These coefficients are nearly identical to those shown in Figure 6, suggesting that these results are in fact driven by the rescheduling of hydrocodone, and not other factors which may have affected opioid use more broadly.

[61] This variable combines oxycodone, morphine, hydromorphone, meperidine (pethidine), methadone, oxymorphone, codeine, dihydrocodeine, and levorphanol.

[62] Figure 2 suggests that this is unlikely, as there is a weak correlation between hydrocodone and overall opioid consumption.

A closely related concern is that these results are not driven by the hydrocodone rescheduling, but rather by relative improvements in economic conditions in areas which have historically consumed large amounts of oxycodone (e.g., OxyContin). If oxycodone consumption is highly correlated with hydrocodone consumption, then the regression coefficients may simply be picking up this pattern. Note that the lack of lack of substitution towards oxycodone following the rescheduling does not rule out this possibility. I present two pieces of evidence which mitigate this concern. First, Figure A14 demonstrates that the consumption of hydrocodone and oxycodone is largely uncorrelated. Each circle represents a zip3, weighted by population, and plots the per capita consumption of oxycodone (x-axis) against the per capita consumption of hydrocodone (y-axis) pooled from 2010-2013. The correlation coefficient is shown in the top right corner of the figure. Second, Figure A15 presents regression coefficients analogous to those in Figure A13, except only controlling for oxycodone prescribing instead of all non-hydrocodone opioid prescribing. Controlling for oxycodone prescribing makes almost no difference, which is not surprising given the lack of correlation between hydrocodone and oxycodone prescribing shown in Figure A14.

Next, I present evidence on possible confoundment from tramadol. As noted earlier, tramadol was first added to Schedule IV of the Controlled Substances Act in 2014. It is therefore possible that the regression results are picking up the effects of the tramadol scheduling rather than the rescheduling of HCPs. While the lack of any significant changes in tramadol prescribing following the HCP rescheduling (Appendix Figure A5) casts doubt on this explanation, I investigate this issue further here. First, I examine the extent to which HCP and tramadol prescribing were correlated prior to the rescheduling. Figure A16 displays a scatter plot of the number of HCP claims per 100 enrollees against the number of tramadol claims per enrollee in each zip3 during 2013. These series are fairly highly correlated ($\rho = 0.63$), which underscores the importance of explicitly including a measure of tramadol claims in my regressions as a robustness test. I do this by adding interactions of pre-period tramadol consumption with time fixed effects into my primary regression equations. The results from this exercise are shown in Figure A17, and indicate that the inclusion of these interactions has essentially no impact on the results. This indicates that it was indeed the rescheduling of HCPs, as opposed to the scheduling of tramadol, that caused the observed changes in the labor market.

Another alternative explanation for these findings is that areas which have historically consumed large quantities of opioids are also areas which were particularly affected by the Great Recession, and the trend break observed in the regressions is simply due to differentially slow recovery from the Recession in these areas. The fact that the trend break occurs around the time of the rescheduling is mere coincidence. However, the findings in Figures A13 and A15 cast doubt on this explanation. These figures show that the trend break occurs specifically in areas with higher reliance on hydrocodone even after controlling for the consumption of other opioids. It is difficult to rationalize why the Recession would differentially affect areas which consumed large amounts of hydrocodone as opposed to other non-hydrocodone opioids.

Furthermore, I present direct evidence that the differential labor market trends in the pre-period do not appear to be driven by the Recession. The crosswalks to map between counties and zip3s do not exist prior to 2010, which means I cannot use the county-by-month LAUS data to include

a longer pre-period in the regressions. However, state-by-year level labor force participation and employment-to-population ratio data exists in the ACS starting in 2005. I present regression results using this data in Figure A9. Panels (A) and (C) show that the differential trends in labor market conditions pre-dated the Recession, going back to at least 2005. Furthermore, the regression coefficients do not appear to change significantly around the time of the Recession, implying that areas which were more reliant on hydrocodone were not differentially affected by the Recession. Panels (B) and (D) show the regression results explicitly controlling for state-specific linear pre-trends. Despite the loss of spatial variation in exposure to treatment that occurs from using state-level data, the results using the ACS data broadly mirror the main regression results presented in Figure 6.

As discussed in section 4, my primary treatment variable pools data on hydrocodone shipments from 2010-2013, the 4 years prior to the rescheduling. The purpose of this variable is to measure the extent to which different areas were likely to be affected by the rescheduling, and pooling over several years should act to reduce measurement error in this variable. However, the specific choice of years to include is somewhat arbitrary. If there are significant differences in the relative ordering of hydrocodone consumption across years in the pre-period, then the regressions may be sensitive to the particular choice of years included. In order to investigate whether this will likely have any effects on the results, I first rank each zip3 in terms of per capita hydrocodone consumption separately for 2010 and 2013 and plot these ranks against each other in Figure A18. The red line represents the 45 degree line while each black circle represents a specific zip3, with more populous zip3s indicated by larger circles. While there are some exceptions, most zip3s fall very close to the 45 degree line, indicating that there is little over-time shuffling in the ranking of hydrocodone consumption over the pre-period. This minimizes the concern that the specific choice of years to include in the variable definition will influence the results. Further illustrating this point, Figure A19 presents the main regression results using several different alternative definitions of the treatment variable. The blue line shows the baseline regression coefficients from equation 1 in blue for the EPR in panel (A) and LFP rate in panel (B). The other colored lines show the same regression coefficients, but defining the treatment variable as the (normalized) per capita hydrocodone shipment for only a single year. The results are similar regardless of which year is used. The fact that using only data from as far back as 2010 gives similar results to those using data from 2013 alone should lessen concerns of the results being driven by mean reversion.

I explore the sensitivity of the regressions to the inclusion of additional control variables in Figure A20. Specifically, I consider the implementation of PDMPs and pill mill laws. As detailed in Horwitz et al. (2018), there are oftentimes significant differences in the coding of PDMP dates across studies. In order to account for this policy uncertainty I examine the sensitivity of my main labor market results to the inclusion of several different sets of laws: (1) the existence of an electronic PDMP (Horwitz et al., 2018), (2) the existence of a "modern" PMDP (Horwitz et al., 2018), (3) the existence of a mandatory access (MA) PDMP (Mallatt, 2018), (4) alternate legal coding of MA-PDMPs from (Sacks et al., 2019), and (5) the existence of a pill mill law from (Mallatt, 2018). The non-parametric regression results for each of these specifications are shown in Appendix Figure A20.

This figure demonstrates that, while in principle we may be concerned about confoundment from

other opioid-related policies, these concerns do not appear to manifest in practice. The regression coefficients are tightly concentrated around those from the baseline specification regardless of the precise legal coding chosen.

Next, I explore the sensitivity of the main labor market findings to labor demand shocks. I construct a Bartik instrument by interacting a state's industry concentration at baseline with the national growth of the industry.[63] Appendix Figure A21 shows the baseline coefficients for the EPR (panel (A)) and LFP rate (panel (B)) in blue, while the coefficients including this Bartik instrument are shown in black. These results are nearly identical, suggesting that the labor market findings are not the result of demand shocks.

Finally, I present regression results including a richer set of fixed effects. Specifically, in Figure A22 I show the regression coefficients from equations 1 (panels (A) and (C)) and 2 (panels (B) and (D)) with the inclusion of Census region-by-time fixed effects.[64] Results for the LFP rate and EPR are shown in the top two and bottom two panels, respectively. By including these fixed effects I net out any variation in labor market conditions across regions over time, relying only on within region variation. The regression coefficients here are almost identical to those shown in Figure 6, indicating that the pattern of results is not just picking up changing economic conditions across different regions of the county coincident with the rescheduling.

[63] Specifically, I construct the instrument following the process outlined in Maestas, Mullen and Powell (2013). Let $\hat{L}_{st} = \sum_k [\frac{\text{National Employment in Industry k at time t}}{\text{National Employment in Industry } k \text{ at time } t-1} \times (\text{State } s \text{ Employment in Industry } k \text{ at time } t-1)]$. The Bartik instrument is then $IV_{st} = \frac{\hat{L}_{st} - L_{s,t-1}}{L_{s,t-1}}$ where $L_{s,t}$ is actual employment.

[64] Census regions are geographic groupings of states. There are four regions in the US.

Figure A1: Geographic Variation in Hydrocodone Distribution, Consumption for 2013

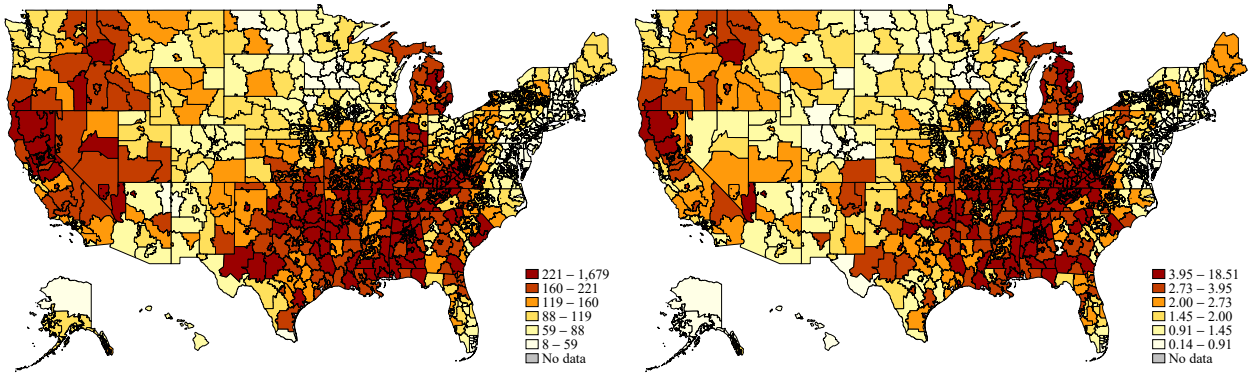

(A) Milligrams Per Person, ARCOS **(B)** Days Supplied Per Person, Medicare Part D

Note: These maps show different measures of the consumption of hydrocodone in ARCOS and Medicare Part D. Panel (A) shows the per capita distribution of hydrocodone in milligrams at the zip3 level in 2013 as reported in ARCOS. Panel (B) shows the per capita number of hydrocodone days supplied in Medicare Part D over the same time frame.

Table A1: Coefficients for Regressions of ARCOS Opioid Measures on Medicare Part D Measures

	(1) Hydrocodone	(2) Codeine	(3) Morphine	(4) Fentanyl	(5) Methadone	(6) Oxycodone	(7) Hydromorphone
	40.87***	190.4***	135.7***	2.812***	170.8***	109.8***	62.22***
	(34.98)	(18.89)	(30.58)	(20.06)	(10.62)	(34.01)	(32.32)
Observations	883	883	883	882	882	883	882
R^2	0.581	0.288	0.515	0.314	0.114	0.568	0.543
Corr. Coef.	0.76	0.54	0.72	0.56	0.34	0.75	0.74

Note: This table presents the coefficients from regressions of the per capita distribution of each drug from ARCOS on the corresponding measure from Medicare Part D. Data are from 2013, as this is the only pre-period year available in Part D.

Table A2: Demographic Summary Statistics by Quartile of Hydro_z^{pre}

	1st	2nd	3rd	4th
Log(Population)	12.42	12.26	12.33	12.24
% Male	49.23	49.7	49.68	49.38
% White	78.45	84.53	83.12	82.37
% Black	11.26	7.25	10.29	11.88
% Hispanic	14.49	11.89	13.72	8.46
% Over 65	15.19	15.85	15.69	16.08
% Under 20	24.81	25.51	25.96	25.44

Note: This table shows demographic summary statistics broken down by quartile of the treatment variable. Variables are measured annually and are taken from the Census Bureau.

Figure A2: Counties Identified in the CPS and ACS

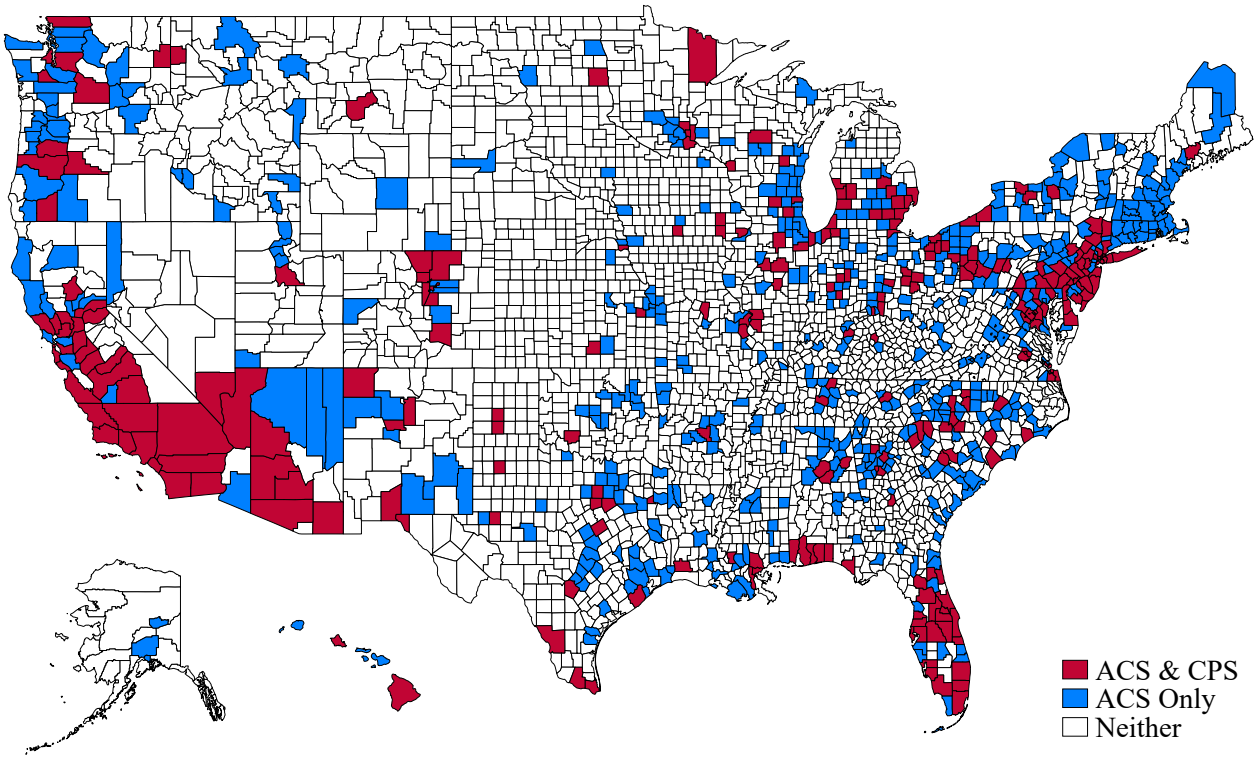

Note: This map illustrates which counties are identified in the ACS and CPS. The counties shaded red are those which are available in the CPS and ACS, while the counties shaded blue are available in the ACS but not CPS. Counties in white are identified in the LAUS, but neither the ACS nor CPS.

Table A3: Placebo Regressions: Impact of One Standard Deviation Increase in Pre-Period Hydrocodone Exposure

	(1) Log(Population)	(2) % Male	(3) % White	(4) % Black	(5) % Hispanic	(6) % Over 65	(7) % Under 20
Hydro x Post	-0.00168	-0.00129	-0.00641	0.00301	-0.00516	0.00420	0.00466
	(-0.96)	(-0.44)	(-1.47)	(1.09)	(-1.86)	(1.10)	(1.02)
Observations	6174	6174	6174	6174	6174	6174	6174
Mean	12.29	49.50	82.05	10.24	12.25	15.69	25.43

Note: This table shows θ coefficients from a series of placebo regressions of the form $Y_{zt} = \lambda_z + \psi_t + \sum_Z \phi_Z \cdot 1(z = Z) \cdot t + \theta \cdot 1(t > Oct.'14) \times \text{Hydro}_z^{pre} + \eta X_{zt} + \nu_{zt}$. Dependent variables are measured annually at the zip3 level and are taken from the Census Bureau.

Figure A3: Zip3 and County Maps

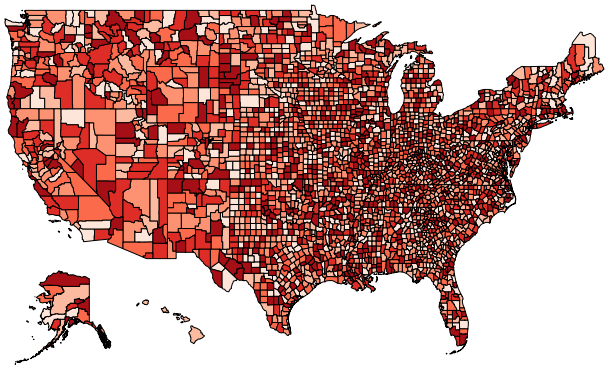

(A) US Counties

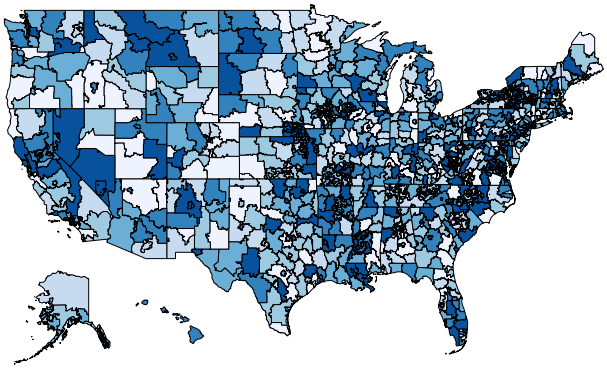

(B) US Zip3s

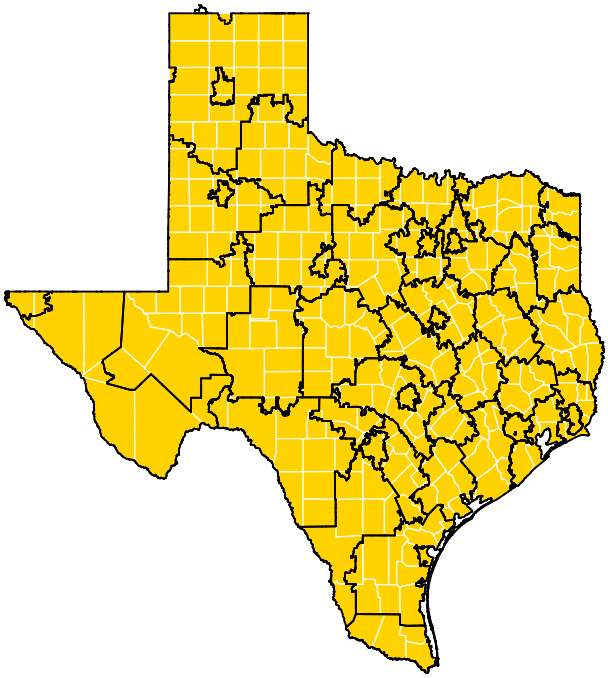

(C) Texas Zip3s and Counties

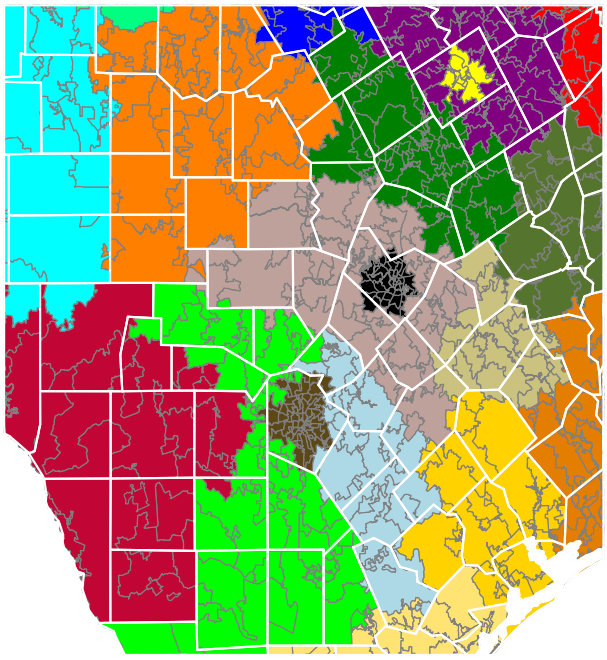

(D) Central Texas Zip5s, Zip3s, and Counties

Note: These maps illustrate the differences between counties, 5-digit zip codes, and 3-digit zip codes. Panel (A) shows counties in varying shades of red, while panel (B) shows zip3s in varying shades of blue. Panel (C) overlays zip3 and county boundaries for the state of Texas, with county boundaries shown in white and zip3 boundaries in black. Panel (D) shows the relationship between 5-digit zip codes (shown in gray), zip3s (different colors), and counties (outlined in white) for the central Texas region.

Figure A4: Hydrocodone Distribution: Dynamics by 1st and 4th Quartiles of Hydro$_z^{pre}$

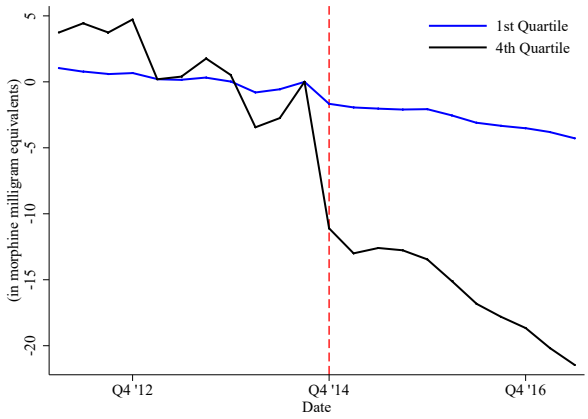

(A) Raw Hydro

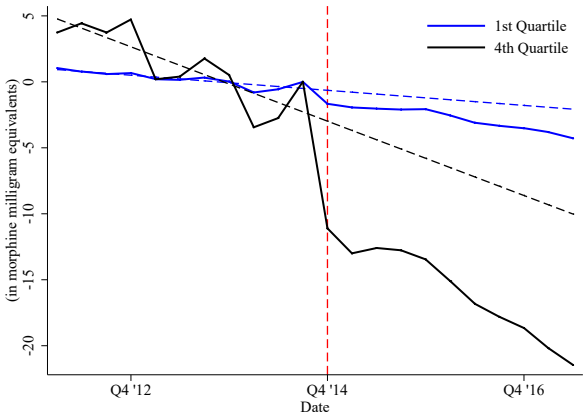

(B) Raw Hydro (with fitted trends)

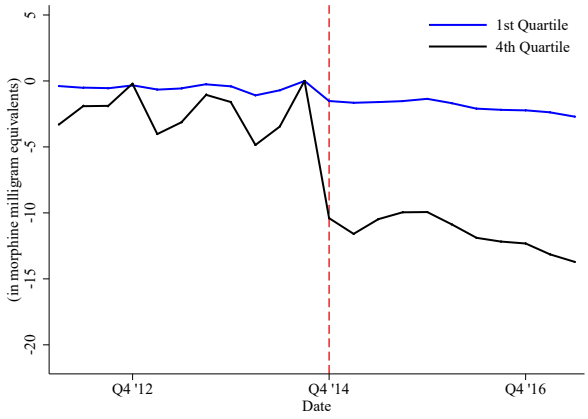

(C) Hydro Detrended

Note: These figures show time series for the distribution of hydrocodone per capita broken down by quartile of Hydro$_z^{pre}$. The first quartile is shown in blue in each panel while the fourth quartile is shown in black. Panel (A) shows the raw data normalized to zero in the third quarter of 2014. Panel (B) superimposes linear time trends fit based on the pre-period, while Panel (C) shows the data after subtracting out linear trends fitted based on the pre-period, analogous to regression equation 2.

Figure A5: Hydrocodone, Codeine, and Tramadol Distribution: Time Series and Regression Estimates

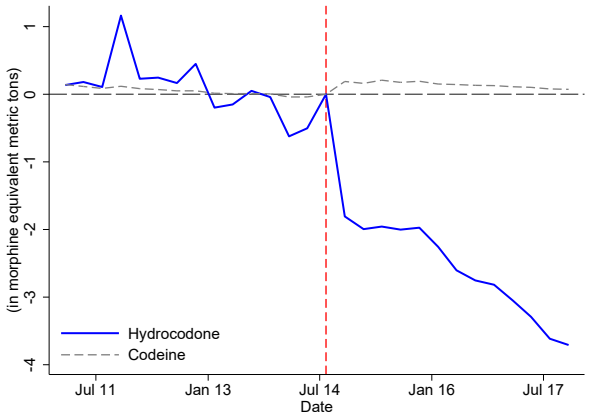

(A) ARCOS: Time Series

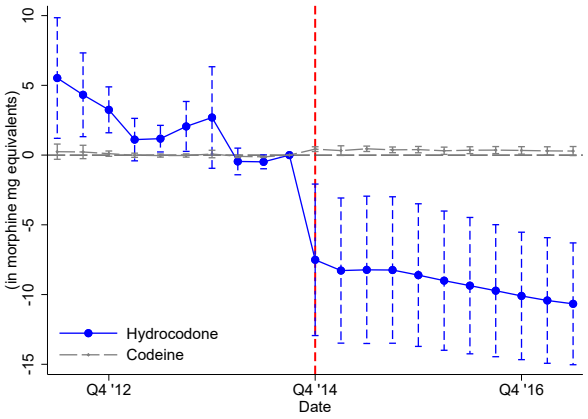

(B) ARCOS: Regression Coefficients

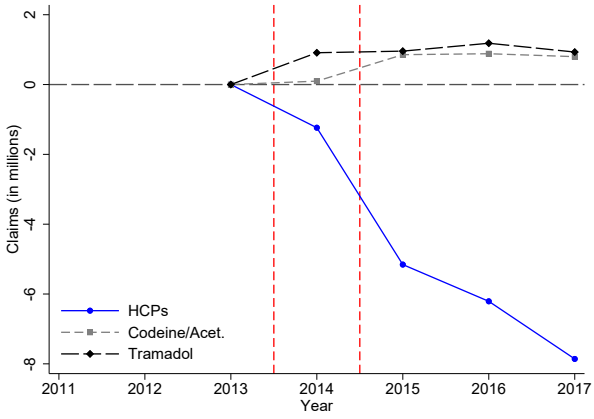

(C) Part D: Time Series

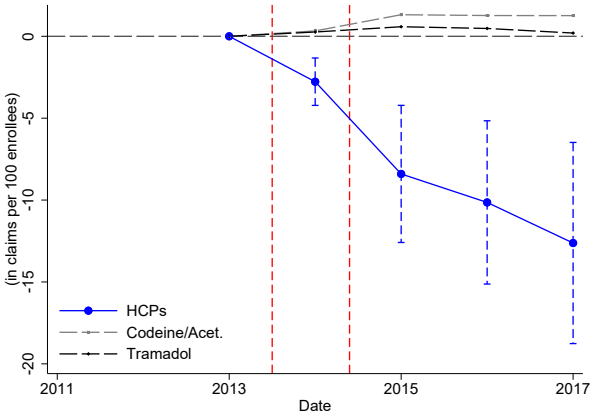

(D) Part D: Regression Coefficients

Note: Panel (A) shows the aggregate quarterly distribution of hydrocodone (blue solid line) and codeine (gray dashed line) from 2011-2017. Data are normalized to zero in the third quarter of 2014. Panel (B) shows regression coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equation 1 for hydrocodone (in blue) and codeine (in gray), both converted to morphine milligram equivalents. Note that Panel (B) replicates Figure 3 Panel (A). Panel (C) shows the annual claim count in millions for hydrocodone combination products (in blue), codeine with acetaminophen (in gray), and tramadol (in black) in Medicare Part D. Data are normalized to zero in 2013. Panel (D) shows regression coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equation 1 using claims from from Medicare Part D. Confidence intervals for codeine with acetaminophen and tramadol are omitted.

Figure A6: Tramadol and Hydrocodone Distribution Over Time

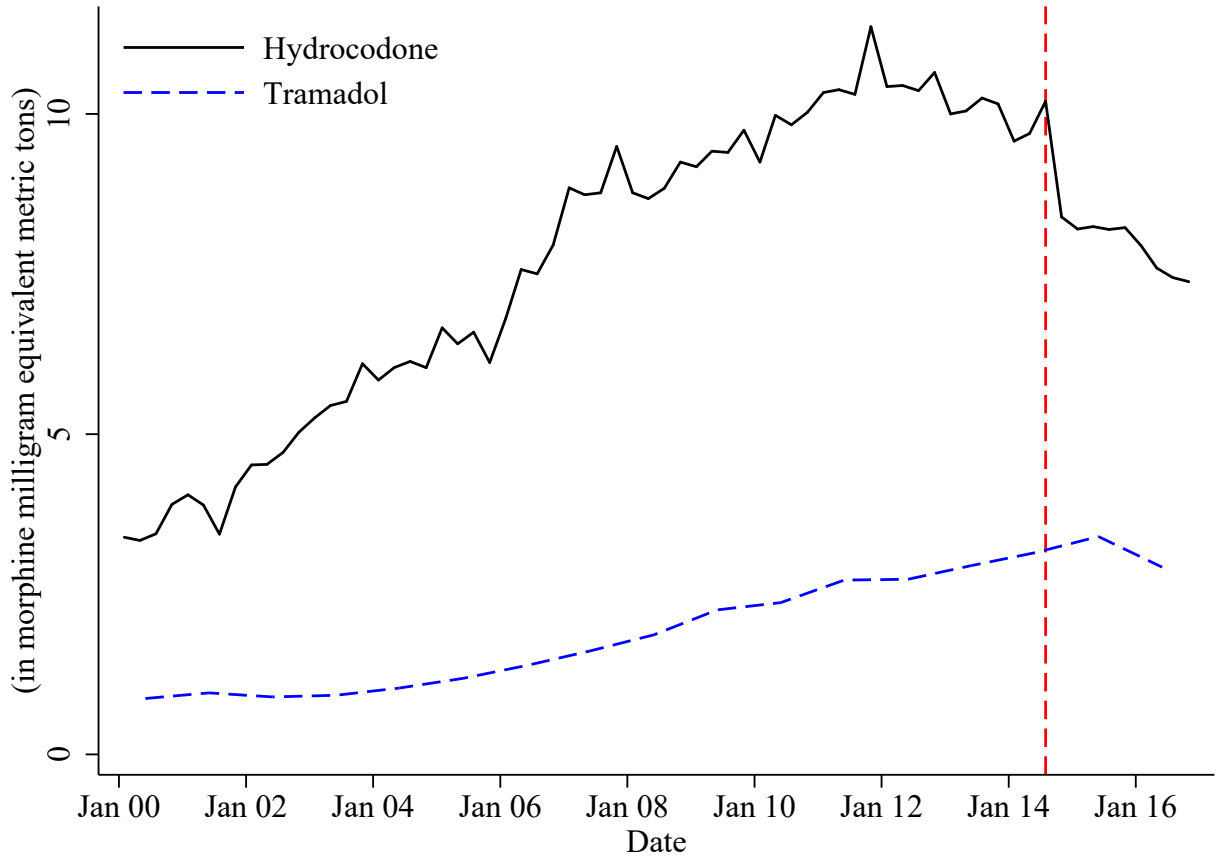

Note: This figure shows the quarterly distribution of hydrocodone (black solid line) and Tramadol (blue dashed line) from 2000 to 2016. Data on the distribution of hydrocodone is taken from ARCOS while the Tramadol data are taken from Figure (3) of US Food and Drug Administration (2018) which uses data from IQVIA. The Tramadol line is the annual distribution divided by four (quarterly data were not published).

Figure A7: Labor Force Participation Rate and Employment-to-Population Ratio: Dynamics by 1st and 4th Quartiles of Hydro$_z^{pre}$

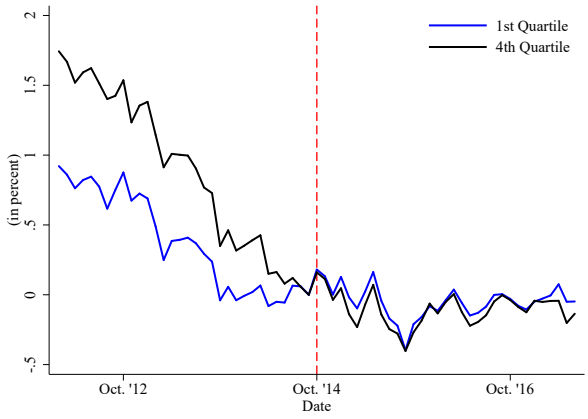

(A) Raw LFP

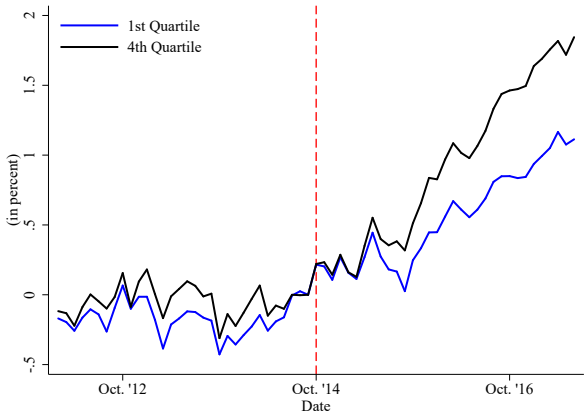

(B) LFP (Detrended)

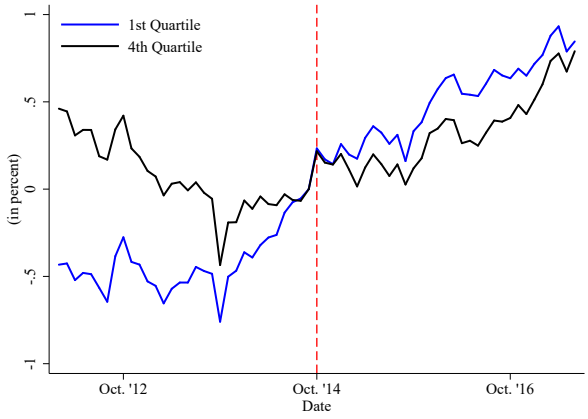

(C) Raw EPR

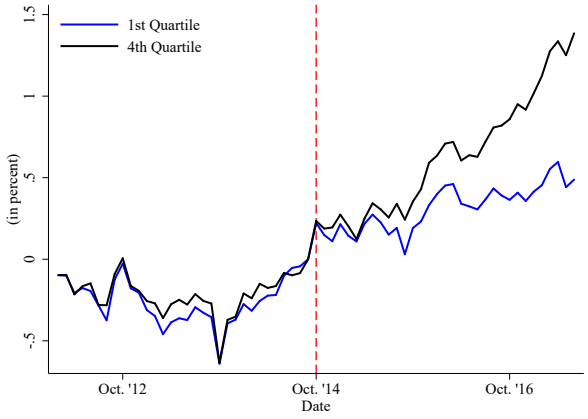

(D) EPR (Detrended)

Note: These figures show time series for the LFP rate and EPR broken down by quartile of Hydro$_z^{pre}$. The first quartile is shown in blue in each panel while the fourth quartile is shown in black. Panels (A) and (B) show the labor force participation rate, while panels (C) and (D) show the employment-to-population ratio. Panels (A) and (C) show the data with September 2014 normalized to zero, while panels (B) and (D) show the data after subtracting out linear trends fit based on the pre-period, analogous to regression equation 2.

Figure A8: Unemployment-to-Population Ratio Regression Estimates: Impact of One Standard Deviation Increase in Hydro_z^{pre}

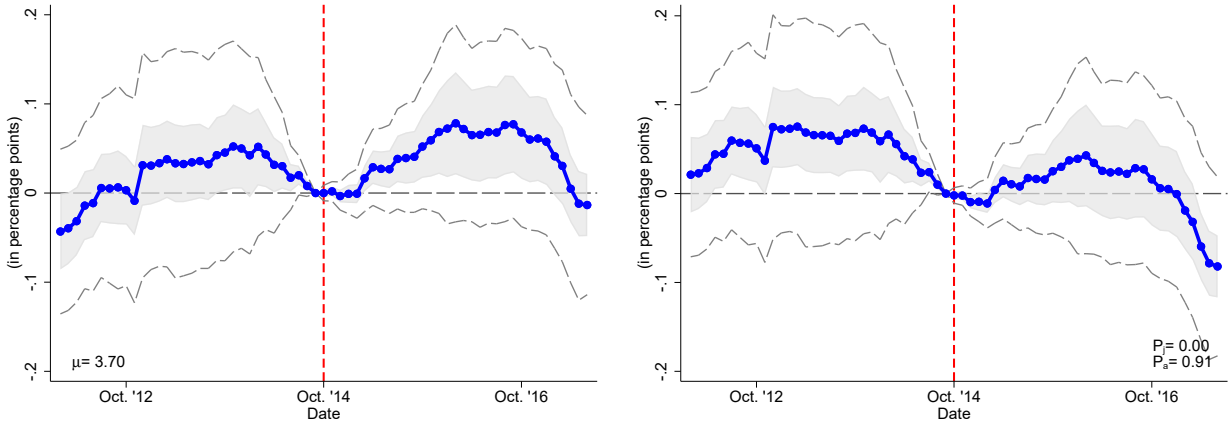

(A) Non-Parametric (equation 1) **(B)** Parametric (equation 2)

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equations 1 (panel (A)) and 2 (panel (B)) with the unemployment-to-population ratio as the dependent variable. The dependent variable is measured at the zip3-by-month level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption.

Figure A9: Labor Force Participation Rate and Employment-to-Population Ratio Regression Estimates Using Data from the ACS: Impact of One Standard Deviation Increase in Hydro_z^{pre}

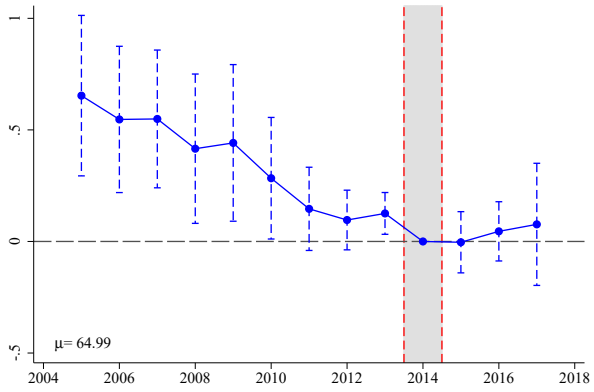

(A) Non-Parametric LFP (equation 1)

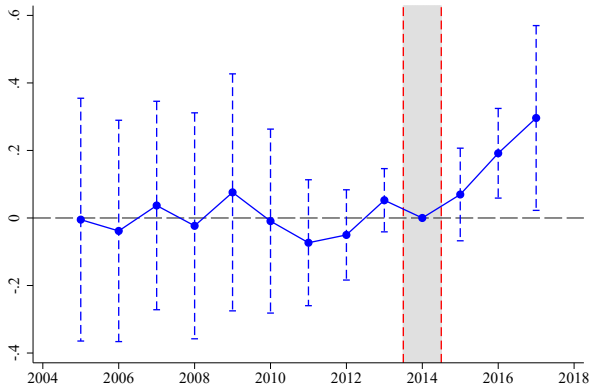

(B) Parametric LFP (equation 2)

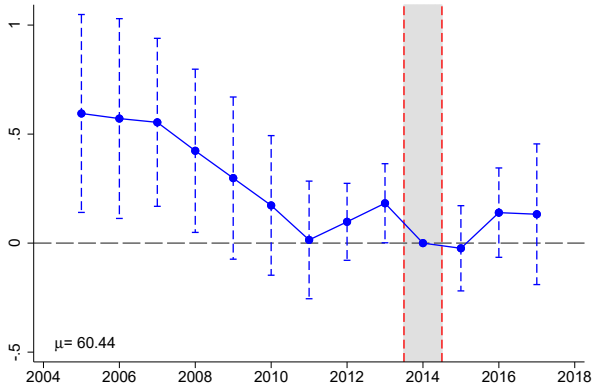

(C) Non-Parametric EPR (equation 1)

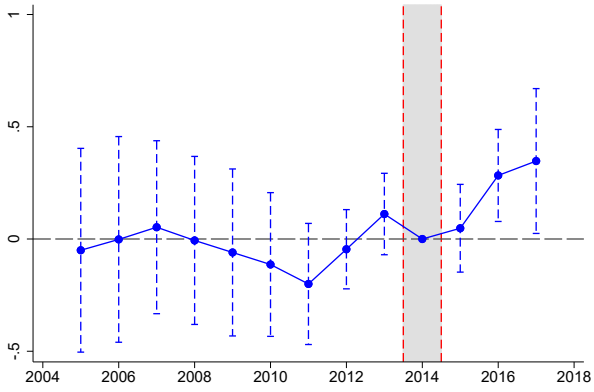

(D) Parametric EPR (equation 2)

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ year interactions from regression equations 1 (panels (A) and (C)) and 2 (panel (B) and (D)). Panels (A) and (B) show the results for the labor force participation rate, while panels (C) and (D) show the results for the employment-to-population ratio. The dependent variables taken from the ACS and are measured at the state-by-year level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption.

Figure A10: Counterfactual Paths of the Labor Force Participation Rate and Employment-to-Population Ratio

(A) Labor Force Participation Rate

(B) Employment-to-Population Ratio

Note: This figure shows the predicted and counterfactual labor force participation rate (panel (A)) and employment-to-population ratio (panel (B)). The counterfactual estimates use the regression coefficients from equation 2, but assume that each zip3 has $\text{Hydro}_z^{pre} = 0$. The black line shows the predicted values with the actual value of Hydro_z^{pre}, while the blue dashed line shows the counterfactual.

Figure A11: Counterfactual Path of the Unemployment-to-Population Ratio

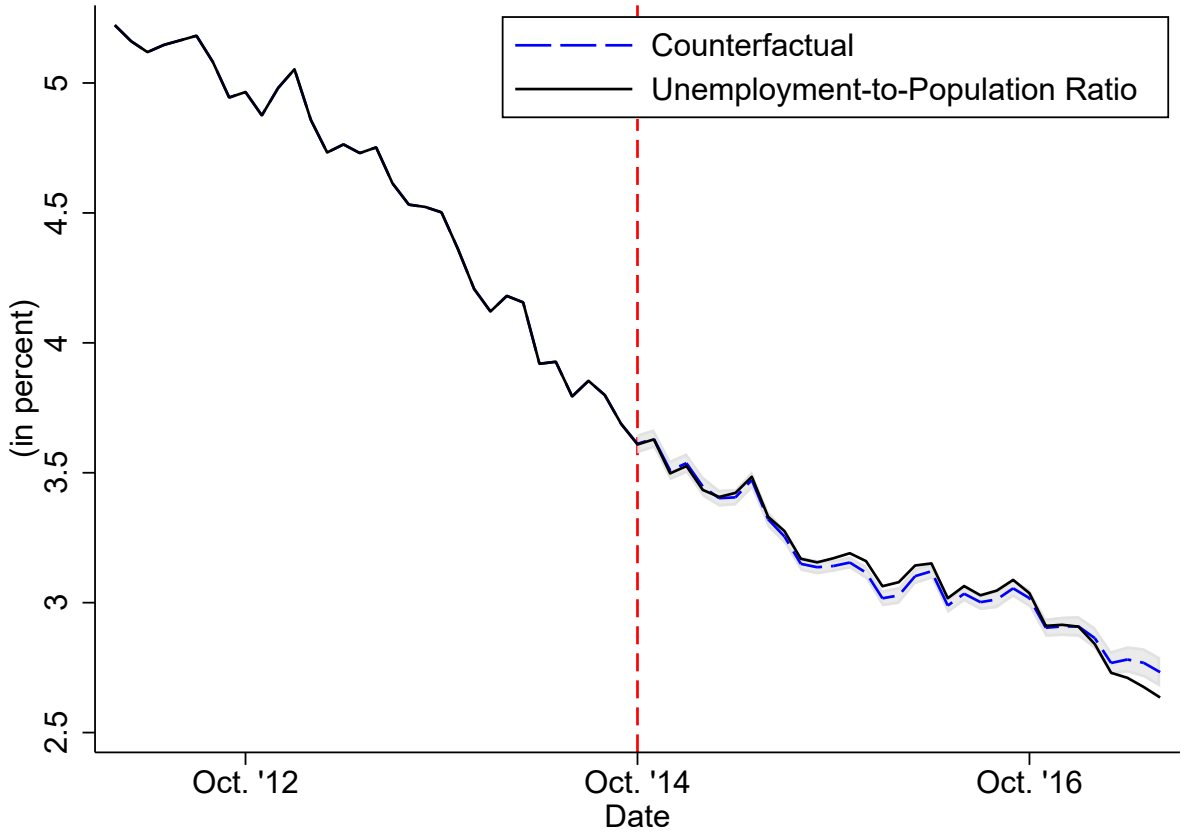

Note: This figure shows the predicted and counterfactual unemployment-to-population ratio. The counterfactual estimates use the regression coefficients from equation 2, but assume that each zip3 has $\text{Hydro}_z^{pre} = 0$. The black line shows the predicted values with the actual value of Hydro_z^{pre}, while the blue dashed line shows the counterfactual. The 95 percent confidence interval for the counterfactual is indicated by the shaded region.

Figure A12: Alternate Treatment Variable: Medicare Part D Hydrocodone Days Supplied

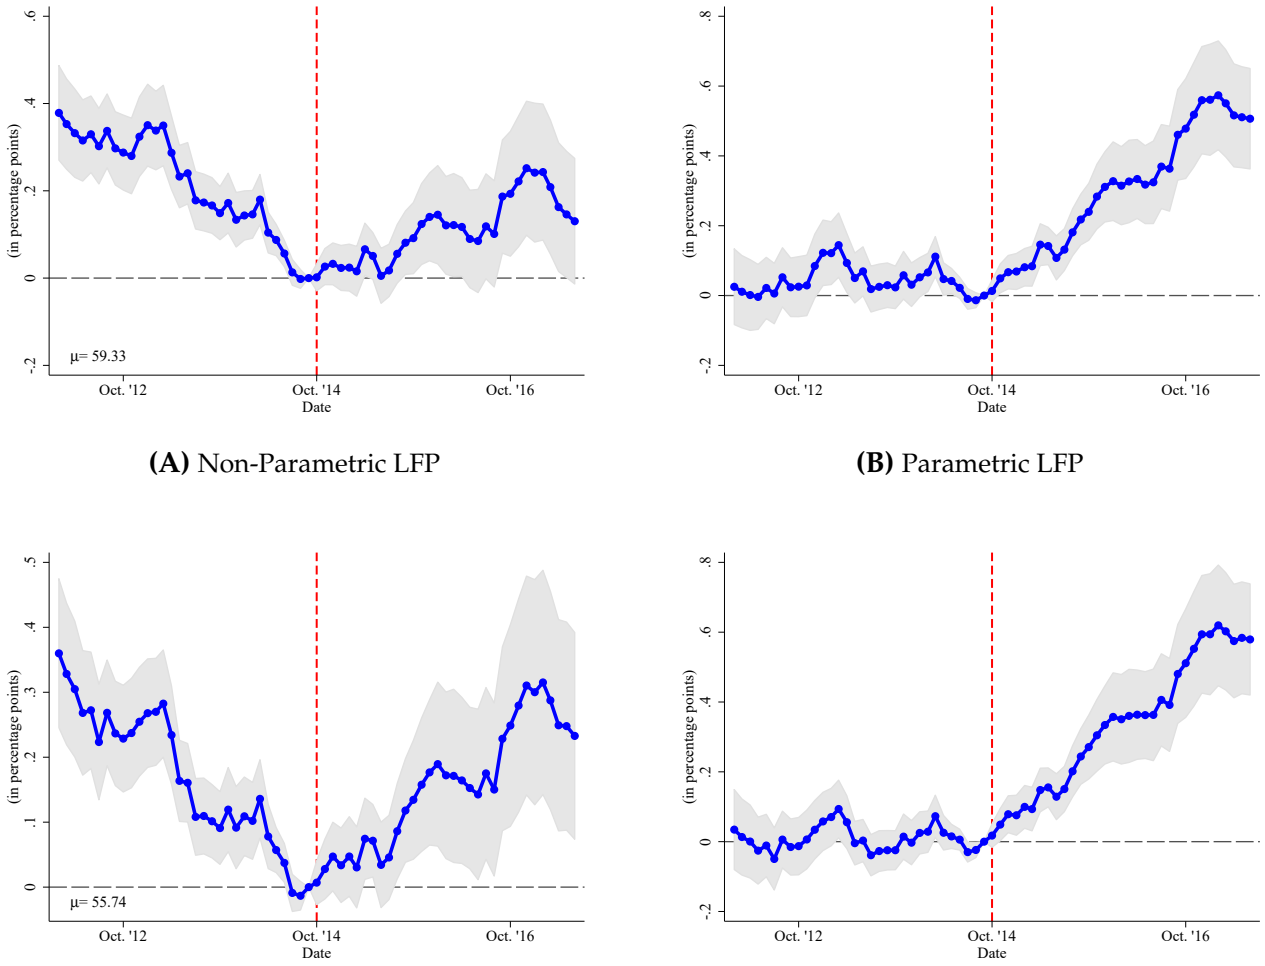

(A) Non-Parametric LFP

(B) Parametric LFP

(C) Non-Parametric EPR

(D) Parametric EPR

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equations 1 (panels (A) and (C)) and 2 (panel (B) and (D)). In these regressions the measure of pre-period hydrocodone consumption comes from Medicare Part D as opposed to ARCOS. Panels (A) and (B) show the results for the labor force participation rate, while panels (C) and (D) show the results for the employment-to-population ratio. The dependent variables are measured at the county-by-month level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption.

Figure A13: Regressions Controlling for Other Opioids: Impact of One Standard Deviation increase in Hydro_z^{pre}

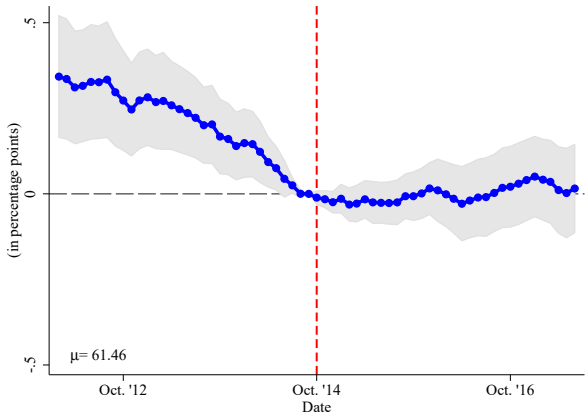

(A) Non-Parametric LFP

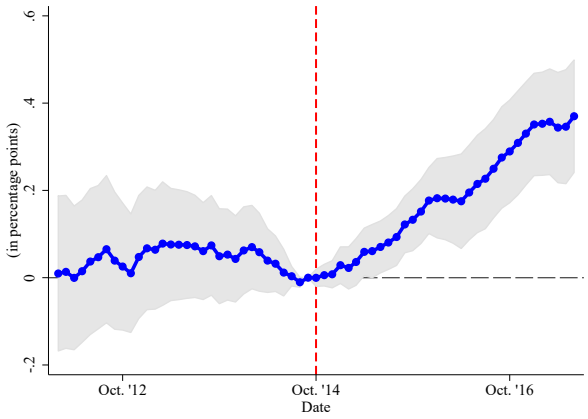

(B) Parametric LFP

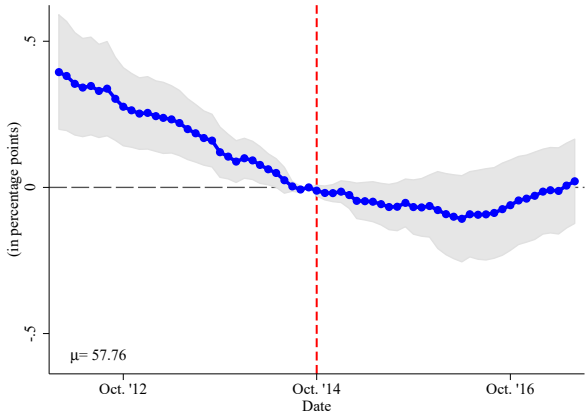

(C) Non-Parametric EPR

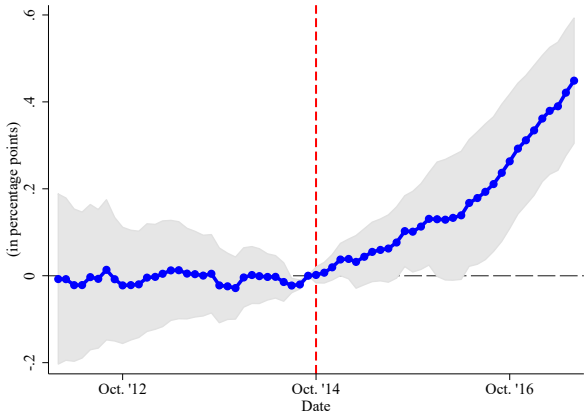

(D) Parametric EPR

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equations 1 (panels (A) and (C)) and 2 (panel (B) and (D)). These regressions explicitly control for the interaction of time fixed effects with non-hydrocodone opioids. Panels (A) and (B) show the results for the labor force participation rate, while panels (C) and (D) show the results for the employment-to-population ratio. The dependent variables are measured at the zip3-by-month level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption.

Figure A14: Correlation Between Hydrocodone and Oxycodone Shipments

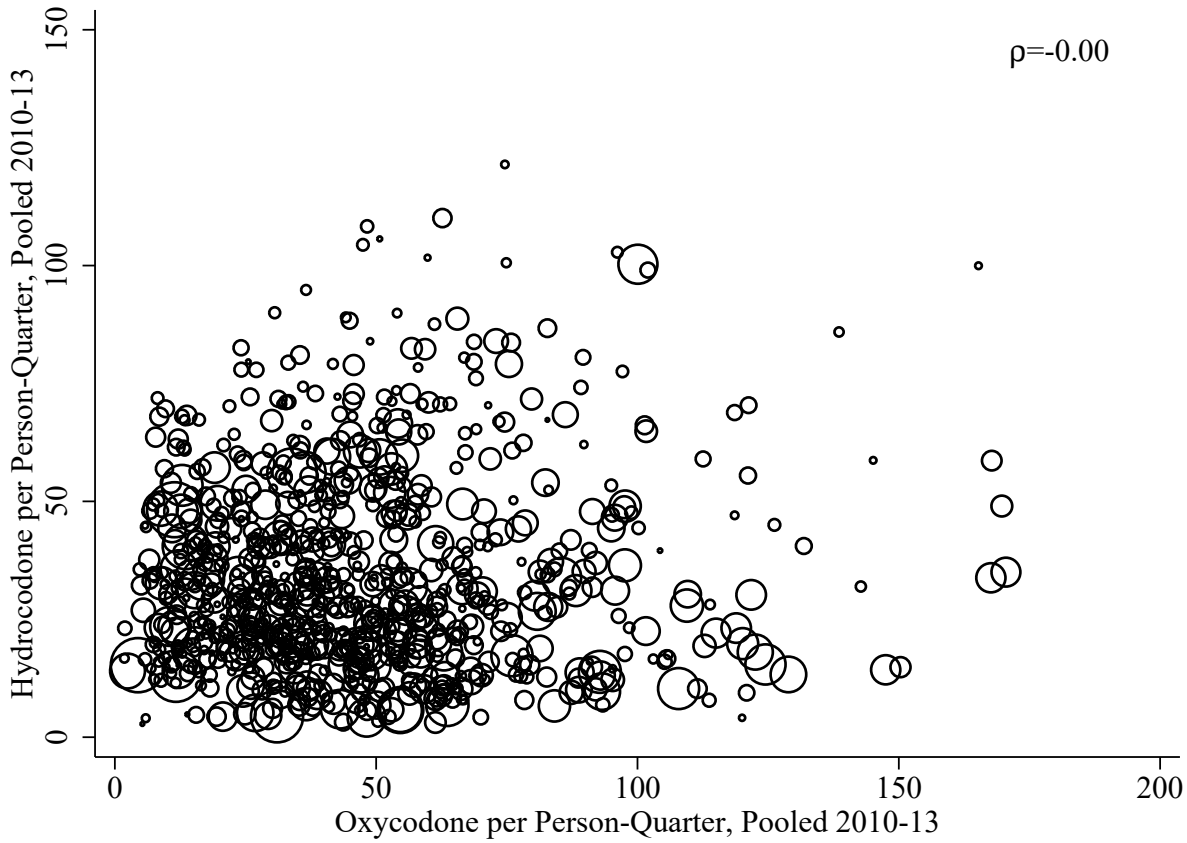

Note: This figure shows a scatter plot of each zip3's hydrocodone consumption per capita (y-axis) versus its oxycodone consumption per capita (x-axis). Data are pooled from 2010-2013 and exclude zip3s above the 99th percentile in either distribution. The size of each circle represents the relative population size of the zip3.

Figure A15: Regressions Controlling for Oxycodone Distribution

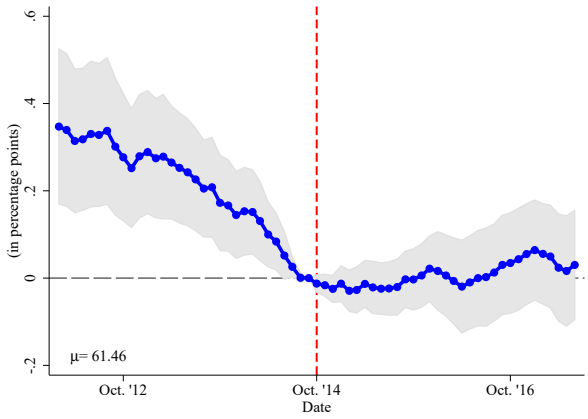

(A) Non-Parametric LFP

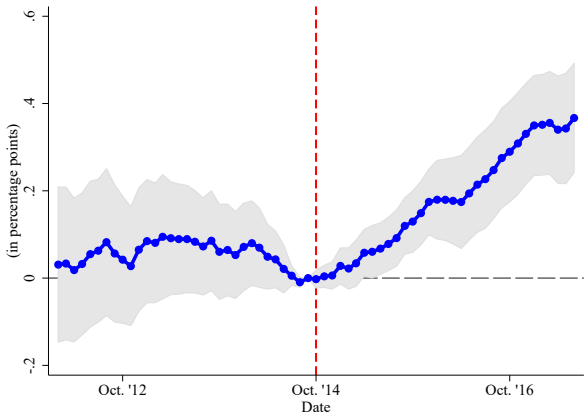

(B) Parametric LFP

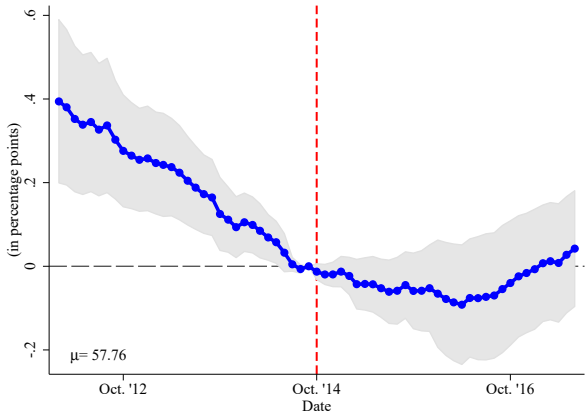

(C) Non-Parametric EPR

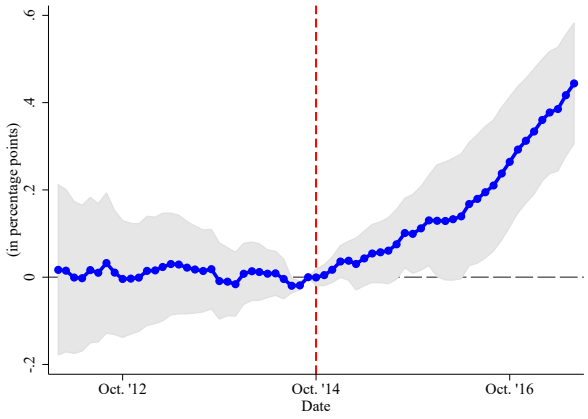

(D) Parametric EPR

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equations 1 (panels (A) and (C)) and 2 (panel (B) and (D)). These regressions explicitly control for the interaction of time fixed effects with a measure of pre-period consumption of oxycodone. Panels (A) and (B) show the results for the labor force participation rate, while panels (C) and (D) show the results for the employment-to-population ratio. The dependent variables are measured at the zip3-by-month level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption.

Figure A16: HCP and Tramadol Claims per 100 Medicare Part D Enrollees

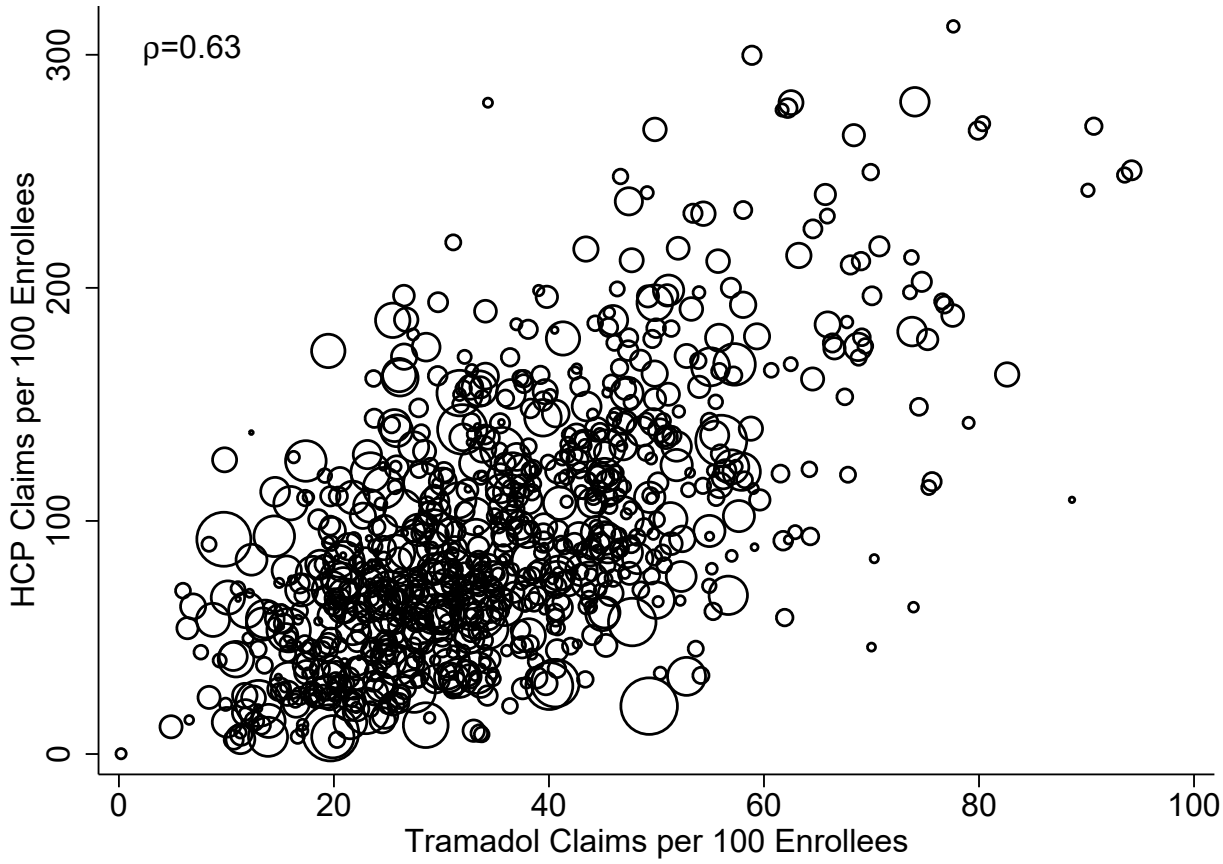

Note: This figure displays the number of HCP claims per 100 Medicare Part D enrollees on the y-axis against the number of tramadol claims per 100 enrollees in 2013. Each circle represents a zip3 weighted by the number of enrollees. The correlation coefficient between the two series is reported in the top left corner.

Figure A17: Regressions Controlling for Tramadol: Impact of One Standard Deviation increase in Hydro_z^{pre}

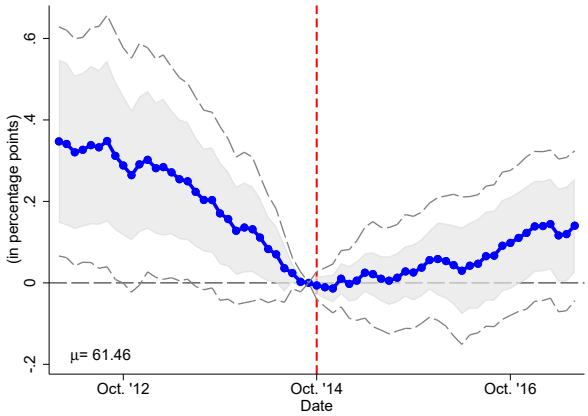

(A) Non-Parametric LFP

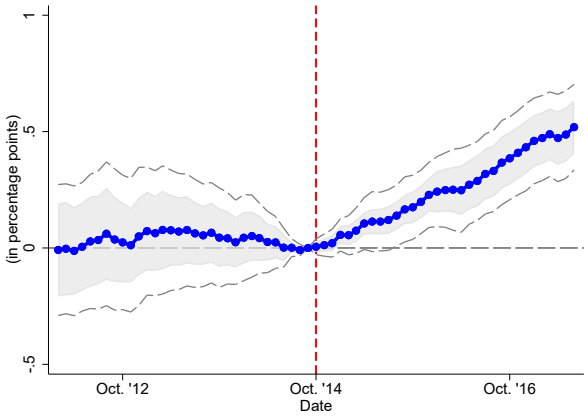

(B) Parametric LFP

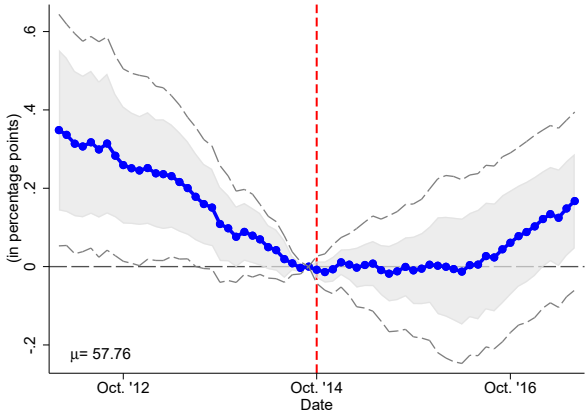

(C) Non-Parametric EPR

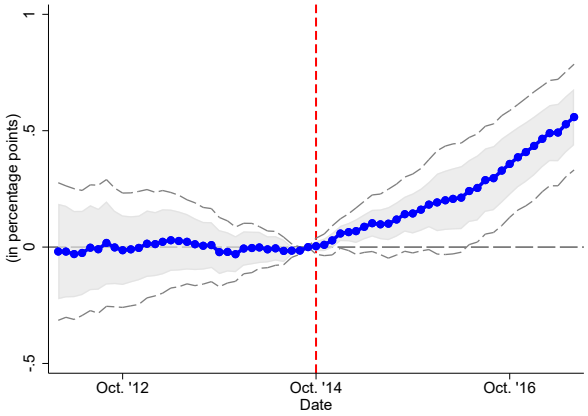

(D) Parametric EPR

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equations 1 (panels (A) and (C)) and 2 (panel (B) and (D)). These regressions explicitly control for the interaction of time fixed effects with non-hydrocodone opioids. Panels (A) and (B) show the results for the labor force participation rate, while panels (C) and (D) show the results for the employment-to-population ratio. The dependent variables are measured at the zip3-by-month level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption. The 95 percent confidence intervals using standard errors clustered at the zip3 level and state level are shown by the shaded region and dashed gray lines, respectively.

Figure A18: Rank in the Hydrocodone Distribution Over Time

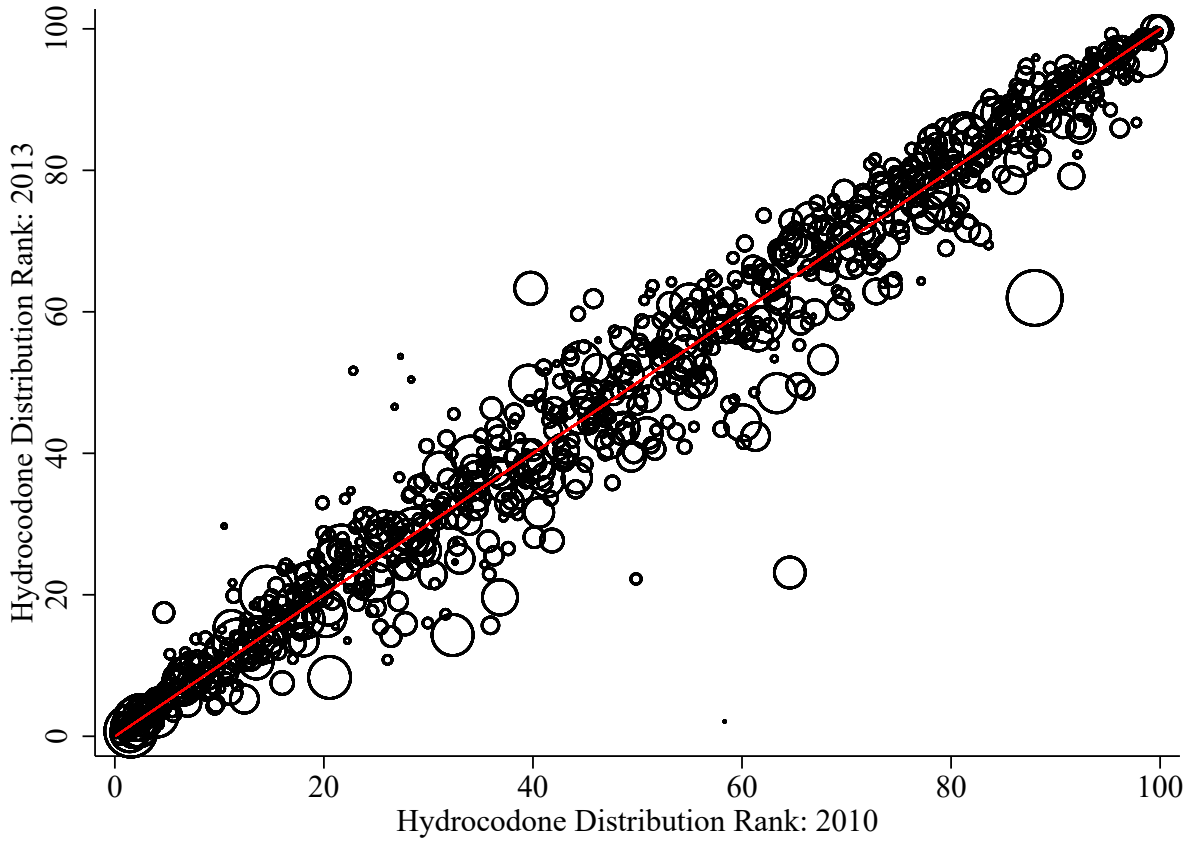

Note: This figure shows a scatter plot of each zip3s rank in the hydrocodone distribution in 2010 versus 2013. The rank is normalized between 0 and 100, with 100 indicating the zip3 with the highest per capita consumption of hydrocodone.

Figure A19: Alternate Definitions of the Treatment Variable

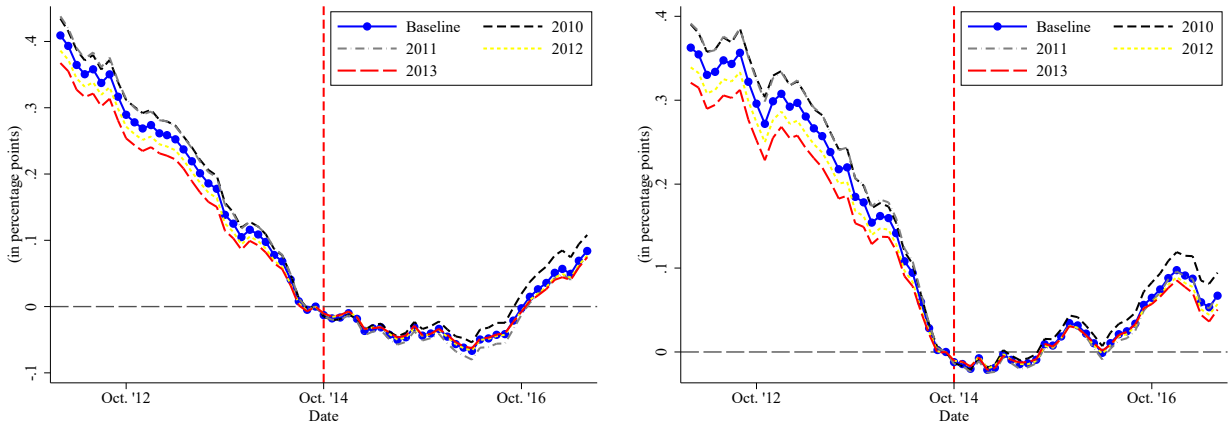

(A) Employment-to-Population Ratio

(B) Labor Force Participation Rate

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equation 1 with the employment-to-population ratio (panel (A)) and labor force participation rate (panel (B)) as the dependent variables. The blue coefficients correspond to the baseline specification pooling data from 2010-2013 for the treatment variable. The other colors indicate the coefficients defining the treatment variable by the per capita distribution of hydrocodone in a single year. The dependent variable is measured at the zip3-by-month level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption.

Figure A20: Regressions Controlling for Other Opioid-Related Policies: Impact of One Standard Deviation increase in Hydro_z^{pre}

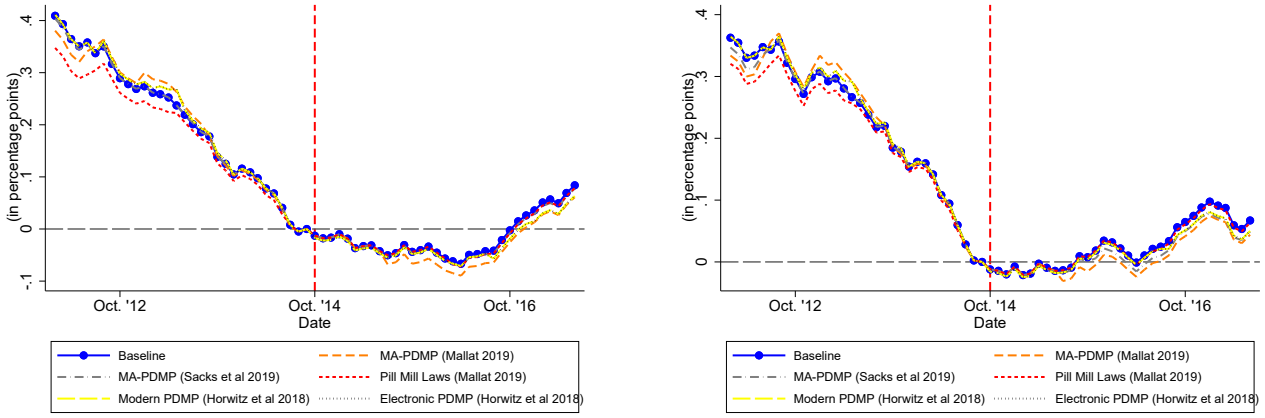

(A) Employment-to-Population Ratio **(B)** Labor Force Participation Rate

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equation 1 with the employment-to-population ratio (panel (A)) and labor force participation rate (panel (B)) as the dependent variables. The blue coefficients correspond to the baseline specification. The orange lines add indicators for MA-PDMPs as coded by Mallatt (2018). The gray lines include indicators for MA-PDMPs as coded by Sacks et al. (2019). The red lines include pill mill laws from Mallatt (2018). The yellow and dotted black lines indicate a modern or electronic PDMP as coded by Horwitz et al. (2018), respectively. The dependent variables are measured at the zip3-by-month level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption.

Figure A21: Regressions Controlling for Bartik Instrument: Impact of One Standard Deviation increase in Hydro_z^{pre}

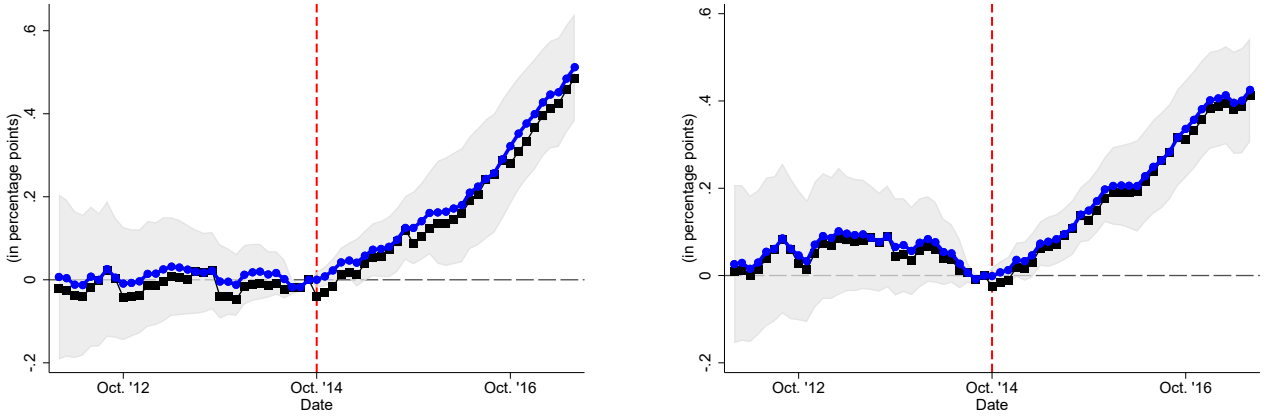

(A) Employment-to-Population Ratio **(B)** Labor Force Participation Rate

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equation 2 with the employment-to-population ratio (panel (A)) and labor force participation rate (panel (B)) as the dependent variables. The blue coefficients correspond to the baseline specification. The black coefficients correspond to an identical regression which also includes the Bartik instrument described in the text. The dependent variables are measured at the zip3-by-month level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption.

Figure A22: Labor Force Participation Rate and Employment-to-Population Ratio Regression Estimates with Region-by-Time Fixed Effects: Impact of One Standard Deviation Increase in Hydro_z^{pre}

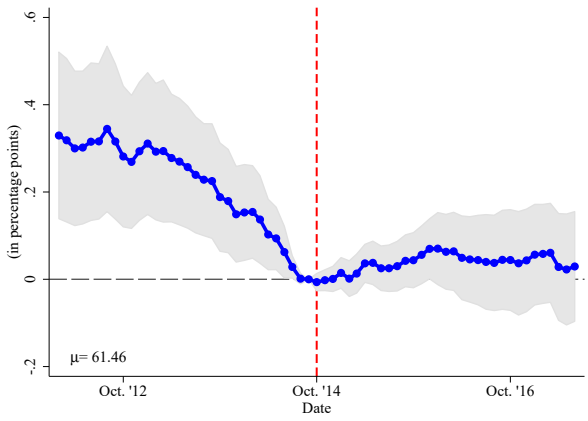

(A) Non-Parametric LFP w/ Region-by-Time FE

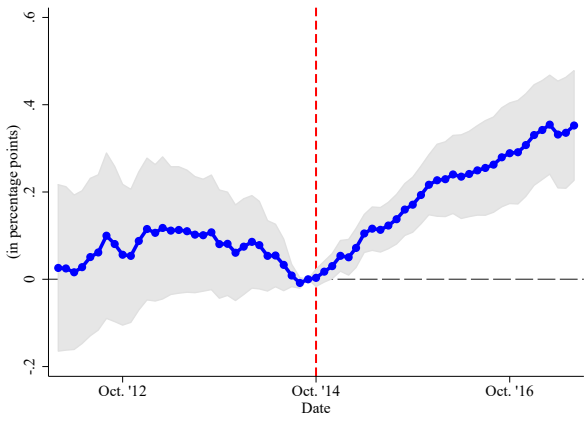

(B) Parametric LFP w/ Region-by-Time FE

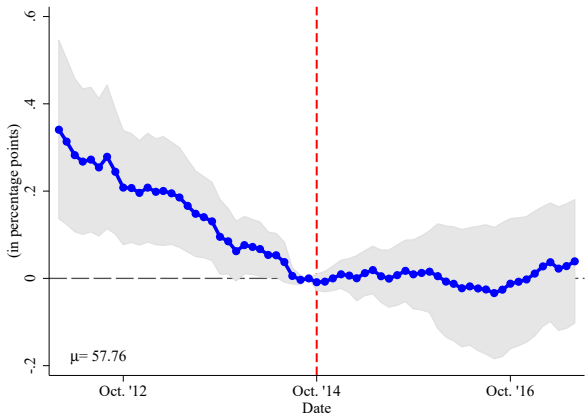

(C) Non-Parametric EPR w/ Region-by-Time FE

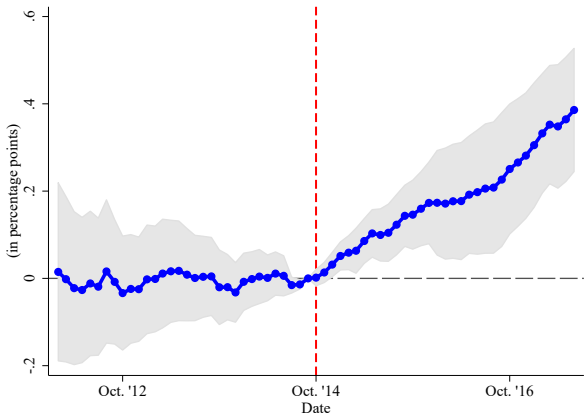

(D) Parametric EPR w/ Region-by-Time FE

Note: These figures show the coefficients on the $\text{Hydro}_z^{pre} \times$ time interactions from regression equations 1 (panels (A) and (C)) and 2 (panels (B) and (D)) with the inclusion of Census region-by-time fixed effects. Panels (A) and (B) show the results for the labor force participation rate, while panels (C) and (D) show the results for the employment-to-population ratio. The dependent variables are measured at the zip3-by-month level. The coefficients can be interpreted as the effect of a one standard deviation increase in pre-period hydrocodone consumption.

Table A4: Comparison of Estimates of the Impact of Opioids on Labor Market Behaviors

Paper	Identification	Direction	Magnitude
Current paper	Policy: Rescheduling of HCPs	(+)	10% decrease in HCPs $\Longrightarrow$ .32% (0.2 pp) increase in LFP
Park and Powell (2021)	Policy: Reformulation of OxyContin	(+)	10% decrease in exposure to reformulation (i.e., less heroin) $\Longrightarrow$ .18 pp increase in EPR
Kilby (2015)	Policy: PDMPs	(-)	10% decrease in opioid use $\Longrightarrow$ 7.4% increase in days missed
Kaestner and Ziedan (2019)	Policy: PDMPs	(-)	5-20% decrease in opioids $\Longrightarrow$ 1-2% reduction in EPR
Franco, Wagner and Whaley (2021)	Policy: PDMPs	(+)	10% decrease in opioids $\Longrightarrow$ 15.6% decrease in workplace absences
Deiana and Giua (2018)	Policy: PDMPs	NA	no effect
Currie, Jin and Schnell (2019)	IV: RX for young with RX for elderly	(-)	10% decrease in prescribing $\Longrightarrow$ 0.45 pp decrease in EPR for women, no effects for men
Harris et al. (2017)	IV: presence of high prescriber	(+)	10% decrease in prescriptions $\Longrightarrow$ 0.56 pp increase in LFP
Savych, Neumark and Lea (2018)	IV: local prescribing patterns	(+)	5 pp decrease in patients receiving long-term opioid RX $\Longrightarrow$ 12.6% decrease in duration of temporary disability; no effects for short term opioid RX
Laird and Nielsen (2016)	Movers: migration from low to high prescribing areas	(-)	10% decrease in local prescriber rate $\Longrightarrow$ 1.5 pp increase in probability of being in the labor force
Aliprantis, Fee and Schweitzer (2019)	Descriptive	(+)	10% decrease in opioid prescriptions $\Longrightarrow$ 0.15-0.47 pp increase in LFP
Krueger (2017)	Descriptive	(+)	10% decrease in opioid prescriptions $\Longrightarrow$ 0.11-0.14 pp increase in LFP

Note: This table shows estimates of the effects of opioids on the labor market from several different papers. The first column lists the authors of the paper. The second provides a brief description of the the identification strategy in each paper. The third column indicates whether the paper finds a positive (+) or negative (-) impact of an opioid reduction on the labor market behaviors studied. Many of these papers examine the effects of increased opioid access, so I rescale their estimates to indicate the implied effect sign of an opioid decrease. The last column indicates the effect magnitude.

B APPENDIX: Model

The purpose of this model is to show, in the simplest framework possible, that changes in opioid access can plausibly increase or decrease labor supply. An important insight from the model is that if opioid consumption is sufficiently efficacious at pain management, opioids could increase the effective amount of time available for both work and leisure. However, it is possible that the presence of side-effects can outweigh these benefits and reduce the effective time endowment. Likewise, opioids carry a risk of addiction which reduces work capacity. Depending on the relative magnitudes of these tradeoffs, restrictions on opioids can generate positive or negative spillovers on the labor market.

Utility is a function of consumption (C), leisure (R), and pain (P). Agents receive an exogenous pain signal (p_i) in period $t = 0$ and decide whether to consume opioids $O \in \{0, 1\}$ to treat their pain. Actual experienced pain is given by

$$P = P(p_i, O)$$

Since p_i is exogenous, I simplify this to $P = P_i(O)$. Opioids are assumed to decrease pain, i.e., $P_i(O = 1) < P_i(O = 0)$. Agents face a time constraint in which their time endowment T_i is a function of both pain, opioid consumption, and a random component ϵ_i. That is,

$$T_i = g(P_i(O), O, \epsilon_i)$$

where $\frac{\partial g}{\partial P_i} < 0$. Essentially, time spent in pain is unavailable for labor or leisure. Opioids affect T_i through their effect on P_i, but can also directly affect T_i. For example, time spent under the side effects of opioids is also unavailable for labor or leisure. The term ϵ_i represents the stochastic risk of addiction from consuming opioids which decreases T_i. The sign of the difference

$$\mathbb{E}[g(P_i(O = 1), O = 1, \epsilon_i)] - g(P_i(O = 0), O = 0, 0) \tag{4}$$

depends on the risk of addiction from opioids and the tradeoff between pain relief increasing the time endowment and side-effects reducing the time endowment.

The time constraint is then

$$L + R \leq T_i \tag{5}$$

where L is labor. In period $t = 1$, agents maximize their utility by choosing C, L, and R subject to a budget constraint

$$pO + C \leq wL + y \tag{6}$$

where y is non-labor income, p is the price of opioids, and w is the wage rate. At this point the decision of whether to consume opioids has already been made and the time endowment T_i is known.

The agent's problem in period $t = 1$ is

$$\max_{L,R,C} u(C, R, P_i(O)) \text{ subject to } 5, 6$$

Assume that $u(C, R, P_i) = \tilde{u}(C, R) - f(P_i(O))$, where f is increasing in P_i. Assuming a Cobb-Douglas

functional form for $\tilde{u}(C, R)$ and substituting 5 for R we can set up the following Lagrangian

$$\mathcal{L} = C^a(T_i - L)^{1-a} - f(P_i(O)) + \lambda[wL + y - pO - C]$$

with FOCs

$$[C] : a(\frac{T_i - L}{C})^{1-a} - \lambda = 0 \tag{7}$$

$$[L] : (1-a)(\frac{C}{T_i - L})^a(-1) + w\lambda = 0 \tag{8}$$

$$[\lambda] : wL + y - pO - C = 0 \tag{9}$$

Combining the FOCs for C and L we obtain $C = Aw(T_i - L)$, where $A = \frac{a}{1-a}$. Plugging this into the last FOC we solve for optimal labor supply

$$L = \frac{AwT_i + pO - y}{w + Aw} \tag{10}$$

which then allows us to solve for optimal consumption

$$C = \frac{Aw}{w + Aw}(wT_i - pO + y) \tag{11}$$

and finally optimal leisure

$$R = \frac{1}{w + AW}(wT_i - pO + y) \tag{12}$$

Knowing what they will choose for L, C, and R in period $t = 1$, the agent can choose whether to consume opioids by comparing their utility from each decision. Let the superscript * denote optimal decisions when $O = 1$ and ** denote optimal decisions when $O = 0$. Then the agent chooses to consume opioids in period $t = 0$ if and only if

$$\mathbb{E}[u(C^*, R^*, P_i(O = 1))] \geq u(C^{**}, R^{**}, P_i(O = 0))$$

Plugging in the functional forms and rearranging terms, the agent consumes opioids if and only if

$$\underbrace{f(P_i(O = 0)) - f(P_i(O = 1))}_{\text{utility gains from decreased pain}} \geq \underbrace{w(T_i^{**} - \mathbb{E}[T_i^*]) - p}_{\substack{\text{utility losses from forgone consumption} \\ \text{and leisure minus cost of opioids}}} \tag{13}$$

The left hand side of 13 will always be positive since P_i is decreasing in O and f is increasing in P_i. The right hand side could be positive or negative depending on how $\mathbb{E}[T_i^*]$ varies with opioid consumption. For example, opioid consumption reduces time spent in pain thus increasing the effective time endowment. However, time spent under the side effects of opioids and the risk of addiction (ϵ_i) both decrease it.

Consider an individual in the model who choses to take opioids. When this choice is disallowed, the individual's optimal choice of labor could increase or decrease based on how their effective time

endowment changes (equations 4 and 10). That I find increases in employment and labor supply in response to the rescheduling suggests that the increased time endowment due to decreased time spent under the side effects of opioids and the eliminated risk of addiction outweighs the decreased time endowment from increased time spent in pain.[65]

Note that this model does not allow for substitution toward the illicit market. To the extent that illicit opioids are close substitutes for legal opioids (in terms of benefits and harms) and are easily accessible, then we'd expect the impact of the policy change to be muted. However, if illicit opioids carry a higher risk of addiction or more sever side effects, then restricted access to the legal market would exacerbate labor market problems.

Likewise, the model does not consider substitution toward over-the-counter (OTC) painkillers (e.g., Tylenol or Ibuprofen). If these drugs have similar medical efficacy but fewer side effects and lower risk of addiction, then we would expect increases in labor supply. On the other hand, if OTC drugs are less effective at treating pain then we would expect decreased labor supply.

[65]Since agents are fully informed in the model, any change to opioid access is utility decreasing even if it increases the time endowment. However, this would change if agents were misinformed about the distribution of ϵ_i.